
The Logtalk Handbook

Release v3.100.0

Paulo Moura

May 14, 2026

CONTENTS

1	User Manual	1
1.1	Declarative object-oriented programming	1
1.2	Main features	2
1.2.1	Integration of logic and object-oriented programming	2
1.2.2	Integration of event-driven and object-oriented programming	2
1.2.3	Support for component-based programming	3
1.2.4	Support for both prototype and class-based systems	3
1.2.5	Support for multiple object hierarchies	3
1.2.6	Separation between interface and implementation	3
1.2.7	Private, protected and public inheritance	3
1.2.8	Private, protected and public object predicates	4
1.2.9	Parametric objects	4
1.2.10	High level multi-threading programming support	4
1.2.11	Smooth learning curve	4
1.2.12	Compatibility with most Prolog systems and the ISO standard	4
1.2.13	Performance	4
1.2.14	Logtalk scope	5
1.3	Nomenclature	6
1.3.1	Prolog nomenclature	6
1.3.2	Smalltalk nomenclature	8
1.3.3	C++ nomenclature	9
1.3.4	Java nomenclature	11
1.3.5	Python nomenclature	12
1.4	Messages	14
1.4.1	Operators used in message-sending	14
1.4.2	Sending a message to an object	14
1.4.3	Delegating a message to an object	15
1.4.4	Sending a message to <i>self</i>	15
1.4.5	Broadcasting	15
1.4.6	Calling imported and inherited predicates	15
1.4.7	Message sending and event generation	16
1.4.8	Sending a message from a module	17
1.4.9	Message sending performance	17
1.5	Objects	17
1.5.1	Objects, prototypes, classes, and instances	17
1.5.2	Defining a new object	18
1.5.3	Parametric objects	21
1.5.4	Finding defined objects	23
1.5.5	Creating a new object at runtime	23
1.5.6	Abolishing an existing object	24

1.5.7	Object directives	24
1.5.8	Object relationships	26
1.5.9	Object properties	27
1.5.10	Built-in objects	29
1.5.11	Multi-threading applications	30
1.6	Protocols	30
1.6.1	Defining a new protocol	30
1.6.2	Finding defined protocols	31
1.6.3	Creating a new protocol at runtime	31
1.6.4	Abolishing an existing protocol	32
1.6.5	Protocol directives	32
1.6.6	Protocol relationships	32
1.6.7	Protocol properties	33
1.6.8	Implementing protocols	34
1.6.9	Built-in protocols	35
1.6.10	Multi-threading applications	35
1.7	Categories	35
1.7.1	Defining a new category	36
1.7.2	Hot patching	38
1.7.3	Finding defined categories	39
1.7.4	Creating a new category at runtime	39
1.7.5	Abolishing an existing category	40
1.7.6	Category directives	40
1.7.7	Category relationships	41
1.7.8	Category properties	42
1.7.9	Importing categories	44
1.7.10	Calling category predicates	45
1.7.11	Parametric categories	46
1.7.12	Built-in categories	46
1.7.13	Multi-threading applications	46
1.8	Predicates	46
1.8.1	Reserved predicate names	47
1.8.2	Declaring predicates	47
1.8.3	Defining predicates	58
1.8.4	Definite clause grammar rules	64
1.8.5	Built-in methods	67
1.8.6	Predicate properties	71
1.8.7	Finding declared predicates	73
1.8.8	Calling Prolog predicates	74
1.8.9	Defining Prolog multifile predicates	78
1.8.10	Asserting and retracting Prolog predicates	78
1.9	Inheritance	80
1.9.1	Protocol inheritance	80
1.9.2	Implementation inheritance	81
1.9.3	Public, protected, and private inheritance	84
1.9.4	Multiple inheritance	85
1.9.5	Composition versus multiple inheritance	85
1.10	Event-driven programming	85
1.10.1	Definitions	86
1.10.2	Event generation	87
1.10.3	Communicating events to monitors	87
1.10.4	Performance concerns	87
1.10.5	Monitor semantics	88
1.10.6	Activation order of monitors	88

1.10.7	Event handling	88
1.11	Multi-threading programming	91
1.11.1	Enabling multi-threading support	91
1.11.2	Enabling objects to make multi-threading calls	91
1.11.3	Multi-threading built-in predicates	91
1.11.4	One-way asynchronous calls	94
1.11.5	Asynchronous calls and synchronized predicates	94
1.11.6	Synchronizing threads through notifications	95
1.11.7	Threaded engines	96
1.11.8	Multi-threading performance	97
1.12	Error handling	97
1.12.1	Raising Exceptions	98
1.12.2	Type-checking	98
1.12.3	Expected terms	99
1.12.4	Compiler warnings and errors	99
1.12.5	Runtime errors	102
1.13	Reflection	102
1.13.1	Structural reflection	103
1.13.2	Behavioral reflection	103
1.14	Writing and running applications	104
1.14.1	Starting Logtalk	104
1.14.2	Running parallel Logtalk processes	104
1.14.3	Source files	105
1.14.4	Multi-pass compiler	106
1.14.5	Compiling and loading your applications	107
1.14.6	Compiler errors, warnings, and comments	108
1.14.7	Loader files	108
1.14.8	Libraries of source files	109
1.14.9	Settings files	110
1.14.10	Compiler linter	111
1.14.11	Compiler flags	111
1.14.12	Reloading source files	119
1.14.13	Batch processing	119
1.14.14	Optimizing performance	119
1.14.15	Portable applications	120
1.14.16	Conditional compilation	120
1.14.17	Avoiding common errors	120
1.14.18	Coding style guidelines	121
1.15	Printing messages and asking questions	121
1.15.1	Printing messages	122
1.15.2	Message tokenization	123
1.15.3	Meta-messages	124
1.15.4	Defining message prefixes and output streams	124
1.15.5	Defining message prefixes and output files	124
1.15.6	Intercepting messages	125
1.15.7	Asking questions	126
1.15.8	Intercepting questions	127
1.15.9	Multi-threading applications	127
1.16	Term and goal expansion	128
1.16.1	Defining expansions	128
1.16.2	Expanding grammar rules	130
1.16.3	Bypassing expansions	130
1.16.4	Hook objects	130
1.16.5	Virtual source file terms and loading context	132

1.16.6	Default compiler expansion workflow	133
1.16.7	User defined expansion workflows	133
1.16.8	Using Prolog defined expansions	133
1.16.9	Debugging expansions	134
1.17	Documenting	135
1.17.1	Entity documenting directives	136
1.17.2	Predicate documenting directives	137
1.17.3	Describing predicates	138
1.17.4	Documenting predicate exceptions	139
1.17.5	Processing and viewing documenting files	140
1.17.6	Inline formatting in comments text	141
1.17.7	Diagrams	142
1.18	Debugging	142
1.18.1	Compiling source files in debug mode	142
1.18.2	Procedure box model	143
1.18.3	Activating the debugger	144
1.18.4	Defining breakpoints	144
1.18.5	Defining log points	148
1.18.6	Tracing program execution	149
1.18.7	Debugging using breakpoints	150
1.18.8	Debugging commands	150
1.18.9	Customizing term writing	152
1.18.10	Context-switching calls	153
1.18.11	Debugging messages	154
1.18.12	Using the term-expansion mechanism for debugging	156
1.18.13	Ports profiling	156
1.18.14	Debug and trace events	156
1.18.15	Source-level debugger	157
1.19	Performance	157
1.19.1	Source code compilation modes	158
1.19.2	Source code compilation order	158
1.19.3	Local predicate calls	158
1.19.4	Calls to imported or inherited predicates	158
1.19.5	Calls to module predicates	158
1.19.6	Messages	158
1.19.7	Automatic expansion of built-in meta-predicates	159
1.19.8	Inlining	159
1.19.9	Generated code simplification and optimizations	159
1.19.10	Size of the generated code	160
1.19.11	Circular references	160
1.19.12	Debug mode overhead	160
1.19.13	Other considerations	160
1.20	Installing Logtalk	161
1.20.1	Hardware and software requirements	161
1.20.2	Logtalk installers	162
1.20.3	Source distribution	162
1.20.4	Distribution overview	162
1.21	Prolog integration and migration	165
1.21.1	Source files with both Prolog code and Logtalk code	165
1.21.2	Encapsulating plain Prolog code in objects	166
1.21.3	Converting Prolog modules into objects	167
1.21.4	Compiling Prolog modules as objects	168
1.21.5	Dealing with proprietary Prolog directives and predicates	170
1.21.6	Calling Prolog module predicates	171

1.21.7	Loading converted Prolog applications	171
2	Reference Manual	173
2.1	Grammar	173
2.1.1	Entities	173
2.1.2	Object definition	173
2.1.3	Category definition	174
2.1.4	Protocol definition	174
2.1.5	Entity relations	175
2.1.6	Entity identifiers	178
2.1.7	Source files	179
2.1.8	Source file names	179
2.1.9	Terms	179
2.1.10	Directives	180
2.1.11	Clauses and goals	188
2.1.12	Lambda expressions	189
2.1.13	Entity properties	189
2.1.14	Predicate properties	193
2.1.15	Compiler flags	194
2.2	Control constructs	194
2.2.1	Message sending	194
2.2.2	Message delegation	197
2.2.3	Calling imported and inherited predicates	198
2.2.4	Calling predicates in <i>this</i>	199
2.2.5	Calling external predicates	200
2.2.6	Context switching calls	202
2.3	Directives	204
2.3.1	Source file directives	204
2.3.2	Conditional compilation directives	210
2.3.3	Entity directives	214
2.3.4	Predicate directives	227
2.4	Built-in predicates	246
2.4.1	Enumerating objects, categories and protocols	246
2.4.2	Enumerating objects, categories and protocols properties	248
2.4.3	Creating new objects, categories and protocols	252
2.4.4	Abolishing objects, categories and protocols	257
2.4.5	Objects, categories, and protocols relations	259
2.4.6	Event handling	269
2.4.7	Multi-threading	272
2.4.8	Multi-threading engines	283
2.4.9	Compiling and loading source files	292
2.4.10	Flags	307
2.4.11	Linter	310
2.5	Built-in methods	311
2.5.1	Logic and control	311
2.5.2	Execution context	314
2.5.3	Reflection	319
2.5.4	Database	323
2.5.5	Meta-calls	332
2.5.6	Error handling	337
2.5.7	All solutions	354
2.5.8	Event handling	360
2.5.9	Message forwarding	361
2.5.10	Definite clause grammar rules	362

2.5.11	Term and goal expansion	368
2.5.12	Coinduction hooks	372
2.5.13	Message printing	374
2.5.14	Question asking	380
3	Tutorial	383
3.1	List predicates	383
3.1.1	Defining a list object	383
3.1.2	Defining a list protocol	384
3.1.3	Summary	386
3.2	Dynamic object attributes	386
3.2.1	Defining a category	386
3.2.2	Importing the category	387
3.2.3	Summary	388
3.3	A reflective class-based system	388
3.3.1	Defining the base classes	388
3.3.2	Summary	389
3.4	Profiling programs	389
3.4.1	Messages as events	390
3.4.2	Profilers as monitors	390
3.4.3	Summary	392
4	FAQ	393
4.1	General	393
4.1.1	Why are all versions of Logtalk numbered 2.x or 3.x?	393
4.1.2	Why do I need a Prolog compiler to use Logtalk?	393
4.1.3	Is the Logtalk implementation based on Prolog modules?	393
4.1.4	Does the Logtalk implementation use term-expansion?	394
4.2	Compatibility	394
4.2.1	What are the backend Prolog compiler requirements to run Logtalk?	394
4.2.2	Can I use constraint-based packages with Logtalk?	394
4.2.3	Can I use Logtalk objects and Prolog modules at the same time?	394
4.3	Installation	394
4.3.1	The integration scripts/shortcuts are not working!	394
4.3.2	I get errors when starting up Logtalk after upgrading to the latest version!	395
4.4	Portability	395
4.4.1	Are my Logtalk applications portable across Prolog compilers?	395
4.4.2	Are my Logtalk applications portable across operating systems?	395
4.5	Programming	395
4.5.1	Should I use prototypes or classes in my application?	396
4.5.2	Can I use both classes and prototypes in the same application?	396
4.5.3	Can I mix classes and prototypes in the same hierarchy?	396
4.5.4	Can I use a protocol or a category with both prototypes and classes?	396
4.5.5	What support is provided in Logtalk for defining and using components?	396
4.5.6	What support is provided in Logtalk for reflective programming?	396
4.6	Troubleshooting	396
4.6.1	Using compiler options on calls to the Logtalk compiling and loading predicates does not work!	397
4.6.2	Gecko-based browsers (e.g., Firefox) show non-rendered HTML entities when browsing XML documenting files!	397
4.6.3	Compiling a source file results in errors or warnings but the Logtalk compiler reports a successful compilation with zero errors and zero warnings!	397
4.7	Usability	397
4.7.1	Is there a shortcut for compiling and loading source files?	397

4.7.2	Is there an equivalent directive to the <code>ensure_loaded/1</code> Prolog directive?	398
4.7.3	Are there shortcuts for the <code>make</code> functionality?	398
4.8	Deployment	398
4.8.1	Can I create standalone applications with Logtalk?	398
4.9	Performance	398
4.9.1	Is Logtalk implemented as a meta-interpreter?	398
4.9.2	What kind of code Logtalk generates when compiling objects? Dynamic code? Static code?	399
4.9.3	How about message-sending performance? Does Logtalk use static binding or dynamic binding?	399
4.9.4	Which Prolog-dependent factors are most crucial for good Logtalk performance?	399
4.9.5	How does Logtalk performance compare with plain Prolog and with Prolog modules?	399
4.10	Licensing	399
4.10.1	What's the Logtalk distribution license?	400
4.10.2	Can Logtalk be used in commercial applications?	400
4.10.3	What's the final license for a combination of Logtalk with a Prolog compiler?	400
4.11	Support	400
4.11.1	Are there professional consulting, training and supporting services?	400
5	Developer Tools	401
5.1	Overview	401
5.1.1	Loading the developer tools	402
5.1.2	Tools documentation	402
5.1.3	Tools common flags	402
5.1.4	Tools requirements	403
5.2	<code>asdf</code>	405
5.3	<code>assertions</code>	405
5.3.1	API documentation	405
5.3.2	Loading	405
5.3.3	Testing	405
5.3.4	Adding assertions to your source code	406
5.3.5	Automatically adding file and line context information to assertions	406
5.3.6	Suppressing assertion calls from source code	406
5.3.7	Redirecting assertion failure messages	406
5.3.8	Converting assertion failures into errors	407
5.4	<code>code_metrics</code>	407
5.4.1	API documentation	407
5.4.2	Loading	408
5.4.3	Testing	408
5.4.4	Available metrics	408
5.4.5	Coupling metrics	408
5.4.6	Lack of cohesion metric	409
5.4.7	WMC metric	410
5.4.8	RFC metric	410
5.4.9	Cognitive complexity metric	411
5.4.10	Lines metric	411
5.4.11	Maintainability index metric	411
5.4.12	Halstead metric	412
5.4.13	UPN metric	412
5.4.14	Cyclomatic complexity metric	413
5.4.15	Usage	413
5.4.16	Excluding code from analysis	413
5.4.17	Defining new metrics	414
5.4.18	Third-party tools	414

5.4.19	Applying metrics to Prolog modules	414
5.4.20	Applying metrics to plain Prolog code	414
5.5	dead_code_scanner	415
5.5.1	API documentation	415
5.5.2	Loading	415
5.5.3	Testing	415
5.5.4	Usage	416
5.5.5	Excluding code from analysis	417
5.5.6	Machine-readable summaries	417
5.5.7	Machine-readable preflight warnings	418
5.5.8	Finding classes and confidence	418
5.5.9	SARIF reports	419
5.5.10	Integration with the make tool	419
5.5.11	Caveats	419
5.5.12	Scanning Prolog modules	420
5.5.13	Scanning plain Prolog files	420
5.5.14	Known issues	420
5.6	debug_messages	420
5.6.1	API documentation	421
5.6.2	Loading	421
5.6.3	Testing	421
5.6.4	Usage	421
5.7	debugger	422
5.7.1	API documentation	422
5.7.2	Loading	422
5.7.3	Testing	423
5.7.4	Usage	423
5.7.5	Alternative debugger tools	423
5.7.6	Known issues	424
5.8	diagrams	424
5.8.1	Requirements	425
5.8.2	API documentation	425
5.8.3	Loading	425
5.8.4	Testing	426
5.8.5	Supported diagrams	426
5.8.6	Graph elements	427
5.8.7	Supported graph languages	428
5.8.8	Customization	429
5.8.9	Linking diagrams	434
5.8.10	Creating diagrams for Prolog module applications	435
5.8.11	Creating diagrams for plain Prolog files	435
5.8.12	Other notes	436
5.9	doclet	437
5.9.1	API documentation	437
5.9.2	Loading	437
5.9.3	Automating running doclets	437
5.9.4	Integration with the make tool	437
5.10	help	438
5.10.1	Requirements	438
5.10.2	Customization	439
5.10.3	API documentation	439
5.10.4	Loading	439
5.10.5	Testing	440
5.10.6	Supported operating-systems	440

5.10.7	Usage	440
5.10.8	Known issues	441
5.11	issue_creator	442
5.11.1	Requirements	442
5.11.2	Loading	442
5.11.3	Usage	442
5.11.4	Known issues	443
5.12	lgtdoc	443
5.12.1	Requirements	443
5.12.2	API documentation	444
5.12.3	Loading	444
5.12.4	Testing	444
5.12.5	Documenting source code	444
5.12.6	Generating documentation	445
5.12.7	Documentation linter checks	446
5.13	lgtunit	446
5.13.1	Main files	446
5.13.2	API documentation	447
5.13.3	Loading	447
5.13.4	Testing	447
5.13.5	Writing and running tests	447
5.13.6	Automating running tests	449
5.13.7	Parametric test objects	450
5.13.8	Test dialects	451
5.13.9	User-defined test dialects	453
5.13.10	QuickCheck	454
5.13.11	Skipping tests	459
5.13.12	Selecting tests	459
5.13.13	Checking test goal results	460
5.13.14	Testing local predicates	461
5.13.15	Testing non-deterministic predicates	461
5.13.16	Testing generators	462
5.13.17	Testing input/output predicates	462
5.13.18	Suppressing tested predicates output	463
5.13.19	Tests with timeout limits	464
5.13.20	Setup and cleanup goals	465
5.13.21	Test annotations	465
5.13.22	Test execution times and memory usage	465
5.13.23	Working with test data files	466
5.13.24	Flaky tests	466
5.13.25	Mocking	467
5.13.26	Debugging messages in tests	468
5.13.27	Debugging failed tests	468
5.13.28	Code coverage	469
5.13.29	Utility predicates	471
5.13.30	Exporting test results	472
5.13.31	Generating Allure reports	474
5.13.32	Exporting code coverage results	475
5.13.33	Automatically creating bug reports at issue trackers	477
5.13.34	Minimizing test results output	477
5.13.35	Help with warnings	477
5.13.36	Known issues	477
5.14	linter	477
5.14.1	Main linter checks	478

5.14.2	Help on linter warnings	479
5.14.3	Extending the linter	479
5.14.4	Linting Prolog modules	479
5.14.5	Linting plain Prolog files	479
5.15	linter_reporter	480
5.15.1	API documentation	480
5.15.2	Loading	480
5.15.3	Testing	481
5.15.4	Usage	481
5.15.5	Options	481
5.16	make	481
5.16.1	API documentation	482
5.16.2	Help with warnings	482
5.16.3	Known issues	482
5.17	mutation_testing	482
5.17.1	API documentation	482
5.17.2	Loading	482
5.17.3	Testing	483
5.17.4	Features	483
5.17.5	Default mutators	483
5.17.6	Usage	484
5.17.7	Limitations	485
5.17.8	Options	485
5.17.9	Sampling semantics	486
5.17.10	Defining new mutators	487
5.18	packs	488
5.18.1	Requirements	488
5.18.2	API documentation	489
5.18.3	Loading	489
5.18.4	Testing	489
5.18.5	Usage	490
5.18.6	Programmatic query API	490
5.18.7	Registries and packs storage	492
5.18.8	Virtual environments	492
5.18.9	Lock files	494
5.18.10	Registry specification	494
5.18.11	Registry handling	495
5.18.12	Registry development	496
5.18.13	Pack specification	497
5.18.14	Encrypted packs	498
5.18.15	Signed packs	498
5.18.16	Pack URLs and Single Sign-On	499
5.18.17	Multiple pack versions	499
5.18.18	Pack dependencies	500
5.18.19	Pack portability	500
5.18.20	Pack development	501
5.18.21	Pack handling	502
5.18.22	Pack documentation	504
5.18.23	Pinning registries and packs	504
5.18.24	Testing packs	505
5.18.25	Security considerations	505
5.18.26	Best practices	506
5.18.27	Installing Prolog packs	506
5.18.28	Help with warnings	506

5.18.29	Known issues	507
5.19	ports_profiler	507
5.19.1	API documentation	507
5.19.2	Loading	508
5.19.3	Testing	508
5.19.4	Compiling source files for port profiling	508
5.19.5	Generating profiling data	508
5.19.6	Printing profiling data reports	508
5.19.7	Interpreting profiling data	510
5.19.8	Profiling Prolog modules	510
5.19.9	Profiling plain Prolog code	511
5.19.10	Known issues	511
5.20	profiler	511
5.20.1	Loading	512
5.20.2	Testing	512
5.20.3	Supported backend Prolog compilers	512
5.20.4	Compiling source code for profiling	513
5.21	sarif	513
5.21.1	API documentation	513
5.21.2	Loading	513
5.21.3	Testing	513
5.21.4	Usage	513
5.22	sbom	514
5.22.1	API documentation	514
5.22.2	Loading	514
5.22.3	Testing	514
5.22.4	Usage	515
5.22.5	Known issues	519
5.23	tool_diagnostics	519
5.23.1	API documentation	520
5.23.2	Loading	520
5.23.3	Protocol overview	520
5.23.4	Term conventions	520
5.23.5	Usage	522
5.24	tutor	522
5.24.1	API documentation	523
5.24.2	Loading	523
5.24.3	Usage	523
5.25	wrapper	523
5.25.1	API documentation	524
5.25.2	Loading	524
5.25.3	Workflows	524
5.25.4	Customization	524
5.25.5	Current limitations	525
6	Libraries	527
6.1	Overview	527
6.1.1	Library documentation	529
6.1.2	Loading libraries	529
6.1.3	Testing libraries	529
6.1.4	Credits	529
6.1.5	Other notes	530
6.2	adaptive_boosting_classifier	530
6.2.1	API documentation	530

6.2.2	Loading	530
6.2.3	Testing	530
6.2.4	Features	530
6.2.5	Options	531
6.2.6	Classifier representation	531
6.2.7	References	531
6.2.8	Usage	531
6.3	agglomerative_clusterer	533
6.3.1	API documentation	533
6.3.2	Loading	533
6.3.3	Testing	534
6.3.4	Features	534
6.3.5	Options	534
6.3.6	Clusterer representation	535
6.3.7	Diagnostics	535
6.4	amqp	535
6.4.1	API documentation	536
6.4.2	Loading	536
6.4.3	Testing	536
6.4.4	Features	536
6.4.5	Usage	537
6.4.6	Connection Pooling	540
6.4.7	Binary Frame Encoding	542
6.4.8	Data Types	542
6.4.9	Float/Double Encoding/Decoding	543
6.4.10	Error Handling	544
6.4.11	Comparison with STOMP	544
6.4.12	Future Work	545
6.4.13	Known Limitations and Issues	545
6.4.14	AMQP 1.0 vs AMQP 0-9-1	545
6.5	anomaly_detection_protocols	546
6.5.1	Export header format	546
6.5.2	API documentation	546
6.5.3	Loading	546
6.5.4	Testing	547
6.5.5	Test datasets	547
6.6	application	547
6.6.1	API documentation	548
6.6.2	Loading	548
6.6.3	Testing	548
6.6.4	Usage	548
6.7	apriori_pattern_miner	550
6.7.1	API documentation	550
6.7.2	Loading	550
6.7.3	Testing	551
6.7.4	Features	551
6.7.5	Options	551
6.7.6	Pattern miner representation	552
6.7.7	References	552
6.8	arbitrary	552
6.8.1	API documentation	552
6.8.2	Loading	552
6.8.3	Testing	552
6.8.4	Pre-defined types	553

6.8.5	Usage	553
6.8.6	Defining new generators and shrinkers	553
6.8.7	Scoped generators and shrinkers	556
6.8.8	Reproducing sequences of arbitrary terms	557
6.8.9	Default size of generated terms	557
6.8.10	Known issues	558
6.9	arrangements	558
6.9.1	API documentation	558
6.9.2	Loading	558
6.9.3	Testing	559
6.10	assignvars	559
6.10.1	API documentation	559
6.10.2	Loading	559
6.10.3	Testing	559
6.11	avro	559
6.11.1	API documentation	560
6.11.2	Loading	560
6.11.3	Testing	560
6.11.4	Schema representation	560
6.11.5	Data representation	560
6.11.6	Schema Examples	560
6.11.7	Encoding	561
6.11.8	Decoding	561
6.12	base32	561
6.12.1	API documentation	562
6.12.2	Loading	562
6.12.3	Testing	562
6.12.4	Encoding	562
6.12.5	Decoding	562
6.13	base58	563
6.13.1	API documentation	563
6.13.2	Loading	563
6.13.3	Testing	563
6.13.4	Encoding	563
6.13.5	Decoding	564
6.14	base64	564
6.14.1	API documentation	564
6.14.2	Loading	564
6.14.3	Testing	564
6.14.4	Encoding	565
6.14.5	Decoding	565
6.15	base85	566
6.15.1	API documentation	566
6.15.2	Loading	566
6.15.3	Testing	566
6.15.4	Encoding	566
6.15.5	Decoding	567
6.16	basic_types	567
6.16.1	API documentation	567
6.16.2	Loading	567
6.16.3	Testing	567
6.17	bayesian_ridge_regression	568
6.17.1	API documentation	568
6.17.2	Loading	568

6.17.3	Testing	568
6.17.4	Features	568
6.17.5	Regressor representation	569
6.17.6	Diagnostics syntax	569
6.17.7	Options	570
6.18	borda_ranker	571
6.18.1	API documentation	571
6.18.2	Loading	571
6.18.3	Testing	571
6.18.4	Features	571
6.18.5	Scoring semantics	572
6.18.6	Usage	572
6.18.7	Diagnostics syntax	573
6.18.8	Options	574
6.18.9	Ranker representation	574
6.18.10	References	574
6.19	bradley_terry_ranker	575
6.19.1	API documentation	575
6.19.2	Loading	575
6.19.3	Testing	575
6.19.4	Features	575
6.19.5	Dataset requirements	576
6.19.6	Usage	576
6.19.7	Diagnostics syntax	576
6.19.8	Options	578
6.19.9	Ranker representation	578
6.19.10	References	578
6.20	c45_classifier	579
6.20.1	API documentation	579
6.20.2	Loading	579
6.20.3	Testing	579
6.20.4	Implemented features	579
6.20.5	Learned tree representation	580
6.20.6	Limitations	580
6.20.7	References	580
6.20.8	Usage	580
6.21	cartesian_products	581
6.21.1	API documentation	581
6.21.2	Loading	581
6.21.3	Testing	581
6.22	cbor	582
6.22.1	Representation	582
6.22.2	Encoding	583
6.22.3	Decoding	583
6.22.4	API documentation	583
6.22.5	Loading	583
6.22.6	Testing	583
6.23	character_sets	583
6.23.1	API documentation	584
6.23.2	Loading	585
6.23.3	Testing	585
6.23.4	Usage	585
6.24	ccsds_frames	585
6.24.1	Available entities	585

6.24.2	Representation	586
6.24.3	Parsing and generating	587
6.24.4	Types	587
6.24.5	Facade helpers	588
6.24.6	API documentation	589
6.24.7	Loading	589
6.24.8	Testing	589
6.25	ccsds_link_profiles	589
6.25.1	Representation	589
6.25.2	Public API	590
6.25.3	Parsing and generating	590
6.25.4	Packet helpers	590
6.25.5	API documentation	591
6.25.6	Loading	591
6.25.7	Testing	591
6.26	ccsds_packet_services	591
6.26.1	Representation	592
6.26.2	Packet-zone helpers	593
6.26.3	Telemetry transfer frame helpers	594
6.26.4	Advanced orbiting systems helpers	594
6.26.5	Cross-frame packet reassembly	594
6.26.6	Scope	596
6.26.7	API documentation	596
6.26.8	Loading	596
6.26.9	Testing	596
6.27	ccsds_packetization	596
6.27.1	Representation	597
6.27.2	Packetization	597
6.27.3	Events	598
6.27.4	Idle packets	598
6.27.5	API documentation	598
6.27.6	Loading	598
6.27.7	Testing	598
6.28	ccsds_packets	598
6.28.1	Packet Structure	599
6.28.2	Representation	599
6.28.3	Parsing	599
6.28.4	Generating	600
6.28.5	Accessor Predicates	600
6.28.6	Types and arbitrary generators	601
6.28.7	API documentation	601
6.28.8	Loading	601
6.28.9	Testing	602
6.29	ccsds_tc_services	602
6.29.1	Representation	602
6.29.2	Scope	604
6.29.3	Examples	604
6.29.4	API documentation	605
6.29.5	Loading	605
6.29.6	Testing	605
6.30	ccsds_time_codes	605
6.30.1	Available entities	605
6.30.2	Representation	606
6.30.3	Parsing and generating	606

6.30.4	Unix time conversion	607
6.30.5	Types and arbitrary generators	608
6.30.6	API documentation	609
6.30.7	Loading	609
6.30.8	Testing	609
6.31	ccsds_time_fields	609
6.31.1	Representation	609
6.31.2	Parsing and generating	610
6.31.3	Descriptor helpers	610
6.31.4	API documentation	611
6.31.5	Loading	611
6.31.6	Testing	611
6.32	classification_protocols	611
6.32.1	Shared category	611
6.32.2	Diagnostics	611
6.32.3	Export header format	612
6.32.4	API documentation	612
6.32.5	Loading	612
6.32.6	Testing	612
6.32.7	Test datasets	612
6.33	clo_span_pattern_miner	613
6.33.1	API documentation	613
6.33.2	Loading	613
6.33.3	Testing	613
6.33.4	Features	613
6.33.5	Options	614
6.33.6	Pattern miner representation	614
6.33.7	References	614
6.34	clustering_protocols	614
6.34.1	API documentation	615
6.34.2	Loading	615
6.34.3	Testing	615
6.34.4	Test datasets	615
6.35	colley_ranker	616
6.35.1	API documentation	616
6.35.2	Loading	616
6.35.3	Testing	616
6.35.4	Features	617
6.35.5	Scoring semantics	617
6.35.6	Usage	617
6.35.7	Diagnostics syntax	618
6.35.8	Ranker representation	619
6.35.9	References	619
6.36	combinations	619
6.36.1	API documentation	619
6.36.2	Loading	620
6.36.3	Testing	620
6.37	command_line_options	620
6.37.1	API documentation	620
6.37.2	Loading	620
6.37.3	Testing	620
6.37.4	Terminology	621
6.37.5	Defining option objects	621
6.37.6	Validating option definitions	621

6.37.7	Example option objects	622
6.37.8	Parsing command-line arguments	623
6.37.9	Generating help text	623
6.37.10	Help options	624
6.37.11	Parse options	624
6.37.12	Notes and tips	624
6.38	core	624
6.38.1	API documentation	624
6.38.2	Loading	625
6.38.3	Testing	625
6.39	copeland_ranker	625
6.39.1	API documentation	625
6.39.2	Loading	625
6.39.3	Testing	625
6.39.4	Features	625
6.39.5	Scoring semantics	626
6.39.6	Usage	626
6.39.7	Diagnostics syntax	627
6.39.8	Options	628
6.39.9	Ranker representation	628
6.39.10	References	628
6.40	coroutining	628
6.40.1	API documentation	629
6.40.2	Loading	629
6.40.3	Testing	629
6.40.4	Usage	629
6.41	crs_projections	629
6.41.1	API documentation	629
6.41.2	Loading	629
6.41.3	Testing	630
6.41.4	Supported coordinate reference systems	630
6.41.5	Usage	630
6.41.6	Notes	632
6.42	csv	632
6.42.1	API documentation	632
6.42.2	Loading	632
6.42.3	Testing	633
6.42.4	Usage	633
6.43	cuid2	634
6.43.1	API documentation	635
6.43.2	Loading	635
6.43.3	Testing	635
6.43.4	Usage	635
6.44	cusum_anomaly_detector	635
6.44.1	API documentation	636
6.44.2	Loading	636
6.44.3	Testing	636
6.44.4	Features	636
6.44.5	Options	637
6.44.6	Detector representation	637
6.44.7	Notes	638
6.45	datalog	638
6.45.1	API documentation	638
6.45.2	Loading	638

6.45.3	Testing	638
6.45.4	Scope	639
6.45.5	Public API overview	639
6.45.6	Basic usage	640
6.45.7	Limitations	640
6.46	dates	640
6.46.1	API documentation	641
6.46.2	Loading	641
6.46.3	Testing	641
6.47	dates_tz	642
6.47.1	API documentation	642
6.47.2	Loading	642
6.47.3	Testing	642
6.47.4	Usage	642
6.47.5	Examples	642
6.47.6	Notes	643
6.48	dbscan_clusterer	644
6.48.1	API documentation	644
6.48.2	Loading	644
6.48.3	Testing	644
6.48.4	Features	644
6.48.5	Options	645
6.48.6	Clusterer representation	645
6.48.7	References	645
6.49	dependents	645
6.49.1	API documentation	646
6.49.2	Loading	646
6.50	dequeues	646
6.50.1	API documentation	646
6.50.2	Loading	646
6.50.3	Testing	646
6.50.4	Usage	646
6.51	derangements	647
6.51.1	API documentation	647
6.51.2	Loading	648
6.51.3	Testing	648
6.52	dictionaries	648
6.52.1	API documentation	648
6.52.2	Loading	648
6.52.3	Testing	648
6.52.4	Usage	648
6.52.5	Credits	649
6.53	dif	649
6.53.1	API documentation	650
6.53.2	Loading	650
6.53.3	Testing	650
6.53.4	Usage	650
6.54	dimension_reduction_protocols	650
6.54.1	API documentation	651
6.54.2	Loading	651
6.54.3	Common options	651
6.54.4	Testing	651
6.54.5	Test datasets	651
6.55	eclat_pattern_miner	652

6.55.1	API documentation	652
6.55.2	Loading	652
6.55.3	Testing	652
6.55.4	Features	652
6.55.5	Options	653
6.55.6	Pattern miner representation	653
6.55.7	References	653
6.56	edcg	653
6.56.1	API documentation	654
6.56.2	Loading	654
6.56.3	Testing	654
6.56.4	Usage	654
6.56.5	Introduction	655
6.56.6	Syntax	655
6.56.7	Declaration of Predicates	656
6.56.8	Declaration of Accumulators	656
6.56.9	Declaration of Passed Arguments	656
6.56.10	Additional documentation	657
6.57	elastic_net_regression	657
6.57.1	API documentation	657
6.57.2	Loading	657
6.57.3	Testing	657
6.57.4	Features	657
6.57.5	Regressor representation	658
6.57.6	Diagnostics syntax	658
6.57.7	Options	659
6.58	elo_ranker	659
6.58.1	API documentation	660
6.58.2	Loading	660
6.58.3	Testing	660
6.58.4	Features	660
6.58.5	Rating semantics	660
6.58.6	Usage	661
6.58.7	Options	662
6.58.8	Ranker representation	662
6.58.9	References	662
6.59	events	663
6.59.1	API documentation	663
6.59.2	Loading	663
6.60	ewma_anomaly_detector	663
6.60.1	API documentation	663
6.60.2	Loading	664
6.60.3	Testing	664
6.60.4	Features	664
6.60.5	Options	665
6.60.6	Detector representation	665
6.60.7	Notes	666
6.61	expand_library_alias_paths	666
6.61.1	API documentation	666
6.61.2	Loading	666
6.61.3	Usage	667
6.62	expecteds	667
6.62.1	API documentation	667
6.62.2	Loading	667

6.62.3	Testing	667
6.62.4	Usage	667
6.62.5	See also	671
6.63	format	671
6.63.1	API documentation	671
6.63.2	Loading	671
6.63.3	Testing	671
6.63.4	Usage	671
6.63.5	Portability	672
6.64	fp_growth_pattern_miner	672
6.64.1	API documentation	672
6.64.2	Loading	672
6.64.3	Testing	672
6.64.4	Features	673
6.64.5	Options	673
6.64.6	Pattern miner representation	673
6.64.7	References	674
6.65	frequent_pattern_mining_protocols	674
6.65.1	API documentation	674
6.65.2	Loading	674
6.65.3	Testing	674
6.65.4	Test datasets	674
6.66	gaussian_mixture_clusterer	675
6.66.1	API documentation	675
6.66.2	Loading	675
6.66.3	Testing	675
6.66.4	Features	675
6.66.5	Gaussian-Mixture-Specific Prediction API	676
6.66.6	Options	676
6.66.7	Clusterer representation	676
6.67	gaussian_process_regression	677
6.67.1	API documentation	677
6.67.2	Loading	677
6.67.3	Testing	677
6.67.4	Features	677
6.67.5	Regressor representation	678
6.67.6	Prediction API	678
6.67.7	Diagnostics syntax	679
6.67.8	Options	680
6.68	genint	681
6.68.1	API documentation	681
6.68.2	Loading	681
6.68.3	Testing	681
6.68.4	Usage	681
6.69	gensym	681
6.69.1	API documentation	682
6.69.2	Loading	682
6.69.3	Testing	682
6.69.4	Usage	682
6.70	geojson	682
6.70.1	API documentation	683
6.70.2	Loading	683
6.70.3	Testing	683
6.70.4	Representation	683

6.70.5	Examples	684
6.70.6	Notes	684
6.71	geospatial	685
6.71.1	API documentation	685
6.71.2	Loading	685
6.71.3	Testing	685
6.71.4	Available predicates	685
6.71.5	Notes	686
6.71.6	Usage	686
6.72	git	690
6.72.1	API documentation	690
6.72.2	Loading	690
6.72.3	Testing	690
6.72.4	Usage	690
6.73	glicko2_ranker	691
6.73.1	API documentation	691
6.73.2	Loading	691
6.73.3	Testing	691
6.73.4	Features	692
6.73.5	Rating semantics	692
6.73.6	Usage	692
6.73.7	Options	693
6.73.8	Diagnostics syntax	694
6.73.9	Ranker representation	694
6.73.10	References	694
6.74	glicko2_periodic_ranker	694
6.74.1	API documentation	695
6.74.2	Loading	695
6.74.3	Testing	695
6.74.4	Features	695
6.74.5	Rating semantics	696
6.74.6	Usage	696
6.74.7	Options	697
6.74.8	Diagnostics syntax	697
6.74.9	Ranker representation	698
6.74.10	References	698
6.75	gradient_boosting_classifier	698
6.75.1	API documentation	698
6.75.2	Loading	698
6.75.3	Testing	699
6.75.4	Features	699
6.75.5	Options	699
6.75.6	Usage	699
6.75.7	Classifier representation	700
6.75.8	References	700
6.76	gradient_boosting_regression	701
6.76.1	API documentation	701
6.76.2	Loading	701
6.76.3	Testing	701
6.76.4	Features	701
6.76.5	Regressor representation	702
6.76.6	Diagnostics syntax	702
6.76.7	Options	702
6.77	grammars	703

6.77.1	API documentation	703
6.77.2	Loading	703
6.77.3	Testing	703
6.77.4	Usage	704
6.78	graphs	704
6.78.1	API documentation	704
6.78.2	Loading	704
6.78.3	Testing	704
6.78.4	Usage	705
6.79	gsp_pattern_miner	705
6.79.1	API documentation	706
6.79.2	Loading	706
6.79.3	Testing	706
6.79.4	Features	706
6.79.5	Options	706
6.79.6	Pattern miner representation	707
6.79.7	References	707
6.80	hashes	707
6.80.1	API documentation	709
6.80.2	Loading	709
6.80.3	Testing	709
6.80.4	Examples	709
6.81	hdbscan_clusterer	710
6.81.1	API documentation	711
6.81.2	Loading	711
6.81.3	Testing	711
6.81.4	Features	711
6.81.5	Options	711
6.81.6	Clusterer representation	712
6.82	heaps	712
6.82.1	API documentation	712
6.82.2	Loading	712
6.82.3	Testing	712
6.82.4	Usage	713
6.82.5	Credits	713
6.83	hierarchical_clustering	713
6.83.1	API documentation	713
6.83.2	Loading	713
6.83.3	Testing	713
6.83.4	Features	714
6.83.5	Options	714
6.83.6	Clusterer representation	714
6.83.7	Additional API	715
6.83.8	Diagnostics	715
6.84	hierarchies	715
6.84.1	API documentation	715
6.84.2	Loading	716
6.84.3	Testing	716
6.85	hmac	716
6.85.1	API documentation	717
6.85.2	Loading	717
6.85.3	Testing	717
6.85.4	Examples	717
6.86	hodge_rank	717

6.86.1	API documentation	718
6.86.2	Loading	718
6.86.3	Testing	718
6.86.4	Features	718
6.86.5	Scoring semantics	718
6.86.6	Residual semantics	719
6.86.7	Usage	719
6.86.8	Diagnostics syntax	719
6.86.9	Ranker representation	720
6.86.10	References	720
6.87	hook_flows	720
6.87.1	API documentation	720
6.87.2	Loading	720
6.87.3	Testing	720
6.87.4	Usage	721
6.88	hook_objects	721
6.88.1	API documentation	721
6.88.2	Loading	721
6.88.3	Testing	722
6.88.4	Usage	722
6.89	html	725
6.89.1	API documentation	725
6.89.2	Loading	725
6.89.3	Testing	726
6.89.4	Generating a HTML document	726
6.89.5	Generating a HTML fragment	726
6.89.6	Working with callbacks to generate content	727
6.89.7	Working with custom elements	727
6.90	ica_projection	727
6.90.1	API documentation	728
6.90.2	Loading	728
6.90.3	Testing	728
6.90.4	Features	728
6.90.5	Options	728
6.90.6	Usage	729
6.90.7	Dimension reducer representation	730
6.90.8	References	730
6.91	ids	730
6.91.1	API documentation	730
6.91.2	Loading	730
6.91.3	Testing	731
6.91.4	Usage	731
6.92	ieee_754	731
6.92.1	Loading	731
6.92.2	Testing	732
6.92.3	Object model	732
6.92.4	Value representation	732
6.92.5	Public API	733
6.92.6	Semantics	733
6.93	intervals	734
6.93.1	Practical limits	734
6.93.2	API documentation	734
6.93.3	Loading	735
6.93.4	Testing	735

6.93.5	Examples	735
6.94	iqr_anomaly_detector	737
6.94.1	API documentation	737
6.94.2	Loading	737
6.94.3	Testing	737
6.94.4	Features	737
6.94.5	Options	738
6.94.6	Detector representation	739
6.94.7	Notes	739
6.95	iso_639	740
6.95.1	API documentation	740
6.95.2	Loading	740
6.95.3	Testing	740
6.96	iso_3166	740
6.96.1	API documentation	741
6.96.2	Loading	741
6.96.3	Testing	741
6.97	iso_4217	741
6.97.1	API documentation	741
6.97.2	Loading	741
6.97.3	Testing	741
6.98	iso_9362	742
6.98.1	API documentation	742
6.98.2	Loading	742
6.98.3	Testing	742
6.99	iso_13616	742
6.99.1	API documentation	743
6.99.2	Loading	743
6.99.3	Testing	743
6.100	isolation_forest_anomaly_detector	743
6.100.1	API documentation	743
6.100.2	Loading	744
6.100.3	Testing	744
6.100.4	Implemented features	744
6.100.5	Options	744
6.100.6	Detector representation	745
6.100.7	Limitations	745
6.100.8	References	745
6.100.9	Usage	746
6.101	java	747
6.101.1	API documentation	747
6.101.2	Loading	747
6.101.3	Testing	747
6.101.4	Usage	747
6.101.5	Known issues	748
6.102	json	748
6.102.1	API documentation	748
6.102.2	Loading	748
6.102.3	Testing	748
6.102.4	Representation	748
6.102.5	Encoding	751
6.102.6	Decoding	751
6.102.7	Known issues	751
6.103	json_ld	751

6.103.1 API documentation	752
6.103.2 Loading	752
6.103.3 Testing	752
6.103.4 Parsing	752
6.103.5 Generating	752
6.103.6 Expansion	752
6.103.7 Compaction	753
6.103.8 Flattening	753
6.103.9 Supported Features	753
6.103.10 Representation	754
6.104 json_lines	754
6.104.1 API documentation	755
6.104.2 Loading	755
6.104.3 Testing	755
6.104.4 Representation	755
6.104.5 Encoding	757
6.104.6 Decoding	758
6.104.7 Known issues	758
6.105 json_pointer	758
6.105.1 API documentation	758
6.105.2 Loading	758
6.105.3 Testing	758
6.105.4 Representation	759
6.105.5 Examples	759
6.106 json_rpc	759
6.106.1 API documentation	759
6.106.2 Loading	759
6.106.3 Testing	760
6.106.4 Usage	760
6.106.5 API summary	762
6.107 json_schema	764
6.107.1 API documentation	764
6.107.2 Loading	764
6.107.3 Testing	764
6.107.4 Schema Parsing	764
6.107.5 Validation	764
6.107.6 Supported Keywords	765
6.107.7 Known Limitations	766
6.108 kcenters_clusterer	767
6.108.1 API documentation	767
6.108.2 Loading	767
6.108.3 Testing	767
6.108.4 Features	767
6.108.5 Options	768
6.108.6 Diagnostics	768
6.108.7 Clusterer representation	768
6.108.8 References	768
6.109 kemeny_young_ranker	769
6.109.1 API documentation	769
6.109.2 Loading	769
6.109.3 Testing	769
6.109.4 Features	769
6.109.5 Usage	770
6.109.6 Options	771

6.109.7	Scoring semantics	771
6.109.8	Diagnostics syntax	771
6.109.9	Ranker representation	771
6.109.10	Notes	772
6.109.1	References	772
6.110	kernel_svm_classifier	772
6.110.1	API documentation	772
6.110.2	Loading	772
6.110.3	Testing	772
6.110.4	Features	773
6.110.5	Options	773
6.110.6	Usage	773
6.110.7	Classifier representation	774
6.110.8	References	774
6.111	kernel_pca_projection	774
6.111.1	API documentation	775
6.111.2	Loading	775
6.111.3	Testing	775
6.111.4	Features	775
6.111.5	Options	775
6.111.6	Usage	776
6.111.7	Dimension reducer representation	777
6.111.8	References	777
6.112	kmeans_clusterer	777
6.112.1	API documentation	778
6.112.2	Loading	778
6.112.3	Testing	778
6.112.4	Features	778
6.112.5	Options	778
6.112.6	Diagnostics	779
6.112.7	Clusterer representation	779
6.112.8	References	779
6.113	kmedians_clusterer	779
6.113.1	API documentation	780
6.113.2	Loading	780
6.113.3	Testing	780
6.113.4	Features	780
6.113.5	Options	780
6.113.6	Diagnostics	781
6.113.7	Clusterer representation	781
6.113.8	References	781
6.114	kmedoids_clusterer	781
6.114.1	API documentation	782
6.114.2	Loading	782
6.114.3	Testing	782
6.114.4	Features	782
6.114.5	Options	782
6.114.6	Diagnostics	783
6.114.7	Clusterer representation	783
6.114.8	References	783
6.115	kmodes_clusterer	783
6.115.1	API documentation	784
6.115.2	Loading	784
6.115.3	Testing	784

6.115.4	Features	784
6.115.5	Options	784
6.115.6	Diagnostics	785
6.115.7	Clusterer representation	785
6.115.8	References	785
6.116	knn_classifier	785
6.116.1	API documentation	786
6.116.2	Loading	786
6.116.3	Testing	786
6.116.4	Features	786
6.116.5	Usage	786
6.116.6	Options	788
6.116.7	Classifier representation	788
6.116.8	References	788
6.117	knn_distance_anomaly_detector	788
6.117.1	API documentation	789
6.117.2	Loading	789
6.117.3	Testing	789
6.117.4	Features	789
6.117.5	Options	789
6.117.6	Detector representation	790
6.117.7	References	790
6.118	knn_regression	790
6.118.1	API documentation	791
6.118.2	Loading	791
6.118.3	Testing	791
6.118.4	Features	791
6.118.5	Regressor representation	792
6.118.6	Diagnostics syntax	792
6.118.7	Options	792
6.119	kprototypes_clusterer	793
6.119.1	API documentation	793
6.119.2	Loading	793
6.119.3	Testing	793
6.119.4	Features	793
6.119.5	Options	794
6.119.6	Distance Function	794
6.119.7	Clusterer representation	794
6.119.8	References	795
6.120	ksuid	795
6.120.1	API documentation	795
6.120.2	Loading	795
6.120.3	Testing	795
6.120.4	Usage	796
6.121	lasso_regression	796
6.121.1	API documentation	796
6.121.2	Loading	796
6.121.3	Testing	796
6.121.4	Features	797
6.121.5	Regressor representation	797
6.121.6	Diagnostics syntax	797
6.121.7	Options	798
6.122	lda_classifier	798
6.122.1	API documentation	799

6.122.2 Loading	799
6.122.3 Testing	799
6.122.4 Features	799
6.122.5 Options	799
6.122.6 Usage	799
6.122.7 Classifier representation	800
6.122.8 References	800
6.123 lda_projection	801
6.123.1 API documentation	801
6.123.2 Loading	801
6.123.3 Testing	801
6.123.4 Features	801
6.123.5 Options	802
6.123.6 Usage	802
6.123.7 Dimension reducer representation	803
6.123.8 References	803
6.124 linda	803
6.124.1 API documentation	804
6.124.2 Loading	804
6.124.3 Testing	804
6.124.4 Usage	804
6.124.5 Examples	806
6.124.6 API Summary	807
6.124.7 Known issues	808
6.124.8 See also	808
6.125 linear_algebra	808
6.125.1 API documentation	809
6.125.2 Loading	809
6.125.3 Testing	809
6.126 linear_regression	809
6.126.1 API documentation	809
6.126.2 Loading	809
6.126.3 Testing	809
6.126.4 Features	810
6.126.5 Regressor representation	810
6.126.6 Diagnostics syntax	810
6.126.7 Options	811
6.127 linear_svm_classifier	811
6.127.1 API documentation	812
6.127.2 Loading	812
6.127.3 Testing	812
6.127.4 Features	812
6.127.5 Options	813
6.127.6 Usage	813
6.127.7 Classifier representation	813
6.127.8 References	814
6.128 listing	814
6.128.1 API documentation	814
6.128.2 Loading	814
6.128.3 Testing	814
6.128.4 Usage	814
6.129 lof_anomaly_detector	815
6.129.1 API documentation	816
6.129.2 Loading	816

6.129.3	Testing	816
6.129.4	Features	816
6.129.5	Options	817
6.129.6	Detector representation	817
6.129.7	References	817
6.130	logging	818
6.130.1	API documentation	818
6.130.2	Loading	818
6.131	logistic_regression_classifier	818
6.131.1	API documentation	818
6.131.2	Loading	818
6.131.3	Testing	818
6.131.4	Features	819
6.131.5	Options	819
6.131.6	Usage	819
6.131.7	Classifier representation	820
6.131.8	References	820
6.132	loops	821
6.132.1	API documentation	821
6.132.2	Loading	821
6.132.3	Testing	821
6.132.4	Usage	821
6.133	massey_ranker	821
6.133.1	API documentation	822
6.133.2	Loading	822
6.133.3	Testing	822
6.133.4	Features	822
6.133.5	Scoring semantics	822
6.133.6	Usage	823
6.133.7	Diagnostics syntax	824
6.133.8	Ranker representation	824
6.133.9	References	824
6.134	mcp_server	824
6.134.1	API documentation	825
6.134.2	Loading	825
6.134.3	Testing	825
6.134.4	Usage	825
6.134.5	Error handling	832
6.134.6	Protocol	832
6.134.7	Supported MCP methods	833
6.135	memcached	833
6.135.1	API documentation	833
6.135.2	Loading	833
6.135.3	Testing	833
6.135.4	Protocol Version	834
6.135.5	Features	834
6.135.6	Usage	834
6.135.7	API Summary	837
6.135.8	Error Handling	838
6.135.9	Key Format	838
6.135.10	Expiration Times	839
6.136	message_pack	839
6.136.1	API documentation	839
6.136.2	Loading	839

6.136.3	Testing	839
6.136.4	Representation	839
6.136.5	Notes	840
6.137	meta	841
6.137.1	API documentation	841
6.137.2	Loading	841
6.137.3	Testing	841
6.137.4	Usage	841
6.138	meta_compiler	841
6.138.1	API documentation	841
6.138.2	Loading	841
6.138.3	Testing	842
6.138.4	Usage	842
6.138.5	Known issues	842
6.139	mime_types	842
6.139.1	API documentation	843
6.139.2	Loading	843
6.139.3	Testing	843
6.139.4	Usage	843
6.140	modified_z_score_anomaly_detector	843
6.140.1	API documentation	844
6.140.2	Loading	844
6.140.3	Testing	844
6.140.4	Features	844
6.140.5	Options	845
6.140.6	Detector representation	845
6.140.7	Notes	846
6.141	mutations	846
6.141.1	API documentation	847
6.141.2	Loading	847
6.141.3	Testing	847
6.141.4	Usage	847
6.142	multisets	848
6.142.1	API documentation	848
6.142.2	Loading	848
6.142.3	Testing	848
6.143	nanoid	848
6.143.1	API documentation	849
6.143.2	Loading	849
6.143.3	Testing	849
6.143.4	Usage	849
6.144	naive_bayes_classifier	850
6.144.1	API documentation	850
6.144.2	Loading	850
6.144.3	Testing	850
6.144.4	Features	850
6.144.5	Usage	850
6.144.6	Classifier representation	851
6.144.7	References	852
6.145	nearest_centroid_classifier	852
6.145.1	API documentation	852
6.145.2	Loading	852
6.145.3	Testing	852
6.145.4	Features	853

6.145.5 Usage	853
6.145.6 Options	854
6.145.7 Classifier representation	854
6.145.8 References	855
6.146 nested_dictionaries	855
6.146.1 API documentation	855
6.146.2 Loading	855
6.146.3 Testing	855
6.146.4 Usage	855
6.146.5 Curly term representation	856
6.147 Raw Sentence Representation	856
6.148 Typed data/2 Results	857
6.149 Date Rule	857
6.150 Current Scope	857
6.151 nmf_projection	858
6.151.1 API documentation	858
6.151.2 Loading	858
6.151.3 Testing	858
6.151.4 Features	858
6.151.5 Options	859
6.151.6 Usage	859
6.151.7 Dimension reducer representation	860
6.151.8 References	860
6.152 optics_clusterer	860
6.152.1 API documentation	860
6.152.2 Loading	860
6.152.3 Testing	860
6.152.4 Features	861
6.152.5 Options	861
6.152.6 Clusterer representation	861
6.153 optionals	862
6.153.1 API documentation	862
6.153.2 Loading	862
6.153.3 Testing	862
6.153.4 Usage	862
6.153.5 See also	865
6.154 options	865
6.154.1 API documentation	865
6.154.2 Loading	865
6.154.3 Testing	865
6.154.4 Usage	865
6.155 os	866
6.155.1 API documentation	866
6.155.2 Loading	866
6.155.3 Testing	867
6.155.4 Known issues	867
6.156 partitions	867
6.156.1 API documentation	868
6.156.2 Loading	868
6.156.3 Testing	868
6.157 pattern_mining_protocols	868
6.157.1 API documentation	868
6.157.2 Loading	868
6.157.3 Testing	868

6.157.4	Common options	869
6.157.5	Diagnostics	869
6.157.6	Related support libraries	869
6.158	pca_projection	870
6.158.1	API documentation	870
6.158.2	Loading	870
6.158.3	Testing	870
6.158.4	Features	870
6.158.5	Options	870
6.158.6	Usage	871
6.158.7	Dimension reducer representation	871
6.158.8	References	872
6.159	permutations	872
6.159.1	API documentation	872
6.159.2	Loading	873
6.159.3	Testing	873
6.160	plackett_luce_ranker	873
6.160.1	API documentation	873
6.160.2	Loading	873
6.160.3	Testing	873
6.160.4	Features	874
6.160.5	Dataset requirements	874
6.160.6	Usage	874
6.160.7	Diagnostics syntax	875
6.160.8	Options	876
6.160.9	Ranker representation	876
6.160.10	See also	876
6.161	plackett_luce_last_ranker	877
6.161.1	API documentation	877
6.161.2	Loading	877
6.161.3	Testing	877
6.161.4	Features	877
6.161.5	Dataset requirements	878
6.161.6	Usage	878
6.161.7	Diagnostics syntax	878
6.161.8	Options	880
6.161.9	Ranker representation	880
6.161.10	See also	880
6.161.11	References	880
6.162	pls_projection	880
6.162.1	API documentation	881
6.162.2	Loading	881
6.162.3	Testing	881
6.162.4	Features	881
6.162.5	Options	881
6.162.6	Usage	882
6.162.7	Dimension reducer representation	882
6.162.8	References	883
6.163	prefix_span_pattern_miner	883
6.163.1	API documentation	883
6.163.2	Loading	883
6.163.3	Testing	883
6.163.4	Features	883
6.163.5	Options	884

6.163.6	Pattern miner representation	884
6.163.7	References	884
6.164	probabilistic_pca_projection	884
6.164.1	API documentation	885
6.164.2	Loading	885
6.164.3	Testing	885
6.164.4	Features	885
6.164.5	Options	885
6.164.6	Usage	886
6.164.7	Dimension reducer representation	886
6.164.8	References	887
6.165	process	887
6.165.1	API documentation	887
6.165.2	Loading	887
6.165.3	Testing	887
6.165.4	Usage	888
6.166	protobuf	888
6.166.1	API documentation	888
6.166.2	Loading	889
6.166.3	Testing	889
6.166.4	Protocol Buffers Overview	889
6.166.5	Schema Representation	889
6.166.6	Data Representation	890
6.166.7	Encoding	891
6.166.8	Decoding	892
6.166.9	Wire Format Details	892
6.166.10	Binary Format Compatibility	893
6.166.11	Examples and Test Files	893
6.166.12	Further Reading	894
6.167	qda_classifier	894
6.167.1	API documentation	894
6.167.2	Loading	894
6.167.3	Testing	894
6.167.4	Features	894
6.167.5	Options	895
6.167.6	Usage	895
6.167.7	Classifier representation	895
6.167.8	References	896
6.168	queues	896
6.168.1	API documentation	896
6.168.2	Loading	896
6.168.3	Testing	896
6.168.4	Usage	896
6.169	random	897
6.169.1	API documentation	897
6.169.2	Loading	897
6.169.3	Testing	897
6.169.4	Algorithms	897
6.169.5	Usage	898
6.170	random_forest_classifier	899
6.170.1	API documentation	899
6.170.2	Loading	899
6.170.3	Testing	899
6.170.4	Features	899

6.170.5 Options	900
6.170.6 Classifier representation	900
6.170.7 References	900
6.170.8 Usage	900
6.171 random_forest_regression	902
6.171.1 API documentation	902
6.171.2 Loading	902
6.171.3 Testing	902
6.171.4 Features	903
6.171.5 Regressor representation	903
6.171.6 Diagnostics syntax	903
6.171.7 Options	904
6.172 random_projection	904
6.172.1 API documentation	905
6.172.2 Loading	905
6.172.3 Testing	905
6.172.4 Features	905
6.172.5 Options	905
6.172.6 Usage	906
6.172.7 Dimension reducer representation	906
6.172.8 References	907
6.173 rank_centrality	907
6.173.1 API documentation	907
6.173.2 Loading	907
6.173.3 Testing	907
6.173.4 Features	908
6.173.5 Dataset requirements	908
6.173.6 Usage	908
6.173.7 Diagnostics syntax	909
6.173.8 Options	910
6.173.9 Ranker representation	910
6.173.10 References	911
6.174 ranked_pairs	911
6.174.1 API documentation	911
6.174.2 Loading	911
6.174.3 Testing	911
6.174.4 Features	912
6.174.5 Usage	912
6.174.6 Diagnostics syntax	914
6.174.7 Options	914
6.174.8 Scoring semantics	914
6.174.9 Ranker representation	915
6.174.10 See also	915
6.175 ranking_protocols	915
6.175.1 Protocol requirements	915
6.175.2 Shared categories	916
6.175.3 Diagnostics	916
6.175.4 Export header format	916
6.175.5 API documentation	917
6.175.6 Loading	917
6.175.7 Testing	917
6.175.8 Test datasets	917
6.176 reader	918
6.176.1 API documentation	918

6.176.2	Loading	918
6.176.3	Testing	919
6.177	recorded_database	919
6.177.1	API documentation	919
6.177.2	Loading	919
6.177.3	Testing	919
6.177.4	Usage	919
6.177.5	Known issues	920
6.178	redis	920
6.178.1	API documentation	920
6.178.2	Loading	920
6.178.3	Testing	920
6.178.4	Supported Redis Features	920
6.178.5	Credits	924
6.178.6	Known issues	924
6.179	regression_protocols	924
6.179.1	API documentation	924
6.179.2	Loading	924
6.179.3	Testing	925
6.179.4	Test datasets	925
6.180	regression_tree	925
6.180.1	API documentation	926
6.180.2	Loading	926
6.180.3	Testing	926
6.180.4	Export header format	926
6.180.5	Features	926
6.180.6	Regressor representation	927
6.180.7	Diagnostics syntax	927
6.180.8	Options	928
6.181	regularized_bradley_terry_ranker	928
6.181.1	API documentation	928
6.181.2	Loading	928
6.181.3	Testing	929
6.181.4	Features	929
6.181.5	Dataset requirements	929
6.181.6	Usage	929
6.181.7	Diagnostics syntax	930
6.181.8	Options	931
6.181.9	Ranker representation	931
6.181.10	References	932
6.182	ridge_regression	932
6.182.1	API documentation	932
6.182.2	Loading	932
6.182.3	Testing	932
6.182.4	Features	932
6.182.5	Regressor representation	933
6.182.6	Diagnostics syntax	933
6.182.7	Options	934
6.183	schulze_ranker	934
6.183.1	API documentation	934
6.183.2	Loading	934
6.183.3	Testing	935
6.183.4	Features	935
6.183.5	Options	935

6.183.6	Strongest paths	935
6.183.7	Diagnostics syntax	936
6.183.8	Ranker representation	936
6.184	sequential_pattern_mining_protocols	936
6.184.1	API documentation	936
6.184.2	Loading	936
6.184.3	Testing	937
6.184.4	Test datasets	937
6.185	sets	937
6.185.1	API documentation	938
6.185.2	Loading	938
6.185.3	Testing	938
6.185.4	Usage	938
6.185.5	Credits	939
6.186	sgd_classifier	939
6.186.1	API documentation	940
6.186.2	Loading	940
6.186.3	Testing	940
6.186.4	Features	940
6.186.5	Options	940
6.186.6	Usage	941
6.186.7	Classifier representation	941
6.186.8	References	942
6.187	snowflakeid	942
6.187.1	API documentation	942
6.187.2	Loading	942
6.187.3	Testing	942
6.187.4	Usage	942
6.188	sockets	943
6.188.1	Design rationale	943
6.188.2	API documentation	943
6.188.3	Loading	943
6.188.4	Testing	944
6.188.5	Usage	944
6.188.6	API Summary	944
6.188.7	Backend-specific notes	945
6.188.8	Known issues	945
6.189	spade_pattern_miner	945
6.189.1	API documentation	945
6.189.2	Loading	946
6.189.3	Testing	946
6.189.4	Features	946
6.189.5	Options	946
6.189.6	Pattern miner representation	947
6.189.7	References	947
6.190	simulated_annealing	947
6.190.1	API documentation	947
6.190.2	Loading	947
6.190.3	Testing	948
6.190.4	Features	948
6.190.5	Defining a problem	949
6.190.6	Options	949
6.190.7	Run statistics	950
6.190.8	Usage	950

6.191	statistics	952
6.191.1	API documentation	952
6.191.2	Loading	952
6.191.3	Testing	952
6.191.4	API overview	952
6.192	stemming	954
6.192.1	API documentation	955
6.192.2	Loading	955
6.192.3	Testing	955
6.192.4	Usage	955
6.192.5	Algorithms	956
6.192.6	Choosing an Algorithm	957
6.192.7	Known Limitations	957
6.193	stomp	957
6.193.1	Protocol Version	957
6.193.2	Features	957
6.193.3	API documentation	958
6.193.4	Loading	958
6.193.5	Testing	958
6.193.6	Usage	958
6.193.7	API Summary	961
6.193.8	Connection Options	962
6.193.9	Send Options	963
6.193.10	Subscribe Options	963
6.193.11	Receive Options	963
6.193.12	Error Handling	963
6.194	string_distance	963
6.194.1	API documentation	964
6.194.2	Loading	964
6.194.3	Testing	964
6.194.4	Algorithms	964
6.195	strings	965
6.195.1	API documentation	965
6.195.2	Loading	965
6.195.3	Testing	965
6.195.4	Predicates	965
6.196	subsequences	966
6.196.1	API documentation	967
6.196.2	Loading	967
6.196.3	Testing	967
6.197	term_io	967
6.197.1	API documentation	967
6.197.2	Loading	967
6.197.3	Testing	967
6.198	thurstone_mosteller_ranker	968
6.198.1	API documentation	968
6.198.2	Loading	968
6.198.3	Testing	968
6.198.4	Features	968
6.198.5	Scoring semantics	969
6.198.6	Options	969
6.198.7	Diagnostics syntax	969
6.198.8	Ranker representation	969
6.199	time_scales	970

6.199.1	API documentation	970
6.199.2	Loading	970
6.199.3	Testing	970
6.199.4	Features	970
6.199.5	Limitations	971
6.199.6	Representation	971
6.199.7	Examples	971
6.200	timeout	973
6.200.1	API documentation	973
6.200.2	Loading	973
6.200.3	Testing	973
6.200.4	Known issues	973
6.201	tle_orbits	973
6.201.1	API documentation	974
6.201.2	Loading	974
6.201.3	Testing	974
6.201.4	Representation	974
6.201.5	Public API	974
6.201.6	Examples	975
6.201.7	Notes	976
6.202	toml	976
6.202.1	API documentation	976
6.202.2	Loading	976
6.202.3	Testing	976
6.202.4	Status	977
6.202.5	Representation	977
6.202.6	Examples	978
6.202.7	Notes	978
6.203	toon	978
6.203.1	API documentation	978
6.203.2	Loading	978
6.203.3	Testing	979
6.203.4	Representation	979
6.203.5	TOON format features	980
6.203.6	Usage examples	980
6.204	truncated_svd_projection	981
6.204.1	API documentation	981
6.204.2	Loading	981
6.204.3	Testing	981
6.204.4	Features	981
6.204.5	Options	982
6.204.6	Usage	982
6.204.7	Dimension reducer representation	983
6.204.8	References	983
6.205	tsv	983
6.205.1	API documentation	984
6.205.2	Loading	984
6.205.3	Testing	984
6.205.4	Usage	984
6.206	types	985
6.206.1	API documentation	985
6.206.2	Loading	985
6.206.3	Testing	985
6.206.4	Type-checking	985

6.206.5	Defining new types	986
6.206.6	Examples	986
6.207	tzif	986
6.207.1	API documentation	987
6.207.2	Loading	987
6.207.3	Testing	987
6.207.4	Representation	987
6.207.5	Zone identifier validation	988
6.207.6	Local Queries	988
6.207.7	Examples	989
6.208	ulid	990
6.208.1	API documentation	990
6.208.2	Loading	990
6.208.3	Testing	991
6.208.4	Generating ULIDs	991
6.208.5	Type-checking ULIDs	992
6.209	unicode_data	992
6.209.1	Authors	992
6.209.2	License	992
6.209.3	Website	992
6.209.4	Description	992
6.209.5	Requirements	993
6.209.6	Usage	993
6.209.7	Known issues	993
6.209.8	Acknowledgements	993
6.209.9	Files and API Summary	993
6.210	union_find	998
6.210.1	API documentation	999
6.210.2	Loading	999
6.210.3	Testing	999
6.210.4	Usage	999
6.211	url	1000
6.211.1	API documentation	1001
6.211.2	Loading	1001
6.211.3	Testing	1001
6.212	uuid	1001
6.212.1	API documentation	1002
6.212.2	Loading	1002
6.212.3	Testing	1002
6.212.4	Generating version 1 UUIDs	1002
6.212.5	Generating version 4 UUIDs	1003
6.212.6	Generating version 3 UUIDs	1003
6.212.7	Generating version 5 UUIDs	1004
6.212.8	Generating version 7 UUIDs	1004
6.212.9	Generating the Nil and Max UUIDs	1005
6.213	validations	1005
6.213.1	API documentation	1005
6.213.2	Loading	1005
6.213.3	Testing	1005
6.213.4	Usage	1006
6.213.5	Comparison with the expecteds library	1007
6.213.6	See also	1008
6.214	wkt_wkb	1008
6.214.1	API documentation	1009

6.214.2	Loading	1009
6.214.3	Testing	1009
6.214.4	Representation	1009
6.214.5	Sources and sinks	1010
6.214.6	Examples	1010
6.214.7	Notes	1011
6.215	yaml	1011
6.215.1	API documentation	1011
6.215.2	Loading	1011
6.215.3	Testing	1012
6.215.4	API	1012
6.215.5	Data Representation	1013
6.215.6	Examples	1013
6.215.7	Features	1014
6.215.8	Supported YAML Features	1015
6.215.9	Limitations	1016
6.215.10	Error Handling	1016
6.215.11	File Organization	1017
6.215.12	Test Files	1017
6.215.13	Performance Considerations	1017
6.215.14	Integration with Other Libraries	1017
6.215.15	Future Enhancements	1017
6.215.16	See Also	1018
6.216	z_score_anomaly_detector	1018
6.216.1	API documentation	1018
6.216.2	Loading	1018
6.216.3	Testing	1018
6.216.4	Features	1018
6.216.5	Options	1019
6.216.6	Detector representation	1020
6.216.7	Notes	1020
6.217	zipper	1021
6.217.1	API documentation	1021
6.217.2	Loading	1021
6.217.3	Testing	1021
7	Ports	1023
7.1	fcube	1023
7.1.1	API documentation	1024
7.1.2	Loading	1024
7.1.3	Testing	1024
7.2	metagol	1024
7.2.1	API documentation	1025
7.2.2	Loading	1025
7.2.3	Testing	1025
7.3	toychr	1025
7.3.1	API documentation	1026
7.3.2	Loading	1026
7.3.3	Testing	1026
8	Contributions	1027
8.1	flags	1027
8.1.1	API documentation	1027
8.1.2	Loading	1027

8.1.3	Testing	1027
8.2	iso8601	1028
8.2.1	API documentation	1028
8.2.2	Loading	1028
8.2.3	Testing	1029
8.3	pddl_parser	1029
8.3.1	API documentation	1029
8.3.2	Loading	1029
8.3.3	Testing	1029
8.4	verdi_neruda	1030
8.5	xml_parser	1032
8.5.1	API documentation	1032
8.5.2	Loading	1033
8.5.3	Testing	1033
8.5.4	Known issues	1033
9	Glossary	1035
	Bibliography	1045
	Index	1049

USER MANUAL

1.1 Declarative object-oriented programming

Logtalk is a *declarative object-oriented logic programming language*. This means that Logtalk shares key concepts with other object-oriented programming languages but abstracts and reinterprets these concepts in the context of declarative logic programming.

The key concepts in *declarative* object-oriented programming are *encapsulation* and *reuse patterns*. Notably, the concept of *mutable state*, which is an *imperative* concept, is not a significant concept in *declarative* object-oriented programming. Declarative object-oriented programming concepts can be materialized in both logic and functional languages. In this section, we focus only on declarative object-oriented logic programming.

The first key generalization of object-oriented programming concepts is the concept of *object* itself. What an object encapsulates depends on the *base programming paradigm* where we apply object-oriented programming concepts. When these concepts are applied to an imperative language, where mutable state and destructive assignment are central, objects naturally encapsulate and abstract mutable state, providing disciplined access and modification. When these concepts are applied to a declarative logic language such as Prolog, objects naturally encapsulate predicates. Therefore, an object can be seen as a *theory*, expressed by a set of related predicates. Theories are usually static, and thus Logtalk objects are static by default. This contrasts with imperative object-oriented languages where usually classes are static and objects are dynamic. This view of an object as a set of predicates also forgoes a distinction between *data* and *procedures* that is central to imperative object-oriented languages but moot in declarative, *homoiconic* logic languages.

The second key generalization concerns the relation between objects and other entities such as protocols (interfaces) and ancestor objects. The idea is that entity relations define *reuse patterns* and the *roles* played by the participating entities. A common reuse pattern is *inheritance*. In this case, an entity inherits, and thus reuses, resources from an ancestor entity. In a reuse pattern, each participating entity plays a specific *role*. The same entity, however, can play multiple roles depending on its relations with other entities. For example, an object can play the role of a class for its instances, the role of a subclass for its superclasses, and the role of an instance for its metaclass. Another common reuse pattern is *protocol implementation*. In this case, an object implementing a protocol reuses its predicate declarations by providing an implementation for those predicates and exposing those predicates to its clients. An essential consequence of this generalization is that protocols, objects, and categories are first-class entities, while e.g. *prototype*, *parent*, *class*, *instance*, *metaclass*, *subclass*, *superclass*, or *ancestor* are just *roles* that an object can play. When sending a message to an object, the corresponding predicate *declaration* and predicate *definition* lookup procedures (reuse patterns) depend on the role or roles that the object plays (see the *Inheritance* section for details). Another consequence of this generalization is that a language can provide multiple reuse patterns instead of selecting a set of patterns and supporting this set as a design choice that excludes other reuse patterns. For example, most imperative object-oriented languages are either class-based or prototype-based. In contrast, Logtalk supports both classes and prototypes by providing the corresponding reuse patterns using objects as first-class entities capable of playing multiple roles.

1.2 Main features

Several years ago, I decided that the best way to learn object-oriented programming was to build my own object-oriented language. Prolog being always my favorite language, I chose to extend it with object-oriented capabilities. Strong motivation also comes from my frustration with Prolog shortcomings for writing large applications. Eventually this work led to the Logtalk programming language as we know it today. The first system to use the name Logtalk appeared in February 1995. At that time, Logtalk was mainly an experiment in computational reflection with a rudimentary runtime and no compiler. Based on feedback by users and on the author's subsequent work, the name was retained and Logtalk was created as a full programming language focusing on using object-oriented concepts for code encapsulation and reuse. Development started in January 1998 with the first public alpha version released in July 1998. The first stable release (2.0) was published in February 1999. Development of the third generation of Logtalk started in 2012 with the first public alpha version in August 2012 and the first stable release (3.0.0) in January 2015.

Logtalk provides the following features:

1.2.1 Integration of logic and object-oriented programming

Logtalk tries to bring together the main advantages of these two programming paradigms. On one hand, the object orientation allows us to work with the same set of entities in the successive phases of application development, giving us a way of organizing and encapsulating the knowledge of each entity within a given domain. On the other hand, logic programming allows us to represent, in a declarative way, the knowledge we have of each entity. Together, these two advantages allow us to minimize the distance between an application and its problem domain, making the writing and maintenance of programming easier and more productive.

From a pragmatic perspective, Logtalk objects provide Prolog with the possibility of defining several namespaces, instead of the traditional Prolog single database, addressing some of the needs of large software projects.

1.2.2 Integration of event-driven and object-oriented programming

Event-driven programming enables the building of reactive systems, where computing which takes place at each moment is a result of the observation of occurring events. This integration complements object-oriented programming, in which each computing is initiated by the explicit sending of a message to an object. The user dynamically defines what events are to be observed and establishes monitors for these events. This is especially useful when representing relationships between objects that imply constraints in the state of participating objects [Rumbaugh87], [Rumbaugh88], [Fornarino_et_al_89], [Razek92]. Other common uses are reflective applications like code debugging or profiling [Maes87]. Predicates can be implicitly called when a spied event occurs, allowing programming solutions which minimize object coupling. In addition, events provide support for behavioral reflection and can be used to implement the concepts of *pointcut* and *advice* found in Aspect-Oriented Programming.

1.2.3 Support for component-based programming

Predicates can be encapsulated inside *categories* which can be imported by any object, without any code duplication and irrespective of object hierarchies. A category is a first-class encapsulation entity, at the same level as objects and protocols, which can be used as a component when building new objects. Thus, objects may be defined through composition of categories, which act as fine-grained units of code reuse. Categories may also extend existing objects. Categories can be used to implement *mixins* and *aspects*. Categories allow for code reuse between unrelated objects, independent of hierarchy relations, in the same vein as protocols allow for interface reuse.

1.2.4 Support for both prototype and class-based systems

Almost any (if not all) object-oriented languages available today are either class-based or prototype-based [Lieberman86], with a strong predominance of class-based languages. Logtalk provides support for both hierarchy types. That is, we can have both prototype and class hierarchies in the same application. Prototypes solve a problem of class-based systems where we sometimes have to define a class that will have only one instance in order to reuse a piece of code. Classes solve a dual problem in prototype based systems where it is not possible to encapsulate some code to be reused by other objects but not by the encapsulating object. Stand-alone objects, that is, objects that do not belong to any hierarchy, are a convenient solution to encapsulate code that will be reused by several unrelated objects.

1.2.5 Support for multiple object hierarchies

Languages like Smalltalk-80 [Goldberg83], Objective-C [Cox86] and Java [Joy_et_al_00] define a single hierarchy rooted in a class usually named `Object`. This makes it easy to ensure that all objects share a common behavior but also tends to result in lengthy hierarchies where it is difficult to express objects which represent exceptions to default behavior. In Logtalk we can have multiple, independent, object hierarchies. Some of them can be prototype-based while others can be class-based. Furthermore, stand-alone objects provide a simple way to encapsulate utility predicates that do not need or fit in an object hierarchy.

1.2.6 Separation between interface and implementation

This is an expected (should we say standard ?) feature of almost any modern programming language. Logtalk provides support for separating interface from implementation in a flexible way: predicate directives can be contained in an object, a category or a protocol (first-order entities in Logtalk) or can be spread in both objects, categories and protocols.

1.2.7 Private, protected and public inheritance

Logtalk supports private, protected and public inheritance in a similar way to C++ [Stroustrup86], enabling us to restrict the scope of inherited, imported or implemented predicates (by default inheritance is public).

1.2.8 Private, protected and public object predicates

Logtalk supports data hiding by implementing private, protected and public object predicates in a way similar to C++ [Stroustrup86]. Private predicates can only be called from the container object. Protected predicates can be called by the container object or by the container descendants. Public predicates can be called from any object.

1.2.9 Parametric objects

Object names can be compound terms (instead of atoms), providing a way to parameterize object predicates. Parametric objects are implemented in a similar way to L&O [McCabe92], OL(P) [Fromherz93] or SICStus Objects [SICStus95] (however, access to parameter values is done via a built-in method instead of making the parameters scope global over the whole object). Parametric objects allows us to treat any predicate clause as defining an *instantiation* of a parametric object. Thus, a parametric object allows us to encapsulate and associate any number of predicates with a compound term.

1.2.10 High level multi-threading programming support

High level multi-threading programming is available when running Logtalk with selected back-end Prolog compilers, allowing objects to support both synchronous and asynchronous messages. Logtalk allows programmers to take advantage of modern multi-processor and multi-core computers without bothering with the details of creating and destroying threads, implement thread communication, or synchronizing threads.

1.2.11 Smooth learning curve

Logtalk has a smooth learning curve, by adopting standard Prolog syntax and by enabling an incremental learning and use of most of its features.

1.2.12 Compatibility with most Prolog systems and the ISO standard

The Logtalk system has been designed to be compatible with most Prolog compilers and, in particular, with the ISO Prolog standard [ISO95]. It runs in almost any computer system with a modern Prolog compiler.

1.2.13 Performance

The current Logtalk implementation works as a trans-compiler: Logtalk source files are first compiled to Prolog source files, which are then compiled by the chosen Prolog compiler. Therefore, Logtalk performance necessarily depends on the *backend Prolog compiler*. The Logtalk compiler preserves the programmers choices when writing efficient code that takes advantage of tail recursion and first-argument indexing.

As an object-oriented language, Logtalk can use both *static binding* and *dynamic binding* for matching messages and methods. Furthermore, Logtalk entities (objects, protocols, and categories) are independently compiled, allowing for a very flexible programming development. Entities can be edited, compiled, and loaded at runtime, without necessarily implying recompilation of all related entities.

When dynamic binding is used, the Logtalk runtime engine implements caching of *message lookups* (including messages to *self* and *super* calls), ensuring a performance level close to what could be achieved when using static binding.

For more detailed information on performance, see its dedicated *section*.

1.2.14 Logtalk scope

Logtalk, being a superset of Prolog, shares with it the same preferred areas of application but also extends them with those areas where object-oriented features provide an advantage compared to plain Prolog. Among these areas we have:

Logic and object-oriented programming teaching and researching

Logtalk smooth learning curve, combined with support for both prototype and class-based programming, protocols, components or aspects via category-based composition, and other advanced object-oriented features allow a smooth introduction to object-oriented programming to people with a background in Prolog programming. The distribution of Logtalk source code using an open-source license provides a framework for people to learn and then modify to try out new ideas on object-oriented programming research. In addition, the Logtalk distribution includes plenty of programming examples that can be used in the classroom for teaching logic and object-oriented programming concepts.

Structured knowledge representations and knowledge-based systems

Logtalk objects, coupled with event-driven programming features, enable easy implementation of frame-like systems and similar structured knowledge representations.

Blackboard systems, agent-based systems, and systems with complex object relationships

Logtalk support for event-driven programming can provide a basis for the dynamic and reactive nature of blackboard type applications.

Highly portable applications

Logtalk is compatible with most modern Prolog systems that support official and de facto standards. Used as a way to provide Prolog with namespaces, it avoids the porting problems of most Prolog module systems. Platform, operating system, or compiler specific code can be isolated from the rest of the code by encapsulating it in objects with well-defined interfaces.

Alternative to a Prolog module system

Logtalk can be used as an alternative to a Prolog compiler module system. Most Prolog applications that use modules can be converted into Logtalk applications, improving portability across Prolog systems and taking advantage of the stronger encapsulation and reuse framework provided by Logtalk object-oriented features.

Integration with other programming languages

Logtalk support for most key object-oriented features helps users integrating Prolog with object-oriented languages like C++, Java, or Smalltalk by facilitating a high-level mapping between the two languages.

1.3 Nomenclature

Depending on your logic programming and object-oriented programming background (or lack of it), you may find Logtalk nomenclature either familiar or at odds with the terms used in other languages. In addition, being a superset of Prolog, terms such as *predicate* and *method* are often used interchangeably. Logtalk inherits most of its nomenclature from Prolog and Smalltalk.

Note that the same terms can have different meanings in different languages. A good example is *class*. The support for meta-classes in e.g. Smalltalk translates to a concept of class that is different in key aspects from the concept of class in e.g. Java or C++. Other terms that can have different meanings are *delegation* and *forwarding*. There are also cases where the same concept is found under different names in some languages (e.g., *self* and *this*) but that can also mean different concepts in Logtalk and other languages. Always be aware of these differences and be cautious with assumptions carried from other programming languages.

In this section, we map nomenclatures from Prolog and popular OOP languages such as Smalltalk, C++, Java, and Python to the Logtalk nomenclature. The Logtalk distribution includes several examples of how to implement common concepts found in other languages, complementing the information in this section. This Handbook also features a [Prolog interoperability section](#) and an extensive [glossary](#) providing the exact meaning of the names commonly used in Logtalk programming.

1.3.1 Prolog nomenclature

Being a superset of Prolog, Logtalk inherits its nomenclature. But Logtalk also aims to fix several Prolog shortcomings, thus introducing new concepts and refining existing Prolog concepts. Logtalk object-oriented nature also introduces names and concepts that are not common when discussing logic programming semantics. We mention here the most relevant ones, notably those where semantics or common practice differ. Further details can be found elsewhere in this Handbook.

arbitrary goals as directives

Although not ISO Prolog Core standard compliant, several Prolog systems accept using arbitrary goal as directives. This is not supported in Logtalk source files. Always use an [initialization/1](#) directive to wrap those goals. This ensures that any initialization goals, which often have side-effects, are only called if the source file is successfully compiled and loaded.

calling a predicate

Sending a [message](#) to an object is similar to *calling a goal* with the difference that the actual predicate that is called is determined not just by the message *term* but also by the object receiving the message and possibly its ancestors. This is also different from calling a Prolog module predicate: a message may result e.g. in calling a predicate [inherited](#) by the object but calling a module predicate requires the predicate to exist in (or be reexported by) the module.

closed-world assumption semantics

Logtalk provides clear closed-world assumption semantics: messages or calls for declared but undefined predicates fail. Messages or calls for unknown (i.e., not declared) predicates throw an error. Crucially, this semantics applies to both *static* and *dynamic* predicates. But in Prolog workarounds are required to have a static predicate being known by the runtime without it being also defined (so that calling it would fail instead of throwing a predicate existence error).

compiling and loading source files

Logtalk provides its own built-in predicates for [compiling and loading](#) source files. It also provides convenient top-level interpreter [shorthands](#) for these and other frequent operations. In general, the traditional Prolog built-in predicates and top-level interpreter shorthands cannot be used to load Logtalk source files.

debugging

In most (if not all) Prolog systems, debugging support is a built-in feature made available using a

set of built-in predicates like `trace/0` and `spy/1`. But in Logtalk the *default debugger* is a regular application, implemented using a public *reflection API*. This means that the debugger must be explicitly loaded (either automatically from a *settings file* at startup or from the top-level). It also means that the debugger can be easily extended or replaced by an alternative application.

directive operators

Some Prolog systems declare directive names as operators (e.g., `dynamic`, `multifile`, ...). This is not required by the ISO Prolog Core standard. It's a practice that should be avoided as it makes code non-portable.

encapsulation

Logtalk enforces encapsulation of object predicates, generating a permission error when a predicate is not within the scope of the caller. In contrast, most Prolog module systems allow any module predicate to be called by using explicit qualification, even if not exported. Worse, some Prolog systems also allow defining clauses for a module predicate outside the module, without declaring the predicate as `multifile`, by simply writing clauses with explicit module-qualified heads.

entity loading

When using Prolog modules, `use_module/1-2` (or equivalent) directives both load the module files and declare that the (implicitly or explicitly) imported predicates can be used with implicit module qualification. But Logtalk separates entity (object, protocol, category, or module) predicate *usage* declarations (via `uses/1` and `uses/2` or its own `use_module/1` and `use_module/2` directives) from *loading* goals (using the `logtalk_load/1` and `logtalk_load/2` predicates), called using an explicit and disciplined approach from *loader files*.

flags scope

The `set_logtalk_flag/2` **directive** is always local to the entity or source file that contains it. Only calls to the `set_logtalk_flag/2` **predicate** set the global default value for a flag. This distinction is lacking in Prolog (where directives usually have a global scope) and Prolog modules (where some flags are local to modules in some systems and global in other systems).

meta-predicate call semantics

Logtalk provides consistent *meta-predicate* call semantics: meta-arguments are always called in the meta-predicate calling context. This contrasts with Prolog module meta-predicates where the semantics of implicitly qualified calls are different from explicitly qualified calls.

operators scope

Operators declared inside an entity (object, protocol, or category) are local to the entity. But operators defined in a source file but outside an entity are global for compatibility with existing Prolog code.

predicates scope

In plain Prolog, all predicates are visible. In a Prolog module, a predicate can be exported or local. In Logtalk, a predicate can be *public*, *protected*, *private*, or *local*.

predicate declaration

Logtalk provides a clear distinction between *declaring* a predicate and *defining* a predicate. This is a fundamental requirement for the concept of *protocol* (aka interface) in Logtalk: we must be able to *declare* a predicate without necessarily *defining* it. This clear distinction is missing in Prolog and Prolog modules. Notably, it's a compiler error for a module to try to export a predicate that it does not define.

predicate loading conflicts

Logtalk does not use predicate import/export semantics. Thus, there are never conflicts when loading entities (objects, protocols, or categories) that declare the same public predicates. But attempting to load two Prolog modules that export the same predicate results in a conflict, usually a compilation error (this is especially problematic when the `use_module/1` directive is used; e.g. adding a new exported predicate can break applications that use the module but not the new predicate).

1.3.2 Smalltalk nomenclature

The Logtalk name originates from a combination of the Prolog and Smalltalk names. Smalltalk had a significant influence on the design of Logtalk and thus inherits some of its ideas and nomenclature. The following list relates the most commonly used Smalltalk terms with their Logtalk counterparts.

abstract class

Similar to Smalltalk, an abstract class is just a class not meant to be instantiated by not understanding a message to create instances.

assignment statement

Logtalk, as a superset of Prolog, uses *logic variables* and *unification* and thus provides no equivalent to the Smalltalk assignment statement.

block

Logtalk supports lambda expressions and meta-predicates, which can be used to provide similar functionality to Smalltalk blocks.

class

In Logtalk, *class* is just a *role* that an object can play. This is similar to Smalltalk, where classes are also objects.

class method

Class methods in Logtalk are simply instance methods declared and defined in the class metaclass.

class variable

Logtalk objects, which can play the roles of class and instance, encapsulate predicates, not state. Class variables, which in Smalltalk are really shared instance variables, can be emulated in a class by defining a predicate locally instead of defining it in the class instances.

inheritance

While Smalltalk only supports single inheritance, Logtalk supports single inheritance, multiple inheritance, and multiple instantiation.

instance

While in Smalltalk every object is an *instance* of some class, objects in Logtalk can play different roles, including the role of a prototype where the concepts of instance and class don't apply. Moreover, instances can be either created dynamically or defined statically.

instance method

Instance methods in Logtalk are simply predicates declared and defined in a class and thus inherited by the class instances.

instance variable

Logtalk being a *declarative* language, objects encapsulate a set of predicates instead of encapsulating *state*. But different objects may provide different definitions of the same predicates. Mutable internal state as in Smalltalk can be emulated by using dynamic predicates.

message

Similar to Smalltalk, a *message* is a request for an operation, which is interpreted in Logtalk as a logic query, asking for the construction of a proof that something is true.

message selector

Logtalk uses the predicate template (i.e., the predicate callable term with all its arguments unbound) as a message selector. The actual type of the message arguments is not considered. Like Smalltalk, Logtalk uses *single dispatch* on the message receiver.

metaclass

Metaclasses are optional in Logtalk (except for a root class) and can be shared by several classes. When metaclasses are used, infinite regression is simply avoided by making a class an instance of itself.

method

Same as in Smalltalk, a *method* is the actual code (i.e., predicate definition) that is run to answer a message. Logtalk uses the words *method* and *predicate* interchangeably.

method categories

There is no support in Logtalk for partitioning the methods of an object into different categories. The Logtalk concept of *category* (a first-class entity) was, however, partially inspired by Smalltalk method categories.

object

Unlike Smalltalk, where *everything* is an object, Logtalk language constructs include both *terms* (as in Prolog representing e.g. numbers and structures) and three first-class entities: objects, protocols, and categories.

pool variables*

Logtalk, as a superset of Prolog, uses *predicates* with no distinction between *variables* and *methods*. Categories can be used to share a set of predicate definitions between any number of objects.

protocol

In Smalltalk, an object *protocol* is the set of messages it understands. The same concept applies in Logtalk. But Logtalk also supports protocols as first-class entities where a protocol can be implemented by multiple objects and an object can implement multiple protocols.

self

Logtalk uses the same definition of *self* found in Smalltalk: the object that received the message being processed. Note, however, that *self* is not a keyword in Logtalk but is implicit in the `(::)/1` message to *self* control construct.

subclass

Same definition in Logtalk.

super

As in Smalltalk, the idea of *super* is to allow calling an inherited predicate (that is usually being redefined). Note, however, that *super* is not a keyword in Logtalk, which provides instead a `(^ ^)/1` *super* call control construct.

superclass

Same definition in Logtalk. But while in Smalltalk a class can only have a single superclass, Logtalk support for multiple inheritance allows a class to have multiple superclasses.

1.3.3 C++ nomenclature

There are several C++ glossaries available on the Internet. The list that follows relates the most commonly used C++ terms with their Logtalk equivalents.

abstract class

Logtalk uses an *operational* definition of abstract class: any class that does not inherit a method for creating new instances can be considered an abstract class. Moreover, Logtalk supports *inter-faces/protocols*, which are often a better way to provide the functionality of C++ abstract classes.

base class

Logtalk uses the term *superclass* with the same meaning.

data member

Logtalk uses *predicates* for representing both behavior and data.

constructor function

There are no special methods for creating new objects in Logtalk. Instead, Logtalk provides a built-in

predicate, *create_object/4*, which can be used as a building block to define more sophisticated object creation predicates.

derived class

Logtalk uses the term *subclass* with the same meaning.

destructor function

There are no special methods for deleting new objects in Logtalk. Instead, Logtalk provides a built-in predicate, *abolish_object/1*, which is often used to define more sophisticated object deletion predicates.

friend function

Not supported in Logtalk. Nevertheless, see the User Manual section on *meta-predicates*.

instance

In Logtalk, an instance can be either created dynamically at runtime or defined statically in a source file in the same way as classes.

member

Logtalk uses the term *predicate*.

member function

Logtalk uses *predicates* for representing both behavior and data.

namespace

Logtalk does not support multiple identifier namespaces. All Logtalk entity identifiers share the same namespace (Logtalk entities are objects, categories, and protocols).

nested class

Logtalk does not support nested classes.

static member

Logtalk does not support a static keyword. But the equivalent of static members can be declared in a class metaclass.

template

Logtalk supports *parametric objects*, which allows you to get the similar functionality of templates at runtime.

this

Logtalk uses the built-in context method *self/1* for retrieving the instance that received the message being processed. Logtalk also provides a *this/1* method but for returning the class containing the method being executed. Why the name clashes? Well, the notion of *self* was inherited from Smalltalk, which predates C++.

virtual member function

There is no virtual keyword in Logtalk. Any inherited or imported predicate can be redefined (either overridden or specialized). Logtalk can use *static binding* or *dynamic binding* for locating both method declarations and method definitions. Moreover, methods that are declared but not defined simply fail when called (as per the *closed-world assumption*).

1.3.4 Java nomenclature

There are several Java glossaries available on the Internet. The list that follows relates the most commonly used Java terms with their Logtalk equivalents.

abstract class

Logtalk uses an *operational* definition of abstract class: any class that does not inherit a method for creating new instances is an abstract class. I.e. there is no abstract keyword in Logtalk.

abstract method

In Logtalk, you may simply declare a method (*predicate*) in a class without defining it, leaving its definition to some descendant subclass.

assertion

There is no assertion keyword in Logtalk. Assertions are supported using Logtalk compilation hooks and developer tools.

class

Logtalk objects can play the role of classes, instances, or protocols (depending on their relations with other objects).

extends

There is no extends keyword in Logtalk. Class inheritance is indicated using *specialization relations*. Moreover, the *extends relation* is used in Logtalk to indicate protocol, category, or prototype extension.

interface

Logtalk uses the term *protocol* with a similar meaning. But note that Logtalk objects and categories declared as implementing a protocol are not required to provide definitions for the declared predicates (*closed-world assumption*).

callback method

Logtalk supports *event-driven programming*, the most common usage context of callback methods. Callback methods can also be implemented using *meta-predicates*.

constructor

There are no special methods for creating new objects in Logtalk. Instead, Logtalk provides a built-in predicate, *create_object/4*, which is often used to define more sophisticated object creation predicates.

final

There is no final keyword in Logtalk. Predicates can always be redeclared and redefined in subclasses (and instances!).

inner class

Inner classes are not supported in Logtalk.

instance

In Logtalk, an instance can be either created dynamically at runtime or defined statically in a source file in the same way as classes.

method

Logtalk uses the term *predicate* interchangeably with the term *method*.

method call

Logtalk usually uses the expression *message-sending* for method calls, true to its Smalltalk heritage.

method signature

Logtalk selects the method/predicate to execute in order to answer a method call based only on the method name and number of arguments. Logtalk (and Prolog) are not typed languages in the same sense as Java.

package

There is no concept of packages in Logtalk. All Logtalk entities (objects, protocols, categories) share a single namespace. But Logtalk does support a concept of *library* that allows grouping of entities whose source files share a common path prefix.

reflection

Logtalk features a *white box* API supporting *structural* reflection about *entity contents*, a *black box* API supporting *behavioral* reflection about *object protocols*, and an *events* API for reasoning about messages exchanged at runtime.

static

There is no static keyword in Logtalk. See the entries below on *static method* and *static variable*.

static method

Static methods may be implemented in Logtalk by using a *metaclass* for the class and defining the static methods in the metaclass. I.e. static methods are simply instance methods of the class metaclass.

static variable

Static variables are *shared instance variables* and can simply be both declared and defined in a class. The built-in database methods can be used to implement destructive updates if necessary by accessing and updating a single clause of a dynamic predicate stored in the class.

super

Instead of a super keyword, Logtalk provides a super operator and control construct, $(\hat{\hat{}})/1$, for calling overridden methods.

synchronized

Logtalk supports *multi-threading programming* in selected Prolog compilers, including a *synchronized/1* predicate directive. Logtalk allows you to synchronize a predicate or a set of predicates using per-predicate or per-predicate-set *mutexes*.

this

Logtalk uses the built-in context method *self/1* for retrieving the instance that received the message being processed. Logtalk also provides a *this/1* method but for returning the class containing the method being executed. Why the name clashes? Well, the notion of *self* was inherited from Smalltalk, which predates C++.

1.3.5 Python nomenclature

The list that follows relates the commonly used Python concepts with their Logtalk equivalents.

abstract class

Logtalk uses a different definition of abstract class: a class that does not inherit a method for creating new instances. Notably, the presence of *abstract methods* (i.e., predicates that are declared but not defined) does not make a class abstract.

abstract method

Logtalk uses the term *predicate* interchangeably with *method*. Predicates can be declared without also being defined in an object (or category).

class

Logtalk objects can play the role of classes, instances, or protocols (depending on their relations with other objects).

dictionary

There is no native, built-in associative data type. But the library provides several implementations of a dictionary protocol.

function

The closest equivalent is a predicate defined in user, a pseudo-object for predicates not defined in regular objects, and thus callable from anywhere without requiring a scope directive.

function object

Predicate calls (goals) can be passed or returned from other predicates and unified with other terms (e.g., variables).

import path

Logtalk uses the term *library* to refer to a directory of source files and supports defining aliases (symbolic names) to library paths to abstract the actual locations.

lambda

Logtalk natively supports lambda expressions.

list

Lists are compound terms with native syntax support.

list comprehensions

There is no native, built-in support for list comprehensions. But the standard `findall/3` predicate can be used to construct a list by calling a goal that generates the list elements.

loader

Logtalk uses the term *loader* to refer to source files whose main or sole purpose is to load other source files.

loop

There are no native loop control constructs based on a counter. But the library provides implementations of several loop predicates.

metaclass

Logtalk objects play the role of metaclasses when instantiated by objects that play the role of classes.

method

Logtalk uses the terms *method* and *predicate* interchangeably. Predicates can be defined in objects (and categories). The value of *self* is implicit, unlike in Python where it is the first parameter of any method.

method resolution order

Logtalk uses a depth-first algorithm to lookup method (predicate) declarations and definitions. It's possible to use predicate *aliases* to access predicate declarations and definitions other than the first ones found by the lookup algorithm.

object

Objects are first-class entities that can play multiple roles, including prototype, class, instance, and metaclass.

package

Logtalk uses the term *library* to refer to a directory of source files defining objects, categories, and protocols.

set

There is no native, built-in set type. But the library provides set implementations.

string

The interpretation of text between double-quotes depends on the `double_quotes` flag. Depending on this flag, double-quoted text can be interpreted as a list of characters, a list of character codes, or an atom. Some backend Prolog compilers allow double-quoted text to be interpreted as a string in the Python sense.

tuple

Compound terms can be used to represent tuples of any complexity.

variable

Logtalk works with *logical variables*, which are close to the mathematical concept of variables and distinct from variables in imperative or imperative-based OOP languages where they are symbolic names for memory locations. Logical variables can be *unified* with any term, including other variables.

while loop

The built-in forall/2 predicate implements a *generate-and-test* loop.

1.4 Messages

Messages allow us to ask an object to prove a goal and must always match a declared predicate within the scope of the *sender* object. Note that sending a message is fundamentally different from calling a predicate. When calling a predicate, the caller decides implicitly which predicate definition will be executed. When sending a message, it is the receiving object, not the sender, that decides which predicate definition (if any) will be called to answer the message. The predicate definition that is used to answer a message depends on the relations between the object and its imported categories and ancestor objects (if any). See the *Inheritance* section for details on the predicate declaration and predicate definition lookup procedures.

When a message corresponds to a *meta-predicate*, the meta-arguments are always called in the context of the object (or category) sending the message.

Logtalk uses nomenclature similar to other object-oriented programming languages such as Smalltalk. Therefore, the terms *query* and *message* are used interchangeably when referring to a declared predicate that is part of an object interface. Likewise, the terms *predicate* and *method* are used interchangeably when referring to the predicate definition (inside an object or category) that is called to answer a message.

1.4.1 Operators used in message-sending

Logtalk declares the following operators for the message-sending control constructs:

```
:- op(600, xfy, ::).  
:- op(600,  fy, ::).  
:- op(600,  fy, ^^).
```

It is assumed that these operators remain active (once the Logtalk compiler and runtime files are loaded) until the end of the Prolog session (this is the usual behavior of most Prolog compilers). Note that these operator definitions are compatible with the predefined operators in the Prolog ISO standard.

1.4.2 Sending a message to an object

Sending a message to an object is accomplished by using the `(::)/2` control construct:

```
..., Object::Message, ...
```

The message must match a public predicate declared for the receiving object. The message may also correspond to a protected or private predicate if the *sender* matches the predicate scope container. If the predicate is declared but not defined, the message simply fails (as per the *closed-world assumption*).

1.4.3 Delegating a message to an object

It is also possible to send a message to an object while preserving the original *sender* and *meta-call context* by using the `[]/1` delegation control construct:

```
..., [Object::Message], ....
```

This control construct can only be used within objects and categories (at the top-level interpreter, the *sender* is always the pseudo-object user so using this control construct would be equivalent to using the `(::)/2` message-sending control construct).

1.4.4 Sending a message to *self*

While defining a predicate, we sometimes need to send a message to *self*, i.e., to the same object that has received the original message. This is done in Logtalk through the `(::)/1` control construct:

```
..., ::Message, ....
```

The message must match either a public or protected predicate declared for the receiving object or a private predicate within the scope of the *sender* otherwise an error will be thrown. If the message is sent from inside a category or if we are using private inheritance, then the message may also match a private predicate. Again, if the predicate is declared but not defined, the message simply fails (as per the *closed-world assumption*).

1.4.5 Broadcasting

In the Logtalk context, *broadcasting* is interpreted as the sending of several messages to the same object. This can be achieved by using the message-sending control construct described above. However, for convenience, Logtalk implements an extended syntax for message-sending that may improve program readability in some cases. This extended syntax uses the `(,)/2`, `(;)/2`, and `(->)/2` control constructs (plus the `(*->)/2` soft-cut control construct when provided by the backend Prolog compiler). For example, if we wish to send several messages to the same object, we can write:

```
| ?- Object::(Message1, Message2, ...).
```

This is semantically equivalent to:

```
| ?- Object::Message1, Object::Message2, ... .
```

This extended syntax may also be used with the `(::)/1` message-sending control construct.

1.4.6 Calling imported and inherited predicates

When redefining a predicate, sometimes we need to call the inherited definition in the new code. This functionality, introduced by the Smalltalk language through the *super* primitive, is available in Logtalk using the `(^^)/1` control construct:

```
..., ^^Predicate, ....
```

Most of the time we will use this control construct by instantiating the pattern:

```
Predicate :-  
    ...,                % do something  
    ^^Predicate,        % call inherited definition  
    ... .               % do something more
```

This control construct is generalized in Logtalk where it may be used to call any imported or inherited predicate definition. This control construct may be used within objects and categories. When combined with *static binding*, this control construct allows imported and inherited predicates to be called with the same performance as local predicates. As with the message-sending control constructs, the $(^^)/1$ call simply fails when the predicate is declared but not defined (as per the *closed-world assumption*).

1.4.7 Message sending and event generation

Assuming the *events* flag is set to allow for the object (or category) sending a message using the $(::)/2$ control construct, two events are generated, one before and one after the message execution. Messages that are sent using the $(::)/1$ (message to *self*) control construct or the $(\wedge\wedge)/1$ super mechanism described above do not generate any events. The rationale behind this distinction is that messages to *self* and *super* calls are only used internally in the definition of methods or to execute additional messages with the same target object (represented by *self*). In other words, events are only generated when using an object's public interface; they cannot be used to break object encapsulation.

If we need to generate events for a public message sent to *self*, then we just need to write something like:

```
Predicate :-  
    ...,  
    % get self reference  
    self(Self),  
    % send a message to self using (::)/2  
    Self::Message,  
    ... .
```

If we also need the sender of the message to be other than the object containing the predicate definition, we can write:

```
Predicate :-  
    ...,  
    % send a message to self using (::)/2  
    % sender will be the pseudo-object user  
    self(Self),  
    {Self::Message},  
    ... .
```

When events are not used, it is possible to turn off event generation globally or on a per-entity basis by using the events compiler flag to optimize message-sending performance (see the *Event-driven programming* section for more details).

1.4.8 Sending a message from a module

Messages can be sent to objects from within Prolog modules. Depending on the backend support for goal-expansion and on the *optimize* flag being turned on, the messages will use static binding when possible. This optimization requires the object to be compiled and loaded before the module. Note that the module can be user. This is usually the case when sending the message from the top-level interpreter. Thus, the same conditions apply in this case. Note that loading Prolog modules using Prolog directives or built-in predicates necessarily limits the range of possible optimizations for messages sent from the modules.

Warning

If you want to benchmark the performance of a message-sending goal at the top-level interpreter, be careful to check first if the goal is pre-compiled to use static binding; otherwise you will also be benchmarking the Logtalk compiler itself.

1.4.9 Message sending performance

For a detailed discussion on message-sending performance, see the *Performance* section.

1.5 Objects

The main goal of Logtalk objects is the encapsulation and reuse of predicates. Instead of a single database containing all your code, Logtalk objects provide separated namespaces or databases, allowing the partitioning of code into more manageable parts. Logtalk is a *declarative programming language* and does not aim to bring some sort of new dynamic state change concept to Logic Programming or Prolog.

Logtalk, defines two built-in objects, *user* and *logtalk*, which are described at the end of this section.

1.5.1 Objects, prototypes, classes, and instances

There are only three kinds of encapsulation entities in Logtalk: *objects*, *protocols*, and *categories*. Logtalk uses the term *object* in a broad sense. The terms *prototype*, *parent*, *class*, *subclass*, *superclass*, *metaclass*, and *instance* always designate an object. Different names are used to emphasize the *role* played by an object in a particular context. I.e. we use a term other than object when we want to make the relationship with other objects explicit. For example, an object with an *instantiation* relation with another object plays the role of an *instance*, while the instantiated object plays the role of a *class*; an object with a *specialization* relation with another object plays the role of a *subclass*, while the specialized object plays the role of a *superclass*; an object with an *extension* relation with another object plays the role of a *prototype*, the same for the extended object. A *stand-alone* object, i.e. an object with no relations with other objects, is always interpreted as a prototype. In Logtalk, entity relations essentially define *patterns* of code reuse. An entity is compiled according to the roles it plays.

Logtalk allows you to work from standalone objects to any kind of hierarchy, either class-based or prototype-based. You may use single or multiple inheritance, use or forgo metaclasses, implement reflective designs, use parametric objects, and take advantage of protocols and categories (think components).

Prototypes

Prototypes are either self-defined objects or objects defined as extensions to other prototypes with whom they share common properties. Prototypes are ideal for representing one-of-a-kind objects. Prototypes usually represent concrete objects in the application domain. When linking prototypes using *extension* relations, Logtalk uses the term *prototype hierarchies* although most authors prefer to use the term *hierarchy* only with class generalization/specialization relations. In the context of logic programming, prototypes are often the ideal replacement for modules.

Classes

Classes are used to represent abstractions of common properties of sets of objects. Classes often provide an ideal structuring solution when you want to express hierarchies of abstractions or work with many similar objects. Classes are used indirectly through *instantiation*. Contrary to most object-oriented programming languages, instances can be created both dynamically at runtime or defined in a source file like other objects. Using classes requires defining at least one *metaclass*, as explained below.

1.5.2 Defining a new object

We can define a new object in the same way we write Prolog code: by using a text editor. Logtalk source files may contain one or more objects, categories, or protocols. If you prefer to define each entity in its own source file, it is recommended that the file be named after the object. By default, all Logtalk source files use the extension `.lgt` but this is optional and can be set in the adapter files. Intermediate Prolog source files (generated by the Logtalk compiler) have, by default, a `_lgt` suffix and a `.pl` extension. Again, this can be set to match the needs of a particular Prolog compiler in the corresponding adapter file. For instance, we may define an object named `vehicle` and save it in a `vehicle.lgt` source file, which will be compiled to a `vehicle_lgt.pl` Prolog file (depending on the *backend compiler*, the names of the intermediate Prolog files may include a directory hash and a process identifier to prevent file name clashes when embedding Logtalk applications or running parallel Logtalk processes).

Object names can be atoms or compound terms (when defining parametric objects, see below). Objects, categories, and protocols share the same name space: we cannot have an object with the same name as a protocol or a category.

Object code (directives and predicates) is textually encapsulated by using two Logtalk directives: *object/1-5* and *end_object/0*. The simplest object will be one that is self-contained, not depending on any other Logtalk entity:

```
:- object(Object).  
    ...  
:- end_object.
```

If an object implements one or more protocols then the opening directive will be:

```
:- object(Object,  
    implements([Protocol1, Protocol2, ...])).  
    ...  
:- end_object.
```

An object can import one or more categories:

```
:- object(Object,  
    imports([Category1, Category2, ...])).
```

(continues on next page)

(continued from previous page)

```
...
:- end_object.
```

If an object both implements protocols and imports categories, then we will write:

```
:- object(Object,
    implements([Protocol1, Protocol2, ...]),
    imports([Category1, Category2, ...])).
...
:- end_object.
```

In object-oriented programming, objects are usually organized in hierarchies that enable interface and code sharing by inheritance. In Logtalk, we can construct prototype-based hierarchies by writing:

```
:- object(Prototype,
    extends(Parent)).
...
:- end_object.
```

We can also have class-based hierarchies by defining instantiation and specialization relations between objects. To define an object as a class instance we will write:

```
:- object(Object,
    instantiates(Class)).
...
:- end_object.
```

A class may specialize another class, its superclass:

```
:- object(Class,
    specializes(Superclass)).
...
:- end_object.
```

If we are defining a reflexive system where every class is also an instance, we will probably be using the following pattern:

```
:- object(Class,
    instantiates(Metaclass),
    specializes(Superclass)).
...
:- end_object.
```

In short, an object can be a *stand-alone* object or be part of an object hierarchy. The hierarchy can be prototype-based (defined by extending other objects) or class-based (with instantiation and specialization relations). An object may also implement one or more protocols or import one or more categories.

A *stand-alone* object (i.e., an object with no extension, instantiation, or specialization relations with other objects) always plays the role of a prototype, that is, a self-describing object. If we want to use classes and instances, then we will need to specify at least one instantiation or specialization relation. The best way to do this is to define a set of objects that provide the basis of a reflective system [Cointe87], [Moura94]. For example:

```
% avoid the inevitable unknown entity warnings as in a
% reflective system there will always be references to
% an entity that will be defined after the reference

:- set_logtalk_flag(unknown_entities, silent).

% default root of the inheritance graph
% providing predicates common to all objects

:- object(object,
    instantiates(class)).
    ...
:- end_object.

% default metaclass for all classes providing
% predicates common to all instantiable classes

:- object(class,
    instantiates(class),
    specializes(abstract_class)).
    ...
:- end_object.

% default metaclass for all abstract classes
% providing predicates common to all classes

:- object(abstract_class,
    instantiates(class),
    specializes(object)).
    ...
:- end_object.
```

Note that with these instantiation and specialization relations, `object`, `class`, and `abstract_class` are, at the same time, classes and instances of some class. In addition, each object inherits its own predicates and the predicates of the other two objects without any inheritance loop.

When a full-blown reflective system solution is not needed, the above scheme can be simplified by making an object an instance of itself, i.e. by making a class its own metaclass. For example:

```
:- object(class,
    instantiates(class)).
    ...
:- end_object.
```

We can use, in the same application, both prototype and class-based hierarchies (and freely exchange messages between all objects). We cannot, however, mix the two types of hierarchies by, e.g., specializing an object that extends another object in this current Logtalk version.

Logtalk also supports public, protected, and private inheritance. See the [inheritance](#) section for details.

1.5.3 Parametric objects

Parametric objects have a compound term as identifier where all the arguments of the compound term are variables. These variables can be bound when sending a message or become bound when a message to the object succeeds, thus acting as *object parameters*. The object predicates can be coded to depend on those parameters, which are *logical variables* shared by all object predicates. Parametric objects can also be used to associate a set of predicates to terms that share a common functor and arity.

As a consequence of parameters being logical variables, when defining, for example, a parametric object `foo/1`, `foo(bar)` and `foo(baz)` are different parameterizations of the same object, not distinct objects.

Accessing object parameters

Object parameters can be accessed using *parameter variables* or built-in execution context methods. Parameter variables is the recommended solution to access object parameters. Although they introduce a concept of entity global variables, their unique syntax (`_ParameterName_`) avoids conflicts, makes them easily recognizable, and distinguishes them from other named anonymous variables. For example:

```
:- object(foo(_Bar_, _Baz_, ...)).

...
bar(_Bar_).

baz :-
    baz(_Baz_),
    ... .
```

Note that using parameter variables doesn't change the fact that entity parameters are logical variables. Parameter variables simplify code maintenance by allowing parameters to be added, reordered, or removed without having to specify or update parameter indexes.

Logtalk provides also a *parameter/2* built-in local method to access individual parameters:

```
:- object(foo(_Bar, _Baz, ...)).

...
bar(Bar) :-
    parameter(1, Bar).

baz :-
    parameter(2, Baz),
    baz(Baz),
    ... .
```

An alternative solution is to use the built-in local method *this/1*, which allows access to all parameters with a single call. For example:

```
:- object(foo(_Bar, _Baz, ...)).

...
baz :-
    this(foo(_, Baz, ...)),
    baz(Baz),
    ... .
```

Both solutions are equally efficient as calls to the methods `this/1` and `parameter/2` are usually compiled inline into a clause head unification. The drawback of this second solution is that we must check all calls of `this/1` if we change the object name. Note that we can't use these method with the message-sending operators `((:)/2`, `((:)/1`, or `(^^)/1`).

When storing a parametric object in its own source file, the convention is to name the file after the object, with the object arity appended. For instance, when defining an object named `sort(Type)`, we may save it in a `sort_1.lgt` text file. This way it is easy to avoid file name clashes when saving Logtalk entities that have the same functor but different arity.

Parametric object proxies

Compound terms with the same functor and with the same number of arguments as a parametric object identifier may act as *proxies* to a parametric object. Proxies may be stored on the database as Prolog facts and be used to represent different instantiations of a parametric object identifier. For example:

```
:- object(circle(_Id_, _Radius_, _Color_)).  
    :- public(area/1).  
    ...  
:- end_object.  
  
% parametric object proxies:  
circle('#1', 1.23, blue).  
circle('#2', 3.71, yellow).  
circle('#3', 0.39, green).  
circle('#4', 5.74, black).  
circle('#5', 8.32, cyan).
```

Logtalk provides a convenient notation for accessing proxies represented as Prolog facts when sending a message:

```
..., {Proxy}::Message, ...
```

For example, using the `circle/3` parametric object above, we can compute a list with the areas of all circles using the following goal:

```
| ?- findall(Area, {circle(_, _, _)}::area(Area), Areas).  
  
Areas = [4.75291, 43.2412, 0.477836, 103.508, 217.468].
```

In this context, the proxy argument is proved as a plain Prolog goal. If successful, the message is sent to the corresponding parametric object. Typically, the proof allows retrieval of parameter instantiations. This construct can either be used with a proxy argument that is sufficiently instantiated in order to unify with a single Prolog fact or with a proxy argument that unifies with several facts on backtracking.

1.5.4 Finding defined objects

We can find, by backtracking, all defined objects by calling the *current_object/1* built-in predicate with an unbound argument:

```
| ?- current_object(Object).
Object = logtalk ;
Object = user ;
...
```

This predicate can also be used to test if an object is defined by calling it with a valid object identifier (an atom or a compound term).

1.5.5 Creating a new object at runtime

An object can be dynamically created at runtime by using the *create_object/4* built-in predicate:

```
| ?- create_object(Object, Relations, Directives, Clauses).
```

The first argument should be either a variable or the name of the new object (a Prolog atom or compound term, which must not match any existing entity name). The remaining three arguments correspond to the relations described in the opening object directive and to the object code contents (directives and clauses).

For example, the call:

```
| ?- create_object(
    foo,
    [extends(bar)],
    [public(foo/1)],
    [foo(1), foo(2)]
).
```

is equivalent to compiling and loading the object:

```
:- object(foo,
    extends(bar)).

:- dynamic.

:- public(foo/1).
foo(1).
foo(2).

:- end_object.
```

If we need to create a lot of (dynamic) objects at runtime, then it is best to define a metaclass or a prototype with a predicate that will call this built-in predicate to make new objects. This predicate may provide automatic object name generation, name checking, and accept object initialization options.

1.5.6 Abolishing an existing object

Dynamic objects can be abolished using the *abolish_object/1* built-in predicate:

```
| ?- abolish_object(Object).
```

The argument must be an identifier of a defined dynamic object; otherwise an error will be thrown.

1.5.7 Object directives

Object directives are used to set initialization goals, define object properties, document an object dependencies on other Logtalk entities, and load the contents of files into an object.

Object initialization

We can define a goal to be executed as soon as an object is (compiled and) loaded to memory with the *initialization/1* directive:

```
:- initialization(Goal).
```

The argument can be any valid Logtalk goal. For example, a call to a local predicate:

```
:- object(foo).

    :- initialization(init).
    :- private(init/0).

    init :-
        ... .

    ...

:- end_object.
```

Or a message to another object:

```
:- object(assembler).

    :- initialization(control::start).
    ...

:- end_object.
```

Another common initialization goal is a message to *self* in order to call an inherited or imported predicate. For example, assuming that we have a monitor category defining a *reset/0* predicate, we could write:

```
:- object(profiler,
    imports(monitor)).

    :- initialization(::reset).
    ...

:- end_object.
```

Note, however, that descendant objects do not inherit initialization directives. In this context, *self* denotes the object that contains the directive. Also note that object initialization does not necessarily mean setting an object dynamic state.

Dynamic objects

Similar to Prolog predicates, an object can be either static or dynamic. An object created during the execution of a program is always dynamic. An object defined in a file can be either dynamic or static. Dynamic objects are declared by using the *dynamic/0* directive in the object source code:

```
:- dynamic.
```

The directive must precede any predicate directives or clauses. Please be aware that using dynamic code results in a performance hit when compared to static code. We should only use dynamic objects when these need to be abolished during program execution. In addition, note that we can declare and define dynamic predicates within a static object.

Object documentation

An object can be documented with arbitrary user-defined information by using the *info/1* entity directive. See the *Documenting* section for details.

Loading files into an object

The *include/1* directive can be used to load the contents of a file into an object. A typical usage scenario is to load a plain Prolog file into an object, thus providing a simple way to encapsulate its contents. For example, assume a *cities.pl* file defining facts for a *city/4* predicate. We could define a wrapper for this database by writing:

```
:- object(cities).

    :- public(city/4).

    :- include(dbs('cities.pl')).

:- end_object.
```

The *include/1* directive can also be used when creating an object dynamically. For example:

```
| ?- create_object(cities, [], [public(city/4), include(dbs('cities.pl'))], []).
```

Declaring object aliases

The *uses/1* directive can be used to declare object aliases. The typical uses of this directive include shortening long object names, working consistently with specific parameterizations of parametric objects, and simplifying experimenting with different object implementations of the same protocol when using explicit message-sending.

1.5.8 Object relationships

Logtalk provides six sets of built-in predicates that enable us to query the system about the relationships that an object has with other entities.

The *instantiates_class/2-3* built-in predicates can be used to query all instantiation relations:

```
| ?- instantiates_class(Instance, Class).
```

or, if we also want to know the instantiation scope:

```
| ?- instantiates_class(Instance, Class, Scope).
```

Specialization relations can be found by using the *specializes_class/2-3* built-in predicates:

```
| ?- specializes_class(Class, Superclass).
```

or, if we also want to know the specialization scope:

```
| ?- specializes_class(Class, Superclass, Scope).
```

For prototypes, we can query extension relations using the *extends_object/2-3* built-in predicates:

```
| ?- extends_object(Object, Parent).
```

or, if we also want to know the extension scope:

```
| ?- extends_object(Object, Parent, Scope).
```

In order to find which objects import which categories, we can use the *imports_category/2-3* built-in predicates:

```
| ?- imports_category(Object, Category).
```

or, if we also want to know the importation scope:

```
| ?- imports_category(Object, Category, Scope).
```

To find which objects implements which protocols, we can use the *implements_protocol/2-3* and *conforms_to_protocol/2-3* built-in predicates:

```
| ?- implements_protocol(Object, Protocol, Scope).
```

or, if we also want to consider inherited protocols:

```
| ?- conforms_to_protocol(Object, Protocol, Scope).
```

Note that, if we use an unbound first argument, we will need to use the *current_object/1* built-in predicate to ensure that the entity returned is an object and not a category.

To find which objects are explicitly complemented by categories, we can use the *complements_object/2* built-in predicate:

```
| ?- complements_object(Category, Object).
```

Note that more than one category may explicitly complement a single object, and a single category can complement several objects.

1.5.9 Object properties

We can find the properties of defined objects by calling the built-in predicate *object_property/2*:

```
| ?- object_property(Object, Property).
```

The following object properties are supported:

static

The object is static

dynamic

The object is dynamic (and thus can be abolished in runtime by calling the *abolish_object/1* built-in predicate)

built_in

The object is a built-in object (and thus always available)

threaded

The object supports/makes multi-threading calls

file(Path)

Absolute path of the source file defining the object (if applicable)

file(Basename, Directory)

Basename and directory of the source file defining the object (if applicable); Directory always ends with a /

lines(BeginLine, EndLine)

Source file begin and end lines of the object definition (if applicable)

directive(BeginLine, EndLine)

Source file begin and end lines of the object opening directive (if applicable)

context_switching_calls

The object supports context-switching calls (i.e., can be used with the *(<<)/2* debugging control construct)

dynamic_declarations

The object supports dynamic declarations of predicates

events

Messages sent from the object generate events

source_data

Source data available for the object

info(List)

List of the compound term representation of the object *info/1* pairs with the key as functor.

complements(Permission)

The object supports complementing categories with the specified permission (allow or restrict)

complements

The object supports complementing categories

public(Resources)

List of public predicates and operators declared by the object

protected(Resources)

List of protected predicates and operators declared by the object

private(Resources)

List of private predicates and operators declared by the object

declares(Predicate, Properties)

List of *properties* for a predicate declared by the object

defines(Predicate, Properties)

List of *properties* for a predicate defined by the object

includes(Predicate, Entity, Properties)

List of *properties* for an object multifile predicate that are defined in the specified entity (the properties include `number_of_clauses(Number)`, `number_of_rules(Number)`, `include(File)`, `lines(Start,End)`, and `line_count(Start)` with `Start-End` being the line range of the first multifile predicate clause)

provides(Predicate, Entity, Properties)

List of *properties* for other entity multifile predicates that are defined in the object (the properties include `number_of_clauses(Number)`, `number_of_rules(Number)`, `include(File)`, `lines(Start,End)`, and `line_count(Start)` with `Start-End` being the line range of the first multifile predicate clause)

references(Reference, Properties)

List of *properties* for other references to entities in calls to the execution-context built-in methods and in directives for multifile predicates and multifile non-terminals that are found in the object (the properties include `in(Context)`, `non_terminal(NonTerminal)`, `include(File)`, `lines(Start,End)`, and `line_count(Start)` with `Start-End` being the line range of the first reference, a directive, a predicate clause, or a non-terminal grammar rule; the possible values for `Context` are multifile, dynamic, discontinuous, meta_predicate, meta_non_terminal, and clause); `Reference` can be either `Entity` or `Entity::Functor/Arity`

alias(Entity, Properties)

List of *properties* for an *entity alias* declared by the object (the properties include `object` in case of an object alias, `module` in case of a module alias, `for(Original)`, `lines(Start,End)`, and `line_count(Start)` with `Start-End` being the line range of the `uses/1` or `use_module/1` directive)

alias(Predicate, Properties)

List of *properties* for a *predicate alias* declared by the object (the properties include `predicate`, `for(Original)`, `from(Entity)`, `non_terminal(NonTerminal)`, `lines(Start,End)`, and `line_count(Start)` with `Start-End` being the line range of the alias directive)

calls(Call, Properties)

List of *properties* for predicate calls made by the object (`Call` is either a predicate indicator or a control construct such as `(:)/1-2` or `(^^)/1` with a predicate indicator as argument; note that `Call` may not be ground in case of a call to a control construct where its argument is only known at runtime; the properties include `caller(Caller)`, `alias(Alias)`, `non_terminal(NonTerminal)`, `lines(Start,End)`, `line_count(Start)` with `Caller`, `Alias`, and `NonTerminal` being predicate indicators and `Start-End` being the line range of the predicate clause or directive making the call)

updates(Predicate, Properties)

List of *properties* for dynamic predicate updates (and also access using the `clause/2` predicate) made by the object (`Predicate` is either a predicate indicator or a control construct such as `(:)/1-2` or `(:)/2` with a predicate indicator as argument; note that `Predicate` may not be ground in case of a control construct argument only known at runtime; the properties include `updater(Updater)`, `alias(Alias)`, `non_terminal(NonTerminal)`, `lines(Start,End)`, and `line_count(Start)` with `Updater` being a (possibly multifile) predicate indicator, `Alias` and `NonTerminal` being predicate indicators, and `Start-End` being the line range of the predicate clause or directive updating the predicate)

number_of_clauses(Number)

Total number of predicate clauses defined in the object at compilation time (includes both user-defined clauses and auxiliary clauses generated by the compiler or by the *expansion hooks* but does not include clauses for multifile predicates defined for other entities or clauses for the object own multifile

predicates contributed by other entities)

number_of_rules(Number)

Total number of predicate rules defined in the object at compilation time (includes both user-defined rules and auxiliary rules generated by the compiler or by the *expansion hooks* but does not include rules for multifile predicates defined for other entities or rules for the object own multifile predicates contributed by other entities)

number_of_user_clauses(Number)

Total number of user-defined predicate clauses defined in the object at compilation time (does not include clauses for multifile predicates defined for other entities or clauses for the object own multifile predicates contributed by other entities)

number_of_user_rules(Number)

Total number of user-defined predicate rules defined in the object at compilation time (does not include rules for multifile predicates defined for other entities or rules for the object own multifile predicates contributed by other entities)

debugging

The object is compiled in debug mode

module

The object resulted from the compilation of a Prolog module

When a predicate is called from an `initialization/1` directive, the argument of the `caller/1` property is `(:-)/1`.

Some properties such as line numbers are only available when the object is defined in a source file compiled with the *source_data* flag turned on. Moreover, line numbers are only supported in *backend Prolog compilers* that provide access to the start line of a read term. When such support is not available, the value `-1` is returned for the start and end lines.

The properties that return the number of clauses (rules) report the clauses (rules) *textually defined in the object* for both multifile and non-multifile predicates. Thus, these numbers exclude clauses (rules) for multifile predicates *contributed* by other entities.

1.5.10 Built-in objects

Logtalk defines some built-in objects that are always available for any application.

The built-in pseudo-object user

The built-in *user* pseudo-object virtually contains all user predicate definitions not encapsulated in a Logtalk entity (or a Prolog module for backends supporting a module system). These predicates are assumed to be implicitly declared public. Messages sent from this pseudo-object, which includes messages sent from the top-level interpreter, generate events when the default value of the *events* flag is set to allow. Defining complementing categories for this pseudo-object is not supported.

With some of the *backend Prolog compilers* that support a module system, it is possible to load (the) Logtalk (compiler/runtime) into a module other than the pseudo-module *user*. In this case, the Logtalk pseudo-object *user* virtually contains all user predicate definitions defined in the module where Logtalk was loaded.

The built-in object logtalk

The built-in `logtalk` object provides *message printing* predicates, *question asking* predicates, *debug and trace event* predicates, predicates for accessing the internal database of loaded files and their properties, halt predicates, and also a set of low-level utility predicates normally used when defining hook objects. Consult its API documentation for details.

1.5.11 Multi-threading applications

When writing multi-threading applications, user-defined predicates calling built-in predicates such as `create_object/4` and `abolish_object/1` may need to be declared synchronized in order to avoid race conditions.

1.6 Protocols

Protocols enable the separation between interface and implementation: several objects can implement the same protocol, and an object can implement several protocols. Protocols may contain only predicate declarations. In some languages the term *interface* is used with a similar meaning. Logtalk allows predicate declarations of any scope within protocols, contrary to some languages that only allow public declarations.

Logtalk defines three built-in protocols, *monitoring*, *expanding*, and *forwarding*, which are described at the end of this section.

1.6.1 Defining a new protocol

We can define a new protocol in the same way we write Prolog code: by using a text editor. Logtalk source files may contain one or more objects, categories, or protocols. If you prefer to define each entity in its own source file, it is recommended that the file be named after the protocol. By default, all Logtalk source files use the extension `.lgt` but this is optional and can be set in the adapter files. Intermediate Prolog source files (generated by the Logtalk compiler) have, by default, a `_lgt` suffix and a `.pl` extension. Again, this can be set to match the needs of a particular Prolog compiler in the corresponding adapter file. For example, we may define a protocol named `listp` and save it in a `listp.lgt` source file that will be compiled to a `listp_lgt.pl` Prolog file (depending on the backend compiler, the names of the intermediate Prolog files may include a directory hash and a process identifier to prevent file name clashes when embedding Logtalk applications or running parallel Logtalk processes).

Protocol names must be atoms. Objects, categories, and protocols share the same namespace: we cannot have a protocol with the same name as an object or a category.

Protocol directives are textually encapsulated by using two Logtalk directives: *protocol/1-2* and *end_protocol/0*. The most simple protocol will be one that is self-contained, not depending on any other Logtalk entity:

```
:- protocol(Protocol).  
    ...  
:- end_protocol.
```

If a protocol extends one or more protocols, then the opening directive will be:

```
:- protocol(Protocol,  
    extends([Protocol1, Protocol2, ...])).
```

(continues on next page)

(continued from previous page)

```
...
:- end_protocol.
```

In order to maximize protocol reuse, all predicates specified in a protocol should relate to the same functionality. Therefore, the only recommended use of protocol extension is when you need both a minimal protocol and an extended version of the same protocol with additional, convenient predicates.

1.6.2 Finding defined protocols

We can find, by backtracking, all defined protocols by using the `current_protocol/1` built-in predicate with an unbound argument:

```
| ?- current_protocol(Protocol).
```

This predicate can also be used to test if a protocol is defined by calling it with a valid protocol identifier (an atom).

1.6.3 Creating a new protocol at runtime

We can create a new (dynamic) protocol at runtime by calling the Logtalk built-in predicate `create_protocol/3`:

```
| ?- create_protocol(Protocol, Relations, Directives).
```

The first argument should be either a variable or the name of the new protocol (a Prolog atom, which must not match an existing entity name). The remaining two arguments correspond to the relations described in the opening protocol directive and to the protocol directives.

For instance, the call:

```
| ?- create_protocol(ppp, [extends(qqq)], [public([foo/1, bar/1]))].
```

is equivalent to compiling and loading the protocol:

```
:- protocol(ppp,
    extends(qqq)).

:- dynamic.

:- public([foo/1, bar/1]).

:- end_protocol.
```

If we need to create a lot of (dynamic) protocols at runtime, then it is best to define a metaclass or a prototype with a predicate that will call this built-in predicate in order to provide more sophisticated behavior.

1.6.4 Abolishing an existing protocol

Dynamic protocols can be abolished using the *abolish_protocol/1* built-in predicate:

```
| ?- abolish_protocol(Protocol).
```

The argument must be an identifier of a defined dynamic protocol; otherwise an error will be thrown.

1.6.5 Protocol directives

Protocol directives are used to define protocol properties and documentation.

Dynamic protocols

As usually happens with Prolog code, a protocol can be either static or dynamic. A protocol created during the execution of a program is always dynamic. A protocol defined in a file can be either dynamic or static. Dynamic protocols are declared by using the *dynamic/0* directive in the protocol source code:

```
:- dynamic.
```

The directive must precede any predicate directives. Please be aware that using dynamic code results in a performance hit when compared to static code. We should only use dynamic protocols when these need to be abolished during program execution.

Protocol documentation

A protocol can be documented with arbitrary user-defined information by using the *info/1* entity directive. See the *Documenting* section for details.

Loading files into a protocol

The *include/1* directive can be used to load the contents of a file into a protocol. See the *Objects* section for an example of using this directive.

1.6.6 Protocol relationships

Logtalk provides two sets of built-in predicates that enable us to query the system about the relationships that a protocol has with other entities.

The *extends_protocol/2-3* built-in predicates return all pairs of protocols so that the first one extends the second:

```
| ?- extends_protocol(Protocol1, Protocol2).
```

or, if we also want to know the extension scope:

```
| ?- extends_protocol(Protocol1, Protocol2, Scope).
```

To find which objects or categories implement which protocols, we can call the *implements_protocol/2-3* built-in predicates:

```
| ?- implements_protocol(ObjectOrCategory, Protocol).
```

or, if we also want to know the implementation scope:

```
| ?- implements_protocol(ObjectOrCategory, Protocol, Scope).
```

Note that, if we use a non-instantiated variable for the first argument, we will need to use the *current_object/1* or *current_category/1* built-in predicates to identify the kind of entity returned.

1.6.7 Protocol properties

We can find the properties of defined protocols by calling the *protocol_property/2* built-in predicate:

```
| ?- protocol_property(Protocol, Property).
```

A protocol may have the property `static`, `dynamic`, or `built_in`. Dynamic protocols can be abolished in runtime by calling the *abolish_protocol/1* built-in predicate. Depending on the *backend Prolog compiler*, a protocol may have additional properties related to the source file where it is defined.

The following protocol properties are supported:

static

The protocol is static

dynamic

The protocol is dynamic (and thus can be abolished in runtime by calling the *abolish_category/1* built-in predicate)

built_in

The protocol is a built-in protocol (and thus always available)

source_data

Source data available for the protocol

info(List)

List of the compound term representation of the object *info/1* pairs with the key as functor.

file(Path)

Absolute path of the source file defining the protocol (if applicable)

file(Basename, Directory)

Basename and directory of the source file defining the protocol (if applicable); Directory always ends with a /

lines(BeginLine, EndLine)

Source file begin and end lines of the protocol definition (if applicable)

directive(BeginLine, EndLine)

Source file begin and end lines of the protocol opening directive (if applicable)

public(Resources)

List of public predicates and operators declared by the protocol

protected(Resources)

List of protected predicates and operators declared by the protocol

private(Resources)

List of private predicates and operators declared by the protocol

declares(Predicate, Properties)

List of *properties* for a predicate declared by the protocol

alias(Predicate, Properties)

List of *properties* for a *predicate alias* declared by the protocol (the properties include `for(Original)`, `from(Entity)`, `non_terminal(NonTerminal)`, and `line_count(Line)` with `Line` being the begin line of the alias directive)

Some of the properties, such as line numbers, are only available when the protocol is defined in a source file compiled with the *source_data* flag turned on.

1.6.8 Implementing protocols

Any number of objects or categories can implement a protocol. The syntax is very simple:

```
:- object(Object,
    implements(Protocol)).
    ...
:- end_object.
```

or, in the case of a category:

```
:- category(Object,
    implements(Protocol)).
    ...
:- end_category.
```

To make all public predicates declared via an implemented protocol protected or to make all public and protected predicates private we prefix the protocol's name with the corresponding keyword. For instance:

```
:- object(Object,
    implements(private::Protocol)).
    ...
:- end_object.
```

or:

```
:- object(Object,
    implements(protected::Protocol)).
    ...
:- end_object.
```

Omitting the scope keyword is equivalent to writing:

```
:- object(Object,
    implements(public::Protocol)).
    ...
:- end_object.
```

The same rules apply to protocols implemented by categories.

1.6.9 Built-in protocols

Logtalk defines a set of built-in protocols that are always available for any application.

The built-in protocol `expanding`

The built-in `expanding` protocol declares the *term_expansion/2* and *goal_expansion/2* predicates. See the description of the *hook* compiler flag for more details.

The built-in protocol `monitoring`

The built-in `monitoring` protocol declares the *before/3* and *after/3* public event handler predicates. See the *Event-driven programming* section for more details.

The built-in protocol `forwarding`

The built-in `forwarding` protocol declares the *forward/1* user-defined message forwarding handler, which is automatically called (if defined) by the runtime for any message that the receiving object does not understand. See also the `[]/1` control construct.

1.6.10 Multi-threading applications

When writing multi-threading applications, user-defined predicates calling built-in predicates such as `create_protocol/3` and `abolish_protocol/1` may need to be declared synchronized in order to avoid race conditions.

1.7 Categories

Categories are *fine-grained units of code reuse* and can be regarded as a dual concept of protocols. Categories provide a way to encapsulate a set of related predicate declarations and definitions that do not represent a complete object and that only make sense when composed with other predicates. Categories may also be used to break a complex object into functional units. A category can be imported by several objects (without code duplication), including objects participating in prototype or class-based hierarchies. This concept of categories shares some ideas with Smalltalk-80 functional categories [Goldberg83], Flavors mix-ins [Moon86] (without necessarily implying multi-inheritance), and Objective-C categories [Cox86]. Categories may also *complement* existing objects, thus providing a *hot patching* mechanism inspired by the Objective-C categories functionality.

Logtalk defines a built-in category, `core_messages`, which is described at the end of this section.

1.7.1 Defining a new category

We can define a new category in the same way we write Prolog code: by using a text editor. Logtalk source files may contain one or more objects, categories, or protocols. If you prefer to define each entity in its own source file, it is recommended that the file be named after the category. By default, all Logtalk source files use the extension `.lgt` but this is optional and can be set in the adapter files. Intermediate Prolog source files (generated by the Logtalk compiler) have, by default, a `_lgt` suffix and a `.pl` extension. Again, this can be set to match the needs of a particular Prolog compiler in the corresponding adapter file. For example, we may define a category named `documenting` and save it in a `documenting.lgt` source file that will be compiled to a `documenting_lgt.pl` Prolog file (depending on the *backend compiler*, the names of the intermediate Prolog files may include a directory hash and a process identifier to prevent file name clashes when embedding Logtalk applications or running parallel Logtalk processes).

Category names can be atoms or compound terms (when defining parametric categories). Objects, categories, and protocols share the same name space: we cannot have a category with the same name as an object or a protocol.

Category code (directives and predicates) is textually encapsulated by using two Logtalk directives: *category/1-4* and *end_category/0*. The most simple category will be one that is self-contained, not depending on any other Logtalk entity:

```
:- category(Category).
...
:- end_category.
```

If a category implements one or more protocols, then the opening directive will be:

```
:- category(Category,
    implements([Protocol1, Protocol2, ...])).
...
:- end_category.
```

A category may be defined as a composition of other categories by writing:

```
:- category(Category,
    extends([Category1, Category2, ...])).
...
:- end_category.
```

This feature should only be used when extending a category without breaking its functional cohesion (for example, when a modified version of a category is needed for importing on several unrelated objects). The preferred way of composing several categories is by importing them into an object. When a category overrides a predicate defined in an extended category, the overridden definition can still be called by using the `(^ ^)/1` control construct.

Categories cannot inherit from objects. In addition, categories cannot define clauses for dynamic predicates. This restriction applies because a category can be imported by several objects and because we cannot use the database handling built-in methods with categories (messages can only be sent to objects). A consequence of this restriction is that a category cannot declare a predicate (or non-terminal) as both multifile and dynamic. However, categories may contain declarations for dynamic predicates, and they can contain predicates that handle dynamic predicates. For example:

```
:- category(attributes).

:- public(attribute/2).
```

(continues on next page)

(continued from previous page)

```

:- public(set_attribute/2).
:- public(del_attribute/2).

:- private(attribute_/2).
:- dynamic(attribute_/2).

attribute(Attribute, Value) :-
    % called in the context of "self"
    ::attribute_(Attribute, Value).

set_attribute(Attribute, Value) :-
    % retract old clauses in "self"
    ::retractall(attribute_(Attribute, _)),
    % assert new clause in "self"
    ::assertz(attribute_(Attribute, Value)).

del_attribute(Attribute, Value) :-
    % retract clause in "self"
    ::retract(attribute_(Attribute, Value)).

:- end_category.

```

Each object importing this category will have its own `attribute_/2` private and dynamic predicate. The predicates `attribute/2`, `set_attribute/2`, and `del_attribute/2` always access and modify the dynamic predicate contained in the object receiving the corresponding messages (i.e., *self*). But it's also possible to define predicates that handle dynamic predicates in the context of *this* instead of *self*. For example:

```

:- category(attributes).

:- public(attribute/2).
:- public(set_attribute/2).
:- public(del_attribute/2).

:- private(attribute_/2).
:- dynamic(attribute_/2).

attribute(Attribute, Value) :-
    % call in the context of "this"
    attribute_(Attribute, Value).

set_attribute(Attribute, Value) :-
    % retract old clauses in "this"
    retractall(attribute_(Attribute, _)),
    % asserts clause in "this"
    assertz(attribute_(Attribute, Value)).

del_attribute(Attribute, Value) :-
    % retract clause in "this"
    retract(attribute_(Attribute, Value)).

:- end_category.

```

When defining a category that declares and handles dynamic predicates, working in the context of *this* ties

those dynamic predicates to the object importing the category, while working in the context of *self* allows each object inheriting from the object that imports the category to have its own set of clauses for those dynamic predicates.

1.7.2 Hot patching

A category may also explicitly complement one or more existing objects, thus providing *hot patching* functionality inspired by Objective-C categories:

```
:- category(Category,
    complements([Object1, Object2, ....])).
...
:- end_category.
```

This allows us to add missing directives (e.g., to define *aliases* for complemented object predicates), replace broken predicate definitions, add new predicates, and add protocols and categories to existing objects without requiring access or modifications to their source code. Common scenarios are adding logging or debugging predicates to a set of objects. Complemented objects need to be compiled with the *complements* compiler flag set *allow* (to allow both patching and adding functionality) or *restrict* (to allow only adding new functionality). A complementing category takes preference over a previously loaded complementing category for the same object, thus allowing patching a previous patch if necessary.

When replacing a predicate definition, it is possible to call the overridden definition in the object from the new definition in the category by using the *(@)/1* control construct. This construct is only meaningful when used within categories and requires a compile-time bound goal argument, which is called in *this* (i.e., in the context of the complemented object or the object importing a category). As an example, consider the following object:

```
:- object(bird).

    :- set_logtalk_flag(complements, allow).

    :- public(make_sound/0).
    make_sound :-
        write('Chirp, chirp!'), nl.

:- end_object.
```

We can use the *(@)/1* control construct to wrap the original *make_sound/0* predicate definition by writing:

```
:- category(logging,
    complements(bird)).

    make_sound :-
        write('Started making sound...'), nl,
        @make_sound,
        write('... finished making sound.'), nl.

:- end_category.
```

After loading the object and the category, calling the *make_sound/0* predicate will result in the following output:

```
| ?- bird::make_sound.

Started making sound...
Chirp, chirp!
... finished making sound.
yes
```

Note that *super calls* from predicates defined in complementing categories lookup inherited definitions as if the calls were made from the complemented object instead of the category ancestors. This allows more comprehensive object patching. But it also means that, if you want to patch an object so that it imports a category that extends another category and uses super calls to access the extended category predicates, you will need to define a (possibly empty) complementing category that extends the category that you want to add.

An unfortunate consequence of allowing an object to be patched at runtime using a complementing category is that it disables the use of *static binding* optimizations for messages sent to the complemented object, as it can always be later patched, thus rendering the static binding optimizations invalid.

Another important caveat is that, while a complementing category can replace a predicate definition, local callers of the replaced predicate will still call the non-patched version of the predicate. This is a consequence of the lack of a portable solution at the *backend Prolog compiler* level for replacing static predicate definitions.

1.7.3 Finding defined categories

We can find, by backtracking, all defined categories by using the *current_category/1* built-in predicate with an unbound argument:

```
| ?- current_category(Category).
```

This predicate can also be used to test if a category is defined by calling it with a valid category identifier (an atom or a compound term).

1.7.4 Creating a new category at runtime

A category can be dynamically created at runtime by using the *create_category/4* built-in predicate:

```
| ?- create_category(Category, Relations, Directives, Clauses).
```

The first argument should be either a variable or the name of the new category (a Prolog atom, which must not match with an existing entity name). The remaining three arguments correspond to the relations described in the opening category directive and to the category code contents (directives and clauses).

For example, the call:

```
| ?- create_category(
    ccc,
    [implements(ppp)],
    [private(bar/1)],
    [(foo(X):-bar(X)), bar(1), bar(2)]
).
```

is equivalent to compiling and loading the category:

```
:- category(ccc,
    implements(ppp)).

:- dynamic.

:- private(bar/1).

foo(X) :-
    bar(X).

bar(1).
bar(2).

:- end_category.
```

If we need to create a lot of (dynamic) categories at runtime, then it is best to define a metaclass or a prototype with a predicate that will call this built-in predicate in order to provide more sophisticated behavior.

1.7.5 Abolishing an existing category

Dynamic categories can be abolished using the `abolish_category/1` built-in predicate:

```
| ?- abolish_category(Category).
```

The argument must be an identifier of a defined dynamic category; otherwise, an error will be thrown.

1.7.6 Category directives

Category directives are used to define category properties, document category dependencies on other Logtalk entities, and load the contents of files into a category.

Dynamic categories

As usually happens with Prolog code, a category can be either static or dynamic. A category created during the execution of a program is always dynamic. A category defined in a file can be either dynamic or static. Dynamic categories are declared by using the `dynamic/0` directive in the category source code:

```
:- dynamic.
```

The directive must precede any predicate directives or clauses. Please be aware that using dynamic code results in a performance hit when compared to static code. We should only use dynamic categories when these need to be abolished during program execution.

Category documentation

A category can be documented with arbitrary user-defined information by using the *info/1* entity directive. See the *Documenting* section for details.

Loading files into a category

The *include/1* directive can be used to load the contents of a file into a category. See the *Objects* section for an example of using this directive.

Declaring object aliases

The *uses/1* directive can be used to declare object aliases. The typical uses of this directive are to shorten long object names and to simplify experimenting with different object implementations of the same protocol when using explicit message-sending.

1.7.7 Category relationships

Logtalk provides two sets of built-in predicates that enable us to query the system about the relationships that a category has with other entities.

The built-in predicates *implements_protocol/2-3* and *conforms_to_protocol/2-3* allow us to find which categories implement which protocols:

```
| ?- implements_protocol(Category, Protocol, Scope).
```

or, if we also want to consider inherited protocols:

```
| ?- conforms_to_protocol(Category, Protocol, Scope).
```

Note that, if we use an unbound first argument, we will need to use the *current_category/1* built-in predicate to ensure that the returned entity is a category and not an object.

To find which objects import which categories, we can use the *imports_category/2-3* built-in predicates:

```
| ?- imports_category(Object, Category).
```

or, if we also want to know the importation scope:

```
| ?- imports_category(Object, Category, Scope).
```

Note that a category may be imported by several objects.

To find which categories extend other categories, we can use the *extends_category/2-3* built-in predicates:

```
| ?- extends_category(Category1, Category2).
```

or, if we also want to know the extension scope:

```
| ?- extends_category(Category1, Category2, Scope).
```

Note that a category may be extended by several categories.

To find which categories explicitly complement existing objects we can use the *complements_object/2* built-in predicate:

```
| ?- complements_object(Category, Object).
```

Note that a category may explicitly complement several objects.

1.7.8 Category properties

We can find the properties of defined categories by calling the built-in predicate *category_property/2*:

```
| ?- category_property(Category, Property).
```

The following category properties are supported:

static

The category is static

dynamic

The category is dynamic (and thus can be abolished in runtime by calling the *abolish_category/1* built-in predicate)

built_in

The category is a built-in category (and thus always available)

file(Path)

Absolute path of the source file defining the category (if applicable)

file(BaseName, Directory)

BaseName and directory of the source file defining the category (if applicable); Directory always ends with a /

lines(BeginLine, EndLine)

Source file begin and end lines of the category definition (if applicable)

directive(BeginLine, EndLine)

Source file begin and end lines of the category opening directive (if applicable)

events

Messages sent from the category generate events

source_data

Source data available for the category

info(List)

List of the compound term representation of the object *info/1* pairs with the key as functor.

public(Resources)

List of public predicates and operators declared by the category

protected(Resources)

List of protected predicates and operators declared by the category

private(Resources)

List of private predicates and operators declared by the category

declares(Predicate, Properties)

List of *properties* for a predicate declared by the category

defines(Predicate, Properties)

List of *properties* for a predicate defined by the category

includes(Predicate, Entity, Properties)

List of *properties* for an object multifile predicate that are defined in the specified entity (the properties include `number_of_clauses(Number)`, `number_of_rules(Number)`, `lines(Start,End)`, and `line_count(Start)` with Start-End being the line range of the first multifile predicate clause)

provides(Predicate, Entity, Properties)

List of *properties* for other entity multifile predicate that are defined in the category (the properties include `number_of_clauses(Number)`, `number_of_rules(Number)`, `lines(Start,End)`, and `line_count(Start)` with Start-End being the line range of the first multifile predicate clause)

references(Reference, Properties)

List of *properties* for other references to entities in calls to the execution-context built-in methods and in directives for multifile predicates and multifile non-terminals that are found in the category (the properties include `in(Context)`, `non_terminal(NonTerminal)`, `include(File)`, `lines(Start,End)`, and `line_count(Start)` with Start-End being the line range of the first reference, a directive, a predicate clause, or a non-terminal grammar rule; the possible values for Context are multifile, dynamic, discontinuous, meta_predicate, meta_non_terminal, and clause); Reference can be either Entity or `Entity::Functor/Arity`

alias(Entity, Properties)

List of *properties* for an *entity alias* declared by the object (the properties include `object` in case of an object alias, `module` in case of a module alias, `for(Original)`, `lines(Start,End)`, and `line_count(Start)` with Start-End being the line range of the `uses/1` or `use_module/1` directive)

alias(Predicate, Properties)

List of *properties* for a *predicate alias* declared by the category (the properties include `predicate`, `for(Original)`, `from(Entity)`, `non_terminal(NonTerminal)`, `lines(Start,End)`, and `line_count(Start)` with Start-End being the line range of the alias directive)

calls(Call, Properties)

List of *properties* for predicate calls made by the category (Call is either a predicate indicator or a control construct such as `(:)/1-2` or `(^)/1` with a predicate indicator as argument; note that Call may not be ground in case of a call to a control construct where its argument is only known at runtime; the properties include `caller(Caller)`, `alias(Alias)`, `non_terminal(NonTerminal)`, `lines(Start,End)`, and `line_count(Start)` with Caller, Alias, and NonTerminal being predicate indicators and Start-End being the line range of the predicate clause or directive making the call)

updates(Predicate, Properties)

List of *properties* for dynamic predicate updates (and also access using the `clause/2` predicate) made by the object (Predicate is either a predicate indicator or a control construct such as `(:)/1-2` or `(:)/2` with a predicate indicator as argument; note that Predicate may not be ground in case of a control construct argument only known at runtime; the properties include `updater(Updater)`, `alias(Alias)`, `non_terminal(NonTerminal)`, `lines(Start,End)`, and `line_count(Start)` with Updater being a (possibly multifile) predicate indicator, Alias and NonTerminal being predicate indicators, and Start-End being the line range of the predicate clause or directive updating the predicate)

number_of_clauses(Number)

Total number of predicate clauses defined in the category (includes both user-defined clauses and auxiliary clauses generated by the compiler or by the *expansion hooks* but does not include clauses for multifile predicates defined for other entities or clauses for the category own multifile predicates contributed by other entities)

number_of_rules(Number)

Total number of predicate rules defined in the category (includes both user-defined rules and auxiliary rules generated by the compiler or by the *expansion hooks* but does not include rules for multifile predicates defined for other entities or rules for the category own multifile predicates contributed by other entities)

number_of_user_clauses(Number)

Total number of user-defined predicate clauses defined in the category (does not include clauses for multifile predicates defined for other entities or clauses for the category own multifile predicates contributed by other entities)

number_of_user_rules(Number)

Total number of user-defined predicate rules defined in the category (does not include rules for multifile predicates defined for other entities or rules for the category own multifile predicates contributed by other entities)

Some properties, such as line numbers, are only available when the category is defined in a source file compiled with the *source_data* flag turned on. Moreover, line numbers are only supported in *backend Prolog compilers* that provide access to the start line of a read term. When such support is not available, the value -1 is returned for the start and end lines.

The properties that return the number of clauses (rules) report the clauses (rules) *textually defined in the object* for both multifile and non-multifile predicates. Thus, these numbers exclude clauses (rules) for multifile predicates *contributed* by other entities.

1.7.9 Importing categories

Any number of objects can import a category. In addition, an object may import any number of categories. The syntax is very simple:

```
:- object(Object,
    imports([Category1, Category2, ...])).
...
:- end_object.
```

To make all public predicates imported via a category protected, or to make all public and protected predicates private, we prefix the category's name with the corresponding keyword:

```
:- object(Object,
    imports(private::Category)).
...
:- end_object.
```

or:

```
:- object(Object,
    imports(protected::Category)).
...
:- end_object.
```

Omitting the scope keyword is equivalent to writing:

```
:- object(Object,
    imports(public::Category)).
...
:- end_object.
```

1.7.10 Calling category predicates

Category predicates can be called from within an object by sending a message to *self* or using a *super* call. Consider the following category:

```
:- category(output).

   :- public(out/1).

   out(X) :-
       write(X), nl.

:- end_category.
```

The predicate `out/1` can be called from within an object importing the category by simply sending a message to *self*. For example:

```
:- object(worker,
   imports(output)).

   ...
   do(Task) :-
       execute(Task, Result),
       ::out(Result).
   ...

:- end_object.
```

This is the recommended way of calling a category predicate that can be specialized/overridden in a descendant object, as the predicate definition lookup will start from *self*.

A direct call to a predicate definition found in an imported category can be made using the `(^^)/1` control construct. For example:

```
:- object(worker,
   imports(output)).

   ...
   do(Task) :-
       execute(Task, Result),
       ^^out(Result).
   ...

:- end_object.
```

This alternative should only be used when the user knows a priori that the category predicates will not be specialized or redefined by descendant objects of the object importing the category. Its advantage is that, when the *optimize* flag is turned on, the Logtalk compiler will try to optimize the calls by using *static binding*. When *dynamic binding* is used due to e.g. the lack of sufficient information at compilation time, the performance is similar to calling the category predicate using a message to *self* (in both cases, a predicate lookup caching mechanism is used).

1.7.11 Parametric categories

Category predicates can be parameterized in the same way as object predicates by using a compound term as the category identifier where all the arguments of the compound term are variables. These variables, the *category parameters*, are *logical variables* and can be accessed by calling the `parameter/2` or `this/1` built-in local methods in the category predicate clauses or by using *parameter variables*.

As a consequence of parameters being logical variables, when defining, for example, a parametric category named `foo/1`, `foo(bar)` and `foo(baz)` are different parameterizations of the same category, not distinct categories.

Category parameter values can be defined by the importing objects. For example:

```
:- object(speech(Season, Event),
    imports([dress(Season), speech(Event)])),
    ...
:- end_object.
```

Note that access to category parameters is only possible from within the category. In particular, calls to the `this/1` built-in local method from category predicates always access the importing object identifier (and thus object parameters, not category parameters).

1.7.12 Built-in categories

Logtalk defines a built-in category that is always available for any application.

The built-in category `core_messages`

The built-in `core_messages` category provides default translations for all compiler and runtime printed messages, such as warnings and errors. It does not define any public predicates.

1.7.13 Multi-threading applications

When writing multi-threading applications, user-defined predicates calling built-in predicates such as `create_category/4` and `abolish_category/1` may need to be declared synchronized in order to avoid race conditions.

1.8 Predicates

Predicate directives and clauses can be encapsulated inside objects and categories. Protocols can only contain predicate directives. From the point of view of a traditional imperative object-oriented language, predicates allow both object state and object behavior to be represented. Mutable object state can be represented using dynamic object predicates but should only be used when strictly necessary, as it breaks declarative semantics.

1.8.1 Reserved predicate names

For practical and performance reasons, some predicate names have a fixed interpretation. These predicates are declared in the built-in protocols. They are: *goal_expansion/2* and *term_expansion/2*, declared in the *expanding* protocol; *before/3* and *after/3*, declared in the *monitoring* protocol; and *forward/1*, declared in the *forwarding* protocol. By default, the compiler prints a warning when a definition for one of these predicates is found but the reference to the corresponding built-in protocol is missing.

1.8.2 Declaring predicates

Logtalk provides a clear distinction between *declaring a predicate* and *defining a predicate* and thus clear *closed-world assumption* semantics. Messages or calls for declared but undefined predicates fail. Messages or calls for unknown (i.e., non-declared) predicates throw an error. Note that this is a fundamental requirement for supporting *protocols*: we must be able to declare a predicate without necessarily defining it.

All object (or category) predicates that we want to access from other objects (or categories) must be explicitly declared. A predicate declaration must contain, at least, a *scope* directive. Other directives may be used to document the predicate or to ensure proper compilation of the predicate clauses.

Scope directives

A predicate scope directive specifies *from where* the predicate can be called, i.e. its *visibility*. Predicates can be *public*, *protected*, *private*, or *local*. Public predicates can be called from any object. Protected predicates can only be called from the container object or from a container descendant. Private predicates can only be called from the container object. Predicates are local when they are not declared in a scope directive. Local predicates, like private predicates, can only be called from the container object (or category), but they are *invisible* to the reflection built-in methods (*current_predicate/1* and *predicate_property/2*) and to the message error handling mechanisms (i.e., sending a message corresponding to a local predicate results in a *predicate_declaration* existence error instead of a scope error).

The scope declarations are made using the directives *public/1*, *protected/1*, and *private/1*. For example:

```
:- public(init/1).

:- protected(valid_init_option/1).

:- private(process_init_options/1).
```

If a predicate does not have a (local or inherited) scope declaration, it is assumed that the predicate is local. Note that we do not need to write scope declarations for all defined predicates. One exception is local dynamic predicates: declaring them as private predicates may allow the Logtalk compiler to generate optimized code for asserting and retracting clauses.

Note that a predicate scope directive doesn't specify *where* a predicate is, or can be, defined. For example, a private predicate can only be called from an object holding its scope directive. But it can be defined in descendant objects. A typical example is an object playing the role of a class defining a private (possibly dynamic) predicate for its descendant instances. Only the class can call (and possibly assert/retract clauses for) the predicate, but its clauses can be found/defined in the instances themselves.

Scope directives may also be used to declare grammar rule non-terminals and operators. For example:

```
:- public(url//1).

:- public(op(800, fx, tag)).
```

Note that, in the case of operators, the operator definitions don't become global when the entity containing the directives is compiled and loaded. This prevents an application from breaking when, for example, an updated third-party library adds new operators. It also allows loading entities that provide conflicting operator definitions. Here the usual programming idiom is to copy the operator definitions to a `uses/2` directive. For example, the *lgtunit* tool makes available a `(=~/2)` predicate (for approximate float equality) that is intended to be used as an infix operator:

```
:- uses(lgtunit, [
    op(700, xfx, =~/2), (=~/2)/2
]).
```

Thus, in practice, the solution to use library entity operators in client entities is the same for using library entity predicates with implicit message-sending.

Mode directive

Often predicates can only be called using specific argument patterns. The valid arguments and instantiation modes of those arguments can be documented using the *mode/2* directive (in the case of *non-terminals*, there's also a *mode_non_terminal/2* directive). For example:

```
:- mode(member(?term, ?list), zero_or_more).
```

The first directive argument describes a valid *calling mode*. The minimum information will be the instantiation mode of each argument. The first four possible values are described in the ISO Prolog Core standard [ISO95] (but with the meaning of the `-` instantiation mode redefined here). The remaining two can also be found in use in some Prolog systems.

- `+`
Argument must be instantiated (but not necessarily ground).
- `-`
Argument should be a free (non-instantiated) variable. When bound, the call will unify the computed term with the given argument.
- `?`
Argument can either be instantiated or free.
- `@`
Argument will not be further instantiated (modified).
- `++`
Argument must be ground.
- `--`
Argument must be unbound. Used mainly when returning an opaque term (e.g., a stream handle).

Note that the `+` and `@` instantiation modes have the same meaning for atomic arguments. E.g. you can write either `+atom` or `@atom` but the first alternative is preferred.

The ISO `-` instantiation mode is equivalent to the Logtalk `--` mode, allowing the use `-` to document predicates with output arguments that don't require those arguments to be unbound at call time and also accept bound arguments without throwing an exception.

These six mode atoms are also declared as prefix operators by the Logtalk compiler. This makes it possible to include type information for each argument as in the example above. Some possible type values are: `event`, `object`, `category`, `protocol`, `callable`, `term`, `nonvar`, `var`, `atomic`, `atom`, `number`, `integer`, `float`, `compound`, and `list`. The first four are Logtalk specific. The remaining are common Prolog types. We can also use our

own types that can be either atoms or ground compound terms. See the [types](#) library documentation for an extensive list of pre-defined types that cover most common use cases.

The second directive argument documents the *number of proofs*, but not necessarily distinct solutions, for the specified mode. As an example, the `member(X, [1,1,1,1])` goal has only one distinct solution but four proofs for that solution. Note that different modes for the same predicate often have different determinism. The possible values are:

zero

Predicate always fails (e.g., the `false/0` standard predicate).

one

Predicate always succeeds once (e.g., the `flush_output/0` standard predicate).

zero_or_one

Predicate either fails or succeeds (e.g., the `atom/1` standard predicate).

zero_or_more

Predicate has zero or more proofs (e.g., the `current_predicate/1` standard predicate).

one_or_more

Predicate has one or more proofs (e.g., the `repeat/0` standard predicate).

zero_or_error

Predicate either fails or throws an error.

one_or_error

Predicate either succeeds once or throws an error (e.g., the `open/3` standard predicate).

zero_or_one_or_error

Predicate succeeds once or fails or throws an error (e.g., the `get_char/1` standard predicate).

zero_or_more_or_error

Predicate may fail or succeed multiple times or throw an error (e.g., the `retract/1` standard predicate).

one_or_more_or_error

Predicate may succeed one or more times or throw an error.

error

Predicate will throw an error (e.g., the `type_error/2` built-in method).

The last six values support documenting that some call modes may throw an error or will throw an error **despite the call arguments complying with the expected types and instantiation modes**. As an example, consider the `open/3` standard predicate:

```
:- mode(open(@source_sink, @io_mode, --stream), one_or_error).
```

In this case, the mode directive tells the user that a valid call can still throw an error (there may be e.g. a permission error opening the specified source or sink).

Notice that using the `zero`, `one`, `zero_or_one`, `zero_or_more`, or `one_or_more` modes is not only for predicates that never throw an exception; they can also be used for any predicate that doesn't throw an exception when the arguments are valid. For example, the `current_predicate/1` standard predicate throws an exception if the argument is neither a variable nor a predicate indicator, but it succeeds zero or more times when its argument is valid:

```
:- mode(current_predicate(?predicate_indicator), zero_or_more).
```

Some predicates have more than one valid mode, thus implying several mode directives. For example, to document the possible use modes of the standard `atom_concat/3` predicate, we would write:

```
:- mode(atom_concat(?atom, ?atom, +atom), one_or_more).
:- mode(atom_concat(+atom, +atom, -atom), one).
```

The first `mode/2` directive specifies that the `atom_concat/3` predicate can be used to split an atom into a prefix and a suffix. The second `mode/2` directive specifies that concatenating two atoms results in a new atom. There are often several alternative `mode/2` directives that can be used to specify a predicate. For example, an alternative to the second `mode/2` directive above would be:

```
:- mode(atom_concat(+atom, +atom, ?atom), zero_or_one).
```

In this case, the same information is provided by both alternatives. But the first alternative is simpler and thus preferred.

Some old Prolog compilers supported some sort of mode directives to improve performance. To the best of my knowledge, there is no modern Prolog compiler supporting this kind of directive for that purpose. The current Logtalk version simply parses this directive for collecting its information for use in the [reflection API](#) (assuming the [source_data](#) flag is turned on). In any case, the use of mode directives is a good starting point for documenting your predicates.

Meta-predicate directive

Some predicates may have arguments that will be called as goals, interpreted as [closures](#) that will be used for constructing goals, or passing meta-arguments to calls to other meta-predicates. To ensure that these goals will be executed in the correct context (i.e., in the [calling context](#), not in the meta-predicate [definition context](#)), we need to use the [meta_predicate/1](#) directive (in the case of *meta non-terminals*, there's also a [meta_non_terminal/1](#) directive). For example:

```
:- meta_predicate(findall(*, 0, *)).
:- meta_predicate(map(2, *, *)).
```

The meta-predicate mode arguments in this directive have the following meaning:

0

Meta-argument that will be called as a goal.

N

Meta-argument that will be a closure used to construct a call by extending it with N arguments. The value of N must be a positive integer.

::

Argument that is context-aware but that will not be called as a goal or a closure. It can contain, however, sub-terms that will be called as goals or closures.

^

Goal that may be existentially quantified (`Vars^Goal`).

Normal argument.

The following meta-predicate mode arguments are for use only when writing backend Prolog [adapter files](#) to deal with proprietary built-in meta-predicates and meta-directives:

/

Predicate indicator (`Name/Arity`), list of predicate indicators, or conjunction of predicate indicators.

//

Non-terminal indicator (`Name//Arity`), list of predicate indicators, or conjunction of predicate indicators.

- [0] List of goals.
- [N] List of closures.
- [/] List of predicate indicators.
- [//] List of non-terminal indicators.

To the best of my knowledge, the use of non-negative integers to specify closures was first introduced on Quintus Prolog for providing information for predicate cross-reference tools.

Note that Logtalk meta-predicate semantics are different from Prolog meta-predicate semantics (assuming a predicate-based module system as common):

- Meta-arguments are always called in the meta-predicate calling context, independent of using explicit or implicit message-sending (to the object defining the meta-predicate when not local). Most Prolog systems have different semantics for explicit versus implicit module qualification.
- Logtalk is not based on a predicate prefixing mechanism. Therefore, the meta-predicate directive is required for any predicate with meta-arguments (including when simply passing the meta-arguments to a call to another meta-predicate). This is usually not required in Prolog systems due to the module prefixing of meta-arguments.
- Sending a message from a meta-predicate definition to call a meta-predicate defined in another object resets the calling context for any passed meta-argument to the object sending the message (including for messages to *self*). Meta-arguments behave differently in Prolog systems due to their module prefixing.
- Logtalk protects from common scenarios where specially crafted meta-predicate definitions are used to break object (and category) encapsulation by changing the meta-arguments passed by client code or trying to subvert the implicit calling context to call client predicates other than the predicates passed as meta-arguments.

Warning

As each Logtalk entity is independently compiled, this directive must be included in every object or category that contains a definition for the described meta-predicate, even if the meta-predicate declaration is inherited from another entity, to ensure proper compilation of meta-arguments.

Discontiguous directive

The clause of an object (or category) predicate may not be contiguous. In that case, we must declare the predicate discontiguous by using the *discontiguous/1* directive:

```
:- discontiguous(foo/1).
```

This is a directive that we should avoid using: it makes your code harder to read, and it is not supported by some Prolog backends.

Warning

As each Logtalk entity is compiled independently of other entities, this directive must be included in every object or category that contains a definition for the described predicate (even if the predicate declaration is inherited from another entity).

Dynamic directive

An object predicate can be static or dynamic. By default, all predicates (and non-terminals) of static objects defined in source files are static. To declare a dynamic predicate (or non-terminal), we use the *dynamic/1* directive. For example:

```
:- dynamic(foo/1).
```

Predicates of objects dynamically created at runtime (using the *create_object/4* built-in predicate) and predicates of dynamic objects defined in source files (using the *dynamic/0* directive) are implicitly dynamic.

Dynamic predicates can be used to represent persistent mutable object state. Note that static objects may declare and define dynamic predicates. Categories can only declare dynamic predicates (with the importing objects holding the predicate definitions).

Warning

As each Logtalk entity is compiled independently from other entities, this directive must be included in every object that contains a definition for the described predicate (even if the predicate declaration is inherited from another object or imported from a category). If we omit the dynamic declaration then the predicate definition will be compiled static.

Operator directive

An object (or category) predicate can be declared as an operator using the familiar *op/3* directive:

```
:- op(Priority, Specifier, Operator).
```

Operators are local to the object (or category) where they are declared. This means that, if you declare a public predicate as an operator, you cannot use operator notation when sending to an object (where the predicate is visible) the respective message (as this would imply visibility of the operator declaration in the context of the *sender* of the message). If you want to declare global operators and, at the same time, use them inside an entity, just write the corresponding directives at the top of your source file, before the entity opening directive.

Note that operators can also be declared using a scope directive. Only these operators are visible to the *current_op/3* reflection method.

When the same operators are used on several entities within the same source file, the corresponding directives must either be repeated in each entity or appear before any entity that uses them. But in the later case, this results in a global scope for the operators. If you prefer the operators to be local to the source file, just *undefine* them at the end of the file. For example:

```
% before any entity that uses the operator
:- op(400, xfx, results).
...

```

(continues on next page)

(continued from previous page)

```
% after all entities that used the operator
:- op(0, xfx, results).
```

Global operators can be declared in the application loader file.

Uses directive

When a predicate makes heavy use of predicates defined on other objects, its predicate clauses can be verbose due to all the necessary message-sending goals. Consider the following example:

```
foo :-
    ...,
    findall(X, list::member(X, L), A),
    list::append(A, B, C),
    list::select(Y, C, R),
    ...
```

Logtalk provides a directive, *uses/2*, which allows us to simplify the code above. One of the usage templates for this directive is:

```
:- uses(Object, [
    Name1/Arity1, Name2/Arity2, ...
]).
```

Rewriting the code above using this directive results in a simplified and more readable predicate definition:

```
:- uses(list, [
    append/3, member/2, select/3
]).

foo :-
    ...,
    findall(X, member(X, L), A),
    append(A, B, C),
    select(Y, C, R),
    ...
```

Logtalk also supports an extended version of this directive that allows the declaration of *predicate aliases* using the notation *Predicate as Alias* (or the alternative notation *Predicate::Alias*). For example:

```
:- uses(btrees, [new/1 as new_btree/1]).
:- uses(queues, [new/1 as new_queue/1]).
```

You may use this extended version for solving conflicts between predicates declared on several *uses/2* directives or just for giving new names to the predicates that will be more meaningful on their using context.

Predicate aliases can also be used to define *predicate shorthands*, simplifying code maintenance. For example:

```
:- uses(pretty_printer, [
    indent(4, Term) as indent(Term)
]).
```

Assuming multiple calls to the shorthand, a change to the indent value will require a change to a single line instead of changing every call.

Another common use of predicate aliases is changing the order of the predicate arguments without using *lambda expressions*. For example:

```
:- uses(meta, [
    fold_left(Closure, Result0, List, Result) as foldl(Closure, List, Result0, Result)
]).
```

See the directive documentation for details and other examples.

The `uses/2` directive allows simpler predicate definitions as long as there are no conflicts between the predicates declared in the directive and the predicates defined in the object (or category) containing the directive. A predicate (or its alias if defined) cannot be listed in more than one `uses/2` directive. In addition, a `uses/2` directive cannot list a predicate (or its alias if defined) that is defined in the object (or category) containing the directive. Any conflicts are reported by Logtalk as compilation errors.

The object identifier argument can also be a *parameter variable* when using the directive in a parametric object or a parametric category. In this case, *dynamic binding* will necessarily be used for all listed predicates (and non-terminals). The parameter variable must be instantiated at runtime when the messages are sent. This feature simplifies experimenting with multiple implementations of the same protocol (for example, to evaluate the performance of each implementation for a particular case). It also simplifies writing tests that check multiple implementations of the same protocol.

An object (or category) can make a predicate listed in a `uses/2` (or `use_module/2`) directive part of its protocol by simply adding a scope directive for the predicate. For example, in the `statistics` library we have:

```
:- public(modes/2).
:- uses(numberlist, [modes/2]).
```

Therefore, a goal such as `sample::modes(Sample, Modes)` implicitly calls `numberlist::modes(Sample, Modes)` without requiring an explicit local definition for the `modes/2` predicate (which would trigger a compilation error).

Alias directive

Logtalk allows the definition of an alternative name for an inherited or imported predicate (or for an inherited or imported grammar rule non-terminal) through the use of the *alias/2* directive:

```
:- alias(Entity, [
    Predicate1 as Alias1,
    Predicate2 as Alias2,
    ...
]).
```

This directive can be used in objects, protocols, or categories. The first argument, `Entity`, must be an entity referenced in the opening directive of the entity containing the `alias/2` directive. It can be an extended or implemented protocol, an imported category, an extended prototype, an instantiated class, or a specialized class. The second argument is a list of pairs of predicate indicators (or grammar rule non-terminal indicators) using the `as` infix operator.

A common use for the `alias/2` directive is to give an alternative name to an inherited predicate in order to improve readability. For example:

```
:- object(square,
    extends(rectangle)).

    :- alias(rectangle, [width/1 as side/1]).

    ...

:- end_object.
```

The directive allows both `width/1` and `side/1` to be used as messages to the object `square`. Thus, using this directive, there is no need to explicitly declare and define a “new” `side/1` predicate. Note that the `alias/2` directive does not rename a predicate, it only provides an alternative, additional name; the original name continues to be available (although it may be masked due to the default inheritance conflict mechanism).

Another common use for this directive is to solve conflicts when two inherited predicates have the same name and arity. We may want to call the predicate that is masked out by the Logtalk lookup algorithm (see the [Inheritance](#) section) or we may need to call both predicates. This is simply accomplished by using the `alias/2` directive to give alternative names to masked-out or conflicting predicates. Consider the following example:

```
:- object(my_data_structure,
    extends(list, set)).

    :- alias(list, [member/2 as list_member/2]).
    :- alias(set, [member/2 as set_member/2]).

    ...

:- end_object.
```

Assuming that both `list` and `set` objects define a `member/2` predicate, without the `alias/2` directives, only the definition of `member/2` predicate in the object `list` would be visible on the object `my_data_structure`, as a result of the application of the Logtalk predicate lookup algorithm. By using the `alias/2` directives, all the following messages would be valid (assuming a public scope for the predicates):

```
% uses list member/2
| ?- my_data_structure::list_member(X, L).

% uses set member/2
| ?- my_data_structure::set_member(X, L).

% uses list member/2
| ?- my_data_structure::member(X, L).
```

When used this way, the `alias/2` directive provides functionality similar to programming constructs of other object-oriented languages that support multi-inheritance (the most notable example probably being the re-naming of inherited features in Eiffel).

Note that the `alias/2` directive never hides a predicate that is visible on the entity containing the directive as a result of the Logtalk lookup algorithm. However, it may be used to make visible a predicate that otherwise would be masked by another predicate, as illustrated in the above example.

The `alias/2` directive may also be used to give access to an inherited predicate, which otherwise would be masked by another inherited predicate, while keeping the original name as follows:

```
:- object(my_data_structure,
    extends(list, set)).

:- alias(list, [member/2 as list_member/2]).
:- alias(set, [member/2 as set_member/2]).

member(X, L) :-
    ^^set_member(X, L).

...

:- end_object.
```

Thus, when sending the message `member/2` to `my_data_structure`, the predicate definition in `set` will be used instead of the one contained in `list`.

Documenting directive

A predicate can be documented with arbitrary user-defined information by using the *info/2* directive:

```
:- info(Name/Arity, List).
```

The second argument is a list of Key is Value terms. See the *Documenting* section for details.

Multifile directive

A predicate can be declared *multifile* by using the *multifile/1* directive:

```
:- multifile(Name/Arity).
```

This allows clauses for a predicate to be defined in several objects and/or categories. This is a directive that should be used with care. It's commonly used in the definition of *hook predicates*. Multifile predicates (and non-terminals) may also be declared dynamic using the same predicate (or non-terminal) notation (multifile predicates are static by default).

Logtalk precludes using a multifile predicate for breaking object encapsulation by checking that the object (or category) declaring the predicate (using a scope directive) defines it also as multifile. This entity is said to contain the *primary declaration* for the multifile predicate. Entities containing primary multifile predicate declarations must always be compiled before entities defining clauses for those multifile predicates. The Logtalk compiler will print a warning if the scope directive is missing. Note also that the *multifile/1* directive is mandatory when defining multifile predicates.

Consider the following simple example:

```
:- object(main).

:- public(a/1).
:- multifile(a/1).
a(1).

:- end_object.
```

After compiling and loading the `main` object, we can define other objects (or categories) that contribute with clauses for the multifile predicate. For example:

```

:- object(other).

    :- multifile(main::a/1).
    main::a(2).
    main::a(X) :-
        b(X).

    b(3).
    b(4).

:- end_object.

```

After compiling and loading the above objects, you can use queries such as:

```

| ?- main::a(X).

X = 1 ;
X = 2 ;
X = 3 ;
X = 4
yes

```

Note that the order of multifile predicate clauses depends on several factors, including loading order and compiler implementation details. Therefore, your code should never assume or rely on a specific order of the multifile predicate clauses.

When a clause of a multifile predicate is a rule, its body is compiled within the context of the object or category defining the clause. This allows clauses for multifile predicates to call local object or category predicates. But the values of the *sender*, *this*, and *self* in the implicit execution context are passed from the clause head to the clause body. This is necessary to ensure that these values are always valid and to allow multifile predicate clauses to be defined in categories. A call to the `parameter/2` execution context methods, however, retrieves parameters of the entity defining the clause, not from the entity for which the clause is defined. The parameters of the entity for which the clause is defined can be accessed by simple unification at the clause head.

Multifile predicate rules should not contain cuts, as these may prevent other clauses for the predicate from being used by callers. The compiler prints by default a warning when a cut is found in a multifile predicate definition.

Local calls to the database methods from multifile predicate clauses defined in an object take place in the object's own database instead of the database of the entity holding the multifile predicate primary declaration. Similarly, local calls to the `expand_term/2` and `expand_goal/2` methods from a multifile predicate clause look for clauses of the `term_expansion/2` and `goal_expansion/2` hook predicates starting from the entity defining the clause instead of the entity holding the multifile predicate primary declaration. Local calls to the `current_predicate/1`, `predicate_property/2`, and `current_op/3` methods from multifile predicate clauses defined in an object also lookup predicates and their properties in the object's own database instead of the database of the entity holding the multifile predicate primary declaration.

Coinductive directive

A predicate can be declared *coinductive* by using the *coinductive/1* directive. For example:

```
:- coinductive(comember/2).
```

Logtalk support for coinductive predicates is experimental and requires a *backend Prolog compiler* with minimal support for cyclic terms. The value of the read-only *coinduction flag* is set to supported for the backend Prolog compilers providing that support.

Synchronized directive

A predicate can be declared *synchronized* by using the *synchronized/1* directive. For example:

```
:- synchronized(write_log_entry/2).  
:- synchronized([produce/1, consume/1]).
```

See the section on *synchronized predicates* for details.

1.8.3 Defining predicates

Object predicates

We define object predicates as we have always defined Prolog predicates, the only difference being that we have four more control structures (the three message-sending operators plus the external call operator) to play with. For example, if we wish to define an object containing common utility list predicates like *append/2* or *member/2* we could write something like:

```
:- object(list).  
  
    :- public(append/3).  
    append([], L, L).  
    append([H| T], L, [H| T2]) :-  
        append(T, L, T2).  
  
    :- public(member/2).  
    member(H, [H| _]).  
    member(H, [_| T]) :-  
        member(H, T).  
  
:- end_object.
```

Note that, abstracting from the opening and closing object directives and the scope directives, what we have written is also valid Prolog code. Calls in a predicate definition body default to the local predicates unless we use the message-sending operators or the external call operator. This simplifies conversion from plain Prolog code to Logtalk objects: often we just need to add the necessary encapsulation and scope directives to the old code.

Category predicates

A category can only contain clauses for static predicates. But there are no restrictions in declaring and calling dynamic predicates from inside a category. Because a category can be imported by multiple objects, dynamic predicates must be called either in the context of *self*, using the *message to self* control structure, *(::)/1*, or in the context of *this* (i.e., in the context of the object importing the category). For example, if we want to define a category implementing attributes using the dynamic database of *self* we could write:

```
:- category(attributes).

:- public(get/2).
:- public(set/2).

:- private(attribute_/2).
:- dynamic(attribute_/2).

get(Var, Value) :-
    ::attribute_(Var, Value).

set(Var, Value) :-
    ::retractall(attribute_(Var, _)),
    ::asserta(attribute_(Var, Value)).

:- end_category.
```

In this case, the *get/2* and *set/2* predicates will always access/update the correct definition, contained in the object receiving the messages.

In alternative, if we want a category implementing attributes using the dynamic database of *this*, we would write instead:

```
:- category(attributes).

:- public(get/2).
:- public(set/2).

:- private(attribute_/2).
:- dynamic(attribute_/2).

get(Var, Value) :-
    attribute_(Var, Value).

set(Var, Value) :-
    retractall(attribute_(Var, _)),
    asserta(attribute_(Var, Value)).

:- end_category.
```

In this case, each object importing the category will have its own clauses for the *attribute_/2* private dynamic predicate.

Meta-predicates

Meta-predicates may be defined inside objects and categories as any other predicate. A meta-predicate is declared using the `meta_predicate/1` directive as described earlier in this section. When defining a meta-predicate, the arguments in the clause heads corresponding to the meta-arguments must be variables. All meta-arguments are called in the context of the object or category calling the meta-predicate. In particular, when sending a message that corresponds to a meta-predicate, the meta-arguments are called in the context of the object or category sending the message.

The most simple example is a meta-predicate with a meta-argument that is called as a goal. E.g. the `ignore/1` built-in predicate could be defined as:

```
:- public(ignore/1).
:- meta_predicate(ignore(0)).

ignore(Goal) :-
    (Goal -> true; true).
```

The `0` in the meta-predicate template tells us that the meta-argument is a goal that will be called by the meta-predicate.

Some meta-predicates have meta-arguments that are not goals but *closures*. Logtalk supports the definition of meta-predicates that are called with closures instead of goals as long as the definition uses the `call/1-N` built-in predicate to call the closure with the additional arguments. A classical example is a list mapping predicate:

```
:- public(map/2).
:- meta_predicate(map(1, *)).

map(_, []).
map(Closure, [Arg| Args]) :-
    call(Closure, Arg),
    map(Closure, Args).
```

Note that in this case the meta-predicate directive specifies that the closure will be extended with exactly one additional argument. When calling a meta-predicate, a closure can correspond to a user-defined predicate, a built-in predicate, a *lambda expression*, or a control construct.

In some cases, it is not a meta-argument but one of its sub-terms that is called as a goal or used as a closure. For example:

```
:- public(call_all/1).
:- meta_predicate(call_all(::)).

call_all([]).
call_all([Goal| Goals]) :-
    call(Goal),
    call_all(Goals).
```

The `::` mode indicator in the meta-predicate template allows the corresponding argument in the meta-predicate definition to be a non-variable term and instructs the compiler to look into the argument sub-terms for goal and closure *meta-variables*.

When a meta-predicate calls another meta-predicate, both predicates require `meta_predicate/1` directives. For example, the `map/2` meta-predicate defined above is usually implemented by exchanging the argument order to take advantage of first-argument indexing:

```

:- meta_predicate(map(1, *)).
map(Closure, List) :-
    map_(List, Closure).

:- meta_predicate(map_(*, 1)).
map_([], _).
map_([Head| Tail], Closure) :-
    call(Closure, Head),
    map_(Tail, Closure).

```

Note that Logtalk, unlike most Prolog module systems, is not based on a predicate prefixing mechanism. Thus, the meta-argument calling context is not part of the meta-argument itself.

Lambda expressions

The use of *lambda expressions* as meta-predicate goal and *closure* arguments often saves writing auxiliary predicates for the sole purpose of calling the meta-predicates. A simple example of a lambda expression is:

```

| ?- meta::map([X,Y]>>(Y is 2*X), [1,2,3], Ys).
Ys = [2,4,6]
yes

```

In this example, a lambda expression, $[X,Y]>>(Y \text{ is } 2*X)$, is used as an argument to the `map/3` list mapping predicate, defined in the library object `meta`, in order to double the elements of a list of integers. Using a lambda expression avoids writing an auxiliary predicate for the sole purpose of doubling the list elements. The *lambda parameters* are represented by the list $[X,Y]$, which is connected to the *lambda goal*, $(Y \text{ is } 2*X)$, by the $(>>)/2$ operator. The `map/3` predicate calls the lambda goal with fresh/unique variables, represented by the X and Y parameters, for each pair of elements of the second and third list arguments.

Currying is supported. I.e. it is possible to write a lambda expression whose goal is another lambda expression. The above example can be rewritten as:

```

| ?- meta::map([X]>>([Y]>>(Y is 2*X)), [1,2,3], Ys).
Ys = [2,4,6]
yes

```

Lambda expressions may also contain *lambda-free variables*. I.e. variables that are global to the lambda expression and shared with the surrounding meta-call context. Consider the following variant of the previous example:

```

| ?- between(1, 3, N), meta::map({N}/[X,Y]>>(Y is N*X), [1,2,3], L).
N = 1, L = [1,2,3] ;
N = 2, L = [2,4,6] ;
N = 3, L = [3,6,9]
yes

```

In this case, the lambda-free variable, N , bound by the `between/3` goal, is fixed across all implicit calls made by the `map/3` goal.

A second example of free variables in a lambda expression using GNU Prolog as the backend compiler:

```

| ?- meta::map({Z}/[X,Y]>>(Z#=X+Y), [1,2,3], Zs).
Z = _#22(3..268435455)

```

(continues on next page)

(continued from previous page)

```
Zs = [_#3(2..268435454),_#66(1..268435453),_#110(0..268435452)]
yes
```

The ISO Prolog construct `{}/1` is used for representing the lambda-free variables as this representation is often associated with set representation. Note that the order of the free variables is of no consequence (on the other hand, a list is used for the lambda parameters as their order does matter).

Both lambda free variables and lambda parameters can be any Prolog term. Consider the following example by Markus Triska:

```
| ?- meta::map([A-B,B-A]>>true, [1-a,2-b,3-c], Zs).
Zs = [a-1,b-2,c-3]
yes
```

Lambda expressions can be used, as expected, in non-deterministic queries, as in the following example using SWI-Prolog as the backend compiler and Markus Triska's CLP(FD) library:

```
| ?- meta::map({Z}/[X,Y]>>(clpfd:(Z#=X+Y)), Xs, Ys).
Xs = [],
Ys = [] ;
Xs = [_G1369],
Ys = [_G1378],
_G1369+_G1378#=Z ;
Xs = [_G1579, _G1582],
Ys = [_G1591, _G1594],
_G1582+_G1594#=Z,
_G1579+_G1591#=Z ;
Xs = [_G1789, _G1792, _G1795],
Ys = [_G1804, _G1807, _G1810],
_G1795+_G1810#=Z,
_G1792+_G1807#=Z,
_G1789+_G1804#=Z ;
...
```

As illustrated by the above examples, lambda expression syntax reuses the ISO Prolog construct `{}/1` and the standard operators `(/)/2` and `(>>)/2`, thus avoiding defining new operators, which is always tricky for a portable system such as Logtalk. The operator `(>>)/2` was chosen as it suggests an arrow, similar to the syntax used in other languages such as OCaml and Haskell to connect lambda parameters with lambda functions. This syntax was also chosen in order to simplify parsing, error checking, and compilation of lambda expressions. The full specification of the lambda expression syntax can be found in the [language grammar](#).

The compiler checks whenever possible that all variables in a lambda expression are either classified as free variables or as lambda parameters. Non-classified variables in a lambda goal (including any anonymous variables) should be regarded as a programming error. The compiler also checks whenever possible if a variable is classified as both a free variable and a lambda parameter. There are a few cases where a variable playing a dual role is intended, but, in general, this also results from a programming error. A third check verifies that no lambda parameter variable is used elsewhere in a clause. Such cases are either programming errors, when the variable appears before the lambda expression, or bad programming style, when the variable is used after the lambda expression. These linter warnings are controlled by the `lambda_variables` flag. Note that the dynamic features of the language and lack of sufficient information at compile-time may prevent the compiler from checking all uses of lambda expressions. To improve linter coverage, compile code using lambda expressions with the `optimize flag` turned on, as that will result in additional cases of meta-arguments being evaluated for possible optimizations.

Warning

Variables listed in lambda parameters must not be shared with other goals in a clause.

An optimizing meta-predicate and lambda expression compiler, based on the *term-expansion mechanism*, is provided as a standard library for practical performance.

A common use of lambda expressions as closure meta-arguments is to workaround closures always being extended by *appending* additional arguments to construct a goal. For example, assume that we want to filter a list of atoms by a given length. We can use the standard `atom_length/2` predicate despite the argument order by writing:

```
filter(Length, Atoms, Filtered) :-
    meta::include({Length}/[Atom]>>atom_length(Atom,Length), Atoms, Filtered).
```

But Logtalk supports a faster alternative by using predicate aliases to change the argument order when calling library or built-in predicates:

```
:- uses(user, [
    atom_length(Atom, Length) as length_atom(Length, Atom)
]).

filter(Length, Atoms, Filtered) :-
    meta::include(length_atom(Length), Atoms, Filtered).
```

In this case, the performance is no longer dependent on compiling away lambda expressions. The resulting code is also easier to read (and thus debug and maintain). But the `uses/2` directive is implicitly defining an auxiliary predicate, which is exactly what we wanted to avoid in the first place by using a lambda expression.

Redefining built-in predicates

Logtalk built-in predicates and Prolog built-in predicates can be redefined inside objects and categories. Although the redefinition of Logtalk built-in predicates should be avoided, the support for redefining Prolog built-in predicates is a practical requirement given the different sets of proprietary built-in predicates provided by backend Prolog systems.

The compiler supports a *redefined_built_ins* flag, whose default value is *silent*, that can be set to *warning* to alert the user of any redefined Logtalk or Prolog built-in predicate.

The redefinition of Prolog built-in predicates can be combined with the *conditional compilation directives* when writing portable applications where some of the supported backends don't provide a built-in predicate found in the other backends. As an example, consider the de facto standard `msort/2` predicate (which sorts a list while keeping duplicates). This predicate is provided as a built-in predicate in most but not all backends. The `list` library object includes the code:

```
:- if(predicate_property(msort(_,_), built_in)).

    msort(List, Sorted) :-
        {msort(List, Sorted)}.

:- else.

    length(List, Length) :-
```

(continues on next page)

(continued from previous page)

```
...
:- endif.
```

I.e. the object will use the built-in predicate when available. Otherwise, it will use the predicate definition provided by the `list` object.

The redefinition of built-in predicates can also be accomplished using *predicate shorthands*. This can be useful when porting code while minimizing the changes. For example, assume that existing code uses the `format/2` de facto standard predicate for writing messages. To convert the code to use the *message printing mechanism*, we could write:

```
:- uses(logtalk, [
    print_message(comment, core, Format+Arguments) as format(Format, Arguments)
]).

process(Crate, Contents) :-
    format('Processing crate ~w...', [Crate]),
    ...,
    format('Filing with ~w...', [Contents]),
    ....
```

The predicate shorthand instructs the compiler to rewrite all `format/2` goals as `logtalk::print_message/3` goals, thus allowing us to reuse the code without changes.

1.8.4 Definite clause grammar rules

Definite clause grammar rules (DCGs) provide a convenient notation to represent the parsing and rewrite rules common of most grammars in Prolog. In Logtalk, definite clause grammar rules can be encapsulated in objects and categories. Currently, the ISO/IEC WG17 group is working on a draft specification for a definite clause grammars Prolog standard. Therefore, in the meantime, Logtalk follows the common practice of Prolog compilers supporting definite clause grammars, extending it to support calling grammar rules contained in categories and objects. A common example of a definite clause grammar is the definition of a set of rules for parsing simple arithmetic expressions:

```
:- object(calculator).

:- public(parse/2).

parse(Expression, Value) :-
    phrase(expr(Value), Expression).

expr(Z) --> term(X), "+", expr(Y), {Z is X + Y}.
expr(Z) --> term(X), "-", expr(Y), {Z is X - Y}.
expr(X) --> term(X).

term(Z) --> number(X), "*", term(Y), {Z is X * Y}.
term(Z) --> number(X), "/", term(Y), {Z is X / Y}.
term(Z) --> number(Z).

number(C) --> "+", number(C).
number(C) --> "-", number(X), {C is -X}.
```

(continues on next page)

(continued from previous page)

```

number(X) --> [C], {0'0 =< C, C =< 0'9, X is C - 0'0}.

:- end_object.

```

After compiling and loading this object, we can test the grammar rules using the `parse/2` message:

```

| ?- calculator::parse("1+2-3*4", Result).

Result = -9
yes

```

The non-terminals can be called from predicates using the private built-in methods `phrase/2` and `phrase/3` as shown in the example above. When we want to use the built-in methods `phrase/2` and `phrase/3`, the non-terminal used as in the first argument must be within the scope of the *sender*. For the above example, assuming that we want the predicate corresponding to the `expr//1` non-terminal to be public, the corresponding scope directive would be:

```

:- public(expr//1).

```

The `//` infix operator used above tells the Logtalk compiler that the scope directive refers to a grammar rule non-terminal, not to a predicate. The idea is that the predicate corresponding to the translation of the `expr//1` non-terminal will have a number of arguments equal to one plus the number of additional arguments necessary for processing the implicit difference list of tokens.

In the body of a grammar rule, we can call rules that are inherited from ancestor objects, imported from categories, or contained in other objects. This is accomplished by using non-terminals as messages. Using a non-terminal as a message to *self* allows us to call grammar rules in categories and ancestor objects. To call grammar rules encapsulated in other objects, we use a non-terminal as a message to those objects. Consider the following example, containing grammar rules for parsing natural language sentences:

```

:- object(sentence,
    imports(determiners, nouns, verbs)).

    :- public(parse/2).

    parse(List, true) :-
        phrase(sentence, List).
    parse(_, false).

    sentence --> noun_phrase, verb_phrase.

    noun_phrase --> ::determiner, ::noun.
    noun_phrase --> ::noun.

    verb_phrase --> ::verb.
    verb_phrase --> ::verb, noun_phrase.

:- end_object.

```

The categories imported by the object would contain the necessary grammar rules for parsing determiners, nouns, and verbs. For example:

```
:- category(determiners).

    :- private(determiner//0).

    determiner --> [the].
    determiner --> [a].

:- end_category.
```

Along with the message-sending operators (`(::)/1`, `(::)/2`, and `(^^)/1`), we may also use other control constructs such as `(<<)/2`, `(\+)/1`, `!/0`, `(;)/2`, `(->)/2`, `{}/1`, `call//1-N`, and `catch/3` in the body of a grammar rule. When using a backend Prolog compiler that supports modules, we may also use the `(:)/2` control construct.

Warning

The semantics of `(\+)/1` and `(->)/2` control constructs in grammar rules with a terminal or a non-terminal in the **first** argument are problematic due to unrestricted lookahead that may or may not be valid depending on the grammar rule implicit arguments. By default, the linter will print warnings for such calls (controlled by the `grammar_rules` flag). Preferably restrict the use of the `(\+)/1` control construct to `{}/1` arguments and the use of the `(->)/2` control construct to `{}/1` test arguments.

In addition, grammar rules may contain meta-calls (a variable taking the place of a non-terminal), which are translated to calls of the built-in method `phrase//1`. The `meta_non_terminal/1` directive allows the declaration of non-terminals that have arguments that are meta-called from grammar rules. For example:

```
:- meta_non_terminal(zero_or_more(1, *)).

zero_or_more(Closure, [Terminal| Terminals]) -->
    call(Closure, Terminal), !, zero_or_more(Closure, Terminals).
zero_or_more(_, []) -->
    [].
```

You may have noticed that Logtalk defines `{}/1` as a control construct for bypassing the compiler when compiling a clause body goal. As exemplified above, this is the same control construct that is used in grammar rules for bypassing the expansion of rule body goals when a rule is converted into a clause. Both control constructs can be combined in order to call a goal from a grammar rule body, while bypassing at the same time the Logtalk compiler. Consider the following example:

```
bar :-
    write('bar predicate called'), nl.

:- object(bypass).

    :- public(foo//0).

    foo --> {{bar}}.

:- end_object.
```

After compiling and loading this code, we may try the following query:

```
| ?- logtalk << phrase(bypass::foo, _, _).
```

```
bar predicate called
```

```
yes
```

This is the expected result, as the expansion of the grammar rule into a clause leaves the {bar} goal untouched, which, in turn, is converted into the goal bar when the clause is compiled. Note that we tested the `bypass::foo/0` non-terminal by calling the `phrase/3` built-in method in the context of the `logtalk` built-in object. This workaround is necessary due to the Prolog backend implementation of the `phrase/3` predicate not being aware of the Logtalk `(:)/2` message-sending control construct semantics.

A grammar rule non-terminal may be declared as dynamic or discontinuous, as any object predicate, using the same `Name//Arity` notation illustrated above for the scope directives. In addition, grammar rule non-terminals can be documented using the [info/2](#) directive, as in the following example:

```
:- public(sentence//0).
```

```
:- info(sentence//0, [
    comment is 'Rewrites sentence into noun and verb phrases.'
]).
```

Note

Future Logtalk versions may compile grammar rules differently from Prolog traditional compilation to prevent name clashes between non-terminals and predicates. Therefore, you should always call non-terminals from predicates using the `phrase/2-3` built-in methods and always call predicates from grammar rules using the `call/1` built-in method. This recommended practice, besides making your code forward compatible with future Logtalk versions, also makes the code more clear. The linter prints warnings when these guidelines are not followed (notably, when a predicate is called as a non-terminal or a non-terminal is called as a predicate).

1.8.5 Built-in methods

Built-in methods are built-in object and category predicates. These include methods to access message execution context, to find sets of solutions, to inspect objects, for database handling, for term and goal expansion, and for printing messages. Some of them are counterparts to standard Prolog built-in predicates that take into account Logtalk semantics. Similar to Prolog built-in predicates, built-in methods cannot be redefined.

Logic and control methods

The `!/0`, `true/0`, `fail/0`, `false/0`, and `repeat/0` standard control constructs and logic predicates are interpreted as built-in public methods and thus can be used as messages to any object. In practice, they are only used as messages when sending multiple messages to the same object (see the section on [message broadcasting](#)).

Execution context methods

Logtalk defines five built-in private methods to access an object execution context. These methods are in the common usage scenarios translated to a single unification performed at compile-time with a clause head context argument. Therefore, they can be freely used without worrying about performance penalties. When called from inside a category, these methods refer to the execution context of the object importing the category. These methods are private and cannot be used as messages to objects.

To find the object that received the message under execution, we may use the `self/1` method. We may also retrieve the object that has sent the message under execution using the `sender/1` method.

The method `this/1` enables us to retrieve the name of the object for which the predicate clause whose body is being executed is defined instead of using the name directly. This helps to avoid breaking the code if we decide to change the object name and forget to change the name references. This method may also be used from within a category. In this case, the method returns the object importing the category on whose behalf the predicate clause is being executed.

Here is a short example including calls to these three object execution context methods:

```
:- object(test).

    :- public(test/0).

    test :-
        this(This),
        write('Calling predicate definition in '),
        writeq(This), nl,
        self(Self),
        write('to answer a message received by '),
        writeq(Self), nl,
        sender(Sender),
        write('that was sent by '),
        writeq(Sender), nl, nl.

:- end_object.

:- object(descendant,
    extends(test)).

:- end_object.
```

After compiling and loading these two objects, we can try the following goal:

```
| ?- descendant::test.

Calling predicate definition in test
to answer a message received by descendant
that was sent by user
yes
```

Note that the goals `self(Self)`, `sender(Sender)`, and `this(This)`, being translated to unifications with the clause head context arguments at compile-time, are effectively removed from the clause body. Therefore, a clause such as:

```

predicate(Arg) :-
    self(Self),
    atom(Arg),
    ... .

```

is compiled with the goal `atom(Arg)` as the first condition on the clause body. As such, the use of these context execution methods does not interfere with the optimizations that some Prolog compilers perform when the first clause body condition is a call to a built-in type-test predicate or a comparison operator.

For parametric objects and categories, the method `parameter/2` enables us to retrieve current parameter values (see the section on *parametric objects* for a detailed description). For example:

```

:- object(block(_Color)).

    :- public(test/0).

    test :-
        parameter(1, Color),
        write('Color parameter value is '),
        writeq(Color), nl.

:- end_object.

```

An alternative to the `parameter/2` predicate is to use *parameter variables*:

```

:- object(block(_Color_)).

    :- public(test/0).

    test :-
        write('Color parameter value is '),
        writeq(_Color_), nl.

:- end_object.

```

After compiling and loading either version of the object, we can try the following goal:

```

| ?- block(blue)::test.

Color parameter value is blue
yes

```

Calls to the `parameter/2` method are translated to a compile-time unification when the second argument is a variable. When the second argument is bound, the calls are translated to a call to the built-in predicate `arg/3`.

When type-checking predicate arguments, it is often useful to include the predicate execution context when reporting an argument error. The `context/1` method provides access to that context. For example, assume a predicate `foo/2` that takes an atom and an integer as arguments. We could type-check the arguments by writing (using the library type object):

```

foo(A, N) :-
    % type-check arguments
    context(Context),

```

(continues on next page)

(continued from previous page)

```

type::check(atom, A, Context),
type::check(integer, N, Context),
% arguments are fine; go ahead
...

```

Error handling and throwing methods

Besides the *catch/3* and *throw/1* methods inherited from Prolog, Logtalk also provides a set of convenience methods to throw standard error/2 exception terms: *instantiation_error/0*, *uninstantiation_error/1*, *type_error/2*, *domain_error/2*, *existence_error/2*, *permission_error/3*, *representation_error/1*, *evaluation_error/1*, *resource_error/1*, *syntax_error/1*, and *system_error/0*. When using these methods, the second argument of the error/2 exception term is bound to the execution context (as it would be provided by the *context/1* method).

Database methods

Logtalk provides a set of built-in methods for *object database* handling similar to the usual database Prolog predicates: *abolish/1*, *asserta/1*, *assertz/1*, *clause/2*, *retract/1*, and *retractall/1*. These methods always operate on the database of the object receiving the corresponding message. When called locally, these predicates take into account any *uses/2* or *use_module/2* directives that refer to the dynamic predicate being handled. For example, in the following object, the clauses for the *data/1* predicate are retracted and asserted in user due to the *uses/2* directive:

```

:- object(an_object).

   :- uses(user, [data/1]).

   :- public(some_predicate/1).
   some_predicate(Arg) :-
       retractall(data(_)),
       assertz(data(Arg)).

:- end_object.

```

When working with dynamic grammar rule non-terminals, you may use the built-in method *expand_term/2* to convert a grammar rule into a clause that can then be used with the database methods.

Logtalk also supports *asserta/2*, *assertz/2*, *clause/3*, and *erase/1* built-in methods when run with a backend that supports the corresponding legacy built-in predicates that work with *clause references*.

Meta-call methods

Logtalk supports the generalized *call/1-N* meta-predicate. This built-in private meta-predicate must be used in the implementation of meta-predicates that work with *closures* instead of goals. In addition, Logtalk supports the built-in private meta-predicates *ignore/1*, *once/1*, and *(\+)/1*. These methods cannot be used as messages to objects.

All solutions methods

The usual all solutions meta-predicates are built-in private methods in Logtalk: *bagof/3*, *findall/3*, *findall/4*, and *setof/3*. There is also a *forall/2* method that implements generate-and-test loops. These methods cannot be used as messages to objects.

Reflection methods

Logtalk provides a comprehensive set of built-in predicates and built-in methods for querying about entities and predicates. Some of the information, however, requires that the source files are compiled with the *source_data* flag turned on.

The *reflection API* supports two different views on entities and their contents, which we may call the *transparent box view* and the *black box view*. In the transparent box view, we look into an entity disregarding how it will be used and returning all information available on it, including predicate declarations and predicate definitions. This view is supported by the entity property built-in predicates. In the black box view, we look into an entity from a usage point of view using built-in methods for inspecting object operators and predicates that are within scope from where we are making the call: *current_op/3*, which returns operator specifications; *predicate_property/2*, which returns predicate properties; and *current_predicate/1*, which enables us to query about user-defined predicate definitions. See below for a more detailed description of these methods.

Definite clause grammar parsing methods and non-terminals

Logtalk supports two definite clause grammar parsing built-in private methods, *phrase/2* and *phrase/3*, with definitions similar to the predicates with the same name found on most Prolog compilers that support definite clause grammars. These methods cannot be used as messages to objects.

Logtalk also supports *phrase//1*, *call//1-N*, and *eos//0* built-in non-terminals. The *call//1-N* non-terminals take a *closure* (which can be a lambda expression) plus zero or more additional arguments and are processed by appending the input list of tokens and the list of remaining tokens to the arguments.

1.8.6 Predicate properties

We can find the properties of visible predicates by calling the *predicate_property/2* built-in method. For example:

```
| ?- bar::predicate_property(foo(_), Property).
```

Note that this method takes into account the predicate's scope declarations. In the above example, the call will only return properties for public predicates.

An object's set of visible predicates is the union of all the predicates declared for the object with all the built-in methods and all the Logtalk and Prolog built-in predicates.

The following predicate properties are supported:

scope(Scope)

The predicate scope (useful for finding the predicate scope with a single call to *predicate_property/2*)

public, protected, private

The predicate scope (useful for testing if a predicate has a specific scope)

static, dynamic

All predicates are either static or dynamic (note, however, that a dynamic predicate can only be abolished if it was dynamically declared)

logtalk, prolog, foreign

A predicate can be defined in Logtalk source code, Prolog code, or in foreign code (e.g., in C)

built_in

The predicate is a built-in predicate

multifile

The predicate is declared multifile (i.e., it can have clauses defined in multiple files or entities)

meta_predicate(Template)

The predicate is declared as a meta-predicate with the specified template

coinductive(Template)

The predicate is declared as a coinductive predicate with the specified template

declared_in(Entity)

The predicate is declared (using a scope directive) in the specified entity

defined_in(Entity)

The predicate definition is looked up in the specified entity (note that this property does not necessarily imply that clauses for the predicate exist in Entity; the predicate can simply be false as per the *closed-world assumption*)

redefined_from(Entity)

The predicate is a redefinition of a predicate definition inherited from the specified entity

non_terminal(NonTerminal//Arity)

The predicate resulted from the compilation of the specified grammar rule non-terminal

alias_of(Predicate)

The predicate (name) is an alias for the specified predicate

alias_declared_in(Entity)

The predicate alias is declared in the specified entity

synchronized

The predicate is declared as synchronized (i.e., it's a deterministic predicate synchronized using a mutex when using a backend Prolog compiler supporting a compatible multi-threading implementation)

Some properties are only available when the entities are defined in source files and when those source files are compiled with the *source_data* flag turned on:

recursive

The predicate definition includes at least one recursive rule

inline

The predicate definition is inlined

auxiliary

The predicate is not user-defined but rather automatically generated by the compiler or the *term-expansion mechanism*

mode(Mode, Solutions)

Instantiation, type, and determinism mode for the predicate (which can have multiple modes)

info(ListOfPairs)

Documentation key-value pairs as specified in the user-defined info/2 directive

number_of_clauses(N)

The number of clauses for the predicate existing at compilation time (note that this property is not updated at runtime when asserting and retracting clauses for dynamic predicates)

number_of_rules(N)

The number of rules for the predicate existing at compilation time (note that this property is not updated at runtime when asserting and retracting clauses for dynamic predicates)

declared_in(Entity, Line)

The predicate is declared (using a scope directive) in the specified entity in a source file at the specified line (if applicable)

defined_in(Entity, Line)

The predicate is defined in the specified entity in a source file at the specified line (if applicable)

redefined_from(Entity, Line)

The predicate is a redefinition of a predicate definition inherited from the specified entity, which is defined in a source file at the specified line (if applicable)

alias_declared_in(Entity, Line)

The *predicate alias* is declared in the specified entity in a source file at the specified line (if applicable)

The properties `declared_in/1-2`, `defined_in/1-2`, and `redefined_from/1-2` do not apply to built-in methods and Logtalk or Prolog built-in predicates. Note that if a predicate is declared in a category imported by the object, it will be the category name — not the object name — that will be returned by the property `declared_in/1`. The same is true for protocol declared predicates.

Some properties, such as line numbers, are only available when the entity holding the predicates is defined in a source file compiled with the *source_data* flag turned on. Moreover, line numbers are only supported in *backend Prolog compilers* that provide access to the start line of a read term. When such support is not available, the value `-1` is returned for the start and end lines.

1.8.7 Finding declared predicates

We can find, by backtracking, all visible user predicates by calling the *current_predicate/1* built-in method. This method takes into account predicate scope declarations. For example, the following call will only return user predicates that are declared public:

```
| ?- some_object::current_predicate(Name/Arity).
```

The predicate property `non_terminal/1` may be used to retrieve all grammar rule non-terminals declared for an object. For example:

```
current_non_terminal(Object, Name//Args) :-
    Object::current_predicate(Name/Arity),
    functor(Predicate, Functor, Arity),
    Object::predicate_property(Predicate, non_terminal(Name//Args)).
```

Usually, the non-terminal and the corresponding predicate share the same functor, but users should not rely on this always being true.

1.8.8 Calling Prolog predicates

Logtalk is designed for both *robustness* and *portability*. In the context of calling Prolog predicates, robustness requires that the compilation of Logtalk source code must not have *accidental* dependencies on Prolog code that happens to be loaded at the time of the compilation. One immediate consequence is that only Prolog *built-in* predicates are visible from within objects and categories. But Prolog systems provide a widely diverse set of built-in predicates, easily raising portability issues. Relying on non-standard predicates is often unavoidable, however, due to the narrow scope of Prolog standards. Logtalk applications may also require calling user-defined Prolog predicates, either in user or in Prolog modules.

Calling Prolog built-in predicates

In predicate clauses and object initialization/1 directives, predicate calls that are not prefixed with a message-sending, super call, or module qualification operator (`::`, `^^`, or `:`), are compiled to either calls to local predicates or as calls to Logtalk/Prolog built-in predicates. A predicate call is compiled as a call to a local predicate if the object (or category) contains a scope directive, a multifile directive, a dynamic directive, or a definition for the called predicate. When that is not the case, the compiler checks if the call corresponds to a Logtalk or Prolog built-in predicate. Consider the following example:

```
foo :-  
    ...,  
    write(bar),  
    ...
```

The call to the `write/1` predicate will be compiled as a call to the corresponding Prolog standard built-in predicate unless the object (or category) containing the above definition also contains a predicate named `write/1` or a directive for the predicate.

When calling non-standard Prolog built-in predicates or using non-standard Prolog arithmetic functions, we may run into portability problems while trying your applications with different backend Prolog compilers. We can use the compiler *portability flag* to generate warnings for calls to non-standard predicates and arithmetic functions. We can also help document those calls using the *uses/2* directive. For example, a few Prolog systems provide an `atom_string/2` non-standard predicate. We can write (in the object or category calling the predicate):

```
:- uses(user, [atom_string/2])
```

This directive is based on the fact that built-in predicates are visible in plain Prolog (i.e., in user). Besides helping to document the dependency on a non-standard built-in predicate, this directive will also silence the compiler portability warning.

Calling Prolog non-standard built-in meta-predicates

Prolog built-in meta-predicates may only be called locally within objects or categories, i.e. they cannot be used as messages. Compiling calls to non-standard, Prolog built-in meta-predicates can be tricky, however, as there is no standard way of checking if a built-in predicate is also a meta-predicate and finding out which are its meta-arguments. But Logtalk supports overriding the original meta-predicate template when not programmatically available or usable. For example, assume a `det_call/1` Prolog built-in meta-predicate that takes a goal as argument. We can add to the object (or category) calling it the directive:

```
:- meta_predicate(user::det_call(0)).
```

Another solution is to explicitly declare non-standard built-in Prolog meta-predicates in the corresponding adapter file using the internal predicate `'$lgt_prolog_meta_predicate'/3`. For example:

```
'$lgt_prolog_meta_predicate'(det_call(_), det_call(0), predicate).
```

The third argument can be either the atom predicate or the atom control_construct, a distinction that is useful when compiling in debug mode.

Calling Prolog foreign predicates

Prolog systems often support defining *foreign* predicates, i.e. predicates defined using languages other than Prolog using a *foreign language interface*. There isn't, however, any standard for defining, making available, and recognizing foreign predicates. From a Logtalk perspective, the two most common scenarios are calling a foreign predicate (from within an object or a category) and making a set of foreign predicates available as part of an object (or category) protocol. Assuming, as this is the most common case, that foreign predicates are globally visible once made available (using a Prolog system-specific loading or linking procedure), we can simply call them as user-defined plain predicates, as explained in the next section. When defining an object (or category) that makes available foreign predicates, the advisable solution is to name the predicates after the object (or category) and then define object (or category) predicates that call the foreign predicates. Most backend adapter files include support for recognizing foreign predicates that allows the Logtalk compiler to inline calls to the predicates (thus avoiding call indirection overheads).

Calling Prolog user-defined plain predicates

User-defined Prolog plain predicates (i.e., predicates that are not defined in a Prolog module) can be called from within objects or categories by sending the corresponding message to user. For example:

```
foo :-
    ...,
    user::bar,
    ...
```

In alternative, we can use the `uses/2` directive and write:

```
:- uses(user, [bar/0]).

foo :-
    ...,
    bar,
    ...
```

Note that user is a pseudo-object in Logtalk containing all predicate definitions that are not encapsulated (either in a Logtalk entity or a Prolog module).

When the Prolog predicate is not a meta-predicate, we can also use the `{}/1` compiler bypass control construct. For example:

```
foo :-
    ...,
    {bar},
    ...
```

But note that in this case the *reflection API* will not record the dependency of the `foo/0` predicate on the Prolog `bar/0` predicate as we are effectively bypassing the compiler.

Calling Prolog module predicates

Prolog module predicates can be called from within objects or categories by using explicit qualification. For example:

```
foo :-
    ...,
    module:bar,
    ...
```

You can also use the `use_module/2` directive to call the module predicates using implicit qualification:

```
:- use_module(module, [bar/0]).

foo :-
    ...,
    bar,
    ...
```

Note that the first argument of the `use_module/2` directive, when used within an object or a category, is a *module name*, not a *file specification* (also be aware that Prolog modules are sometimes defined in files with names that differ from the module names).

As loading a Prolog module varies between Prolog systems, the actual loading directive or goal is preferably done from the application *loader file*. An advantage of this approach is that it contributes to a clean separation between *loading* and *using* a resource, with the loader file being the central point that loads all application resources (complex applications often use a *hierarchy* of loader files, but the main idea remains the same).

As an example, assume that we need to call predicates defined in a CLP(FD) Prolog library, which can be loaded using `library(clpfd)` as the file specification. In the loader file, we would add:

```
:- use_module(library(clpfd), []).
```

Specifying an empty import list is often used to avoid adding the module-exported predicates to plain Prolog. In the objects and categories we can then call the library predicates, using implicit or explicit qualification, as explained. For example:

```
:- object(puzzle).

:- public(puzzle/1).

:- use_module(clpfd, [
    all_different/1, ins/2, label/1,
    (#=)/2, (#\=)/2,
    op(700, xfx, #=), op(700, xfx, #\=)
]).

puzzle([S,E,N,D] + [M,O,R,E] = [M,O,N,E,Y]) :-
    Vars = [S,E,N,D,M,O,R,Y],
    Vars ins 0..9,
    all_different(Vars),
    S*1000 + E*100 + N*10 + D +
    M*1000 + O*100 + R*10 + E #=
    M*10000 + O*1000 + N*100 + E*10 + Y,
    M #\= 0, S #\= 0,
```

(continues on next page)

(continued from previous page)

```
label([M,0,N,E,Y]).

:- end_object.
```

Warning

The actual module code **must** be loaded prior to the compilation of Logtalk source code that uses it. In particular, programmers should not expect that the module be auto-loaded (including when using a backend Prolog compiler that supports an auto-loading mechanism).

The module identifier argument can also be a *parameter variable* when using the directive in a parametric object or a parametric category. In this case, dynamic binding will necessarily be used for all listed predicates (and non-terminals). The parameter variable must be instantiated at runtime when the calls are made.

Logtalk supports the declaration of *predicate aliases* and *predicate shorthands* in `use_module/2` directives used within objects and categories. For example, the ECLiPSe IC Constraint Solvers define a `(:)/2` variable domain operator that clashes with the Logtalk `(:)/2` message-sending operator. We can solve the conflict by writing:

```
:- use_module(ic, [(:)/2 as ins/2]).
```

With this directive, calls to the `ins/2` predicate alias will be automatically compiled by Logtalk to calls to the `(:)/2` predicate in the `ic` module.

Calling Prolog module meta-predicates

The Logtalk library provides implementations of common meta-predicates, which can be used in place of module meta-predicates (e.g., list mapping meta-predicates). If that is not the case, the Logtalk compiler may need help to understand the module meta-predicate templates. Despite some recent progress in standardization of the syntax of `meta_predicate/1` directives and of the `meta_predicate/1` property returned by the `predicate_property/2` reflection predicate, portability is still a major problem. Thus, Logtalk allows the original `meta_predicate/1` directive to be **overridden** with a local directive that Logtalk can make sense of. It also allows providing a `meta_predicate/1` directive when it's missing from the module defining the meta-predicate. Note that Logtalk is not based on a predicate prefixing mechanism as found in module systems. This fundamental difference precludes an automated solution at the Logtalk compiler level.

As an example, assume that you want to call from an object (or a category) a module meta-predicate with the following meta-predicate directive:

```
:- module(foo, [bar/2]).

:- meta_predicate(bar(*, :)).
```

The `:` meta-argument specifier is ambiguous. It tells us that the second argument of the meta-predicate is module sensitive, but it does not tell us *how*. Some legacy module libraries and some Prolog systems use `:` to mean `0` (i.e., a meta-argument that will be meta-called). Some others use `:` for meta-arguments that are not meta-called but that still need to be augmented with module information. Whichever the case, the Logtalk compiler doesn't have enough information to unambiguously parse the directive and correctly compile the meta-arguments in the meta-predicate call. Therefore, the Logtalk compiler will generate an error stating that `:` is not a valid meta-argument specifier when trying to compile a `foo:bar/2` goal. There are two alternative solutions for this problem. The advised solution is to override the meta-predicate directive by writing, inside the object (or category) where the meta-predicate is called:

```
:- meta_predicate(foo:bar(*, *)).
```

or:

```
:- meta_predicate(foo:bar(*, 0)).
```

depending on the true meaning of the second meta-argument. The second alternative, only usable when the meta-argument can be handled as a normal argument, is to simply use the `{}/1` compiler bypass control construct to call the meta-predicate as-is:

```
... :- {foo:bar(..., ...)}, ...
```

The downside of this alternative is that it hides the dependency on the module library from the reflection API and thus from the developer tools.

1.8.9 Defining Prolog multifile predicates

Some Prolog module libraries, e.g. constraint packages, expect clauses for some library predicates to be defined in other modules. This is accomplished by declaring the library predicate *multifile* and by explicitly prefixing predicate clause heads with the library module identifier. For example:

```
:- multifile(clpfd:run_propagator/2).
clpfd:run_propagator(..., ...) :-
    ...
```

Logtalk supports the definition of Prolog module multifile predicates in objects and categories. While the clause head is compiled as-is, the clause body is compiled in the same way as a regular object or category predicate, thus allowing calls to local object or category predicates. For example:

```
:- object(...).

    :- multifile(clpfd:run_propagator/2).
    clpfd:run_propagator(..., ...) :-
        % calls to local object predicates
        ...

:- end_object.
```

The Logtalk compiler will print a warning if the `multifile/1` directive is missing. These multifile predicates may also be declared dynamic using the same `Module:Name/Arity` notation.

1.8.10 Asserting and retracting Prolog predicates

To assert and retract clauses for Prolog dynamic predicates, we can use an explicitly qualified module argument. For example:

```
:- object(...).

    :- dynamic(m:bar/1).

    foo(X) :-
        retractall(m:bar(_)),
```

(continues on next page)

(continued from previous page)

```

    assertz(m:bar(X)),
    ...

:- end_object.

```

In alternative, we can use *use_module/2* directives to declare the module predicates. For example:

```

:- object(...).

:- use_module(m, [bar/1]).
:- dynamic(m:bar/1).

foo(X) :-
    % retract and assert bar/1 clauses in module m
    retractall(bar(_)),
    assertz(bar(X)),
    ...

:- end_object.

```

When the Prolog dynamic predicates are defined in user, the recommended and most portable practice (as not all backends support a module system) is to use a *uses/2* directive:

```

:- object(...).

:- uses(user, [bar/1]).
:- dynamic(user::bar/1).

foo(X) :-
    % retract and assert bar/1 clauses in user
    retractall(bar(_)),
    assertz(bar(X)),
    ...

:- end_object.

```

Note that in the alternatives using *uses/2* or *use_module/2* directives, the argument of the database handling predicates must be known at compile-time. If that is not the case, you must use either an explicitly-qualified argument or the *{}/1* control construct instead. For example:

```

:- object(...).

add(X) :-
    % assert clause X in module m
    assertz(m:X),
    ...

remove(Y) :-
    % retract all clauses in user whose head unifies with Y
    {retractall(Y)},
    ...

:- end_object.

```

1.9 Inheritance

The inheritance mechanisms found in object-oriented programming languages allow the specialization of previously defined objects, avoiding the unnecessary repetition of code and allowing the definition of common functionality for sets of objects. In the context of logic programming, we can interpret inheritance as a form of *theory extension*: an object will virtually contain, besides its own predicates, all the predicates inherited from other objects that are not redefined locally. Inheritance is not, however, the only mechanism for theory extension. Logtalk also supports *composition* using *categories*.

Logtalk uses a depth-first lookup procedure for finding predicate declarations and predicate definitions, as explained below, when a *message* is sent to an object. The lookup procedures locate the entity holding the *predicate declaration* and the entity holding the *predicate definition* using the predicate name and arity. The *alias/2* predicate directive may be used for defining alternative names for inherited predicates, for solving inheritance conflicts, and for giving access to all inherited definitions (thus overriding the default lookup procedure).

The lookup procedures are used when sending a message (using the *(::)/2*, *(:)/1*, and *[]/1* control constructs) and when making *super* calls (using the *(^)/1* control construct). The exact details of the lookup procedures depend on the *role* played by the object receiving the message or making the *super* call, as explained next. The lookup procedures are also used by the *current_predicate/1* and *predicate_property/2* reflection predicates.

1.9.1 Protocol inheritance

Protocol inheritance refers to the inheritance of predicate declarations (*scope directives*). These can be contained in objects, protocols, or categories. Logtalk supports single and multi-inheritance of protocols: an object or a category may implement several protocols, and a protocol may extend several protocols.

Lookup order for prototype hierarchies

The lookup order for predicate declarations is first the object, second the implemented protocols (and the protocols that these may extend), third the imported categories (and the protocols that they may implement), and finally the objects that the object extends (following their declaration order). This lookup is performed in depth-first order. When an object inherits two different declarations for the same predicate, by default, only the first one will be considered.

Lookup order for class hierarchies

The lookup order for predicate declarations is first the object classes (following their declaration order), second the classes implemented protocols (and the protocols that these may extend), third the classes imported categories (and the protocols that they may implement), and finally the superclasses of the object classes. This lookup is performed in depth-first order. If the object inherits two different declarations for the same predicate, by default, only the first one will be considered.

1.9.2 Implementation inheritance

Implementation inheritance refers to the inheritance of predicate definitions. These can be contained in objects or in categories. Logtalk supports multi-inheritance of implementation: an object may import several categories or extend, specialize, or instantiate several objects.

Lookup order for prototype hierarchies

The lookup order for predicate definitions is similar to the lookup for predicate declarations, except that implemented protocols are ignored (as they can only contain predicate directives).

Lookup order for class hierarchies

The lookup order for predicate definitions is similar to the lookup for predicate declarations, except that implemented protocols are ignored (as they can only contain predicate directives) and that the lookup starts at the instance itself (that received the message) before proceeding, if no predicate definition is found there, to the instance classes imported categories and then to the class superclasses.

Redefining inherited predicate definitions

When we define a predicate that is already inherited from an ancestor object or an imported category, the inherited definition is hidden by the new definition. This is called inheritance overriding: a local definition overrides any inherited definitions. For example, assume that we have the following two objects:

```
:- object(root).

    :- public(bar/1).
    bar(root).

    :- public(foo/1).
    foo(root).

:- end_object.

:- object(descendant,
    extends(root)).

    foo(descendant).

:- end_object.
```

After compiling and loading these objects, we can check the overriding behavior by trying the following queries:

```
| ?- root::(bar(Bar), foo(Foo)).

Bar = root
Foo = root
yes
```

(continues on next page)

(continued from previous page)

```
| ?- descendant::(bar(Bar), foo(Foo)).

Bar = root
Foo = descendant
yes
```

However, we can explicitly code other behaviors. Some examples follow.

Specializing inherited predicate definitions

Specialization of inherited definitions: the new definition calls the inherited definition and makes additional calls. This is accomplished by calling the $(\wedge\wedge)/1$ *super call* operator in the new definition. For example, assume a `init/0` predicate that must account for object specific initializations along the inheritance chain:

```
:- object(root).

    :- public(init/0).

    init :-
        write('root init'), nl.

:- end_object.

:- object(descendant,
    extends(root)).

    init :-
        write('descendant init'), nl,
        ^^init.

:- end_object.
```

```
| ?- descendant::init.

descendant init
root init
yes
```

Union of inherited and local predicate definitions

Union of the new with the inherited definitions: all the definitions are taken into account, the calling order being defined by the inheritance mechanisms. This can be accomplished by writing a clause that just calls, using the $(\wedge\wedge)/1$ *super call* operator, the inherited definitions. The relative position of this clause among the other definition clauses sets the calling order for the local and inherited definitions. For example:

```
:- object(root).

    :- public(foo/1).
```

(continues on next page)

(continued from previous page)

```

foo(1).
foo(2).

:- end_object.

:- object(descendant,
    extends(root)).

    foo(3).
    foo(Foo) :-
        ^^foo(Foo).

:- end_object.

```

```

| ?- descendant::foo(Foo).

Foo = 3 ;
Foo = 1 ;
Foo = 2 ;
no

```

Selective inheritance of predicate definitions

Selective inheritance of predicate definitions (also known as differential inheritance) is normally used in the representation of exceptions to inherited default definitions. We can use the *(^^)/1 super call* operator to test and possibly reject some of the inherited definitions. A common example is representing flightless birds:

```

:- object(bird).

    :- public(mode/1).

    mode(walks).
    mode(flies).

:- end_object.

:- object(penguin,
    extends(bird)).

    mode(swims).
    mode(Mode) :-
        ^^mode(Mode),
        Mode \== flies.

:- end_object.

```

```

| ?- penguin::mode(Mode).

```

(continues on next page)

(continued from previous page)

```
Mode = swims ;
Mode = walks ;
no
```

1.9.3 Public, protected, and private inheritance

To make all *public predicates* declared via implemented protocols, imported categories, or ancestor objects *protected predicates* or to make all public and protected predicates *private predicates*, we prefix the entity's name with the corresponding keyword. For example:

```
:- object(Object,
    implements(private::Protocol)).

    % all the Protocol public and protected
    % predicates become private predicates
    % for the Object clients

    ...

:- end_object.
```

or:

```
:- object(Class,
    specializes(protected::Superclass)).

    % all the Superclass public predicates become
    % protected predicates for the Class clients

    ...

:- end_object.
```

Omitting the scope keyword is equivalent to using the public scope keyword. For example:

```
:- object(Object,
    imports(public::Category)).

    ...

:- end_object.
```

This is the same as:

```
:- object(Object,
    imports(Category)).

    ...

:- end_object.
```

This way we ensure backward compatibility with older Logtalk versions and a simplified syntax when protected or private inheritance is not used.

1.9.4 Multiple inheritance

Logtalk supports multiple inheritance by enabling an object to extend, instantiate, or specialize more than one object. Likewise, a protocol may extend multiple protocols, and a category may extend multiple categories. In this case, the depth-first lookup algorithms described above traverse the list of entities per relation from left to right. Consider as an example the following object opening directive:

```
:- object(foo,
    extends((bar, baz))).
```

The lookup procedure will look first into the parent object *bar* and its related entities before looking into the parent object *baz*. The *alias/2* predicate directive can always be used to solve multiple inheritance conflicts. It should also be noted that the multi-inheritance support does not affect performance when we use single inheritance.

1.9.5 Composition versus multiple inheritance

It is not possible to discuss inheritance mechanisms without referring to the long and probably endless debate on single versus multiple inheritance. The single inheritance mechanism can be implemented efficiently, but it imposes several limitations on reusing, even if the multiple characteristics we intend to inherit are orthogonal. On the other hand, the multiple inheritance mechanisms are attractive in their apparent capability of modeling complex situations. However, they include a potential for conflict between inherited definitions whose variety does not allow a single and satisfactory solution for all the cases.

No solution that we might consider satisfactory for all the problems presented by the multiple inheritance mechanisms has been found. From the simplicity of some extensions that use the Prolog search strategy, such as [McCabe92] or [Moss94], to the sophisticated algorithms of CLOS [Bobrow_et_al_88], there is no adequate solution for all the situations. Besides, the use of multiple inheritance carries some complex problems in the domain of software engineering, particularly in the reuse and maintenance of the applications. All these problems are substantially reduced if we preferably use in our software development composition mechanisms instead of specialization mechanisms [Taenzer89]. Multiple inheritance is best used as an analysis and project abstraction, rather than as an implementation technique [Shan_et_al_93]. Note that Logtalk provides first-class support for composition using *categories*.

1.10 Event-driven programming

The addition of event-driven programming capacities to the Logtalk language [Moura94] is based on a simple but powerful idea:

The computations must result not only from message-sending but also from the **observation** of message-sending.

The need to associate computations to the occurrence of events was very early recognized in knowledge representation languages, programming languages [Stefik_et_al_86], [Moon86], operative systems [Tanenbaum87], and graphical user interfaces.

With the integration between object-oriented and event-driven programming, we intend to achieve the following goals:

- Minimize the *coupling* between objects. An object should only contain what is intrinsic to it. If an object observes another object, that means that it should depend only on the public protocol of the object observed and not on the implementation of that protocol.
- Provide a mechanism for building *reflexive systems* in Logtalk based on the dynamic behavior of objects in complement to the reflective information on object predicates and relations.
- Provide a mechanism for easily defining method *pre- and post-conditions* that can be toggled using the *events* compiler flag. The pre- and post-conditions may be defined in the same object containing the methods or distributed between several objects acting as method monitors.
- Provide a *publish-subscribe* mechanism where public messages play the role of events.

1.10.1 Definitions

The words *event* and *monitor* have multiple meanings in computer science. To avoid misunderstandings, we start by defining them in the Logtalk context.

Event

In an object-oriented system, all computations start through message sending. It thus becomes quite natural to declare that the only event that can occur in this kind of system is precisely the sending of a message. An event can thus be represented by the ordered tuple (Object, Message, Sender).

If we consider message processing an indivisible activity, we can interpret the sending of a message and the return of the control to the object that has sent the message as two distinct events. This distinction allows us to have more precise control over a system's dynamic behavior. In Logtalk, these two types of events have been named *before* and *after*, respectively for sending a message and for returning control to the sender. Therefore, we refine our event representation using the ordered tuple (Event, Object, Message, Sender).

The implementation of events in Logtalk enjoys the following properties:

Independence between the two types of events

We can choose to watch only one event type or to process each one of the events associated with a message-sending goal in an independent way.

All events are automatically generated by the message-sending mechanism

The task of generating events is transparently accomplished by the message-sending mechanism. The user only needs to define the events that will be monitored.

The events watched at any moment can be dynamically changed during program execution

The notion of event allows the user not only to have the possibility of observing but also of controlling and modifying an application behavior, namely by dynamically changing the observed events during program execution. It is our goal to provide the user with the possibility of modeling the largest number of situations.

Monitor

Complementary to the notion of event is the notion of monitor. A monitor is an object that is automatically notified by the message-sending mechanism whenever a registered event occurs. Any object that defines the event-handling predicates can play the role of a monitor.

The implementation of monitors in Logtalk enjoys the following properties:

Any object can act as a monitor

The monitor status is a role that any object can perform during its existence. The minimum protocol necessary is declared in the built-in `monitoring` protocol. Strictly speaking, the reference to this protocol is only needed when specializing event handlers. Nevertheless, it is considered good programming practice to always refer to the protocol when defining event handlers.

Unlimited number of monitors for each event

Several monitors can observe the same event for distinct reasons. Therefore, the number of monitors per event is bounded only by the available computing resources.

The monitor status of an object can be dynamically changed at runtime

This property does not imply that an object must be dynamic to act as a monitor (the monitor status of an object is not stored in the object).

Event handlers cannot modify the event arguments

Notably, if the message contains unbound variables, these cannot be bound by the calls to the monitor event handlers.

1.10.2 Event generation

Assuming that the `events` flag is set to allow for the object (or category) sending the messages we want to observe, for each message that is sent using the `(::)/2` control construct, the runtime system automatically generates two events. The first — *before event* — is generated when the message is sent. The second — *after event* — is generated after the message has successfully been executed.

Note that *self* messages (using the `(::)/1` control construct) and *super* calls (using the `(^^)/1` control construct) don't generate events.

1.10.3 Communicating events to monitors

Whenever a spied event occurs, the message-sending mechanism calls the corresponding event handlers directly for all registered monitors. These calls are internally made, thus bypassing the message-sending primitives in order to avoid potential endless loops. The event handlers consist of user definitions for the public predicates declared in the built-in `monitoring` protocol (see below for more details).

1.10.4 Performance concerns

Ideally, the existence of monitored messages should not affect the processing of the remaining messages. On the other hand, for each message that has been sent, the system must verify if its respective event is monitored. Whenever possible, this verification should be performed in constant time and independently of the number of monitored events. The representation of events takes advantage of the first argument indexing performed by most Prolog compilers, which ensure — in the general case — access in constant time.

Event support can be turned off on a per-object (or per-category) basis using the `events` compiler flag. With event support turned off, Logtalk uses optimized code for processing message-sending calls that skips the checking of monitored events, resulting in a small but measurable performance improvement.

1.10.5 Monitor semantics

The established semantics for monitor actions consists of considering its success as a necessary condition so that a message can succeed:

- All actions associated with events of type *before* must succeed so that the message processing can start.
- All actions associated with events of type *after* also have to succeed so that the message itself succeeds. The failure of any action associated with an event of type *after* forces backtracking over the message execution (the failure of a monitor never causes backtracking over the preceding monitor actions).

Note that this is the most general choice. If we require a transparent presence of monitors in a message processing, we just have to define the monitor actions in such a way that they never fail (which is very simple to accomplish).

1.10.6 Activation order of monitors

Ideally, whenever there are several monitors defined for the same event, the calling order should not interfere with the result. However, this is not always possible. In the case of an event of type *before*, the failure of a monitor prevents a message from being sent and prevents the execution of the remaining monitors. In the case of an event of type *after*, a monitor failure will force backtracking over message execution. Different orders of monitor activation can therefore lead to different results if the monitor actions imply object modifications unrecoverable in case of backtracking. Therefore, the order for monitor activation should be assumed as arbitrary. In effect, to assume or to try to impose a specific sequence requires a global knowledge of an application dynamics, which is not always possible. Furthermore, that knowledge can reveal itself as incorrect if there is any change in the execution conditions. Note that, given the independence between monitors, it does not make sense that a failure forces backtracking over the actions previously executed.

1.10.7 Event handling

Logtalk provides three built-in predicates for event handling. These predicates support defining, enumerating, and abolishing events. Applications that use events extensively usually define a set of objects that use these built-in predicates to implement more sophisticated and higher-level behavior.

Defining new events

New events can be defined using the *define_events/5* built-in predicate:

```
| ?- define_events(Event, Object, Message, Sender, Monitor).
```

Note that if any of the *Event*, *Object*, *Message*, and *Sender* arguments is a free variable or contains free variables, this call will define a **set** of matching events.

Abolishing defined events

Events that are no longer needed may be abolished using the *abolish_events/5* built-in predicate:

```
| ?- abolish_events(Event, Object, Message, Sender, Monitor).
```

If called with free variables, this goal will remove all matching events.

Finding defined events

The events that are currently defined can be retrieved using the *current_event/5* built-in predicate:

```
| ?- current_event(Event, Object, Message, Sender, Monitor).
```

Note that this predicate will return **sets** of matching events if some of the returned arguments are free variables or contain free variables.

Defining event handlers

The *monitoring* built-in protocol declares two public predicates, *before/3* and *after/3*, that are automatically called to handle before and after events. Any object that plays the role of monitor must define one or both of these event handler methods:

```
before(Object, Message, Sender) :-
    ...

after(Object, Message, Sender) :-
    ...
```

The arguments in both methods are instantiated by the message-sending mechanism when a monitored event occurs. For example, assume that we want to define a monitor called *tracer* that will track any message sent to an object by printing a descriptive text to the standard output. Its definition could be something like:

```
:- object(tracer,
    % built-in protocol for event handler methods
    implements(monitoring)).

    before(Object, Message, Sender) :-
        write('call: '), writeq(Object),
        write(' <-- '), writeq(Message),
        write(' from '), writeq(Sender), nl.

    after(Object, Message, Sender) :-
        write('exit: '), writeq(Object),
        write(' <-- '), writeq(Message),
        write(' from '), writeq(Sender), nl.

:- end_object.
```

Assume that we also have the following object:

```
:- object(any).  
  
    :- public(bar/1).  
    bar(bar).  
  
    :- public(foo/1).  
    foo(foo).  
  
:- end_object.
```

After compiling and loading both objects and setting the *events* flag to allow, we can start tracing every message sent to any object by calling the *define_events/5* built-in predicate:

```
| ?- set_logtalk_flag(events, allow).  
  
yes  
  
| ?- define_events(_, _, _, _, tracer).  
  
yes
```

From now on, every message sent from user to any object will be traced to the standard output stream:

```
| ?- any::bar(X).  
  
call: any <-- bar(X) from user  
exit: any <-- bar(bar) from user  
X = bar  
  
yes
```

To stop tracing, we can use the *abolish_events/5* built-in predicate:

```
| ?- abolish_events(_, _, _, _, tracer).  
  
yes
```

The *monitoring* protocol declares the event handlers as public predicates. If necessary, *protected or private implementation of the protocol* may be used in order to change the scope of the event handler predicates. Note that the message-sending processing mechanism is able to call the event handlers irrespective of their scope. Nevertheless, the scope of the event handlers may be restricted in order to prevent other objects from calling them.

The pseudo-object *user* can also act as a monitor. This object expects the *before/3* and *after/3* predicates to be defined in the plain Prolog database. To avoid predicate existence errors when setting *user* as a monitor, this object declares the predicates multifile. Thus, any plain Prolog code defining the predicates should include the directives:

```
:- multifile(before/3).  
:- multifile(after/3).
```

1.11 Multi-threading programming

Logtalk provides **experimental** support for multi-threading programming on selected Prolog compilers. Logtalk makes use of the low-level Prolog built-in predicates that implement message queues and interface with POSIX threads and mutexes (or a suitable emulation), providing a small set of high-level predicates and directives that allows programmers to easily take advantage of modern multi-processor and multi-core computers without worrying about the tricky details of creating, synchronizing, or communicating with threads, mutexes, and message queues. Logtalk multi-threading programming integrates with object-oriented programming by providing a *threaded engines* API, enabling objects and categories to prove goals concurrently, and supporting synchronous and asynchronous messages.

1.11.1 Enabling multi-threading support

Multi-threading support may be disabled by default. It can be enabled on the Prolog adapter files of supported compilers by setting the read-only *threads* compiler flag to supported.

1.11.2 Enabling objects to make multi-threading calls

The *threaded/0* object directive is used to enable an object to make multi-threading calls:

```
:- threaded.
```

1.11.3 Multi-threading built-in predicates

Logtalk provides a small set of built-in predicates for multi-threading programming. For simple tasks where you simply want to prove a set of goals, each one in its own thread, Logtalk provides a *threaded/1* built-in predicate. The remaining predicates allow for fine-grained control, including postponing retrieval of thread goal results at a later time, supporting non-deterministic thread goals, and making *one-way* asynchronous calls. Together, these predicates provide high-level support for multi-threading programming, covering most common use cases.

Proving goals concurrently using threads

A set of goals may be proved concurrently by calling the Logtalk built-in predicate *threaded/1*. Each goal in the set runs in its own thread.

When the *threaded/1* predicate argument is a *conjunction* of goals, the predicate call is akin to *and-parallelism*. For example, assume that we want to find all the prime numbers in a given interval, $[N, M]$. We can split the interval into two parts and then span two threads to compute the prime numbers in each sub-interval:

```
prime_numbers(N, M, Primes) :-
    M > N,
    N1 is N + (M - N) // 2,
    N2 is N1 + 1,
    threaded((
        prime_numbers(N2, M, [], Acc),
        prime_numbers(N, N1, Acc, Primes)
    )).
```

(continues on next page)

(continued from previous page)

```
prime_numbers(N, M, Acc, Primes) :-
    ...
```

The `threaded/1` call terminates when the two implicit threads terminate. In a computer with two or more processors (or with a processor with two or more cores), the code above can be expected to provide better computation times when compared with single-threaded code for sufficiently large intervals.

When the `threaded/1` predicate argument is a *disjunction* of goals, the predicate call is akin to *or-parallelism*, here reinterpreted as a set of goals *competing* to find a solution. For example, consider the different methods that we can use to find the roots of real functions. Depending on the function, some methods will be faster than others. Some methods will converge to the solution while others may diverge and never find it. We can try all the methods simultaneously by writing:

```
find_root(Function, A, B, Error, Zero) :-
    threaded((
        bisection::find_root(Function, A, B, Error, Zero)
    ;   newton::find_root(Function, A, B, Error, Zero)
    ;   muller::find_root(Function, A, B, Error, Zero)
    )).
```

The above `threaded/1` goal succeeds when one of the implicit threads succeeds in finding the function root, leading to the termination of all the remaining competing threads.

The `threaded/1` built-in predicate is most useful for lengthy, independent, deterministic computations where the computational costs of each goal outweigh the overhead of the implicit thread creation and management.

Proving goals asynchronously using threads

A goal may be proved asynchronously using a new thread by calling the `threaded_call/1-2` built-in predicate. Calls to this predicate are always true and return immediately (assuming a callable argument). The term representing the goal is copied, not shared with the thread. The thread computes the first solution to the goal, posts it to the implicit message queue of the object from where the `threaded_call/1` predicate was called, and suspends waiting for either a request for an alternative solution or for the program to commit to the current solution.

The results of proving a goal asynchronously in a new thread may be later retrieved by calling the `threaded_exit/1-2` built-in predicate within the same object where the call to the `threaded_call/1` predicate was made. The `threaded_exit/1` calls suspend execution until the results of the `threaded_call/1` calls are sent back to the object message queue.

The `threaded_exit/1` predicate allows us to retrieve alternative solutions through backtracking (if you want to commit to the first solution, you may use the `threaded_once/1-2` predicate instead of the `threaded_call/1` predicate). For example, assuming a lists object implementing the usual `member/2` predicate, we could write:

```
| ?- threaded_call(lists::member(X, [1,2,3])).
X = _G189
yes

| ?- threaded_exit(lists::member(X, [1,2,3])).

X = 1 ;
X = 2 ;
```

(continues on next page)

(continued from previous page)

```
X = 3 ;
no
```

In this case, the `threaded_call/1` and the `threaded_exit/1` calls are made within the pseudo-object user. The implicit thread running the `lists::member/2` goal suspends itself after providing a solution, waiting for a request for an alternative solution; the thread is automatically terminated when the runtime engine detects that backtracking to the `threaded_exit/1` call is no longer possible.

Calls to the `threaded_exit/1` predicate block the caller until the object message queue receives the reply to the asynchronous call. The predicate `threaded_peek/1-2` may be used to check if a reply is already available without removing it from the thread queue. The `threaded_peek/1` predicate call succeeds or fails immediately without blocking the caller. However, keep in mind that repeated use of this predicate is equivalent to polling a message queue, which may hurt performance.

Be careful when using the `threaded_exit/1` predicate inside failure-driven loops. When all the solutions have been found (and the thread generating them is therefore terminated), re-calling the predicate will generate an exception. Note that failing instead of throwing an exception is not an acceptable solution, as it could be misinterpreted as a failure of the `threaded_call/1` argument.

The example in the previous section with prime numbers could be rewritten using the `threaded_call/1` and `threaded_exit/1` predicates:

```
prime_numbers(N, M, Primes) :-
    M > N,
    N1 is N + (M - N) // 2,
    N2 is N1 + 1,
    threaded_call(prime_numbers(N2, M, [], Acc)),
    threaded_call(prime_numbers(N, N1, Acc, Primes)),
    threaded_exit(prime_numbers(N2, M, [], Acc)),
    threaded_exit(prime_numbers(N, N1, Acc, Primes)).

prime_numbers(N, M, Acc, Primes) :-
    ...
```

When using asynchronous calls, the link between a `threaded_exit/1` call and the corresponding `threaded_call/1` call is established using unification. If there are multiple `threaded_call/1` calls for a matching `threaded_exit/1` call, the connection can potentially be established with any of them (this is akin to what happens with tabling). Nevertheless, you can easily use a call *tag* by using the alternative `threaded_call/2`, `threaded_once/2`, and `threaded_exit/2` built-in predicates. For example:

```
?- threaded_call(member(X, [1,2,3]), Tag).

Tag = 1
yes

?- threaded_call(member(X, [1,2,3]), Tag).

Tag = 2
yes

?- threaded_exit(member(X, [1,2,3]), 2).

X = 1 ;
X = 2 ;
```

(continues on next page)

(continued from previous page)

```
X = 3
yes
```

When using these predicates, the tags shall be considered as an opaque term; users shall not rely on its type. Tagged asynchronous calls can be canceled by using the `threaded_cancel/1` predicate.

1.11.4 One-way asynchronous calls

Sometimes we want to prove a goal in a new thread without caring about the results. This may be accomplished by using the built-in predicate `threaded_ignore/1`. For example, assume that we are developing a multi-agent application where an agent may send a “happy birthday” message to another agent. We could write:

```
..., threaded_ignore(agent::happy_birthday), ...
```

The call succeeds with no reply of the goal success, failure, or even exception ever being sent back to the object making the call. Note that this predicate implicitly performs a deterministic call of its argument.

1.11.5 Asynchronous calls and synchronized predicates

Proving a goal asynchronously using a new thread may lead to problems when the goal results in side effects such as input/output operations or modifications to an *object database*. For example, if a new thread is started with the same goal before the first one finished its job, we may end up with mixed output, a corrupted database, or unexpected goal failures. In order to solve this problem, predicates (and grammar rule non-terminals) with side effects can be declared as *synchronized* by using the `synchronized/1` predicate directive. Proving a query to a synchronized predicate (or synchronized non-terminal) is internally protected by a mutex, thus allowing for easy thread synchronization. For example:

```
% ensure thread synchronization
:- synchronized(db_update/1).

db_update(Update) :-
    % predicate with side-effects
    ...
```

A second example: assume an object defining two predicates for writing, respectively, even and odd numbers in a given interval to the standard output. Given a large interval, a goal such as:

```
| ?- threaded_call(obj::odd_numbers(1,100)),
    threaded_call(obj::even_numbers(1,100)).

1 3 2 4 6 8 5 7 10 ...
...
```

will most likely result in a mixed-up output. By declaring the `odd_numbers/2` and `even_numbers/2` predicates synchronized:

```
:- synchronized([
    odd_numbers/2,
    even_numbers/2]).
```

one goal will only start after the other one finished:

```
| ?- threaded_ignore(obj::odd_numbers(1,99)),
   threaded_ignore(obj::even_numbers(1,99)).

1 3 5 7 9 11 ...
...
2 4 6 8 10 12 ...
...
```

Note that, in a more realistic scenario, the two `threaded_ignore/1` calls would be made concurrently from different objects. Using the same synchronized directive for a set of predicates implies that they all use the same mutex, as required for this example.

As each Logtalk entity is independently compiled, this directive must be included in every object or category that contains a definition for the described predicate, even if the predicate declaration is inherited from another entity, in order to ensure proper compilation. Note that a synchronized predicate cannot be declared dynamic. To ensure atomic updates of a dynamic predicate, declare as synchronized the predicate performing the update.

Synchronized predicates may be used as wrappers for messages sent to objects that are not multi-threading aware. For example, assume a log object defining a `write_log_entry/2` predicate that writes log entries to a file, thus using side effects on its implementation. We can specify and define, for example, a `sync_write_log_entry/2` predicate as follows:

```
:- synchronized(sync_write_log_entry/2).

sync_write_log_entry(File, Entry) :-
    log::write_log_entry(File, Entry).
```

and then call the `sync_write_log_entry/2` predicate instead of the `write_log_entry/2` predicate from multi-threaded code.

The synchronization directive may be used when defining objects that may be reused in both single-threaded and multi-threaded Logtalk applications. The directive simply makes calls to the synchronized predicates deterministic when the objects are used in a single-threaded application.

1.11.6 Synchronizing threads through notifications

Declaring a set of predicates as synchronized can only ensure that they are not executed at the same time by different threads. Sometimes we need to suspend a thread not on a synchronization lock but on some condition that must hold true for a thread goal to proceed. I.e. we want a thread goal to be suspended until a condition becomes true instead of simply failing. The built-in predicate `threaded_wait/1` allows us to suspend a predicate execution (running in its own thread) until a notification is received. Notifications are posted using the built-in predicate `threaded_notify/1`. A notification is a Prolog term that a programmer chooses to represent some condition becoming true. Any Prolog term can be used as a notification argument for these predicates. Related calls to the `threaded_wait/1` and `threaded_notify/1` must be made within the same object, *this*, as the object message queue is used internally for posting and retrieving notifications.

Each notification posted by a call to the `threaded_notify/1` predicate is consumed by a single `threaded_wait/1` predicate call (i.e., these predicates implement a peer-to-peer mechanism). Care should be taken to avoid deadlocks when two (or more) threads both wait and post notifications to each other.

1.11.7 Threaded engines

Threaded engines provide an alternative to the multi-threading predicates described in the previous sections. An *engine* is a computing thread whose solutions can be lazily computed and retrieved. In addition, an engine also supports a term queue that allows passing arbitrary terms to the engine.

An engine is created by calling the `threaded_engine_create/3` built-in predicate. For example:

```
| ?- threaded_engine_create(X, member(X, [1,2,3]), worker).
yes
```

The first argument is an *answer template* to be used for retrieving solution bindings. The user can name the engine, as in this example where the atom `worker` is used, or have the runtime generate a name, which should be treated as an opaque term.

Engines are scoped by the object within which the `threaded_engine_create/3` call takes place. Thus, different objects can create engines with the same names with no conflicts. Moreover, engines share the visible predicates of the object creating them.

The engine computes the first solution of its goal argument and suspends waiting for it to be retrieved. Solutions can be retrieved one at a time using the `threaded_engine_next/2` built-in predicate:

```
| ?- threaded_engine_next(worker, X).
X = 1
yes
```

The call blocks until a solution is available and fails if there are no solutions left. After returning a solution, this predicate signals the engine to start computing the next one. Note that this predicate is deterministic. In contrast with the `threaded_exit/1-2` built-in predicates, retrieving the next solution requires calling the predicate again instead of backtracking into its call. For example:

```
collect_all(Engine, [Answer| Answers]) :-
    threaded_engine_next(Engine, Answer),
    !,
    collect_all(Engine, Answers).
collect_all(_, []).
```

There is also a reified alternative version of the predicate, `threaded_engine_next_reified/2`, which returns the(`Answer`), `no`, and `exception(Error)` terms as answers. Using this predicate, collecting all solutions to an engine uses a different programming pattern:

```
... :-
    ...,
    threaded_engine_next_reified(Engine, Reified),
    collect_all_reified(Reified, Engine, Answers),
    ...

collect_all_reified(no, _, []).
collect_all_reified(the(Answer), Engine, [Answer| Answers]) :-
    threaded_engine_next_reified(Engine, Reified),
    collect_all_reified(Reified, Engine, Answers).
```

Engines must be explicitly terminated using the `threaded_engine_destroy/1` built-in predicate:

```
| ?- threaded_engine_destroy(worker).
yes
```

A common usage pattern for engines is to define a recursive predicate that uses the engine term queue to retrieve a task to be performed. For example, assume we define the following predicate:

```
loop :-
    threaded_engine_fetch(Task),
    handle(Task),
    loop.
```

The `threaded_engine_fetch/1` built-in predicate fetches a task for the engine term queue. The engine clients would use the `threaded_engine_post/2` built-in predicate to post tasks into the engine term queue. The engine would be created using the call:

```
| ?- threaded_engine_create(none, loop, worker).

yes
```

The `handle/1` predicate, after performing a task, can use the `threaded_engine_yield/1` built-in predicate to make the task results available for consumption using the `threaded_engine_next/2` and `threaded_engine_next_reified/2` built-in predicates. Blocking semantics are used by these two predicates: the `threaded_engine_yield/1` predicate blocks until the returned solution is consumed, while the `threaded_engine_next/2` predicate blocks until a solution becomes available.

1.11.8 Multi-threading performance

The performance of multi-threading applications is highly dependent on the *backend Prolog compiler*, the operating-system, and the use of *dynamic binding* and dynamic predicates. All compatible backend Prolog compilers that support multi-threading features make use of POSIX threads or *pthreads*. The performance of the underlying pthreads implementation can exhibit significant differences between operating systems. An important point is synchronized access to dynamic predicates. As different threads may try to simultaneously access and update dynamic predicates, these operations may use a lock-free algorithm or be protected by a lock, usually implemented using a mutex. In the latter case, poor mutex lock operating-system performance, combined with a large number of collisions by several threads trying to acquire the same lock, can result in severe performance penalties. Thus, whenever possible, avoid using dynamic predicates and dynamic binding.

1.12 Error handling

Error handling is accomplished in Logtalk by using the standard `catch/3` and `throw/1` predicates [ISO95] together with a set of built-in methods that simplify generating errors decorated with expected context.

Errors thrown by Logtalk have, whenever possible, the following format:

```
error(Error, logtalk(Goal, ExecutionContext))
```

In this exception term, `Goal` is the goal that triggered the error `Error` and `ExecutionContext` is the context in which `Goal` is called. For example:

```
error(
    permission_error(modify,private_predicate,p/0),
    logtalk(foo::abolish(p/0), _)
)
```

Note, however, that `Goal` and `ExecutionContext` can be unbound or only partially instantiated when the corresponding information is not available (e.g., due to compiler optimizations that throw away the necessary error context information). The `ExecutionContext` argument is an opaque term that can be decoded using the `logtalk::execution_context/7` predicate.

1.12.1 Raising Exceptions

The *error handling section* in the reference manual lists a set of convenient built-in methods that generate error/2 exception terms with the expected context argument. For example, instead of manually constructing a type error as in:

```
...,
context(Context),
throw(error(type_error(atom, 42), Context)).
```

we can simply write:

```
...,
type_error(atom, 42).
```

The provided error built-in methods cover all standard error types found in the ISO Prolog Core standard.

1.12.2 Type-checking

One of the most common cases where errors may be generated is when type-checking predicate arguments and input data before processing it. The standard library includes a `type` object that defines an extensive set of types, together with predicates for validating and checking terms. The set of types is user extensible. New types can be defined by adding clauses for the `type/1` and `check/2` multifile predicates. For example, assume that we want to be able to check *temperatures* expressed in Celsius, Fahrenheit, or Kelvin scales. We start by declaring (in an object or category) the new type:

```
:- multifile(type::type/1).
type::type(temperature(_Unit)).
```

Next, we need to define the actual code that would verify that a temperature is valid. As the different scales use a different value for absolute zero, we can write:

```
:- multifile(type::check/2).
type::check(temperature(Unit), Term) :-
    check_temperature(Unit, Term).

% given that temperature has only a lower bound, we make use of the library
% property/2 type to define the necessary test expression for each unit
check_temperature(celsius, Term) :-
    type::check(property(float, [Temperature]>>(Temperature >= -273.15)), Term).
check_temperature(fahrenheit, Term) :-
    type::check(property(float, [Temperature]>>(Temperature >= -459.67)), Term).
check_temperature(kelvin, Term) :-
    type::check(property(float, [Temperature]>>(Temperature >= 0.0)), Term).
```

With this definition, a term is first checked that it is a float value before checking that it is in the expected open interval. But how do we use this new type? If we just want to test if a temperature is valid, we can write:

```
..., type::valid(temperature(celsius), 42.0), ...
```

The `type::valid/2` predicate succeeds or fails depending on the second argument being of the type specified in the first argument. If instead of success or failure we want to generate an error for invalid values, we can use the `type::check/2` predicate instead:

```
..., type::check(temperature(celsius), 42.0), ...
```

If we require an error/2 exception term with the error context, we can use instead the `type::check/3` predicate:

```
...,
context(Context),
type::check(temperature(celsius), 42.0, Context),
...
```

Note that `context/1` calls are inlined and messages to the library type object use *static binding* when compiling with the *optimize flag* turned on, thus enabling efficient type-checking.

1.12.3 Expected terms

Support for representing and handling *expected terms* is provided by the *expecteds* library. Expected terms allow deferring errors to later stages of an application in lieu to raising an exception as soon as an error is detected.

1.12.4 Compiler warnings and errors

The current Logtalk compiler uses the standard `read_term/3` built-in predicate to read and compile a Logtalk source file. This improves the compatibility with *backend Prolog compilers* and their proprietary syntax extensions and standard compliance quirks. But one consequence of this design choice is that invalid Prolog terms or syntax errors may abort the compilation process with limited information given to the user (due to the inherent limitations of the `read_term/3` predicate).

Assuming that all the terms in a source file are valid, there is a set of errors and potential errors, described below, that the compiler will try to detect and report, depending on the used compiler flags (see the *Compiler flags* section of this manual on lint flags for details).

Unknown entities

The Logtalk compiler warns about any referenced entity that is not currently loaded. The warning may reveal a misspelled entity name or just an entity that will be loaded later. Out-of-order loading should be avoided when possible as it prevents some code optimizations, such as *static binding* of messages to methods.

Singleton variables

Singleton variables in a clause are often misspelled variables and, as such, are one of the most common errors when programming in Prolog. Assuming that the *backend Prolog compiler* implementation of the `read_term/3` predicate supports the standard `singletons/1` option, the compiler warns about any singleton variable found while compiling a source file.

Redefinition of Prolog built-in predicates

The Logtalk compiler will warn us of any redefinition of a Prolog built-in predicate inside an object or category. Sometimes the redefinition is intended. In other cases, the user may not be aware that a particular *backend Prolog compiler* may already provide the predicate as a built-in predicate or may want to ensure code portability among several Prolog compilers with different sets of built-in predicates.

Redefinition of Logtalk built-in predicates

Similar to the redefinition of Prolog built-in predicates, the Logtalk compiler will warn us if we try to redefine a Logtalk built-in. But the redefinition will probably be an error in most (if not all) cases.

Redefinition of Logtalk built-in methods

An error will be thrown if we attempt to redefine a Logtalk built-in method inside an entity. The default behavior is to report the error and abort the compilation of the offending entity.

Misspell calls of local predicates

A warning will be reported if Logtalk finds (in the body of a predicate definition) a call to a local predicate that is not defined, built-in (either in Prolog or in Logtalk) or declared dynamic. In most cases these calls are simple misspell errors.

Portability warnings

A warning will be reported if a predicate clause contains a call to a non-standard built-in predicate or arithmetic function. Portability warnings are also reported for non-standard flags or flag values. These warnings often cannot be avoided due to the limited scope of the ISO Prolog standard.

Deprecated elements

A warning will be reported if a deprecated directive, control construct, or predicate is used. These warnings should be fixed as soon as possible, as support for any deprecated features will likely be discontinued in future versions.

Missing directives

A warning will be reported for any missing dynamic, discontinuous, meta-predicate, and public predicate directive.

Duplicated directives

A warning will be reported for any duplicated scope, multifile, dynamic, discontinuous, meta-predicate, and meta-non-terminal directives. Note that conflicting directives for the same predicate are handled as errors, not as duplicated directive warnings.

Duplicated clauses

A warning will be reported for any duplicated entity clauses. This check is computationally heavy, however, and usually turned off by default.

Goals that are always true or false

A warning will be reported for any goal that is always true or false. This is usually caused by typos in the code. For example, writing `X == y` instead of `X == Y`.

Trivial fails

A warning will be reported for any call to a local static predicate with no matching clause.

Suspicious calls

A warning will be reported for calls that are syntactically correct but most likely a semantic error. An example is `(::)/1` calls in clauses that apparently are meant to implement recursive predicate definitions where the user intention is to call the local predicate definition.

Lambda variables

A warning will be reported for *lambda expressions* with unclassified variables (not listed as either *lambda free* or *lambda parameter* variables), for variables playing a dual role (as both lambda free and lambda parameter variables), and for lambda parameters used elsewhere in a clause.

Redefinition of predicates declared in `uses/2` or `use_module/2` directives

An error will be reported for any attempt to define locally a predicate that is already declared in an *uses/2* or *use_module/2* directive.

Other warnings and errors

The Logtalk compiler will throw an error if it finds a predicate clause or a directive that cannot be parsed. The default behavior is to report the error and abort the compilation.

1.12.5 Runtime errors

This section briefly describes runtime errors that result from misuse of Logtalk built-in predicates, built-in methods, or from message-sending. For a complete and detailed description of runtime errors, please consult the Reference Manual.

Logtalk built-in predicates

Most Logtalk built-in predicates check the type and mode of the calling arguments, throwing an exception in case of misuse.

Logtalk built-in methods

Most Logtalk built-in method check the type and mode of the calling arguments, throwing an exception in case of misuse.

Message sending

The message-sending mechanisms always check if the receiver of a message is a defined object and if the message corresponds to a declared predicate within the scope of the sender. The built-in protocol [forwarding](#) declares a predicate, *forward/1*, which is automatically called (if defined) by the runtime for any message that the receiving object does not understand. The usual definition for this error handler is to delegate or forward the message to another object that might be able to answer it:

```
forward(Message) :-  
    % forward the message while preserving the sender  
    [Object::Message].
```

If preserving the original sender is not required, this definition can be simplified to:

```
forward(Message) :-  
    Object::Message.
```

More sophisticated definitions are, of course, possible.

1.13 Reflection

Logtalk provides support for both *structural* and *behavioral* reflection. Structural reflection supports computations over an application structure. Behavioral reflection supports computations over what an application does while running. The structural and behavioral reflection APIs are used by all the [developer tools](#), which are regular applications.

1.13.1 Structural reflection

Structural reflection allows querying the properties of objects, categories, protocols, and predicates. This API provides two views on the structure of an application: a *transparent-box view* and a *black-box view*, described next.

Transparent-box view

The transparent-box view provides a structural view of the contents and properties of entities, predicates, and source files akin to accessing the corresponding source code. I.e. this is the view we use when asking questions such as: *What predicates are declared in this protocol? Which predicates are called by this predicate? Where are clauses for this multifile predicate defined?*

For entities, built-in predicates are provided for *enumerating entities*, *enumerating entity properties* (including entity declared, defined, called, and updated predicates; i.e. full predicate cross-referencing data), and *enumerating entity relations* (for full entity cross-referencing data). For a detailed description of the supported entity properties, see the sections on *object properties*, *protocol properties*, and *category properties*. For examples of querying entity relations, see the sections on *object relations*, *protocol relations*, and *category relations*.

Note

Some entity and predicate properties are only available when the source files are compiled with the `source_data` flag turned on.

The `logtalk` built-in object provides predicates for querying loaded source files and their properties.

Black-box view

The black-box view provides a view that takes into account entity encapsulation and thus only allows querying about predicates and operators that are within the scope of the entity calling the reflection methods. This is the view we use when asking questions such as: *What messages can be sent to this object?*

Built-in methods are provided for querying the *predicates that are declared and can be called or used as messages* and for querying the *predicate properties*. It is also possible to enumerate *entity operators*. See the sections on *finding declared predicates* and on *predicate properties* for more details.

1.13.2 Behavioral reflection

Behavioral reflection provides insight on what an application does when running. Specifically, by observing and acting on the messages being exchanged between objects. See the section on *event-driven programming* for details. There is also a *dependents* library that provides an implementation of Smalltalk dependents mechanism.

For use in debugging tools, there is also a small reflection API providing *trace and debug event predicates* provided by the `logtalk` built-in object.

1.14 Writing and running applications

For successful programming in Logtalk, you need a good working knowledge of Prolog and an understanding of the principles of object-oriented programming. Most guidelines for writing good Prolog code apply as well to Logtalk programming. To those guidelines, you should add the basics of good object-oriented design.

One of the advantages of a system like Logtalk is that it enables us to use the currently available object-oriented methodologies, tools, and metrics [Champaux92] in logic programming. That said, writing applications in Logtalk is similar to writing applications in Prolog: we define new predicates describing what is true about our domain objects, about our problem solution. We encapsulate our predicate directives and definitions inside new objects, categories, and protocols that we create by hand with a text editor or by using the Logtalk built-in predicates. Some of the information collected during the analysis and design phases can be integrated into the objects, categories, and protocols that we define by using the available entity and predicate documenting directives.

1.14.1 Starting Logtalk

We run Logtalk inside a normal Prolog session, after loading the necessary files. Logtalk extends but does not modify your Prolog compiler. We can freely mix Prolog queries with the sending of messages, and our applications can be made of both normal Prolog clauses and object definitions.

Depending on your Logtalk installation, you may use a script or a shortcut to start Logtalk with your chosen Prolog compiler. On POSIX operating-systems, Bash shell integration scripts should be available from the command-line. On Windows, PowerShell integration scripts should be available from the command-line and integration shortcuts should be available from the Start Menu. Scripts are named upon the used backend Prolog compilers.

For example, assuming a POSIX operating-system and GNU Prolog as the backend:

```
$ gplgt
...
```

Depending on your Logtalk installation, you may need to type instead `gplgt.sh`. On Windows, using PowerShell 7.2 or a later version and ECLiPSe as the backend:

```
PS> eclipselgt.ps1
...
```

1.14.2 Running parallel Logtalk processes

Running parallel Logtalk processes is enabled by setting the *clean* flag to on. This is the default flag value in the backend adapter files. With this setting, the intermediate Prolog files generated by the Logtalk compiler include the process identifier in the names, thus preventing file name clashes when running parallel processes. When the flag is turned off, the generated intermediate Prolog file names don't include the process identifier and are kept between runs. This is usually done to avoid repeated recompilation of stable code when developing large applications or when running multiple test sets for performance (by avoiding repeated recompilation of the *lgtunit* tool).

To run parallel Logtalk processes with the *clean* flag turned off, each process must use its own *scratch directory*. This is accomplished by defining the `scratch_directory` library alias to a per-process location **before** loading the compiler/runtime. For example, assuming we're using GNU Prolog as the backend, a possible definition could be:

```
:- multifile(logtalk_library_path/2).
:- dynamic(logtalk_library_path/2).

logtalk_library_path(scratch_directory, Directory) :-
    temporary_name(lgtXXXXXX, Name),
    decompose_file_name(Name, _, Prefix, _),
    atom_concat('/tmp/', Prefix, Directory),
    ( file_exists(Directory) ->
      true
    ; make_directory(Directory)
    ).
```

Assuming the code above is saved in a `parallel_logtalk_processes_setup.pl` file, we would then start Logtalk using:

```
$ gplgt --init-goal "consult('parallel_logtalk_processes_setup.pl')"
```

The details on how to define and load the definition of the `scratch_directory` library alias are, however, backend specific (due to the lack of Prolog standardization) and possibly also operating-system specific (different locations for the temporary directory). The Logtalk library includes a `parallel_logtalk_processes_setup.pl` file with support for selected backends and usage instructions.

1.14.3 Source files

Logtalk source files may define any number of entities (objects, categories, or protocols). Source files may also contain Prolog code interleaved with Logtalk entity definitions. Plain Prolog code is usually copied as-is to the corresponding Prolog output file (except, of course, if subject to the *term-expansion mechanism*). Prolog modules are compiled as objects. The following Prolog directives are processed when read (thus affecting the compilation of the source code that follows): `ensure_loaded/1`, `use_module/1-2`, `op/3`, and `set_prolog_flag/2`. The *initialization/1* directive may be used for defining an initialization goal to be executed when loading a source file.

Logtalk source files can include the text of other files by using the *include/1* directive. Although there is also a standard Prolog `include/1` directive, any occurrences of this directive in a Logtalk source file is handled by the Logtalk compiler, not by the *backend Prolog compiler*, to improve portability.

When writing a Logtalk source file, the following advice applies:

- When practical and when performance is critical, define each entity on its own source file.
- Source file loading order can impact performance (e.g., if an object imports a category defined in a source file loaded after the object source file, no static binding optimizations will be possible).
- Initialization directives that result in the compilation and loading of other source files (e.g., libraries) should preferably be written in the application loader file to ensure the availability of the entities they define when compiling the application source files (thus enabling static binding optimizations).

Note that you can use the `logtalk::loaded_files_topological_sort/1` and `logtalk::loaded_files_topological_sort/2` predicates to find an optimal file loading order to eliminate compilation warnings due to files being loaded before their dependencies. In the case of simple applications with no library dependencies:

```
?- {loader}. % load you application code
...
```

(continues on next page)

(continued from previous page)

```
?- logtalk::loaded_files_topological_sort(Sorted).
Sorted = [...]
```

Use the file basenames from the returned list to update the application driver files (typically, `loader.lgt` and `tester.lgt`).

In more complex applications with external library dependencies, define a library alias for your application (e.g., `my_app`) and use it to get an initial list of your application files:

```
?- {my_app(loader)}. % load you application code
...

?- findall(Path, logtalk::loaded_file_property(Path, library(my_app)), Paths),
   logtalk::loaded_files_topological_sort(Paths, Sorted).
Sorted = [...]
```

Note that the `Sorted` list will include the loader file itself.

Naming conventions

When defining each entity in its own source file, it is recommended that the source file be named after the entity identifier. For parametric objects, the identifier arity can be appended to the identifier functor. By default, all Logtalk source files use the extension `.lgt` but this is optional and can be set in the adapter files. For example, we may define an object named `vehicle` and save it in a `vehicle.lgt` source file. A `sort(_)` parametric object would be saved in a `sort_1.lgt` source file.

Source file text encoding

The text encoding used in a source file may be declared using the [encoding/1](#) directive when running Logtalk with backend Prolog compilers that support multiple encodings (check the [encoding_directive](#) flag in the adapter file of your Prolog compiler).

1.14.4 Multi-pass compiler

Logtalk is implemented using a *multi-pass* compiler. In comparison, some Prolog systems use a multi-pass compiler while others use a single-pass compiler. While there are pros and cons with each solution, the most relevant consequence in this context is for the content of source files. In Logtalk, entities and predicates only become available (for the runtime system) after the source file is successfully compiled and loaded. This may prevent some compiler optimizations, notably [static binding](#), if some of the referred entities are defined in the same source file. On the other hand, the order of predicate directives and predicate definitions is irrelevant. In contrast, in a system implemented using a single-pass compiler, the order of the source file terms can and often is significant for proper and successful compilation. In these systems, predicates may become available for calling as soon as they are compiled, even if the rest of the source file is yet to be compiled.

The Logtalk compiler reads source files using the Prolog standard `read_term/3` predicate. This ensures compatibility with any syntax extensions that the used backend may implement. In the first compiler stage, all source file terms are read, and data about all defined entities, directives, predicates, and grammar rules is collected. Any defined [term-expansion rules](#) are applied to the read terms. Grammar rules are expanded into predicate clauses unless expanded by user-defined term-expansion rules. The second stage compiles all initialization goals and clause bodies, taking advantage of the data collected in the first stage and applying any defined goal-expansion rules. Depending on the compilation mode, the generated code can be instrumented for debugging tools or optimized for performance. Linter checks are performed during these two first stages.

The final step in the second stage is to write the generated intermediate Prolog code into a temporary file. In the third and final stage, this intermediate Prolog file is compiled and loaded by the used backend. These intermediate files are deleted by default after loading (see the [clean](#) flag description for details).

1.14.5 Compiling and loading your applications

Your applications will be made of source files containing your objects, protocols, and categories. The source files can be compiled to disk by calling the [logtalk_compile/1](#) built-in predicate:

```
| ?- logtalk_compile([source_file1, source_file2, ...]).
```

This predicate runs the compiler on each file and, if no fatal errors are found, outputs Prolog source files that can then be consulted or compiled in the usual way by your Prolog compiler.

To compile to disk and also load into memory the source files, we can use the [logtalk_load/1](#) built-in predicate:

```
| ?- logtalk_load([source_file1, source_file2, ...]).
```

This predicate works in the same way as the predicate [logtalk_compile/1](#) but also loads the compiled files into memory.

Both predicates expect a source file name or a list of source file names as an argument. The Logtalk source file name extension, as defined in the adapter file (by default, `.lgt`), can be omitted.

If you have more than a few source files, then you may want to use a [loader file](#) helper file containing the calls to the [logtalk_load/1-2](#) predicates. Consulting or compiling the loader file will then compile and load all your Logtalk entities into memory (see below for details).

With most [backend Prolog compilers](#), you can use the shorthand `{File}` for `logtalk_load(File)` and `{File1, File2, ...}` for `logtalk_load([File1, File2, ...])`. The use these shorthands should be restricted to the Logtalk/Prolog top-level interpreter, as they are not part of the language specification and may be commented out in case of conflicts with backend Prolog compiler features.

The built-in predicate [logtalk_make/0](#) can be used to reload all modified source files. With most backend Prolog compilers, you can also use the `{*}` top-level shortcut. Files are also reloaded when the compilation mode changes. An extended version of this predicate, [logtalk_make/1](#), accepts multiple targets, including `all`, `clean`, `check`, `circular`, `documentation`, `caches`, `debug`, `normal`, and `optimal`. For example, assume that you have loaded your application files and found a bug. You can easily recompile the files in debug mode by using the `logtalk_make(debug)` goal. After debugging and fixing the bug, you can reload the files in normal mode using the `logtalk_make(normal)` or in optimized mode using the `logtalk_make(optimal)` goal. See the predicates documentation for a complete list of targets and top-level shortcuts. In particular, the `logtalk_make(clean)` goal can be specially useful before switching backend Prolog compilers, as the generated intermediate files may not be compatible. The `logtalk_make(caches)` goal is usually used when benchmarking compiler performance improvements.

1.14.6 Compiler errors, warnings, and comments

Following a Prolog tradition inherited from Quintus Prolog, the compiler prefixes (by default) errors with a `!` and warnings with a `*`. For example:

```
!      Existence error: directive object/1 does not exist
!      in directive end_object/0
!      in file /home/jdoe/logtalk/examples/errors/unmatched_directive.lgt at or above line_
↪27

*      No matching clause for goal: baz(a)
*      while compiling object main_include_compiler_warning
*      in file /home/jdoe/logtalk/examples/errors/include_compiler_warning.lgt between_
↪lines 38-39
```

Compiler comments are prefixed by `%`. For example:

```
?- {ack(loader)}.
% [ /home/jdoe/logtalk/examples/ack/ack.lgt loaded ]
% [ /home/jdoe/logtalk/examples/ack/loader.lgt loaded ]
% (0 warnings)
true.
```

1.14.7 Loader files

If you look into the Logtalk distribution, you will notice that most source code directories (e.g., of tools, libraries, and examples) contain a *driver file* that can be used to load all included source files and any required libraries. These loader files are usually named `loader.lgt` or contain the word *loader* in their name. Loader files are ordinary source files and thus compiled and loaded like any source file. By also defining a loader file for your project, you can then load it by simply typing:

```
| ?- {loader}.
```

Another driver file, usually named `tester.lgt` (or containing the word *tester* in its name) is commonly used to load and run tests. By also defining a tester file for your project, you can then run its tests by simply typing:

```
| ?- {tester}.
```

Usually these driver files contain calls to the built-in predicates `set_logtalk_flag/2` (e.g., for setting global, project-specific, flag values) and `logtalk_load/1` or `logtalk_load/2` (for loading project files), wrapped inside a Prolog initialization/1 directive for portability. For instance, if your code is split into three source files named `source1.lgt`, `source2.lgt`, and `source3.lgt`, then the contents of your loader file could be:

```
:- initialization((
    % set project-specific global flags
    set_logtalk_flag(events, allow),
    % load the project source files
    logtalk_load([source1, source2, source3])
)).
```

Another example of directives that are often used in a loader file would be `op/3` directives declaring global operators needed by your project. Loader files are also often used for setting source file-specific compiler

flags (this is useful even when you only have a single source file if you always load it with the same set of compiler flags). For example:

```
:- initialization((
    % set project-specific global flags
    set_logtalk_flag(source_data, off),
    % load the project source files
    logtalk_load(
        [source1, source2, source3],
        % source file-specific flags
        [portability(warning)]),
    logtalk_load(
        [source4, source5],
        % source file-specific flags
        [portability(silent)])
)).
```

To take the best advantage of loader and tester files, define a clause for the multifile and dynamic `logtalk_library_path/2` predicate for the directory containing your source files as explained in the next section.

When your project also uses Prolog module resources, the loader file is also the advised place to load them, preferably without any exports. For example:

```
:- use_module(library(clpfd), []).
...

:- initialization((
    ...
)).
```

Complex projects often use a main loader file that loads the loader files of each of the project components. Thus, loader files provide a central point to understand a project organization and dependencies.

It is worth mentioning a common mistake when writing the first loader files. New users sometimes try to set compiler flags using `logtalk_load/2` when loading a loader file. For example, by writing:

```
| ?- logtalk_load(loader, [optimize(on)]).
```

This will not work as you might expect, as the compiler flags will only be used in the compilation of the `loader.lgt` file itself and will not affect the compilation of files loaded through the `initialization/1` directive contained on the loader file.

1.14.8 Libraries of source files

Logtalk defines a *library* simply as a directory containing source files. Library locations can be specified by defining or asserting clauses for the dynamic and multifile predicate `logtalk_library_path/2`. For example:

```
:- multifile(logtalk_library_path/2).
:- dynamic(logtalk_library_path/2).

logtalk_library_path(shapes, '$LOGTALKUSER/examples/shapes/').
```

The first argument of the predicate is used as an alias for the path on the second argument. Library aliases may also be used on the second argument. For example:

```
:- multifile(logtalk_library_path/2).
:- dynamic(logtalk_library_path/2).

logtalk_library_path(lgtuser, '$LOGTALKUSER/').
logtalk_library_path(examples, lgtuser('examples/')).
logtalk_library_path(viewpoints, examples('viewpoints/')).
```

This allows us to load a library source file without the need to first change the current working directory to the library directory and then back to the original directory. For example, in order to load a `loader.lgt` file, contained in a library named `viewpoints`, we just need to type:

```
| ?- logtalk_load(viewpoints(loader)).
```

The best way to take advantage of this feature is to load at startup a source file containing clauses for the `logtalk_library_path/2` predicate needed for all available libraries (typically, using a *settings file*, as discussed below). This allows us to load library source files or entire libraries without worrying about library paths, improving code portability. The directory paths on the second argument should always end with the path directory separator character. Most backend Prolog compilers allow the use of environment variables in the second argument of the `logtalk_library_path/2` predicate. Use of POSIX relative paths (e.g., `'../'` or `'./'`) for top-level library directories (e.g., `lgtuser` in the example above) is not advised, as different backend Prolog compilers may start with different initial working directories, which may result in portability problems of your loader files.

This *library notation* provides functionality inspired by the `file_search_path/2` mechanism introduced by Quintus Prolog and later adopted by some other Prolog compilers but with a key difference: there is no fragile search mechanism, and the Logtalk make can be used to check for duplicated library aliases. Multiple definitions for the same alias are problematic when using external dependencies, as any third-party update to those dependencies can introduce file name clashes. Note that the potential for these clashes cannot be reliably minimized by a careful ordering of the `logtalk_library_path/2` predicate clauses due to this predicate being multifile and dynamic.

1.14.9 Settings files

Although it is always possible to edit the *backend Prolog compiler* adapter files, the recommended solution to customize compiler flags is to create a `settings.lgt` file in the Logtalk user folder or in the user home folder. Depending on the backend Prolog compiler and the operating-system, is also possible to define per-project settings files by creating a `settings.lgt` file in the project directory and by starting Logtalk from this directory. At startup, Logtalk tries to load a `settings.lgt` file from the following directories, searched in sequence:

- Startup directory (`$LOGTALK_STARTUP_DIRECTORY`)
- Logtalk user directory (`$LOGTALKUSER`)
- User home directory (`$HOME`; `%USERPROFILE%` on Windows if `%HOME%` is not defined)
- Application data directory (`%APPDATA%\Logtalk`; only on Windows)
- Config directory (`$XDG_CONFIG_HOME/logtalk`)
- Default config directory (`$HOME/.config/logtalk/`)

The startup directory is only searched when the read-only *settings file* flag is set to allow. When no settings files are found, Logtalk will use the default compiler flag values set on the backend Prolog compiler adapter files. When limitations of the backend Prolog compiler or the operating-system prevent Logtalk from finding the settings files, these can always be loaded manually after Logtalk startup.

Settings files are normal Logtalk source files (although when automatically loaded by Logtalk they are compiled and loaded silently with any errors being reported but otherwise ignored). The usual contents is an initialization/1 Prolog directive containing calls to the [set_logtalk_flag/2](#) Logtalk built-in predicate and asserting clauses for the [logtalk_library_path/2](#) multifile dynamic predicate. Note that the [set_logtalk_flag/2](#) directive cannot be used as its scope is local to the source file being compiled.

One of the troubles of writing portable applications is the different feature sets of Prolog compilers. Using the Logtalk support for conditional compilation and the [prolog_dialect](#) flag we can write a single settings file that can be used with several [backend Prolog compilers](#):

```
:- if(current_logtalk_flag(prolog_dialect, yap)).

    % YAP specific settings
    ...

:- elif(current_logtalk_flag(prolog_dialect, gnu)).

    % GNU Prolog specific settings
    ...

:- else.

    % generic Prolog settings

:- endif.
```

The Logtalk distribution includes a `samples/settings-sample.lgt` sample file with commented out code snippets for common settings.

1.14.10 Compiler linter

The compiler includes a [linter](#) that checks for a wide range of possible problems in source files. Notably, the compiler checks for unknown entities, unknown predicates, undefined predicates (i.e., predicates that are declared but not defined), missing directives (including missing `dynamic/1` and `meta_predicate/1` directives), redefined built-in predicates, calls to non-portable predicates, singleton variables, goals that are always true or always false (i.e., goals that can be replaced by `true` or `fail`), and trivial fails (i.e., calls to predicates with no match clauses). Most of the linter warnings are controlled by [compiler flags](#). See the next section for details.

1.14.11 Compiler flags

The [logtalk_load/1](#) and [logtalk_compile/1](#) always use the current set of default compiler flags as specified in your settings file and the Logtalk adapter files or changed for the current session using the built-in predicate [set_logtalk_flag/2](#). Although the default flag values cover the usual cases, you may want to use a different set of flag values while compiling or loading some of your Logtalk source files. This can be accomplished by using the [logtalk_load/2](#) or the [logtalk_compile/2](#) built-in predicates. These two predicates accept a list of options affecting how a Logtalk source file is compiled and loaded:

```
| ?- logtalk_compile(Files, Options).
```

or:

```
| ?- logtalk_load(Files, Options).
```

In fact, the `logtalk_load/1` and `logtalk_compile/1` predicates are just shortcuts to the extended versions called with the default compiler flag values. The options are represented by a compound term where the functor is the flag name and the sole argument is the flag value.

We may also change the default flag values from the ones loaded from the adapter file by using the [set_logtalk_flag/2](#) built-in predicate. For example:

```
| ?- set_logtalk_flag(unknown_entities, silent).
```

The current default flags values can be enumerated using the [current_logtalk_flag/2](#) built-in predicate:

```
| ?- current_logtalk_flag(unknown_entities, Value).
```

```
Value = silent  
yes
```

Logtalk also implements a [set_logtalk_flag/2](#) directive, which can be used to set flags within a source file or within an entity. For example:

```
% compile objects in this source file with event support  
:- set_logtalk_flag(events, allow).  
  
:- object(foo).  
  
    % compile this object with support  
    % for dynamic predicate declarations  
    :- set_logtalk_flag(dynamic_declarations, allow).  
    ...  
  
:- end_object.  
  
...
```

Note that the scope of the `set_logtalk_flag/2` directive is local to the entity or to the source file containing it.

Note

Applications should never rely on default flag values for working properly. Whenever the compilation of a source file or an entity requires a specific flag value, the flag should be set explicitly in the entity, in the source file, or in the loader file.

Read-only flags

Some flags have read-only values and thus cannot be changed at runtime. Their values are defined in the Prolog backend *adapter files*. These are:

settings_file

Allows or disables loading of a *settings file* at startup. Possible values are allow, restrict, and deny. The usual default value is allow but it can be changed by editing the adapter file when e.g. embedding Logtalk in a compiled application. With a value of allow, settings files are searched in the startup directory, in the Logtalk user directory, in the user home directory, in the APPDATA if running on Windows, and in the XDG configuration directory. With a value of restrict, the search for the settings files skips the startup directory.

prolog_dialect

Identifier of the *backend Prolog compiler* (an atom). This flag can be used for *conditional compilation* of Prolog compiler specific code.

prolog_version

Version of the *backend Prolog compiler* (a compound term, v(Major, Minor, Patch), whose arguments are integers). This flag availability depends on the Prolog compiler. Checking the value of this flag fails for any Prolog compiler that does not provide access to version data.

prolog_compatible_version

Compatible version of the *backend Prolog compiler* (a compound term, usually with the format @>=(v(Major, Minor, Patch)), whose arguments are integers). This flag availability depends on the Prolog compiler. Checking the value of this flag fails for any Prolog compiler that does not provide access to version data.

unicode

Informs Logtalk if the *backend Prolog compiler* supports the Unicode standard. Possible flag values are unsupported, full (all Unicode planes supported), and bmp (supports only the Basic Multilingual Plane).

encoding_directive

Informs Logtalk if the *backend Prolog compiler* supports the *encoding/1* directive. This directive is used for declaring the text encoding of source files. Possible flag values are unsupported, full (can be used in both Logtalk source files and compiler generated Prolog files), and source (can be used only in Logtalk source files).

tabling

Informs Logtalk if the *backend Prolog compiler* provides tabling programming support. Possible flag values are unsupported and supported.

engines

Informs if the *backend Prolog compiler* provides the required low level multi-threading programming support for Logtalk *threaded engines*. Possible flag values are unsupported and supported.

threads

Informs if the *backend Prolog compiler* provides the required low level multi-threading programming support for all high-level Logtalk *multi-threading features*. Possible flag values are unsupported and supported.

modules

Informs Logtalk if the *backend Prolog compiler* provides suitable module support. Possible flag values are unsupported and supported (independently of this flag, Logtalk provides limited support for compiling Prolog modules as objects).

coinduction

Informs Logtalk if the *backend Prolog compiler* provides the required minimal support for cyclic terms

necessary for working with *coinductive predicates*. Possible flag values are unsupported and supported.

Version flags

version_data(Value)

Read-only flag whose value is the compound term `logtalk(Major,Minor,Patch,Status)`. The first three arguments are integers and the last argument is an atom, possibly empty, representing version status: `aN` for alpha versions, `bN` for beta versions, `rcN` for release candidates (with `N` being a natural number), and `stable` for stable versions. The `version_data` flag is also a de facto standard for Prolog compilers.

Lint flags

linter(Option)

Meta-flag for managing the values of all the linter flags as a group. Possible option values are `on` to set all the individual linter flags to warning, `off` to set all the individual linter flags to silent, and `default` to set all the individual linter flag values to their defaults as defined in the backend adapter files (the usual default). This flag **must** always be defined in the backend adapter files with the value of `default`.

unknown_entities(Option)

Controls the unknown entity warnings, resulting from loading an entity that references some other entity that is not currently loaded. Possible option values are `warning` (the usual default) and `silent`. Note that these warnings are not always avoidable, specially when using reflective designs of class-based hierarchies.

unknown_predicates(Option)

Defines the compiler behavior when unknown messages or calls to unknown predicates (or non-terminals) are found. An unknown message is a message sent to an object that is not part of the object protocol. An unknown predicate is a called predicate that is neither locally declared or defined. Possible option values are `error`, `warning` (the usual default), and `silent` (not recommended).

undefined_predicates(Option)

Defines the compiler behavior when calls to declared but undefined predicates (or non-terminals) are found. Note that these calls will fail at runtime as per closed-world assumption. Possible option values are `error`, `warning` (the usual default), and `silent` (not recommended).

steadfastness(Option)

Controls warnings about *possible* non *steadfast* predicate definitions due to variable aliasing at a clause head and a cut in the clause body. Possible option values are `warning` and `silent` (the usual default due to the possibility of false positives).

portability(Option)

Controls the non-ISO specified Prolog built-in predicate and non-ISO specified Prolog built-in arithmetic function calls warnings plus use of non-standard Prolog flags and/or flag values. Possible option values are `warning` and `silent` (the usual default).

deprecated(Option)

Controls the deprecated predicate warnings. Possible option values are `warning` (the usual default) and `silent`.

missing_directives(Option)

Controls the missing predicate directive warnings. Possible option values are `warning` (the usual default) and `silent` (not recommended).

duplicated_directives(Option)

Controls the duplicated predicate directive warnings. Possible option values are `warning` (the usual

default) and silent (not recommended). Note that conflicting directives for the same predicate are handled as errors, not as duplicated directive warnings.

trivial_goal_fails(*Option*)

Controls the printing of warnings for calls to local static predicates with no matching clauses. Possible option values are warning (the usual default) and silent (not recommended).

always_true_or_false_goals(*Option*)

Controls the printing of warnings for goals that are always true or false. Possible option values are warning (the usual default) and silent (not recommended). A unexpected exception in the goal being checked is also reported.

grammar_rules(*Option*)

Controls the printing of grammar rules related warnings. Possible option values are warning (the usual default) and silent (not recommended).

arithmetic_expressions(*Option*)

Controls the printing of arithmetic expressions related warnings. Possible option values are warning (the usual default) and silent (not recommended).

lambda_variables(*Option*)

Controls the printing of lambda variable related warnings. Possible option values are warning (the usual default) and silent (not recommended).

suspicious_calls(*Option*)

Controls the printing of suspicious call warnings. Possible option values are warning (the usual default) and silent (not recommended).

redefined_built_ins(*Option*)

Controls the Logtalk and Prolog built-in predicate redefinition warnings. Possible option values are warning and silent (the usual default). Warnings about redefined Prolog built-in predicates are often the result of running a Logtalk application on several Prolog compilers, as each Prolog compiler defines its set of built-in predicates.

redefined_operators(*Option*)

Controls the Logtalk and Prolog built-in operator redefinition warnings. Possible option values are warning (the usual default) and silent. Redefining Logtalk operators or standard Prolog operators can break term parsing, causing syntax errors or change how terms are parsed, introducing bugs.

singleton_variables(*Option*)

Controls the singleton variable warnings. Possible option values are warning (the usual default) and silent (not recommended).

naming(*Option*)

Controls warnings about entity, predicate, and variable names per official coding guidelines (which advise using underscores for entity and predicate names and camel case for variable names). Additionally, variable names should not differ only in case. Possible option values are warning and silent (the usual default due to the current limitation to ASCII names and the computational cost of the checks).

duplicated_clauses(*Option*)

Controls warnings of duplicated entity clauses (and duplicated entity grammar rules). Possible option values are warning and silent (the usual default due to the required heavy computations). When the term-expansion mechanism is used and results in duplicated clauses, the reported line numbers are for lines of the original clauses that were expanded.

disjunctions(*Option*)

Controls warnings on clauses where the body is a disjunction. Possible option values are warning (the usual default) and silent. As per coding guidelines, in most cases, these clauses can be rewritten using a clause per disjunction branch for improved code readability.

conditionals(*Option*)

Controls warnings on if-then-else and soft-cut control constructs. Possible option values are warning (the usual default) and silent. Warnings include misuse of cuts, potential bugs in the test part, and missing else part (lack of compliance with coding guidelines).

catchall_catch(*Option*)

Controls warnings on catch/3 goals that catch all exceptions. Possible option values are warning and silent (the usual default). Lack of standardization often makes it tricky or cumbersome to avoid too generic catch/3 goals when writing portable code.

left_recursion(*Option*)

Controls warnings of left-recursion on clauses and grammar rules. Specifically, when the clause or grammar rule head and the leftmost goal in the body are variants. Possible option values are warning (the usual default) and silent.

tail_recursive(*Option*)

Controls warnings of non-tail recursive predicate (and non-terminal) definitions. The lint check does not detect all cases of non-tail recursive predicate definitions, however. Also, definitions that make two or more recursive calls are not reported as usually they cannot be changed to be tail recursive. Possible option values are warning and silent (the usual default).

encodings(*Option*)

Controls the source file text encoding warnings. Possible option values are warning (the usual default) and silent.

general(*Option*)

Controls warnings that are not controlled by a specific flag. Possible option values are warning (the usual default) and silent.

Optional features compilation flags

complements(*Option*)

Allows objects to be compiled with support for complementing categories turned off in order to improve performance and security. Possible option values are allow (allow complementing categories to override local object predicate declarations and definitions), restrict (allow complementing categories to add predicate declarations and definitions to an object but not to override them), and deny (ignore complementing categories; the usual default). This option can be used on a per-object basis. Note that changing this option is of no consequence for objects already compiled and loaded.

dynamic_declarations(*Option*)

Allows objects to be compiled with support for dynamic declaration of new predicates turned off in order to improve performance and security. Possible option values are allow and deny (the usual default). This option can be used on a per-object basis. Note that changing this option is of no consequence for objects already compiled and loaded. This option is only checked when sending an *asserta/1* or *assertz/1* message to an object. Local asserting of new predicates is always allowed.

events(*Option*)

Allows message-sending calls to be compiled with or without *event-driven programming* support. Possible option values are allow and deny (the usual default). Objects (and categories) compiled with this option set to deny use optimized code for message-sending calls that do not trigger events. As such, this option can be used on a per-object (or per-category) basis. Note that changing this option is of no consequence for objects already compiled and loaded.

context_switching_calls(*Option*)

Allows context-switching calls (*((<<)/2*) to be either allowed or denied. Possible option values are allow and deny. The default flag value is allow. Note that changing this option is of no consequence for objects already compiled and loaded.

Backend Prolog compiler and loader flags

underscore_variables(Option)

Controls the interpretation of variables that start with an underscore (excluding the anonymous variable) that occur once in a term as either don't care variables or singleton variables. Possible option values are `dont_care` (the default for all supported backends) and `singletons`. Although a changeable flag, its value is backend dependent and thus expected to be set only in the backend adapter files.

prolog_compiler(Flags)

List of compiler flags for the generated Prolog files. The valid flags are specific to the used Prolog backend compiler. The usual default is the empty list. These flags are passed to the backend Prolog compiler built-in predicate that is responsible for compiling to disk a Prolog file. For Prolog compilers that don't provide separate predicates for compiling and loading a file, use instead the *prolog_loader* flag.

prolog_loader(Flags)

List of loader flags for the generated Prolog files. The valid flags are specific to the used Prolog backend compiler. The usual default is the empty list. These flags are passed to the backend Prolog compiler built-in predicate that is responsible for loading a (compiled) Prolog file.

Other flags

scratch_directory(Directory)

Sets the directory to be used to store the temporary files generated when compiling Logtalk source files. This directory can be specified using an atom or using *library notation*. The directory must always end with a slash. The default value is a sub-directory of the source files directory, either `./lgt_tmp/` or `./lgt_tmp/` (depending on the backend Prolog compiler and operating-system). Relative directories must always start with `./` due to the lack of a portable solution to check if a path is relative or absolute. The default value set on the *backend Prolog compiler* adapter file can be overridden by defining the `scratch_directory` library alias (see the *logtalk_library_path/2* predicate documentation for details).

report(Option)

Controls the default printing of messages. Possible option values are `on` (by usual default, print all messages that are not intercepted by the user), `warnings` (only print warning and error messages that are not intercepted by the user), and `off` (do not print any messages that are not intercepted by the user).

code_prefix(Character)

Enables the definition of prefix for all functors of Prolog code generated by the Logtalk compiler. The option value must be a single character atom. Its default value is `'$'`. Specifying a code prefix provides a way to solve possible conflicts between Logtalk compiled code and other Prolog code. In addition, some Prolog compilers automatically hide predicates whose functor starts with a specific prefix, such as the character `$`. Although this is not a read-only flag, it should only be changed at startup time and **before** loading any source files. When changing this flag (e.g., from a *settings file*), restart with the *clean* flag turned on to ensure that any compiled files using the old `code_prefix` value will be recompiled.

optimize(Option)

Controls the compiler optimizations. Possible option values are `on` (used by default for deployment) and `off` (used by default for development). Compiler optimizations include the use of static binding whenever possible, the removal of redundant calls to `true/0` from predicate clauses, the removal of redundant unifications when compiling grammar rules, and inlining of predicate definitions with a single clause that links to a local predicate, to a plain Prolog built-in (or foreign) predicate, or to a Prolog module predicate with the same arguments. Care should be taken when developing applications

with this flag turned on as changing and reloading a file may render *static binding* optimizations invalid for code defined in other loaded files. Turning on this flag automatically turns off the *debug* flag.

source_data(Option)

Defines how much information is retained when compiling a source file. Possible option values are on (the usual default for development) and off. With this flag set to on, Logtalk will keep the information represented using documenting directives plus source location data (including source file names and line numbers). This information can be retrieved using the *reflection API* and is useful for documenting, debugging, and integration with third-party development tools. This flag can be turned off in order to generate more compact code.

debug(Option)

Controls the compilation of source files in debug mode (the Logtalk default debugger can only be used with files compiled in this mode). Also controls, by default, printing of debug> and debug(Topic) messages. Possible option values are on and off (the usual default). Turning on this flag automatically turns off the *optimize* flag.

reload(Option)

Defines the reloading behavior for source files. Possible option values are skip (skip reloading of already loaded files; this value can be used to get similar functionality to the Prolog directive *ensure_loaded/1* but should be used only with fully debugged code), changed (the usual default; reload files only when they are changed since last loaded, provided that any explicit flags and the compilation mode are the same as before), and always (always reload files).

relative_to(Directory)

Defines a base directory for resolving relative source file paths. The default value is the directory of the source file being compiled.

hook(Object)

Allows the definition of an object (which can be the pseudo-object *user*) implementing the *expanding* built-in protocol. The hook object must be compiled and loaded when this option is used. It's also possible to specify a Prolog module instead of a Logtalk object, but the module must be pre-loaded, and its identifier must be different from any object identifier.

clean(Option)

Controls cleaning of the intermediate Prolog files generated when compiling Logtalk source files. Possible option values are off and on (the usual default). When turned on, intermediate files are deleted after loading, and all source files are recompiled disregarding any existing intermediate files. When turned off, the intermediate files are kept. This is useful when embedding applications, which requires collecting the intermediate code, and when working on large applications to avoid repeated recompilation of stable code. The flag must be turned on when changing compilation modes, changing flags such as *code_prefix*, or when turning on linter flags that are off by default without at the same time making changes to the application source files themselves, as any existing intermediate files would not be recompiled as necessary due to file timestamps not changing.

User-defined flags

Logtalk provides a `create_logtalk_flag/3` predicate that can be used for defining new flags.

1.14.12 Reloading source files

As a general rule, reloading source files should never occur in production code and should be handled with care in development code. Reloading a Logtalk source file usually requires reloading the intermediate Prolog file that is generated by the Logtalk compiler. The problem is that there is no standard behavior for reloading Prolog files. For static predicates, almost all Prolog compilers replace the old definitions with the new ones. However, for dynamic predicates, the behavior depends on the Prolog compiler. Most compilers replace the old definitions, but some of them simply append the new ones, which usually leads to trouble. See the compatibility notes for the backend Prolog compiler you intend to use for more information. There is an additional potential problem when using multi-threading programming. Reloading a threaded object does not recreate from scratch its old message queue, which may still be in use (e.g., threads may be waiting on it).

When using library entities and stable code, you can avoid reloading the corresponding source files (and, therefore, recompiling them) by setting the `reload` compiler flag to skip. For code under development, you can turn off the `clean` flag to avoid recompiling files that have not been modified since the last compilation (assuming that backend Prolog compiler that you are using supports retrieving file modification dates). You can disable deleting the intermediate files generated when compiling source files by changing the default flag value in your settings file, by using the corresponding compiler flag with the compiling and loading built-in predicates, or, for the remaining of a working session, by using the call:

```
| ?- set_logtalk_flag(clean, off).
```

Some caveats that you should be aware of. First, some warnings that might be produced when compiling a source file will not show up if the corresponding object file is up-to-date because the source file is not being (re)compiled. Second, if you are using several Prolog compilers with Logtalk, be sure to perform the first compilation of your source files with the `clean` flag turned off: the intermediate Prolog files generated by the Logtalk compiler may not be compatible across Prolog compilers or even for the same Prolog compiler across operating systems (e.g., due to the use of different character encodings or end-of-line characters).

1.14.13 Batch processing

When doing batch processing, you probably want to turn off the `report` flag to suppress all messages of type banner, comment, comment(_), warning, and warning(_) that are normally printed. Note that error messages and messages providing information requested by the user will still be printed.

1.14.14 Optimizing performance

The default compiler flag settings are appropriate for the **development** but not necessarily for the **deployment** of applications. To minimize the generated code size, turn the `source_data` flag off. To optimize runtime performance, turn on the `optimize` flag. Your chosen backend Prolog compiler may also provide performance-related flags; check its documentation.

Pay special attention to file compilation/loading order. Whenever possible, compile and load your files by taking into account file dependencies. By default, the compiler will print a warning whenever a file references an entity that is not yet loaded. Solving these warnings is key for optimal performance by enabling `static binding` optimizations. For a clear picture of file dependencies, use the `diagrams` tool to generate a file dependency diagram for your application.

Minimize the use of dynamic predicates. Parametric objects can often be used in alternative. When dynamic predicates cannot be avoided, try to make them private. Declaring a dynamic predicate also as a private predicate allows the compiler to optimize local calls to the database methods (e.g., *assertz/1* and *retract/1*) that modify the predicate.

Sending a *message to self* implies *dynamic binding*, but there are often cases where *(::)/1* is misused to call an imported or inherited predicate that is never going to be redefined in a descendant. In these cases, a *super call*, *(^^)/1*, can be used instead with the benefit of often enabling static binding. Most of the guidelines for writing efficient Prolog code also apply to Logtalk code. In particular, define your predicates to take advantage of first-argument indexing. In the case of recursive predicates, define them as tail-recursive predicates whenever possible.

See the *section on performance* for a detailed discussion on Logtalk performance.

1.14.15 Portable applications

Logtalk is compatible with most modern standards-compliant Prolog compilers. However, this does not necessarily imply that your Logtalk applications will have the same level of portability. If possible, you should only use in your applications Logtalk built-in predicates and ISO Prolog-specified built-in predicates and arithmetic functions. If you need to use built-in predicates (or built-in arithmetic functions) that may not be available in other Prolog compilers, you should try to encapsulate the non-portable code in a small number of objects and provide a portable **interface** for that code through the use of Logtalk protocols. An example will be code that access operating-system specific features. The Logtalk compiler can warn you of the use of non-ISO-specified built-in predicates and arithmetic functions by using the *portability* compiler flag.

1.14.16 Conditional compilation

Logtalk supports conditional compilation within source files using the *if/1*, *elif/1*, *else/0*, and *endif/0* directives. This support is similar to the support found in several Prolog systems such as ECLiPSe, GNU Prolog, SICStus Prolog, SWI-Prolog, XSB, and YAP.

1.14.17 Avoiding common errors

Try to write objects and protocol documentation **before** writing any other code; if you are having trouble documenting a predicate, perhaps you need to go back to the design stage.

Try to avoid lengthy hierarchies. Composition is often a better choice over inheritance for defining new objects (Logtalk supports component-based programming through the use of *categories*). In addition, prototype-based hierarchies are semantically simpler than class-based hierarchies.

Dynamic predicates or dynamic entities are sometimes needed, but we should always try to minimize the use of non-logical features such as asserts and retracts.

Since each Logtalk entity is independently compiled, if an object inherits a dynamic or a meta-predicate predicate, then the respective directives must be repeated to ensure a correct compilation.

In general, Logtalk does not verify if a user predicate call/return arguments comply with the declared modes. On the other hand, Logtalk built-in predicates, built-in methods, and message-sending control structures are fully checked for calling mode errors.

Logtalk error handling strongly depends on the ISO compliance of the chosen Prolog compiler. For instance, the error terms that are generated by some Logtalk built-in predicates assume that the Prolog built-in predicates behave as defined in the ISO standard regarding error conditions. In particular, if your Prolog compiler does not support a *read_term/3* built-in predicate compliant with the ISO Prolog Standard definition, then

the current version of the Logtalk compiler may not be able to detect misspelled variables in your source code.

1.14.18 Coding style guidelines

It is suggested that all code between an entity opening and closing directives be indented by one tab stop. When defining entity code, both directives and predicates, Prolog coding style guidelines may be applied. All Logtalk source files, examples, and standard library entities use tabs (the recommended setting is a tab width equivalent to 4 spaces) for laying out code. Closely related entities can be defined in the same source file. However, for the best performance, is often necessary to have an entity per source file. Entities that might be useful in different contexts (such as library entities) are best defined in their own source files.

A detailed coding style guide is available at the Logtalk official website.

1.15 Printing messages and asking questions

Applications, components, and libraries often print all sorts of messages. These include banners, logging, debugging, and computation results messages. But also, in some cases, user interaction messages. However, the authors of applications, components, and libraries often cannot anticipate the context where their software will be used and thus decide which and when messages should be displayed, suppressed, or diverted. Consider the different components in a Logtalk application development and deployment. At the base level, you have the Logtalk compiler and runtime. The compiler writes messages related to e.g. compiling and loading files, compiling entities, and compilation warnings and errors. The runtime may write banner messages or throw execution errors that may result in printing human-level messages. The development environment can be console-based, or you may be using a GUI tool such as PDT. In the latter case, PDT needs to intercept the Logtalk compiler and runtime messages to present the relevant information using its GUI. Then you have all the other components in a typical application. For example, your own libraries and third-party libraries. The libraries may want to print messages on their own, e.g. banners, debugging information, or logging information. As you assemble all your application components, you want to have the final word on which messages are printed, where, and when. Uncontrolled message printing by libraries could potentially disrupt application flow, expose implementation details, spam the user with irrelevant details, or break user interfaces.

The solution is to decouple the calls to print a message from the actual printing of the output text. The same is true for calls to read user input. By decoupling the call to input some data from the actual read of the data, we can easily switch from, for example, a command-line interface to a GUI input dialog or even automate providing the data (e.g., when automating testing of the user interaction).

Logtalk provides a solution based on the *structured message printing mechanism* that was introduced by Quintus Prolog, where it was apparently implemented by Dave Bowen (thanks to Richard O’Keefe for the historical bits). This mechanism gives the programmer full control of message printing, allowing it to filter, rewrite, or redirect any message. Variations of this mechanism can also be found in some Prolog systems, including SICStus Prolog, SWI-Prolog, and YAP. Based on this mechanism, Logtalk introduces an extension that also allows abstracting asking a user for input. Both mechanisms are implemented by the `logtalk` built-in object and described in this section. The message printing mechanism is extensively used by the Logtalk compiler itself and by the developer tools. The question-asking mechanism is used e.g. in the `debugger` tool.

1.15.1 Printing messages

The main predicate for printing a message is `logtalk::print_message/3`. A simple example, using the Logtalk runtime, is:

```
| ?- logtalk::print_message(banner, core, banner).
```

```
Logtalk 3.23.0
```

```
Copyright (c) 1998-2018 Paulo Moura
```

```
yes
```

The first argument of the predicate is the kind of message that we want to print. In this case, we use `banner` to indicate that we are printing a product name and copyright banner. An extensive list of message kinds is supported by default:

banner

banner messages (used e.g. when loading tools or main application components; can be suppressed by setting the `report` flag to `warnings` or `off`)

help

messages printed in reply to the user asking for help (mostly for helping port existing Prolog code)

information and information(Group)

messages usually printed in reply to a user's request for information

silent and silent(Group)

not printed by default (but can be intercepted using the `message_hook/4` predicate)

comment and comment(Group)

useful but usually not essential messages (can be suppressed by setting the `report` flag to `warnings` or `off`)

warning and warning(Group)

warning messages (generated e.g. by the compiler; can be suppressed by turning off the `report` flag)

error and error(Group)

error messages (generated e.g. by the compiler)

debug, debug(Group)

debugging messages (by default, only printed when the `debug` flag is turned on; the `print_message/3` goals for these messages are suppressed by the compiler when the `optimize` flag is turned on)

question, question(Group)

questions to a user

Using a compound term allows easy partitioning of messages of the same kind in different groups. Note that you can define your own alternative message kind identifiers for your own components, together with suitable definitions for their associated prefixes and output streams.

The second argument of `print_message/3` represents the *component* defining the message being printed. In this context, *component* is a generic term that can designate, e.g., a tool, a library, or some sub-system in a large application. In our example, the component name is `core`, identifying the Logtalk compiler/runtime. This argument was introduced to provide multiple namespaces for message terms and thus simplify programming-in-the-large by allowing easy filtering of all messages from a specific component and also avoiding conflicts when two components happen to define the same message term (e.g., `banner`). Users should choose and use a unique name for a component, which usually is the name of the component itself. For example, all messages from the `lgtunit` tool use `lgtunit` for the component argument. The compiler and runtime are interpreted as a single component designated as `core`.

The third argument of `print_message/3` is the message itself, represented by a term. In the above example, the message term is `banner`. Using a term to represent a message instead of a string with the message text itself has significant advantages. Notably, it allows using a compound term for easy parameterization of the message text and simplifies machine processing, localization of applications, and message interception. For example:

```
| ?- logtalk::print_message(comment, core, redefining_entity(object, foo)).

% Redefining object foo
yes
```

1.15.2 Message tokenization

The use of message terms requires a solution for generating the actual text of the messages. This is supported by defining grammar rules for the `logtalk::message_tokens//2` multifile non-terminal, which translates a message term, for a given component, to a list of tokens. For example:

```
:- multifile(logtalk::message_tokens//2).
:- dynamic(logtalk::message_tokens//2).

logtalk::message_tokens(redefining_entity(Type, Entity), core) -->
    ['Redefining ~w ~q'-[Type, Entity], nl].
```

The following tokens can be used when translating a message:

at_same_line

Signals a following part to a multi-part message with no line break in between; this token is ignored when it's not the first in the list of tokens

tab(Expression)

Evaluate the argument as an arithmetic expression and write the resulting number of spaces; this token is ignored when the number of spaces is not positive

nl

Change line in the output stream

flush

Flush the output stream (by calling the `flush_output/1` standard predicate)

Format-Arguments

Format must be an atom and Arguments must be a list of format arguments (the token arguments are passed to a call to the `format/3` de facto standard predicate)

term(Term, Options)

Term can be any term and Options must be a list of valid `write_term/3` output options (the token arguments are passed to a call to the `write_term/3` standard predicate)

ansi(Attributes, Format, Arguments)

Taken from SWI-Prolog; by default, do nothing; can be used for styled output

begin(Kind, Var)

Taken from SWI-Prolog; by default, do nothing; can be used together with `end(Var)` to wrap a sequence of message tokens

end(Var)

Taken from SWI-Prolog; by default, do nothing

The logtalk object also defines public predicates for printing a list of tokens, for hooking into printing an individual token, and for setting default output streams and message prefixes. For example, the SWI-Prolog adapter file uses the print message token hook predicate to enable coloring of messages printed on a console.

1.15.3 Meta-messages

Defining tokenization rules for every message is not always necessary, however. Logtalk defines several *meta-messages* that are handy for simple cases and temporary messages only used during application development, notably debugging messages. See the [Debugging messages](#) section and the [logtalk built-in object](#) remarks section for details.

1.15.4 Defining message prefixes and output streams

The `logtalk::message_prefix_stream/4` hook predicate can be used to define a message line prefix and an output stream for printing messages of a given kind and component. For example:

```
:- multifile(logtalk::message_prefix_stream/4).
:- dynamic(logtalk::message_prefix_stream/4).

logtalk::message_prefix_stream(comment, my_app, '% ', user_output).
logtalk::message_prefix_stream(warning, my_app, '* ', user_error).
```

A single clause at most is expected per message kind and component pair. When this predicate is not defined for a given kind and component pair, the following defaults are used:

```
kind_prefix_stream(banner,      ' ',      user_output).
kind_prefix_stream(help,       ' ',      user_output).
kind_prefix_stream(question,   ' ',      user_output).
kind_prefix_stream(question(_), ' ',      user_output).
kind_prefix_stream(information, '% ',     user_output).
kind_prefix_stream(information(_), '% ',   user_output).
kind_prefix_stream(comment,    '% ',     user_output).
kind_prefix_stream(comment(_), '% ',     user_output).
kind_prefix_stream(warning,    '* ',     user_error).
kind_prefix_stream(warning(_), '* ',     user_error).
kind_prefix_stream(error,      '! ',     user_error).
kind_prefix_stream(error(_),   '! ',     user_error).
kind_prefix_stream(debug,      '>>> ',   user_error).
kind_prefix_stream(debug(_),   '>>> ',   user_error).
```

When the message kind is unknown, information is used instead.

1.15.5 Defining message prefixes and output files

Some applications require copying and saving messages without diverting them from their default stream. For simple cases, this can be accomplished by intercepting the messages using the `logtalk::message_hook/4` multifile hook predicate (see next section). In more complex cases, where messages are already intercepted for a different purpose, it can be tricky to use multiple definitions of the `message_hook/4` predicate as the order of the clauses of a multiple predicate cannot be assumed in general (for all `message_hook/4` predicate definitions to run, all but the last one to be called must fail). Using a single *master* definition is also not ideal as it would result in strong coupling instead of a clean separation of concerns.

The experimental `logtalk::message_prefix_file/6` hook predicate can be used to define a message line prefix and an output file for copying messages of a given kind and component pair. For example:

```
:- multifile(logtalk::message_prefix_file/6).
:- dynamic(logtalk::message_prefix_file/6).

logtalk::message_prefix_file(error, app, '!', 'log.txt', append, []).
logtalk::message_prefix_file(warning, app, '!', 'log.txt', append, []).
```

A single clause at most is expected per message kind and component pair.

This predicate is called by default by the message printing mechanism. Definitions of the `message_hook/4` hook predicate are free to decide if the `logtalk::message_prefix_file/6` predicate should be called and acted upon.

1.15.6 Intercepting messages

Calls to the `logtalk::print_message/3` predicate can be intercepted by defining clauses for the `logtalk::message_hook/4` multifile hook predicate. This predicate can suppress, rewrite, and divert messages.

As a first example, assume that you want to make Logtalk startup less verbose by suppressing printing of the default compiler flag values. This can be easily accomplished by defining the following category in a settings file:

```
:- category(my_terse_logtalk_startup_settings).

:- multifile(logtalk::message_hook/4).
:- dynamic(logtalk::message_hook/4).

logtalk::message_hook(default_flags, comment(settings), core, _).

:- end_category.
```

The printing message mechanism automatically calls the `message_hook/4` hook predicate. When this call succeeds, the mechanism assumes that the message has been successfully handled.

As another example, assume that you want to print all otherwise silent compiler messages:

```
:- category(my_verbose_logtalk_message_settings).

:- multifile(logtalk::message_hook/4).
:- dynamic(logtalk::message_hook/4).

logtalk::message_hook(_Message, silent, core, Tokens) :-
    logtalk::message_prefix_stream(comment, core, Prefix, Stream),
    logtalk::print_message_tokens(Stream, Prefix, Tokens).

logtalk::message_hook(_Message, silent(Key), core, Tokens) :-
    logtalk::message_prefix_stream(comment(Key), core, Prefix, Stream),
    logtalk::print_message_tokens(Stream, Prefix, Tokens).

:- end_category.
```

1.15.7 Asking questions

Logtalk *structured question-asking* mechanism complements the message printing mechanism. It provides an abstraction for the common task of asking a user a question and reading back its reply. By default, this mechanism writes the question, writes a prompt, and reads the answer using the current user input and output streams but allows all steps to be intercepted, filtered, rewritten, and redirected. Two typical examples are using a GUI dialog for asking questions and automatically providing answers to specific questions.

The question-asking mechanism works in tandem with the message printing mechanism, using it to print the question text and a prompt. It provides an asking predicate and a hook predicate, both declared and defined in the logtalk built-in object. The asking predicate, `logtalk::ask_question/5`, is used for asking a question and reading the answer. Assume that we defined the following message tokenization and question prompt and stream:

```
:- category(hitchhikers_guide_to_the_galaxy).

:- multifile(logtalk::message_tokens//2).
:- dynamic(logtalk::message_tokens//2).

% abstract the question text using the atom ultimate_question;
% the second argument, hitchhikers, is the application component
logtalk::message_tokens(ultimate_question, hitchhikers) -->
  ['The answer to the ultimate question of life, the universe and everything is?'-[],_
  ↪nl].

:- multifile(logtalk::question_prompt_stream/4).
:- dynamic(logtalk::question_prompt_stream/4).

% the prompt is specified here instead of being part of the question text
% as it will be repeated if the answer doesn't satisfy the question closure
logtalk::question_prompt_stream(question, hitchhikers, '> ', user_input).

:- end_category.
```

After compiling and loading this category, we can now ask the ultimate question:

```
| ?- logtalk::ask_question(question, hitchhikers, ultimate_question, '==(42), N).

The answer to the ultimate question of life, the universe and everything is?
> 42.

N = 42
yes
```

Note that the fourth argument, `'==(42)` in our example, is a *closure* that is used to check the answers provided by the user. The question is repeated until the goal constructed by extending the closure with the user answer succeeds. For example:

```
| ?- logtalk::ask_question(question, hitchhikers, ultimate_question, '==(42), N).
The answer to the ultimate question of life, the universe and everything is?
> icecream.
> tea.
> 42.
```

(continues on next page)

(continued from previous page)

```
N = 42
yes
```

Practical usage examples of this mechanism can be found, e.g., in the debugger tool where it's used to abstract the user interaction when tracing a goal execution in debug mode.

1.15.8 Intercepting questions

Calls to the `logtalk::ask_question/5` predicate can be intercepted by defining clauses for the `logtalk::question_hook/6` multifile hook predicate. This predicate can suppress, rewrite, and divert questions. For example, assume that we want to automate testing and thus cannot rely on someone manually providing answers:

```
:- category(hitchhikers_fixed_answers).

    :- multifile(logtalk::question_hook/6).
    :- dynamic(logtalk::question_hook/6).

    logtalk::question_hook(ultimate_question, question, hitchhikers, _, _, 42).

:- end_category.
```

After compiling and loading this category, trying the question again will now skip asking the user:

```
| ?- logtalk::ask_question(question, hitchhikers, ultimate_question, '==(42), N).

N = 42
yes
```

In a practical case, the fixed answer would be used for follow-up goals being tested. The question-answer read loop (which calls the question check closure) is not used when a fixed answer is provided using the `logtalk::question_hook/6` predicate thus preventing the creation of endless loops. For example, the following query succeeds:

```
| ?- logtalk::ask_question(question, hitchhikers, ultimate_question, '==(41), N).

N = 42
yes
```

Note that the `logtalk::question_hook/6` predicate takes as argument the closure specified in the `logtalk::ask_question/5` call, allowing a fixed answer to be checked before being returned.

1.15.9 Multi-threading applications

When writing multi-threading applications, user-defined predicates calling methods such as `print_message/3` or `ask_question/5` may need to be declared synchronized in order to avoid race conditions.

1.16 Term and goal expansion

Logtalk supports a *term and goal expansion mechanism* that can be used to define source-to-source transformations. Two common uses are the definition of language extensions and domain-specific languages.

Logtalk improves upon the term-expansion mechanism found on some Prolog systems by providing the user with fine-grained control on *if*, *when*, and *how* expansions are applied. It allows declaring in a source file itself which expansions, if any, will be used when compiling it. It allows declaring which expansions will be used when compiling a file using compile and loading predicate options. It also allows defining a default expansion for all source files. It defines a concept of *hook objects* that can be used as building blocks to create custom and reusable *expansion workflows* with explicit and well-defined semantics. It prevents the simple act of loading expansion rules affecting subsequent compilation of files. It prevents conflicts between groups of expansion rules of different origins. It avoids a set of buggy expansion rules from breaking other sets of expansion rules.

1.16.1 Defining expansions

Term and goal expansions are defined using, respectively, the predicates *term_expansion/2* and *goal_expansion/2*, which are declared in the *expanding* built-in protocol. Note that, unlike Prolog systems also providing these two predicates, they are **not** declared as *multifile predicates* in the protocol. This design decision is key for giving the programmer full control of the expansion process and preventing the problems inflicting most of the Prolog systems that provide a term-expansion mechanism.

An example of an object defining expansion rules:

```
:- object(an_object,
    implements(expanding)).

    term_expansion(ping, pong).
    term_expansion(
        colors,
        [white, yellow, blue, green, read, black]
    ).

    goal_expansion(a, b).
    goal_expansion(b, c).
    goal_expansion(X is Expression, true) :-
        catch(X is Expression, _, fail).

:- end_object.
```

These predicates can be explicitly called using the *expand_term/2* and *expand_goal/2* built-in methods or called automatically by the compiler when compiling a source file (see the section below on *hook objects*).

In the case of source files referenced in *include/1* directives, expansions are only applied automatically when the directives are found in source files, not when used as arguments in the *create_object/4*, *create_protocol/3*, and *create_category/4*, predicates. This restriction prevents inconsistent results when, for example, an expansion is defined for a predicate with clauses found in both an included file and as argument in a call to the *create_object/4* predicate.

Clauses for the *term_expansion/2* predicate are called until one of them succeeds. The returned expansion can be a single term or a list of terms (including the empty list). For example:

```
| ?- an_object::expand_term(ping, Term).

Term = pong
yes

| ?- an_object::expand_term(colors, Colors).

Colors = [white, yellow, blue, green, read, black]
yes
```

When no `term_expansion/2` clause applies, the same term that we are trying to expand is returned:

```
| ?- an_object::expand_term(sounds, Sounds).

Sounds = sounds
yes
```

Clauses for the `goal_expansion/2` predicate are recursively called on the expanded goal until a fixed point is reached. For example:

```
| ?- an_object::expand_goal(a, Goal).

Goal = c
yes

| ?- an_object::expand_goal(X is 3+2*5, Goal).

X = 13,
Goal = true
yes
```

When no `goal_expansion/2` clause applies, the same goal that we are trying to expand is returned:

```
| ?- an_object::expand_goal(3 == 5, Goal).

Goal = (3==5)
yes
```

The goal-expansion mechanism prevents an infinite loop when expanding a goal by checking that a goal to be expanded was not the result from a previous expansion of the same goal. For example, consider the following object:

```
:- object(fixed_point,
implements(expanding)).

goal_expansion(a, b).
goal_expansion(b, c).
goal_expansion(c, (a -> b; c)).

:- end_object.
```

The expansion of the goal `a` results in the goal `(a -> b; c)` with no attempt to further expand the `a`, `b`, and `c` goals as they have already been expanded.

Goal-expansion applies to goal arguments of control constructs, meta-arguments in built-in or user defined

meta-predicates, meta-arguments in local user-defined meta-predicates, meta-arguments in meta-predicate messages when static binding is possible, and `initialization/1`, `if/1`, and `elif/1` directives.

1.16.2 Expanding grammar rules

A common term expansion is the translation of grammar rules into predicate clauses. This transformation is performed automatically by the compiler when a source file entity defines grammar rules. It can also be done explicitly by calling the `expand_term/2` built-in method. For example:

```
| ?- logtalk::expand_term((a --> b, c), Clause).
```

```
Clause = (a(A,B) :- b(A,C), c(C,B))
yes
```

Note that the default translation of grammar rules can be overridden by defining clauses for the `term_expansion/2` predicate.

1.16.3 Bypassing expansions

Terms and goals wrapped by the `{}/1` control construct are not expanded. For example:

```
| ?- an_object::expand_term({ping}, Term).
```

```
Term = {ping}
yes
```

```
| ?- an_object::expand_goal({a}, Goal).
```

```
Goal = {a}
yes
```

This also applies to source file terms and source file goals when using hook objects (discussed next).

1.16.4 Hook objects

Term and goal expansion of a source file during its compilation is performed by using *hook objects*. A hook object is simply an object implementing the `expanding` built-in protocol and defining clauses for the term and goal expansion hook predicates. Hook objects must be compiled and loaded prior to being used to expand a source file.

To compile a source file using a hook object, we can use the `hook` compiler flag in the second argument of the `logtalk_compile/2` and `logtalk_load/2` built-in predicates. For example:

```
| ?- logtalk_load(source_file, [hook(hook_object)]).
...
```

In alternative, we can use a `set_logtalk_flag/2` directive in the source file itself. For example:

```
:- set_logtalk_flag(hook, hook_object).
```

To use multiple hook objects in the same source file, simply write each directive before the block of code that it should handle. For example:

```
:- object(h1,
    implements(expanding)).

    term_expansion((:- public(a/0)), (:- public(b/0))).
    term_expansion(a, b).

:- end_object.
```

```
:- object(h2,
    implements(expanding)).

    term_expansion((:- public(a/0)), (:- public(c/0))).
    term_expansion(a, c).

:- end_object.
```

```
:- set_logtalk_flag(hook, h1).

:- object(s1).

    :- public(a/0).
    a.

:- end_object.

:- set_logtalk_flag(hook, h2).

:- object(s2).

    :- public(a/0).
    a.

:- end_object.
```

```
| ?- {h1, h2, s}.
...

| ?- s1::b.
yes

| ?- s2::c.
yes
```

It is also possible to define a default hook object by defining a global value for the hook flag by calling the [set_logtalk_flag/2](#) predicate. For example:

```
| ?- set_logtalk_flag(hook, hook_object).

yes
```

Note that, due to the `set_logtalk_flag/2` directive being local to a source file, using it to specify a hook object will override any defined default hook object or any hook object specified as a `logtalk_compile/2` or

logtalk_load/2 predicate compiler option for compiling or loading the source file.

Note

Clauses for the `term_expansion/2` and `goal_expansion/2` predicates defined within an object or a category are never used in the compilation of the object or the category itself.

1.16.5 Virtual source file terms and loading context

When using a hook object to expand the terms of a source file, two virtual file terms are generated: `begin_of_file` and `end_of_file`. These terms allow the user to define term-expansions before and after the actual source file terms.

Logtalk also provides a `logtalk_load_context/2` built-in predicate that can be used to access the compilation/loading context when performing expansions. The `logtalk` built-in object also provides a set of predicates that can be useful, notably when adding Logtalk support for language extensions originally developed for Prolog.

As an example of using the virtual terms and the `logtalk_load_context/2` predicate, assume that you want to convert plain Prolog files to Logtalk by wrapping the Prolog code in each file using an object (named after the file) that implements a given protocol. This could be accomplished by defining the following hook object:

```
:- object(wrapper(_Protocol_),
    implements(expanding)).

    term_expansion(begin_of_file, (:- object(Name,implements(_Protocol_))) :-
        logtalk_load_context(file, File),
        os::decompose_file_name(File,_ , Name, _).

    term_expansion(end_of_file, (:- end_object)).

:- end_object.
```

Assuming, e.g., `my_car.pl` and `lease_car.pl` files to be wrapped and a `car_protocol` protocol, we could then load them using:

```
| ?- logtalk_load(
    ['my_car.pl', 'lease_car.pl'],
    [hook(wrapper(car_protocol))])
    ).
yes
```

Note

When a source file also contains plain Prolog directives and predicates, these are term-expanded but not goal-expanded (with the exception of the `initialization/1`, `if/1`, and `elif/1` directives, where the goal argument is expanded to improve code portability across backends).

1.16.6 Default compiler expansion workflow

When *compiling a source file*, the compiler will first try, by default, the source file-specific hook object specified using a `set_logtalk_flag/2` directive, if defined. If that expansion fails, it tries the hook object specified using the `hook/1` compiler option in the `logtalk_compile/2` or `logtalk_load/2` goal that compiles or loads the file, if defined. If that expansion fails, it tries the default hook object, if defined. If that expansion also fails, the compiler tries the Prolog dialect-specific expansion rules found in the *adapter file* (which are used to support non-standard Prolog features).

1.16.7 User defined expansion workflows

Sometimes we have multiple hook objects that we need to combine and use in the compilation of a source file. Logtalk includes a *hook_flows* library that supports two basic expansion workflows: a *pipeline* of hook objects, where the expansion results from a hook object are fed to the next hook object in the pipeline, and a *set* of hook objects, where expansions are tried until one of them succeeds. These workflows are implemented as parametric objects, allowing combining them to implement more sophisticated expansion workflows. There is also a *hook_objects* library that provides convenient hook objects for defining custom expansion workflows. This library includes a hook object that can be used to restore the default expansion workflow used by the compiler.

For example, assuming that you want to apply the Prolog backend-specific expansion rules defined in its adapter file, using the *backend_adapter_hook* library object, passing the resulting terms to your own expansion when compiling a source file, we could use the goal:

```
| ?- logtalk_load(
    source,
    [hook(hook_pipeline([backend_adapter_hook, my_expansion]))]
).
```

As a second example, we can prevent expansion of a source file using the library object *identity_hook* by adding as the first term in a source file the directive:

```
:- set_logtalk_flag(hook, identity_hook).
```

The file will be compiled as-is as any hook object (specified as a compiler option or as a default hook object) and any backend adapter expansion rules are overridden by the directive.

1.16.8 Using Prolog defined expansions

In order to use clauses for the `term_expansion/2` and `goal_expansion/2` predicates defined in plain Prolog, simply specify the pseudo-object `user` as the hook object when compiling source files. When using *backend Prolog compilers* that support a module system, it can also be specified a module containing clauses for the expanding predicates as long as the module name doesn't coincide with an object name. When defining a custom workflow, the library object *prolog_module_hook/1* can be used as a workflow step. For example, assuming a module `functions` defining expansion rules that we want to use:

```
| ?- logtalk_load(
    source,
    [hook(hook_set([prolog_module_hook(functions), my_expansion]))]
).
```

But note that Prolog module libraries may provide definitions of the expansion predicates that are not compatible with the Logtalk compiler. In particular, when setting the hook object to `user`, be aware of any Prolog

library that is loaded, possibly by default or implicitly by the Prolog system, that may be contributing definitions of the expansion predicates. It is usually safer to define a specific hook object for combining multiple expansions in a fully controlled way.

Note

The user object declares `term_expansion/2` and `goal_expansion/2` as multifile and dynamic predicates. This helps in avoiding predicate existence errors when compiling source files with the hook flag set to user as these predicates are only natively declared by some of the supported backend Prolog compilers.

1.16.9 Debugging expansions

The `term_expansion/2` and `goal_expansion/2` predicates can be *debugged* like any other object predicates. Note that expansions can often be manually tested by sending `expand_term/2` and `expand_goal/2` messages to a hook object with the term or goal whose expansion you want to check as argument. An alternative to the debugging tools is to use a *monitor* for the runtime messages that call the predicates. For example, assume a `expansions_debug.lgt` file with the contents:

```
:- initialization(
    define_events(after, edcg, _, _, expansions_debug)
).

:- object(expansions_debug,
    implements(monitors)).

    after(edcg, term_expansion(T,E), _) :-
        writeq(term_expansion(T,E)), nl.

:- end_object.
```

We can use this monitor to help debug the expansion rules of the `edcg` library when applied to the `edcgs` example using the queries:

```
| ?- {expansions_debug}.
...

| ?- set_logtalk_flag(events, allow).
yes

| ?- {edcgs(loader)}.
...
term_expansion(begin_of_file,begin_of_file)
term_expansion((:-object(gemini)),[(:-object(gemini)),(:-op(1200,xfx,-->))])
term_expansion(acc_info(castor,A,B,C,true),[])
term_expansion(pass_info(pollux),[])
term_expansion(pred_info(p,1,[castor,pollux]),[])
term_expansion(pred_info(q,1,[castor,pollux]),[])
term_expansion(pred_info(r,1,[castor,pollux]),[])
term_expansion((p(A)-->B is A+1,q(B),r(B)),(p(A,C,D,E):-B is A+1,q(B,C,F,E),r(B,F,D,E)))
term_expansion((q(A)-->[]),(q(A,B,B,C):-true))
term_expansion((r(A)-->[]),(r(A,B,B,C):-true))
```

(continues on next page)

(continued from previous page)

```
term_expansion(end_of_file,end_of_file)
...
```

This solution does not require compiling the edcg hook object in debug mode or access to its source code (e.g., to modify its expansion rules to emit debug messages). We could also simply use the user pseudo-object as the monitor object:

```
| ?- assertz((
    after(_, term_expansion(T,E), _) :-
        writeq(term_expansion(T,E)), nl
    )).
yes

| ?- define_events(after, edcg, _, Sender, user).
yes
```

Another alternative is to use a pipeline of hook objects with the library `hook_pipeline/1` and `write_to_stream_hook` objects to write the expansion results to a file. For example, using the `unique.lgt` test file from the edcgs library directory:

```
| ?- {hook_flows(loader), hook_objects(loader)}.
...

| ?- open('unique_expanded.lgt', write, Stream),
    logtalk_compile(
        unique,
        [hook(hook_pipeline([edcg,write_to_stream_hook(Stream,[quoted(true)])))]
    ),
    close(Stream).
...
```

The generated `unique_expanded.lgt` file will contain the clauses resulting from the expansion of the EDCG rules found in the `unique.lgt` file by the edcg hook object expansion.

1.17 Documenting

Assuming that the `source_data` flag is turned on, the compiler saves all relevant documenting information collected when compiling a source file. The provided `lgt doc` tool can access this information by using the `reflection` support and generate a documentation file for each compiled entity (object, protocol, or category) in XML format. Contents of the XML file include the entity name, type, and compilation mode (static or dynamic), the entity relations with other entities, and a description of any declared predicates (name, compilation mode, scope, ...). The XML documentation files can be enriched with arbitrary user-defined information, either about an entity or about its predicates, by using documentation directives. The `lgt doc` tool includes POSIX and Windows scripts for converting the XML documentation files to several final formats (such as HTML and PDF).

Logtalk supports two documentation directives for providing arbitrary user-defined information about an entity or a predicate. These two directives complement other directives that also provide important documentation information, such as the `mode/2` and `meta_predicate/1` directives.

1.17.1 Entity documenting directives

Arbitrary user-defined entity information can be represented using the *info/1* directive:

```
:- info([
    Key1 is Value1,
    Key2 is Value2,
    ...
]).
```

In this pattern, keys should be atoms and values should be bound terms. The following keys are predefined and may be processed specially by Logtalk tools:

comment

Comment describing the entity purpose (an atom). End the comment with a period (full stop). As a style guideline, don't use overly long comments. If you need to provide additional details, use the `fails_if` and `remarks` keys.

author

Entity author(s) (an atom or a compound term `{entity}` where `entity` is the name of an XML entity in a user-defined `custom.ent` file).

version

Version number (a `Major:Minor:Patch` compound term) Following the [Semantic Versioning guidelines](#) is strongly advised.

date

Date of last modification in ISO 8601 standard format (Year-Month-Day where Year, Month, and Day are integers).

parameters

Parameter names and descriptions for parametric entities (a list of Name-Description pairs where both names and descriptions are atoms). End the Description with a period (full stop).

parnames

Parameter names for parametric entities (a list of atoms; a simpler version of the previous key, used when parameter descriptions are deemed unnecessary).

copyright

Copyright notice for the entity source code (an atom or a compound term `{entity}` where `entity` is the name of an XML entity defined in a user-defined `custom.ent` file).

license

License terms for the entity source code; usually, just the license name (an atom or a compound term `{entity}` where `entity` is the name of an XML entity in a user-defined `custom.ent` file). License names should, whenever possible, be a license identifier as specified in the [SPDX standard](#).

remarks

List of general remarks about the entity using Topic-Text pairs where both the topic and the text must be atoms. End the Text with a period (full stop).

see_also

List of related entities (using the entity identifiers, which can be atoms or compound terms).

For example:

```
:- info([
    version is 2:1:0,
    author is 'Paulo Moura',
```

(continues on next page)

(continued from previous page)

```

date is 2000-11-20,
comment is 'Building representation.',
diagram is 'UML Class Diagram #312'
]).

```

Use only the keywords that make sense for your application, and remember that you are free to invent your own keywords. All key-value pairs can be retrieved programmatically using the [reflection API](#) and are visible to the [lgt doc](#) tool (which includes them in the generated documentation).

1.17.2 Predicate documenting directives

Arbitrary user-defined predicate information can be represented using the [info/2](#) directive:

```

:- info(Name/Arity, [
    Key1 is Value1,
    Key2 is Value2,
    ...
]).

```

The first argument can also a grammar rule non-terminal indicator, `Name//Arity`. Keys should be atoms. Values should be bound terms. The following keys are predefined and may be processed specially by Logtalk tools:

comment

Comment describing the predicate (or non-terminal) purpose (an atom). End the comment with a period (full stop). As a style guideline, don't use overly long comments. If you need to provide additional details, use the `remarks` key.

fails_if

Comment describing failing conditions for the predicate. As a style guideline, don't use overly long comments. If you need to provide additional details, use the `remarks` key.

arguments

Names and descriptions of predicate arguments for pretty print output (a list of Name-Description pairs where both names and descriptions are atoms). End the Description with a period (full stop).

argnames

Names of predicate arguments for pretty print output (a list of atoms; a simpler version of the previous key, used when argument descriptions are deemed unnecessary).

allocation

Objects where we should define the predicate. Some possible values are `container`, `descendants`, `instances`, `classes`, `subclasses`, and `any`.

redefinition

Describes if a predicate is expected to be redefined and, if so, in what way. Some possible values are `never`, `free`, `specialize`, `call_super_first`, `call_super_last`.

exceptions

List of possible exceptions thrown by the predicate using Description-Exception pairs. The description must be an atom. The exception term must be a ground term.

examples

List of typical predicate call examples using the format Description-Goal-Bindings. The description must be an atom with the goal term sharing variables with the bindings. The variable bindings term uses the format `{Variable = Term, ...}`. When there are no variable bindings, the success or failure

of the predicate call should be represented by the terms {true} or {false}, respectively (you can also use in alternative the terms {yes} or {no}).

remarks

List of general remarks about the predicate using Topic-Text pairs where both the topic and the text must be atoms. End the Text with a period (full stop).

since

Version that added the predicate (Major:Minor:Patch).

see_also

List of related predicates and non-terminals (using the predicate and non-terminal indicators).

For example:

```
:- info(color/1, [
    comment is 'Table of defined colors.',
    argnames is ['Color'],
    constraint is 'Up to four visible colors allowed.',
    examples is [
        'Check that the color blue is defined' - color(blue) - {true}
    ]
]).
```

As with the `info/1` directive, use only the keywords that make sense for your application and remember that you are free to invent your own keywords. All key-value pairs can also be retrieved programmatically using the [reflection API](#) and are visible to the [lgt doc](#) tool (which includes them in the generated documentation).

1.17.3 Describing predicates

The value of the `comment` key, possibly extended with the `remarks` key, should describe a predicate purpose and, when applicable, the circumstances under which a call may fail. Descriptions should be consistent across library and application APIs. Some guidelines:

1. When starting the description with a verb, use the *third-person singular simple present form*. For example, write 'Runs ...', 'Calls ...', 'Compares ...', 'Parses ...', 'Generates ...', 'Converts ...', 'Creates ...', 'Maps ...', 'Merges ...', 'Finds ...', etc.
2. Predicates that are pure logical relations often have descriptions starting with 'True iff ...' or 'True if ...'.
3. Predicates with multiple solutions often have descriptions starting with 'Enumerates, by backtracking, all ...' or 'Enumerates, by backtracking, the ...'.
4. Predicate call failure conditions often have descriptions with one or more sentences starting with 'Fails when ...' or 'Fails if ...'.

If you're not sure how best to describe a predicate, look for examples in the Logtalk libraries and developer tools APIs documentation.

1.17.4 Documenting predicate exceptions

As described above, the `info/2` predicate directive supports an `exceptions` key that allows us to list all exceptions that may occur when calling a predicate. For example:

```
:- info(check_option/1, [
    comment is 'Succeeds if the option is valid. Throws an error otherwise.',
    argnames is ['Option'],
    exceptions is [
        '`Option` is a variable' - instantiation_error,
        '`Option` is neither a variable nor a compound term' - type_error(compound, 'Option
→ '),
        '`Option` is a compound term but not a valid option' - domain_error(option, 'Option
→ ')
    ]
]).
```

When possible, only standard exceptions should be used. See e.g. the [error handling methods](#) section for a full list. The argument names should be the same as those provided in the arguments or `argnames` keys. Exceptions are usually listed starting with instantiation and un instantiation errors, followed by type errors, and then domain errors. These may then be followed by permission, existence, evaluation, representation, or resource errors.

For each exception, use of *controlled language* as found, e.g., in the ISO Prolog Core standard and this Handbook is advised. Some examples:

Instantiation error when one or more arguments cannot be a variable

Argument is a variable

Argument1 and Argument2 are variables

Instantiation error when a closed list with bound elements is required

Argument is a partial list or a list with an element `Element` which is a variable

Uninstantiation error when an argument is not a variable

Argument is not a variable

Type error when an argument is not a variable but also not of the expected type

Argument is neither a variable nor a TYPE

Argument is neither a partial list nor a list

Type error when an element of a list is not a variable but is not of the expected type

An element `Element` of the Argument list is neither a variable nor a TYPE

Domain error when an argument is of the correct type but not in the expected domain

Argument is a TYPE but not a valid DOMAIN

Argument is an integer that is less than zero

Domain error when an element of a list is of the correct type but not in the expected domain

An element `Element` of the Argument list is a TYPE but not a valid DOMAIN

Existence error when an entity of a given kind does not exist

The KIND Argument does not exist

Other classes of errors have a less rigid style. In case of doubt, look for examples in this Handbook, in the APIs documentation, and in standard documents.

1.17.5 Processing and viewing documenting files

The *lgt doc* tool generates an XML documenting file per entity. It can also generate library, directory, entity, and predicate indexes when documenting libraries and directories. For example, assuming the default file-name extensions, a trace object and a sort(*_*) parametric object will result in *trace_0.xml* and *sort_1.xml* XML files.

Each entity XML file contains references to two other files, an XML specification file and a XSLT stylesheet file. The XML specification file can be either a DTD file (*logtalk_entity.dtd*) or an XML Scheme file (*logtalk_entity.xsd*). The XSLT stylesheet file is responsible for converting the XML files to some desired format such as HTML or PDF. The default names for the XML specification file and the XSL stylesheet file are defined by the *lgt doc* tool but can be overridden by passing a list of options to the tool predicates. The *lgt doc/xml* sub-directory in the Logtalk installation directory contains the XML specification files described above, along with several sample XSL stylesheet files and sample scripts for converting XML documenting files to several formats (e.g., reStructuredText, Markdown, HTML, and PDF). For example, assume that you want to generate the API documentation for the *types* library:

```
| ?- {types(loader)}.  
....  
  
| ?- {lgt doc(loader)}.  
....  
  
| ?- lgt doc::library(types).  
...
```

The above queries will result in the creation of a *xml_docs* in your current directory by default. Assuming that we want to generate Sphinx-based documentation and that we are using a POSIX operating-system, the next steps would be:

```
$ cd xml_docs  
$ lgt2rst -s -m
```

The *lgt2rst* script will ask a few questions (project name, author, version, ...). After its completion, the generated HTML files will be found in the *_build/html* directory by default:

```
$ open _build/html/index.html
```

For Windows operating-systems, PowerShell scripts are available. For example, assuming that we want to generate HTML documentation, we could run in a PowerShell window:

```
PS > cd xml_docs  
PS > lgt2html.ps1 -p saxon
```

After completion, the generated HTML files will be found in the *xml_docs* directory by default.

See the *NOTES* file in the tool directory for details, specially on the XSLT processor dependencies. You may use the supplied sample files as a starting point for generating the documentation of your Logtalk applications.

The Logtalk DTD file, *logtalk_entity.dtd*, contains a reference to a user-customizable file, *custom.ent*, which declares XML entities for source code author names, license terms, and copyright strings. After editing the *custom.ent* file to reflect your personal data, you may use the XML entities on *info/1* documenting directives. For example, assuming that the XML entities are named *author*, *license*, and *copyright* we may write:

```
:- info([
    version is 1:1:0,
    author is {author},
    license is {license},
    copyright is {copyright}
]).
```

The entity references are replaced by the value of the corresponding XML entity when the XML documenting files are processed (**not** when they are generated; this notation is just a shortcut to take advantage of XML entities).

The *lgt doc* tool supports a set of options that can be used to control the generation of the XML documentation files. See the tool documentation for details. There is also a *doclet* tool that allows automating the steps required to generate the documentation for an application.

1.17.6 Inline formatting in comments text

Inline formatting in comments text can be accomplished by using Markdown or reStructuredText syntax and converting XML documenting files to Markdown or reStructuredText files (and these, if required, to e.g. HTML, ePub, or PDF formats). Note that Markdown and reStructuredText common syntax elements are enough for most API documentation:

```
Mark italic text with one asterisk.
Mark bold text with two asterisks.
Mark ``monospaced text`` with two backquotes.
```

Rendering this block as markup gives:

Mark *italic text* with one asterisk. Mark **bold text** with two asterisks. Mark `monospaced text` with two backquotes.

As single backquotes have different purposes in Markdown (monospaced text) and reStructuredText (domain- or application-dependent meaning), never use them. This also avoids doubts if there's an inline formatting typo in text meant to be rendered as monospaced text (usually inline code fragments).

1.17.7 Diagrams

The *diagrams* tool supports a wide range of diagrams that can also help in documenting an application. The generated diagrams can include URL links to both source code and API documentation. They can also be linked, connecting, for example, high level diagrams to detail diagrams. These features allow diagrams to be an effective solution for navigating and understanding the structure and implementation of an application. This tool uses the same *reflection API* as the *lgt doc* tool and thus has access to the same source data. See the tool documentation for details.

1.18 Debugging

The Logtalk distribution includes a command-line *debugger* tool implemented as a Logtalk application using the debugging API. It can be loaded at the top-level interpreter by typing:

```
| ?- logtalk_load(debugger(loader)).
```

It can also be loaded automatically at startup time by using a *settings file*.

The *debugger* tool includes the debugging features found in traditional Prolog debuggers. There are some differences, however, between the usual implementation of Prolog debuggers and the current implementation of the Logtalk debugger that you should be aware of. First, unlike most Prolog debuggers, the Logtalk debugger is not a built-in feature but a regular Logtalk application using documented debugging hook predicates. This translates to a different, although similar, set of debugging features when compared with some of the more sophisticated Prolog debuggers. Second, debugging is only possible for entities compiled in debug mode. When compiling an entity in debug mode, Logtalk decorates clauses with source information to allow tracing of the goal execution. Third, the tool provides several types of breakpoints (for pausing and interacting with the debugger) and also log points, while most Prolog systems are limited to traditional predicate spy points.

1.18.1 Compiling source files in debug mode

Compilation of source files in debug mode is controlled by the *debug* compiler flag. The default value for this flag, usually off, is defined in the *backend adapter files*. Its default value may be changed globally at runtime by calling:

```
| ?- set_logtalk_flag(debug, on).
```

Implicitly, this goal also turns off the optimize flag. In alternative, if we want to compile only some source files in debug mode, we may instead write:

```
| ?- logtalk_load([file1, file2, ...], [debug(on)]).
```

The *logtalk_make/1* built-in predicate can also be used to recompile all loaded files (that were compiled without using explicit values for the *debug* and *optimize* compiler flags in a *logtalk_load/2* call or in a *loader file*, if used) in debug mode:

```
| ?- logtalk_make(debug).
```

With most *backend Prolog compilers*, the *{+d}* top-level shortcut can also be used. After debugging, the files can be recompiled in normal or optimized mode using, respectively, the *{+n}* or *{+o}* top-level shortcuts.

Warning

The `clean` compiler flag should be turned on whenever the `debug` flag is turned on at runtime. This is necessary because debug code would not be generated for files previously compiled in normal or optimized mode if there are no changes to the source files.

After loading the debugger, we may check (or enumerate by backtracking), all loaded entities compiled in debug mode as follows:

```
| ?- debugger::debugging(Entity).
```

To compile only a specific entity in debug mode, use the `set_logtalk_flag/2` directive inside the entity. To compile all entities in a source file in debug mode, use the `set_logtalk_flag/2` directive at the beginning of the file.

1.18.2 Procedure box model

Logtalk uses a *procedure box model* similar to those found on most Prolog systems. The traditional Prolog procedure box model defines four ports (*call*, *exit*, *redo*, and *fail*) for describing control flow when calling a predicate:

```
call
    predicate call
exit
    success of a predicate call
redo
    backtracking into a predicate call
fail
    failure of a predicate call
```

Logtalk, as found on some recent Prolog systems, adds a port for dealing with exceptions thrown when calling a predicate:

```
exception
    predicate call throws an exception
```

In addition to the ports described above, Logtalk adds two more ports, *fact* and *rule*, which show the result of the unification of a goal with, respectively, a fact and a rule head:

```
fact
    unification success between a goal and a fact
rule
    unification success between a goal and a rule head
```

Following Prolog tradition, the user may define for which ports the debugger should pause for user interaction by specifying a list of *leashed* ports. Unleashed ports are just printed with no pause for user interaction when tracing. For example:

```
| ?- debugger::leash([call, exit, fail]).
```

Alternatively, the user may use an atom abbreviation for a pre-defined set of ports. For example:

```
| ?- debugger::leash(loose).
```

The abbreviations defined in Logtalk are similar to those defined on some Prolog compilers:

```
none
    []
loose
    [fact, rule, call]
half
    [fact, rule, call, redo]
tight
    [fact, rule, call, redo, fail, exception]
full
    [fact, rule, call, exit, redo, fail, exception]
```

By default, the debugger pauses at every port for user interaction.

1.18.3 Activating the debugger

The `debugger::trace/0` and `debugger::debug/0` predicates implicitly select the debugger tool as the active *debug handler*. If you have additional debug handlers loaded (e.g., the `ports_profiler` tool), those would no longer be active (there can be only one active debug handler at any given time). The `debugger::nodebug/0` predicate implicitly deselects the debugger tool as the active debug handler.

1.18.4 Defining breakpoints

The debugger tool provides the following breakpoint types where the debugger pauses at a leashed port for user interaction:

- **Predicate breakpoints**
Traditional Prolog spy points are defined using a predicate (or a non-terminal) indicator. The debugger pauses for user input at all ports for the predicate (or non-terminal).
- **Clause breakpoints**
Defined using the location of a clause. The debugger pauses for user input at unification ports.
- **Conditional breakpoints**
Defined using the location of a clause and a condition for pausing for user input at an unification port.
- **Hit count breakpoints**
Defined using the location of a clause and an unification count expression as a condition for pausing for user input at an unification port.

- **Triggered breakpoints**

Defined using the location of a clause and another breakpoint that must be hit first as a condition for pausing for user input at an unification port.

- **Context breakpoints**

Defined using execution context and goal templates as a condition for pausing for user input at all ports.

Clause breakpoints are checked when the current goal successfully unifies with a clause head. To simplify their definition, these are specified using the entity identifier instead of the file name (as all entities share a single namespace, an entity can only be defined in a single file) and the first line number of the clause head. But note that only some Prolog backends provide accurate source file term line numbers. Check the [debugger](#) tool documentation for details.

Defining predicate and clause breakpoints

Predicate and clause breakpoints can be defined using the debugger `spy/1` predicate. The argument can be a predicate indicator (`Name/Arity`), a non-terminal indicator (`Name//Arity`), a clause location (expressed as an `Entity-Line` pair), or a list of breakpoints. Predicate and non-terminal indicators can also be qualified with a specific object or category identifiers. For example:

```
| ?- debugger::spy(person-42).

Clause breakpoint added.
yes

| ?- debugger::spy(foo/2).

Predicate breakpoint added.
yes

| ?- debugger::spy(qux::baz/3).

Entity predicate breakpoint added.
yes

| ?- debugger::spy([foo/4, bar//1, agent-99]).

All specified breakpoints added.
yes
```

Note that setting a clause breakpoint implicitly removes any existing conditional breakpoint, triggered breakpoint, or log point for the same clause.

Unconditional clause and predicate breakpoints can be removed by using the debugger `nospy/1` predicate. The argument can also be a list of breakpoints or a non-instantiated variable, in which case all breakpoints will be removed. For example:

```
| ?- debugger::nospy(_).

All matching predicate and clause breakpoints removed.
yes
```

Defining conditional breakpoints

Conditional clause breakpoints are specified using the debugger `spy/3` predicate. The first two arguments are the entity and the line of a clause head. The third argument is the condition, which can be a lambda expression, an unification count expression (see next section), or another breakpoint (see next below).

The supported lambda expressions are `[Count, N, Goal]>>Condition` and `[Goal]>>Condition` where `Count` is the unification count, `N` is the goal invocation number, and `Goal` is the goal that unified with the clause head; `Condition` is called in the context of the user pseudo-object and must not have any side effects. Some examples:

```
| ?- debugger::spy(planet, 76, [weight(m1,_)]>>true).
```

```
Conditional breakpoint added.  
yes
```

Note that setting a conditional breakpoint will remove any existing clause breakpoint or log point for the same location.

Conditional breakpoints can be removed by using the debugger `nospy/3` predicate. For example:

```
| ?- debugger::nospy(planet, _, _).
```

```
All matching conditional breakpoints removed.  
yes
```

Defining hit count breakpoints

Conditional clause breakpoints that depend on the unification count are known as *hit count* clause breakpoints. The debugger pauses at a hit count breakpoint depending on an unification count expression:

- `>(Count)` - break when the unification count is greater than `Count`
- `>=(Count)` - break when the unification count is greater than or equal to `Count`
- `=(Count)` - break when the unification count is equal to `Count`
- `=(Count)` - break when the unification count is less than or equal to `Count`
- `<(Count)` - break when the unification count is less than `Count`
- `mod(M)` - break when the unification count modulo `M` is zero
- `Count` - break when the unification count is greater than or equal to `Count`

For example:

```
| ?- debugger::spy(planet, 41, =(2)).
```

```
Conditional breakpoint added.  
yes
```

Defining triggered breakpoints

Conditional clause breakpoints that depend on other clause breakpoint or on a log point are known as *triggered* clause breakpoints. The debugger only pauses at a triggered breakpoint if the clause breakpoint or log point it depends on is hit first. For example:

```
| ?- debugger::spy(mars, 98, planet-76).
```

```
Triggered breakpoint added.
yes
```

In this case, the debugger will break for user interaction at the unification port for the clause in the source file defining the `mars` object at line 98 if and only if the debugger paused earlier at the unification port for the clause in the source file defining the `planet` category at line 76.

The debugger prints a `^` character at the beginning of the line for easy recognition of triggered breakpoints.

Defining context breakpoints

A context breakpoint is a tuple describing a message execution context and a goal:

```
(Sender, This, Self, Goal)
```

The debugger pauses for user interaction whenever the breakpoint goal and execution context subsume the goal currently being executed and its execution context. The user may establish any number of context breakpoints as necessary. For example, in order to call the debugger whenever a predicate defined on an object named `foo` is called, we may define the following context breakpoint:

```
| ?- debugger::spy(_, foo, _, _).
```

```
Context breakpoint set.
yes
```

For example, we can spy all calls to a `foo/2` predicate with a *bar* atom in the second argument by setting the condition:

```
| ?- debugger::spy(_, _, _, foo(_, bar)).
```

```
Context breakpoint set.
yes
```

The debugger `nospy/4` predicate may be used to remove all matching breakpoints. For example, the call:

```
| ?- debugger::nospy(_, _, foo, _).
```

```
All matching context breakpoints removed.
yes
```

will remove all context breakpoints where the value of *self* is the atom `foo`.

Removing all breakpoints

We can remove all breakpoints by using the debugger `nospyall/0` predicate:

```
| ?- debugger::nospyall.  
All breakpoints removed.  
yes
```

There's also a `reset/0` predicate that can be used to reset the debugger to its default settings and delete all defined breakpoints and log points.

1.18.5 Defining log points

Logtalk log points are similar to clause breakpoints. Therefore, the line number must correspond to the first line of an entity clause. When the debugger reaches a log point at an unification port, it prints a log message and continues without pausing execution for reading a port command. When the log message is an empty atom, the default port output message is printed. When the log message starts with a `%` character, the default port output message is printed, followed by the log message. In these two cases, the debugger prints a `@` character at the beginning of the line for easy recognition of log points output. When the log message is neither empty nor starts with a `%` character, the log message is printed instead of the default port output message. In this case, the message can contain `$KEYWORD` placeholders that are expanded at runtime. The valid keywords are:

- PORT
- ENTITY
- CLAUSE_NUMBER
- FILE
- LINE
- UNIFICATION_COUNT
- INVOCATION_NUMBER
- GOAL
- PREDICATE
- EXECUTION_CONTEXT
- SENDER
- THIS
- SELF
- METACALL_CONTEXT
- COINDUCTION_STACK
- THREAD

Log points are defined using the `log/3` predicate. For example:

```
| ?- debugger::log(agent, 99, '% At the secret headquarters!').  
Log point added.  
yes
```

(continues on next page)

(continued from previous page)

```
| ?- debugger::log(loop, 42, 'Message $PREDICATE from $SENDER at thread $THREAD').
    Log point added.
yes
```

The logging/3 and nolog/3 predicate can be used to, respectively, query and remove log points. There's also a nologall/0 predicate that removes all log points.

Note that setting a log point will remove any existing clause breakpoint for the same location.

1.18.6 Tracing program execution

Logtalk allows tracing of execution for all objects compiled in debug mode. To start the debugger in trace mode, write:

```
| ?- debugger::trace.
yes
```

Next, type the query to be debugged. For example, using the family example in the Logtalk distribution compiled for debugging:

```
| ?- addams::sister(Sister, Sibling).
    Call: (1) sister(_1082,_1104) ?
    Rule: (1) sister(_1082,_1104) ?
    Call: (2) ::female(_1082) ?
    Call: (3) female(_1082) ?
    Fact: (3) female(morticia) ?
    *Exit: (3) female(morticia) ?
    *Exit: (2) ::female(morticia) ?
    ...
```

While tracing, the debugger will pause for user input at each leashed port, printing an informative message. Each trace line starts with the port, followed by the goal invocation number, followed by the goal. The invocation numbers are unique and allow us to correlate the ports used for a goal. In the output above, you can see, for example, that the goal `::female(_1082)` succeeds with the answer `::female(morticia)`. The debugger also provides determinism information by prefixing the exit port with a `*` character when a call succeeds with choice-points pending, thus indicating that there might be alternative solutions for the goal.

Note that breakpoints are ignored when tracing. But when a breakpoint is set for the current predicate or clause, the debugger prints, before the port name and number, a `+` character for predicate breakpoints, a `#` character for clause breakpoints, a `?` character for conditional clause breakpoints, a `^` for triggered breakpoints, and a `*` character for context breakpoints. For example:

```
| ?- debugger::spy(female/2).
yes

| ?- addams::sister(Sister, Sibling).
    Call: (1) sister(_1078,_1100) ?
    Rule: (1) sister(_1078,_1100) ?
    Call: (2) ::female(_1078) ?
    + Call: (3) female(_1078) ?
```

To stop tracing (but still allowing the debugger to pause at the defined breakpoints), write:

```
| ?- debugger::notrace.  
yes
```

1.18.7 Debugging using breakpoints

Tracing a program execution may generate large amounts of debugging data. Debugging using breakpoints allows the user to concentrate on specific points of the code. To start a debugging session using breakpoints, write:

```
| ?- debugger::debug.  
yes
```

For example, assuming the predicate breakpoint we set in the previous section on the `female/1` predicate:

```
| ?- addams::sister(Sister, Sibling).  
+ Call: (3) female(_1078) ?
```

To stop the debugger, write:

```
| ?- debugger::nodebug.  
yes
```

Note that stopping the debugger does not remove any defined breakpoints or log points.

1.18.8 Debugging commands

The debugger pauses for user interaction at leashed ports when tracing and when hitting a breakpoint. The following commands are available:

- c — creep**
go on; you may use the spacebar, return, or enter keys in alternative
- l — leap**
continues execution until the next breakpoint is found
- s — skip**
skips tracing for the current goal; valid at call, redo, and unification ports
- S - Skip**
similar to skip but displaying all intermediate ports unleashed
- q — quasi-skip**
skips tracing until returning to the current goal or reaching a breakpoint; valid at call and redo ports
- r — retry**
retries the current goal but side-effects are not undone; valid at the fail port
- j — jump**
reads invocation number and continues execution until a port is reached for that number

- z — zap**
reads either a port name and continues execution until that port is reached or a negated port name (e.g. `-exit`) and continues execution until a port other than the negated port is reached
- i — ignore**
ignores goal, assumes that it succeeded; valid at call and redo ports
- f — fail**
forces backtracking; may also be used to convert an exception into a failure
- n — nodebug**
turns off debugging
- N — notrace**
turns off tracing
- @ — command; ! can be used in alternative**
reads and executes a query
- b — break**
suspends execution and starts new interpreter; type `end_of_file` to terminate
- a — abort**
returns to top level interpreter
- Q — quit**
quits Logtalk
- p — print**
writes current goal using the `print/1` predicate if available
- d — display**
writes current goal without using operator notation
- w — write**
writes current goal quoting atoms if necessary
- \$ — dollar**
outputs the compiled form of the current goal (for low-level debugging)
- x — context**
prints execution context
- . — file**
prints file, entity, predicate, and line number information at an unification port
- e — exception**
prints exception term thrown by the current goal
- E — raise exception**
reads and throws an exception term
- = — debugging**
prints debugging information
- < — write depth**
sets the write term depth (set to 0 to reset)
- * — add**
adds a context breakpoint for the current goal
- / — remove**
removes a context breakpoint for the current goal

- + — **add**
adds a predicate breakpoint for the current goal
- — **remove**
removes a predicate breakpoint for the current goal
- # — **add**
adds a breakpoint for the current clause
- | — **remove**
removes a breakpoint for the current clause
- h — **condensed help**
prints list of command options
- ? — **extended help**
prints list of command options

1.18.9 Customizing term writing

Debugging complex applications often requires customizing term writing. The available options are limiting the writing depth of large compound terms and using the `p` command at a leashed port. This command uses the `format/3` de facto standard predicate with the `~p` formatting option to delegate writing the term to the `print/1` predicate. But note that some backends don't support this formatting option.

Term write depth

The terms written by the debugger can be quite large depending on the application being debugged. As described in the previous section, the debugger accepts the `<` command to set the maximum write term depth for compound terms. This command requires that the used *backend Prolog compiler* supports the non-standard but common `max_depth/1` option for the `write_term/3` predicate. When the compound term being written is deeply nested, the sub-terms are only written up to the specified depth with the omitted sub-terms replaced usually by `...`. For example:

```
| ?- write_term([0,1,2,3,4,5,6,7,8,9], [max_depth(5)]).  
  
[0,1,2,3,4|...]  
yes
```

The default maximum depth depends on the backend. To print compound terms without a depth limit, set it explicitly to zero if necessary. The `debuggerp::write_max_depth/1` and `debuggerp::set_write_max_depth/1` predicates can be used to query and set the default maximum depth.

Custom term writing

The implicit use of the traditional `print/1` predicate (using the `p` command) and the `portray/1` user-defined hook predicate requires backend Prolog compiler support for these predicates. See the documentation of the backend you intend to use for details. As an example, assuming the following `portray/1` definition:

```
portray(e(V1,V2)) :-  
    format('~q ---> ~q~n', [V1,V2]).
```

Calling the `print/1` predicate with e.g. a `e(x1,x7)` compound term argument will output:

```
| ?- print(e(x1,x7)).
```

```
x1 ---> x7
```

```
yes
```

1.18.10 Context-switching calls

Logtalk provides a debugging control construct, `(<<)/2`, which allows the execution of a query within the context of an object. Common debugging uses include checking an object local predicates (e.g. predicates representing internal dynamic state) and sending a message from within an object. This control construct may also be used to write unit tests.

Consider the following toy example:

```
:- object(broken).

    :- public(a/1).

    a(A) :- b(A, B), c(B).
    b(1, 2). b(2, 4). b(3, 6).
    c(3).

:- end_object.
```

Something is wrong when we try the object public predicate, `a/1`:

```
| ?- broken::a(A).
```

```
no
```

For helping in diagnosing the problem, instead of compiling the object in debug mode and doing a *trace* of the query to check the clauses for the non-public predicates, we can instead simply type:

```
| ?- broken << c(C).
```

```
C = 3
```

```
yes
```

The `(<<)/2` control construct works by switching the execution context to the object in the first argument and then compiling and executing the second argument within that context:

```
| ?- broken << (self(Self), sender(Sender), this(This)).
```

```
Self = broken
```

```
Sender = broken
```

```
This = broken
```

```
yes
```

As exemplified above, the `(<<)/2` control construct allows you to call an object local and private predicates. However, it is important to stress that we are not bypassing or defeating an object predicate scope directives. The calls take place within the context of the specified object, not within the context of the object making the `(<<)/2` call. Thus, the `(<<)/2` control construct implements a form of *execution-context-switching*.

The availability of the (`<<`)/2 control construct is controlled by the `context_switching_calls` compiler flag (its default value is defined in the adapter files of the backend Prolog compilers).

1.18.11 Debugging messages

Messages generated by calls to the `logtalk::print_message/3` predicate where the message kind is either `debug` or `debug(Group)` are only printed, by default, when the `debug` flag is turned on. Moreover, these predicate calls are suppressed by the compiler when the `optimize` flag is turned on. Note that actual printing of debug messages does not require compiling the code in debug mode, only turning on the debug flag.

Meta-messages

To avoid having to define `message_tokens//2` grammar rules for translating each and every debug message, Logtalk provides default tokenization for seven *meta-messages* that cover the most common cases:

@Message

By default, the message is printed as passed to the `write/1` predicate followed by a newline.

Key-Value

By default, the message is printed as `Key: Value` followed by a newline. The key is printed as passed to the `write/1` predicate while the value is printed as passed to the `writeln/1` predicate.

Format+Arguments

By default, the message is printed as passed to the `format/2` predicate.

List

By default, the list items are printed indented, one per line. The items are preceded by a dash and can be @Message, Key-Value, or Format+Arguments messages. If that is not the case, the item is printed as passed to the `writeln/1` predicate.

Title::List

By default, the title is printed, followed by a newline and the indented list items, one per line. The items are printed as in the List meta message.

[Stream,Prefix]>>Goal

By default, call user-defined printing Goal in the context of user. The use of a lambda expression allows passing the message stream and prefix. Printing the prefix is delegated to the goal.

[Stream]>>Goal

By default, call user-defined printing Goal in the context of user. The use of a lambda expression allows passing the message stream.

Some simple examples of using these meta-messages:

```
| ?- logtalk::print_message(debug, core, @'Phase 1 completed').
yes

| ?- logtalk::print_message(debug, core, [Stream]>>write(Stream,foo)).
yes

| ?- set_logtalk_flag(debug, on).
yes

| ?- logtalk::print_message(debug, core, [Stream]>>write(Stream,foo)).
foo
yes
```

(continues on next page)

(continued from previous page)

```
| ?- logtalk::print_message(debug, core, @'Phase 1 completed').
>>> Phase 1 completed
yes

| ?- logtalk::print_message(debug, core, answer-42).
>>> answer: 42
yes

| ?- logtalk::print_message(debug, core, 'Position: <~d,~d>'+[42,23]).
>>> Position: <42,23>
yes

| ?- logtalk::print_message(debug, core, [arthur,ford,marvin]).
>>> - arthur
>>> - ford
>>> - marvin
yes

| ?- logtalk::print_message(debug, core, names::[arthur,ford,marvin]).
>>> names:
>>> - arthur
>>> - ford
>>> - marvin
yes
```

The >>> prefix is the default message prefix for debug messages. It can be redefined using the [logtalk::message_prefix_stream/4](#) hook predicate. For example:

```
:- multifile(logtalk::message_prefix_stream/4).
:- dynamic(logtalk::message_prefix_stream/4).

logtalk::message_prefix_stream(debug, core, '(dbg) ', user_error).
```

Debugging messages only output information by default. These messages can, however, be *intercepted* to perform other actions.

Selective printing of debug messages

By default, all debug messages are either printed or skipped, depending on the [debug](#) and [optimize](#) flags. When the code is not compiled in optimal mode, the [debug_messages](#) tool allows selective enabling of debug messages per [component](#) and per debug group. For example, to enable all debug and debug(Group) messages for the parser component:

```
% upon loading the tool, all messages are disabled by default:
| ?- logtalk_load(debug_messages(loader)).
...

% enable both debug and debug(_) messages:
| ?- debug_messages::enable(parser).
yes
```

To enable only debug(tokenization) messages for the parser component:

```
% first disable any and all enabled messages:
| ?- debug_messages::disable(parser).
yes

% enable only debug(tokenization) messages:
| ?- debug_messages::enable(parser, tokenization).
yes
```

See the tool documentation for more details.

1.18.12 Using the term-expansion mechanism for debugging

The *term-expansion mechanism* can also be used for conditional compilation of debugging goals. For example, the *hook_objects* library provides a `print_goal_hook` object that simplifies printing entity goals before or after calling them by simply prefixing them with an operator. See the library and hook object documentation for details. You can also define your own specialized hook objects for custom debugging tasks.

1.18.13 Ports profiling

The Logtalk distribution includes a *ports_profiler* tool based on the same procedure box model described above. This tool is specially useful for debugging performance issues (e.g., due to lack of determinism or unexpected backtracking). See the tool documentation for details.

1.18.14 Debug and trace events

The debugging API defines two multifile predicates, `logtalk::trace_event/2` and `logtalk::debug_handler/3` for handling trace and debug events. It also provides a `logtalk::debug_handler/1` multifile predicate that allows an object (or a category) to declare itself as a debug handler provider. The Logtalk debugger and *ports_profiler* tools are regular applications that are implemented using this API, which can also be used to implement alternative or new debugging-related tools. See the API documentation for details and the source code of the debugger and *ports_profiler* tools for usage examples.

To define a new debug handler provider, add (to an object or category) clauses for the `debug_handler/1` and `debug_handler/3` predicates. For example:

```
% declare my_debug_handler as a debug handler provider
:- multifile(logtalk::debug_handler/1).
logtalk::debug_handler(my_debug_handler).

% handle debug events
:- multifile(logtalk::debug_handler/3).
logtalk::debug_handler(my_debug_handler, Event, ExCtx) :-
    debug_handler(Event, ExCtx).

debug_handler(fact(Entity,Fact,Clause,File,Line), ExCtx) :-
    ...
debug_handler(rule(Entity,Head,Clause,File,Line), ExCtx) :-
    ...
debug_handler(top_goal(Goal, TGoal), ExCtx) :-
    ...
```

(continues on next page)

(continued from previous page)

```
debug_handler(goal(Goal, TGoal), ExCtx) :-
    ...
```

Your debug handler provider should also either automatically call the `logtalk::activate_debug_handler/1` and `logtalk::deactivate_debug_handler/0` predicate or provide public predicates to simplify calling these predicates. For example:

```
:- public(start/0).
start :-
    logtalk::activate_debug_handler(my_debug_handler).

:- public(stop/0).
stop :-
    logtalk::deactivate_debug_handler.
```

If you only need to define a trace event handler, then simply define clauses for the `logtalk::trace_event/2` multifile predicate:

```
:- multifile(logtalk::trace_event/2).
:- dynamic(logtalk::trace_event/2).

% the Logtalk runtime calls all defined logtalk::trace_event/2 hooks using
% a failure-driven loop; thus we don't have to worry about handling all
% events or failing after handling an event to give other hooks a chance
logtalk::trace_event(fact(Entity, Fact, N, _, _), _) :-
    ...
logtalk::trace_event(rule(Entity, Head, N, _, _), _) :-
    ...
```

1.18.15 Source-level debugger

A minimal source-level debugger is provided by the Logtalk for VSCode extension: when debugging in the integrated terminal using the debugger tool, the current clause (at leashed unification ports) is shown in the active editor window. The extension can also be used with VSCodium. See its documentation for more details.

1.19 Performance

Logtalk is implemented as a *trans-compiler* to Prolog. When compiling predicates, it preserves in the generated Prolog code all cases of first-argument indexing and tail-recursion. In practice, this means that if you know how to write efficient Prolog predicates, you already know the basics of how to write efficient Logtalk predicates.

The Logtalk compiler appends a single argument to the compiled form of all entity predicate clauses. This hidden argument is used to pass the *execution-context* when proving a query. In the common case where a predicate makes no calls to the *execution-context predicates* and *message-sending control constructs* and is neither a meta-predicate nor a coinductive predicate, the execution-context is simply passed between goals. In this case, with most backend Prolog virtual machines, the cost of this extra argument is null or negligible. When the execution-context needs to be accessed (e.g. to fetch the value of *self* for a `(::)/1` call) there may be a small inherent overhead due to the access to the individual arguments of the compound term used to represent the execution-context.

1.19.1 Source code compilation modes

Source code can be compiled in *optimal*, *normal*, or *debug* mode, depending on the *optimize* and *debug* compiler flags. Optimal mode is used when deploying an application, while normal and debug modes are used when developing an application. Compiling code in optimal mode enables several optimizations, notably the use of *static binding* whenever enough information is available at compile-time. In debug mode, most optimizations are turned off, and the code is instrumented to generate *debug events* that enable developer tools such as the *command-line debugger* and the *ports profiler*.

1.19.2 Source code compilation order

Static binding optimizations, notably message sending and super calls, require referenced code to be compiled before the calls so that the calls can be resolved at compile time. The compiler prints warnings when the file compilation/loading order is not ideal (controlled by the *unknown_entities* flag). See the *Source files* section on using the `logtalk::loaded_files_topological_sort/1` and `logtalk::loaded_files_topological_sort/2` predicates to find an optimal file loading order. See the *Source files* section for details. See also the section below on circular references.

1.19.3 Local predicate calls

Local calls to object (or category) predicates have zero overhead in terms of the number of inferences, as expected, compared with local Prolog calls.

1.19.4 Calls to imported or inherited predicates

Assuming the *optimize* flag is turned on and a static predicate, $(\wedge \wedge)/1$ calls have zero overhead in terms of number of inferences.

1.19.5 Calls to module predicates

Local calls from an object (or category) to a module predicate have zero overhead (assuming both the module and the predicate are bound at compile-time).

1.19.6 Messages

Logtalk implements *static binding* and *dynamic binding* for message-sending calls. For dynamic binding, a caching mechanism is used by the runtime. It's useful to measure the performance overhead in *number of logic inferences* compared with plain Prolog and Prolog modules. Note that the number of logic inferences is a metric independent of the chosen backend Prolog compiler. The results for Logtalk 3.17.0 and later versions are:

- Static binding: 0
- Dynamic binding (object bound at compile-time): +1
- Dynamic binding (object bound at runtime): +2

Static binding is the common case with libraries and most application code; it requires compiling code with the *optimize* flag turned on. Dynamic binding numbers are after the first call (i.e. after the generalization of the query is cached). All numbers with the *events* flag set to deny (setting this flag to allow adds an overhead of +5 inferences to the results above; note that this flag can be defined on a per-object basis as needed instead of globally and thus minimizing the performance impact).

The dynamic binding caches assume the used *backend Prolog compiler* does indexing of dynamic predicates. This is a common feature of modern Prolog systems, but the actual details vary from system to system and may have an impact on dynamic binding performance.

Note that messages to *self* (`(::)/1` calls) and messages to an object (`(::)/2` calls) from the top-level interpreter always use dynamic binding, as the object that receives the message is only known at runtime.

Messages sent from Prolog modules may use static binding depending on the used backend Prolog compiler native support for goal-expansion. Consult the Prolog compiler documentation and adapter file notes for details.

Warning

Some Prolog systems provide a `time/1` predicate that also reports the number of inferences. But the reported numbers are often misleading when the predicate is called from the top-level. Besides common top-level bookkeeping operations (e.g., keeping track of goal history or applying goal-expansion) that may influence the inference counting, the Logtalk runtime code for a `(:)/2` top-level goal is necessarily different from the code generated for a `(:)/2` goal from a compiled object, as it requires *runtime* compilation of the goal into the same low-level message-sending primitive (assuming dynamic-binding is also required for the compiled object goal).

1.19.7 Automatic expansion of built-in meta-predicates

The compiler always expands calls to the *forall/2*, *once/1*, and *ignore/1* meta-predicates into equivalent definitions using the negation and conditional control constructs. It also expands calls to the *call/1-N*, *phrase/2*, and *phrase/3* meta-predicates when the first argument is bound. These expansions are performed independently of the `optimize` flag value.

1.19.8 Inlining

When the *optimize* flag is turned on, the Logtalk compiler performs *inlining* of predicate calls whenever possible. This includes calls to Prolog predicates that are either built-in, foreign, or defined in a module (including user). Inlining notably allows wrapping module or foreign predicates using an object without introducing any overhead. In the specific case of the *execution-context predicates*, calls are inlined independently of the `optimize` flag value.

1.19.9 Generated code simplification and optimizations

When the *optimize* flag is turned on, the Logtalk compiler simplifies and optimizes generated clauses (including those resulting from the compilation of grammar rules), by flattening conjunctions, folding left unifications (e.g. generated as a by-product of the compilation of grammar rules), and removing redundant calls to `true/0`.

When using *lambda expressions* and library meta-predicates, use the *meta_compiler* library to avoid most meta-call overheads.

1.19.10 Size of the generated code

The size of the intermediate Prolog code generated by the compiler is proportional to the size of the source code. Assuming that the *term-expansion mechanism* is not used, each predicate clause in the source code is compiled into a single predicate clause. But the Logtalk compiler also generates internal tables for the defined entities, for the entity relations, and for the declared and defined predicates. These tables enable support for fundamental features such as *inheritance* and *reflection*. The size of these tables is proportional to the number of entities, entity relations, and predicate declarations and definitions. When the *source_data* is turned on (the default when *developing* an application), the generated code also includes additional data about the source code, such as entity and predicate positions in a source file. This data enables advanced developer tool functionality. But it is usually not required when *deploying* an application. Thus, turning this flag off is a common setting for minimizing an application footprint.

1.19.11 Circular references

Circular references, i.e. two objects sending messages to each other, are relatively costly and should be avoided if possible as they prevent using static binding for the messages sent from the first loaded object to the second object. The *logtalk_make(circular)* goal (or its {*@*} top-level abbreviation) can be used to scan for circular entity dependencies.

1.19.12 Debug mode overhead

Code compiled in debug mode runs slower, as expected, when compared with normal or optimized mode. The overhead depends on the number of *debug events* generated when running the application. A debug event is simply a pass on a call or unification port of the *procedure box model*. These debug events can be intercepted by defined clauses for the *logtalk::trace_event/2* and *logtalk::debug_handler/3* multifile predicates. With no application (such as a debugger or a port profiler) loaded defining clauses for these predicates, each goal has an overhead of four extra inferences due to the runtime checking for a definition of the hook predicates and a meta-call of the user goal. The clause head unification events result in one or more inferences per goal (depending on the number of clauses whose head unifies with the goal and backtracking). In practice, this overhead translates to code compiled in debug mode running typically ~2x to ~7x slower than code compiled in normal or optimized mode, depending on the application (the exact overhead is proportional to the number of passes on the call and unification ports; deterministic code often results in a relatively larger overhead when compared with code performing significant backtracking).

1.19.13 Other considerations

One aspect of performance that affects both Logtalk and Prolog code is the characteristics of the Prolog VM. The Logtalk distribution includes two examples, *bench* and *benchmarks*, to help evaluate performance with specific backend Prolog systems. A table with benchmark *results* for a subset of the supported systems is also available at the Logtalk website. But note that multiple factors affect the performance of an application. The benchmark examples and their results only provide a partial assessment.

1.20 Installing Logtalk

This page provides an overview of Logtalk installation requirements and instructions and a description of the files contained in the Logtalk distribution. For detailed, up-to-date installation and configuration instructions, please see the `README.md`, `INSTALL.md`, and `CUSTOMIZE.md` files distributed with Logtalk. The broad compatibility of Logtalk, both with Prolog compilers and operating-systems, together with all the possible user scenarios, means that installation can vary from very simple, by running an installer or a couple of scripts, to the need of patching both Logtalk and Prolog compilers to workaround the lack of strong Prolog standards or to cope with the requirements of less common operating-systems.

The preferred installation scenario is to have Logtalk installed in a system-wide location, thus available for all users, and a local copy of user-modifiable files on each user home directory (even when you are the single user of your computer). This scenario allows each user to independently customize Logtalk and to freely modify the provided libraries and programming examples. Logtalk installers, installation shell scripts, and Prolog integration scripts favor this installation scenario, although alternative installation scenarios are always possible. The installers set two environment variables, `LOGTALKHOME` and `LOGTALKUSER`, pointing, respectively, to the Logtalk installation folder and to the Logtalk user folder.

User applications should preferably be kept outside of the Logtalk user folder created by the installation process, as updating Logtalk often results in updating the contents of this folder. If your applications depend on customizations to the distribution files, backup those changes before updating Logtalk.

1.20.1 Hardware and software requirements

Computer and operating system

Logtalk is compatible with almost any computer/operating-system with a modern, standards-compliant, Prolog compiler available.

Prolog compiler

Logtalk requires a *backend Prolog compiler* supporting official and de facto standards. Capabilities needed by Logtalk that are not defined in the official ISO Prolog Core standard include:

- access to predicate properties
- operating-system access predicates
- de facto standard predicates not (yet) specified in the official standard

Logtalk needs access to the predicate property `built_in` to properly compile objects and categories that contain Prolog built-in predicate calls. In addition, some Logtalk built-ins need to know the dynamic/static status of predicates to ensure correct application. The ISO standard for Prolog modules defines a `predicate_property/2` predicate that is already implemented by most Prolog compilers. Note that if these capabilities are not built-in the user cannot easily define them.

For optimal performance, Logtalk requires that the Prolog compiler supports **first-argument indexing** for both static and dynamic code (most modern compilers support this feature).

Since most Prolog compilers are moving closer to the ISO Prolog standard [ISO95], it is advisable that you try to use the most recent version of your favorite Prolog compiler.

1.20.2 Logtalk installers

Logtalk installers are available for macOS, Linux, and Microsoft Windows. Depending on the chosen installer, some tasks (e.g., setting environment variables or integrating Logtalk with some Prolog compilers) may need to be performed manually.

1.20.3 Source distribution

Logtalk sources are available in a tar archive compressed with bzip2, `lgt3xxx.tar.bz2`. You may expand the archive by using a decompressing utility or by typing the following commands at the command-line:

```
% tar -jxvf lgt3xxx.tar.bz2
```

This will create a sub-directory named `lgt3xxx` in your current directory. Almost all files in the Logtalk distribution are text files. Different operating-systems use different end-of-line codes for text files. Ensure that your decompressing utility converts the end-of-lines of all text files to match your operating system.

1.20.4 Distribution overview

In the Logtalk installation directory, you will find the following files and directories:

`ACKNOWLEDGMENTS.md` - List of authors, contributors, sponsors, and open source credits

`BIBLIOGRAPHY.bib` – Logtalk bibliography in BibTeX format

`CITATION.cff` - Information on how to cite Logtalk

`CODE_OF_CONDUCT.md` - Code of conduct for contributors and users posting on support forums

`CUSTOMIZE.md` – Logtalk end-user customization instructions

`INSTALL.md` – Logtalk installation instructions

`LICENSE.txt` – Logtalk user license

`NOTICE.txt` – Logtalk copyright notice

`QUICK_START.md` – Quick start instructions for those that do not like to read manuals

`README.md` – several useful pieces of information

`RELEASE_NOTES.md` – release notes for this version

`UPGRADING.md` – instructions on how to upgrade your programs to the current Logtalk version

`VERSION.txt` – file containing the current Logtalk version number (used for compatibility checking when upgrading Logtalk)

adapters

`NOTES.md` – notes on the provided adapter files

`template.pl` – template adapter file

`...` – specific adapter files

coding

`NOTES.md` – notes on syntax highlighter and text editor support files providing syntax coloring for publishing and editing Logtalk source code

`...` – syntax coloring support files

contributions

NOTES.md – notes on the user-contributed code
... – user-contributed code files

core

NOTES.md – notes on the current status of the compiler and runtime
... – core source files

docs/apis

NOTES.md – notes on the provided documentation for core, library, tools, and contributions entities
index.html – root document for all entities documentation
... – other entity documentation files

docs/handbook

NOTES.md – notes on the provided documentation
bibliography.html – bibliography
glossary.html – glossary
index.html – root document for all documentation
... – other documentation files

docs/man

... – POSIX man pages for the shell scripts

examples

NOTES.md – short description of the provided examples

ack

NOTES.md – example notes and sample queries
loader.lgt – loader utility file for the example objects
... – ack example source files
... – other examples

integration

NOTES.md – notes on scripts for Logtalk integration with Prolog compilers
... – Prolog integration scripts

library

NOTES.md – short description of the library contents
all_loader.lgt – loader utility file for all library entities
... – library source files

paths

NOTES.md – description on how to setup library and example paths
paths.pl – default library and example paths

ports

NOTES.md – description of included ports of third-party software
... – ports

samples

loader-sample.lgt – sample loader file for user applications
settings-sample.lgt – sample file for user-defined Logtalk settings
tester-sample.lgt – sample file for helping to automate running user application unit tests
tests-sample.lgt – sample file for writing user application unit tests

scratch

NOTES.md – notes on the scratch directory

scripts

NOTES.md – notes on scripts for Logtalk user setup, packaging, and installation
... – packaging, installation, and setup scripts

tests

NOTES.md – notes on the current status of the unit tests
... – unit tests for built-in features

tools

NOTES.md – notes on the provided programming tools
... – programming tools

Adapter files

Adapter files provide the glue code between the Logtalk compiler/runtime and a Prolog compiler. Each adapter file contains two sets of predicates: ISO Prolog standard predicates and directives not built-in in the target Prolog compiler and Logtalk specific predicates.

Logtalk already includes ready-to-use adapter files for most academic and commercial Prolog compilers. If an adapter file is not available for the compiler that you intend to use, then you need to build a new one, starting from the included `template.pl` file. Start by making a copy of the template file. Carefully check (or complete if needed) each listed definition. If your Prolog compiler conforms to the ISO standard, this task should only take you a few minutes. In most cases, you can borrow code from the predefined adapter files. If you are unsure that your Prolog compiler provides all the ISO predicates needed by Logtalk, try to run the system by setting the unknown predicate error handler to report as an error any call to a missing predicate. Better yet, switch to a modern, ISO-compliant, Prolog compiler. If you send me your adapter file, with a reference to the target Prolog compiler, maybe I can include it in the next release of Logtalk.

The adapter files specify *default* values for most of the Logtalk *compiler flags*. They also specify values for read-only flags that are used to describe Prolog backend-specific features.

Compiler and runtime

The core sub-directory contains the Prolog and Logtalk source files that implement the Logtalk compiler and the Logtalk runtime. The compiler and the runtime may be split in two (or more) separate files or combined into a single file, depending on the Logtalk release that you are installing.

Library

The Logtalk distribution includes a standard library of useful objects, categories, and protocols. Read the corresponding `NOTES.md` file for details about the library contents.

Examples

The Logtalk distribution includes a large number of programming examples. The sources of each one of these examples can be found included in a subdirectory with the same name, inside the directory `examples`. The majority of these examples include tests and a file named `SCRIPT.txt` with sample calls. Some examples may depend on other examples and library objects to work properly. Read the corresponding `NOTES.md` file for details before running an example.

Logtalk source files

Logtalk source files are text files containing one or more entity definitions (objects, categories, or protocols). The Logtalk source files may also contain plain Prolog code. The extension `.lgt` is normally used. Logtalk compiles these files to plain Prolog by appending to the file name a suffix derived from the extension and by replacing the `.lgt` extension with `.pl` (`.pl` is the default Prolog extension; if your Prolog compiler expects the Prolog source filenames to end with a specific, different extension, you can set it in the corresponding adapter file).

1.21 Prolog integration and migration

This section provides suggestions for integrating and migrating plain Prolog code and Prolog module code to Logtalk. Detailed instructions are provided for encapsulating plain Prolog code in objects, converting Prolog modules into objects, and compiling and reusing Prolog modules as objects from inside Logtalk. An interesting application of the techniques described in this section is a solution for running a Prolog application that uses modules on a Prolog compiler with no module system. The *wrapper* tool can be used to help in migrating Prolog code.

1.21.1 Source files with both Prolog code and Logtalk code

Logtalk source files may contain plain Prolog code intermixed with Logtalk code. The Logtalk compiler simply copies the plain Prolog code as-is to the generated Prolog file. With Prolog modules, it is assumed that the module code starts with a `module/1-2` directive and ends at the end of the file. There is no module ending directive that would allow us to define more than one module per file. In fact, most, if not all, Prolog module systems always define a single module per file. Some of them mandate that the `module/1-2` directive be the first term in a source file. As such, when the Logtalk compiler finds a `module/1-2` directive, it assumes that all code that follows until the end of the file belongs to the module.

1.21.2 Encapsulating plain Prolog code in objects

Most applications consist of several plain Prolog source files, each one defining a few top-level predicates and auxiliary predicates that are not meant to be directly called by the user. Encapsulating plain Prolog code in objects allows us to make clear the different roles of each predicate, to hide implementation details, to prevent auxiliary predicates from being called outside the object, and to take advantage of Logtalk advanced code encapsulating and reusing features. It also simplifies using its developer tools.

Encapsulating Prolog code using Logtalk objects is simple. First, for each source file, add an opening object directive, `object/1-5`, to the beginning of the file and an ending object directive, `end_object/0`, to the end of the file. Choose an object name that reflects the purpose of the source file code (this is a good opportunity for code refactoring if necessary). Second, add `public/1` predicate directives for the top-level predicates that are used directly by the user or called from other source files. Third, we need to be able to call from inside an object predicates defined in other source files/objects. The easiest solution, which has the advantage of not requiring any changes to the predicate definitions, is to use the `uses/2` directive. If your Prolog compiler supports cross-referencing tools, you may use them to help you make sure that all calls to predicates on other source files/objects are listed in the `uses/2` directives. The Logtalk wrapper tool can also help in detecting cross-predicate calls. Compiling the resulting objects with the Logtalk `unknown_predicates` and `portability` flags set to warning will help you identify calls to predicates defined in other converted source files and possible portability issues.

Prolog multifile predicates

Prolog *multifile* predicates are used when clauses for the same predicate are spread among several source files. When encapsulating plain Prolog code that uses multifile predicates, it's often the case that the clauses of the multifile predicates get spread between different objects and categories, but conversion is straightforward. In the Logtalk object (or category) holding the multifile predicate *primary declaration*, add a *predicate scope directive* and a `multifile/1` directive. In all other objects (or categories) defining clauses for the multifile predicate, add a `multifile/1` directive and predicate clauses using the format:

```
:- multifile(Entity::Name/Arity).  
  
Entity::Functor(...) :-  
    ...
```

See the [section](#) on the `multifile/1` predicate directive for more information. An alternative solution is to simply keep the clauses for the multifile predicates as plain Prolog code and define, if necessary, a parametric object to encapsulate all predicates working with the multifile predicate clauses. For example, assume the following `multifile/1` directive:

```
% city(Name, District, Population, Neighbors)  
:- multifile(city/4).
```

We can define a parametric object with `city/4` as its identifier:

```
:- object(city(_Name, _District, _Population, _Neighbors)).  
    % predicates for working with city/4 clauses  
:- end_object.
```

This solution is preferred when the multifile predicates are used to represent large tables of data. See the section on *parametric objects* for more details.

1.21.3 Converting Prolog modules into objects

Converting Prolog modules into objects may allow an application to run on a wider range of Prolog compilers, overcoming portability problems. Some Prolog compilers don't support a module system. Among those Prolog compilers that support a module system, the lack of standardization leads to several issues, notably with semantics, operators, and meta-predicates. In addition, the conversion allows you to take advantage of Logtalk more powerful abstraction and reuse mechanisms, such as separation between interface and implementation, inheritance, parametric objects, and categories. It also allows you to take full advantage of Logtalk developer tools for improved productivity.

Converting a Prolog module into an object is simplified when the directives used in the module are supported by Logtalk (see the listing in the next section). Assuming that this is the case, apply the following steps:

1. Convert the module `module/1` directive into an `object/1` opening object directive, using the module name as the object name. For `module/2` directives apply the same conversion and convert the list of exported predicates into `public/1` predicate directives. Add a closing object directive, `end_object/0`, at the end of the source code.
2. Convert any `export/1` directives into `public/1` predicate directives.
3. Convert any `use_module/1` directives for modules that will not be converted to objects into `use_module/2` directives (see next section), replacing the file spec in the first argument with the module name.
4. Convert any `use_module/1-2` directives referencing other modules also being converted to objects into Logtalk `uses/2` directives.
5. Convert each `reexport/1` directive into a `uses/2` directive and `public/1` predicate directives (see next section).
6. Convert any `meta_predicate/1` directives into Logtalk `meta_predicate/1` directives by replacing the module meta-argument indicator, `:`, with the Logtalk meta-argument indicator `0` for goal meta-arguments. For closure meta-arguments, use an integer denoting the number of additional arguments that will be appended to construct a goal. Arguments that are not meta-arguments are represented by the `*` character. Do not use argument mode indicators such as `?`, `or +`, or `-` as Logtalk supports *mode directives*.
7. Convert any explicit qualified calls to module predicates to messages by replacing the `(:)/2` operator with the `(::)/2` message-sending operator when the referenced modules are also being converted into objects. Calls in the pseudo-module `user` can be encapsulated using the `{}/1` Logtalk external call control construct. You can also use instead a `uses/2` directive where the first argument would be the atom `user` and the second argument a list of all external predicates. This alternative has the advantages of not requiring changes to the code making the predicate calls and of better visibility for the documenting and diagramming tools.
8. If your module uses the database built-in predicates to implement module-local mutable state using dynamic predicates, add both `private/1` and `dynamic/1` directives for each dynamic predicate.
9. If your module declares or defines clauses for multifile module predicates, replace the `(:)/2` functor by `(::)/2` in the `multifile/1` directives and in the clause heads for all modules defining the multifile predicates that are also being converted into objects; if that is not the case, just keep the `multifile/1` directives and the clause heads as-is.
10. Compile the resulting objects with the Logtalk `unknown_predicates`, and `portability` flags set to warning to help you locate possible issues and calls to proprietary Prolog built-in predicates and to predicates defined on other converted modules. In order to improve code portability, check the Logtalk library for possible alternatives to the use of proprietary Prolog built-in predicates.

Before converting your modules to objects, you may try to compile them first as objects (using the `logtalk_compile/1` Logtalk built-in predicates) to help identify any issues that must be dealt with when doing

the conversion to objects. Note that Logtalk supports compiling Prolog files as Logtalk source code without requiring changes to the file name extensions.

1.21.4 Compiling Prolog modules as objects

A possible alternative to porting Prolog code to Logtalk is to compile the Prolog source files using the `logtalk_load/1-2` and `logtalk_compile/1-2` predicates. The Logtalk compiler provides partial support for compiling Prolog modules as Logtalk objects. This support may allow using modules from a backend Prolog system in a different backend Prolog system, although its main purpose is to help in porting existing Prolog code to Logtalk in order to benefit from its extended language features and its developer tools. Why partial support? Although there is an ISO Prolog standard for modules, it is (rightfully) ignored by most implementers and vendors (due to its flaws and deviation from common practice). In addition, there is no de facto standard for module systems, despite otherwise frequent misleading claims. Key system differences include the set of implemented module directives, the directive semantics, the handling of operators, the locality of flags, and the integration of term-expansion mechanisms (when provided). Another potential issue is that, when compiling modules as objects, Logtalk assumes that any referenced module (e.g., using `use_module/1-2` directives) is also being compiled as an object. If that's not the case, the compiled module calls being compiled as message-sending goals will still work for normal predicates but will not work for meta-predicates called using implicit module qualification. The reason is that, unlike in Logtalk, calls to implicitly and explicitly qualified module meta-predicates have different semantics. Follows a discussion of other limitations of this approach that you should be aware of.

Supported module directives

Currently, Logtalk supports the following module directives:

module/1

The module name becomes the object name.

module/2

The module name becomes the object name. The exported predicates become public object predicates. The exported grammar rule non-terminals become public grammar rule non-terminals. The exported operators become public object operators but are not active elsewhere when loading the code.

use_module/2

This directive is compiled as a Logtalk *uses/2* directive in order to ensure correct compilation of the module predicate clauses. The first argument of this directive must be the module **name** (an atom), not a module file specification (the adapter files attempt to use the Prolog dialect level term-expansion mechanism to find the module name from the module file specification). Note that the module is not automatically loaded by Logtalk (as it would be when compiling the directive using Prolog instead of Logtalk; the programmer may also want the specified module to be compiled as an object). The second argument must be a predicate indicator (`Name/Arity`), a grammar rule non-terminal indicator (`Name//Arity`), a operator declaration, or a list of predicate indicators, grammar rule non-terminal indicators, and operator declarations. Predicate aliases can be declared using the notation `Name/Arity` as `Alias/Arity` or, in alternative, the notation `Name/Arity:Alias/Arity`. Similar for non-terminal aliases.

export/1

Exported predicates are compiled as public object predicates. The argument must be a predicate indicator (`Name/Arity`), a grammar rule non-terminal indicator (`Name//Arity`), an operator declaration, or a list of predicate indicators, grammar rule non-terminal indicators, and operator declarations.

reexport/2

Reexported predicates are compiled as public object predicates. The first argument is the module name. The second argument must be a predicate indicator (`Name/Arity`), a grammar rule non-terminal

indicator (Name//Arity), an operator declaration, or a list of predicate indicators, grammar rule non-terminal indicators, and operator declarations. Predicate aliases can be declared using the notation Name/Arity as Alias/Arity or, in alternative, the notation Name/Arity:Alias/Arity. Similar for non-terminal aliases.

meta_predicate/1

Module meta-predicates become object meta-predicates. All meta-predicates must be declared using the [meta_predicate/1](#) directive using Logtalk syntax for normal arguments and meta-arguments. Note that Prolog module meta-predicates and Logtalk meta-predicates don't share the same explicit-qualification calling semantics: in Logtalk, meta-arguments are always called in the context of the [sender](#). Moreover, Logtalk is not based on the predicate-prefixing mechanism common to most Prolog module systems.

A common issue when compiling modules as objects is the use of the atoms `dynamic`, `discontiguous`, and `multifile` as operators in directives. For better portability, avoid this usage. For example, write:

```
:- dynamic([foo/1, bar/2]).
```

instead of:

```
:- dynamic foo/1, bar/2.
```

Another common issue is missing `meta_predicate/1`, `dynamic/1`, `discontiguous/1`, and `multifile/1` predicate directives. The Logtalk compiler supports detection of missing directives (by setting its [missing_directives](#) flag to warning).

When compiling modules as objects, you probably don't need event support turned on. You may use the [events](#) compiler flag to deny in the Logtalk compiling and loading built-in methods for a small performance gain for the compiled code.

Unsupported module directives

The `reexport/1` and `use_module/1` directives are not directly supported by the Logtalk compiler. But most Prolog adapter files provide support for compiling these directives using Logtalk's first stage of its [term-expansion mechanism](#). Nevertheless, these directives can be converted, respectively, into a sequence of `:- use_module/2` and `export/1` directives and `use_module/2` directives by finding which predicates exported by the specified modules are reexported or imported into the module containing the directive. For `use_module/1` directives, finding the names of the imported predicates that are actually used is easy. First, comment out the directive and compile the file (making sure that the [unknown_predicates](#) compiler flag is set to warning). Logtalk will print a warning with a list of predicates that are called but never defined. Second, use this list to replace the `use_module/1` directives by `use_module/2` directives. You should then be able to compile the modified Prolog module as an object.

Modules using a term-expansion mechanism

Although Logtalk supports [term and goal expansion mechanisms](#), the usage semantics are different from similar mechanisms found in some Prolog compilers. In particular, Logtalk does not support defining term and goal expansions clauses in a source file for expanding the source file itself. Logtalk forces a clean separation between expansion clauses and the source files that will be subject to source-to-source expansions by using [hook objects](#). But hook objects also provide a working solution here when the expansion code is separated from the code to be expanded. Logtalk supports using a module as a hook object as long as its name doesn't coincide with the name of an object and that the module uses `term_expansion/2` and `goal_expansion/2` predicates. Assuming that's the case, before attempting to compile the modules as objects, set the default hook object to the module containing the expansion code. For example, if the expansions are stored in a system module:

```
| ?- set_logtalk_flag(hook, system).
...
```

This, however, may not be enough, as expansions may be stored in multiple modules. A common example is to use a module named `prolog` for system expansions and to store the user-defined expansions in `user`. The Logtalk library provides a solution for these scenarios. Using the `hook_flows` library we can select multiple hook objects or hook modules. For example, assuming expansions stored on both user and system modules:

```
| ?- logtalk_load(hook_flows(loader)).
...

| ?- set_logtalk_flag(hook, hook_set([user, system])).
...
```

After these queries, we can try to compile the modules and look for other porting or portability issues. A well-know issue is Prolog module term-expansions calling predicates such as `prolog_load_context/2`, which will always fail when it's the Logtalk compiler instead of the Prolog compiler loading a source file. In some of these cases, it may be possible to rewrite the expansion rules to use the [logtalk_load_context/2](#) predicate instead.

File search paths

Some Prolog systems provide a mechanism for defining file search paths (this mechanism works differently from Logtalk own support for defining library path aliases). When porting Prolog code that defines file search paths, e.g. for finding module libraries, it often helps to load the pristine Prolog application before attempting to compile its source files as Logtalk source files. Depending on the Prolog backend, this may allow the file search paths to be used when compiling modules as objects that use file directives such as `use_module/2`.

1.21.5 Dealing with proprietary Prolog directives and predicates

Most Prolog compilers define proprietary, non-standard directives and predicates that may be used in both plain code and module code. Non-standard Prolog built-in predicates are usually not problematic, as Logtalk is usually able to identify and compile them correctly (but see the notes on built-in meta-predicates for possible caveats). However, Logtalk will generate compilation errors on source files containing proprietary directives unless you first specify how the directives should be handled. Several actions are possible on a per-directive basis: ignoring the directive (i.e., do not copy the directive, although a goal can be proved as a consequence), rewriting and copying the directive to the generated Prolog files, or rewriting and re-compiling the resulting directive. To specify these actions, the adapter files contain clauses for the internal `'$lgt_prolog_term_expansion'/2` predicate. For example, assume that a given Prolog compiler defines a `comment/2` directive for predicates using the format:

```
:- comment(foo/2, "Brief description of the predicate").
```

We can rewrite this predicate into a Logtalk `info/2` directive by defining a suitable clause for the `'$lgt_prolog_term_expansion'/2` predicate:

```
'$lgt_prolog_term_expansion'(
    :- comment(F/A, String)),
    :- info(F/A, [comment is Atom]))
) :-
    atom_codes(Atom, String).
```

This Logtalk feature can be used to allow compilation of legacy Prolog code without the need of changing the sources. When used, it is advisable to set the *portability* compiler flag to warning in order to more easily identify source files that are likely non-portable across Prolog compilers.

A second example, where a proprietary Prolog directive is discarded after triggering a side effect:

```
'$lgt_prolog_term_expansion'(
    :- load_foreign_files(Files,Libs,InitRoutine)),
    []
) :-
    load_foreign_files(Files,Libs,InitRoutine).
```

In this case, although the directive is not copied to the generated Prolog file, the foreign library files are loaded as a side effect of the Logtalk compiler calling the '\$lgt_prolog_term_expansion'/2 hook predicate.

1.21.6 Calling Prolog module predicates

Prolog module predicates can be called from within objects or categories by simply using explicit module qualification, i.e. by writing `Module:Goal` or `Goal@Module` (depending on the module system). Logtalk also supports the use of `use_module/2` directives in objects and categories (with the restriction that the first argument of the directive must be the actual module name and not the module file name or the module file path). In this case, these directives are parsed in a similar way to Logtalk *uses/2* directives, with calls to the specified module predicates being automatically translated to `Module:Goal` calls.

As a general rule, the Prolog modules should be loaded (e.g., in the auxiliary Logtalk loader files) *before* compiling objects that make use of module predicates. Moreover, the Logtalk compiler does not generate code for the automatic loading of modules referenced in `use_module/1-2` directives. This is a consequence of the lack of standardization of these directives, whose first argument can be a module name, a straight file name, or a file name using some kind of library notation, depending on the *backend Prolog compiler*. Worse, modules are sometimes defined in files with names different from the module names, requiring finding, opening, and reading the file in order to find the actual module name.

Logtalk allows you to send a message to a module in order to call one of its predicates. This is usually not advised as it implies a performance penalty when compared to just using the `Module:Call` notation. Moreover, this works only if there is no object with the same name as the module you are targeting. This feature is necessary, however, in order to properly support the compilation of modules containing `use_module/2` directives as objects. If the modules specified in the `use_module/2` directives are not compiled as objects but are instead loaded as-is by Prolog, the exported predicates would need to be called using the `Module:Call` notation but the converted module will be calling them through message-sending. Thus, this feature ensures that, on a module compiled as an object, any predicate calling other module predicates will work as expected, either these other modules are loaded as-is or also compiled as objects.

For more details, see the *Calling Prolog predicates* section.

1.21.7 Loading converted Prolog applications

Logtalk strongly favors and advises users to provide a main *loader file* for applications that explicitly load any required libraries and the application source files. In contrast, Prolog applications often either scatter loading of source files from multiple files or use implicit loading of source files via `use_module/1-2` directives. Due to this frequent ad-hoc approach, it's common to find Prolog applications with duplicated loading directives, and where loading order ignores the dependencies between source files. These issues are easily exposed by the Logtalk linter when compiling Prolog files as Logtalk files. Also common are Prolog files with multiple circular dependencies. While this should not affect the *semantics* of the ported code, it may cause some performance penalties as it prevents the Logtalk compiler from optimizing the message sending goals using

static-binding. It also makes the application architecture more difficult to understand. The definition of explicit loader files provides a good opportunity for sorting out loading order and circular dependencies, with the linter warnings providing hints for possible code refactoring to eliminate these issues. The *diagrams* tool supports directory and file loading and dependency diagrams that are also useful in understanding applications architecture.

REFERENCE MANUAL

2.1 Grammar

The Logtalk grammar is here described using W3C-style Extended Backus-Naur Form syntax. Non-terminal symbols not defined here can be found in the ISO Prolog Core standard. Terminal symbols are represented between double-quotes.

2.1.1 Entities

```
entity ::=
  object
  | category
  | protocol
```

2.1.2 Object definition

```
object ::=
  begin_object_directive ( object_term )* end_object_directive

begin_object_directive ::=
  ":- object(" object_identifier ( "," object_relations )? ")."

end_object_directive ::=
  ":- end_object."

object_relations ::=
  prototype_relations
  | non_prototype_relations

prototype_relations ::=
  prototype_relation
  | prototype_relation "," prototype_relations

prototype_relation ::=
  implements_protocols
  | imports_categories
  | extends_objects
```

(continues on next page)

(continued from previous page)

```
non_prototype_relations ::=
  non_prototype_relation
  | non_prototype_relation "," non_prototype_relations

non_prototype_relation ::=
  implements_protocols
  | imports_categories
  | instantiates_classes
  | specializes_classes
```

2.1.3 Category definition

```
category ::=
  begin_category_directive ( category_term )* end_category_directive

begin_category_directive ::=
  ":- category(" category_identifier ( "," category_relations )? ")."

end_category_directive ::=
  ":- end_category."

category_relations ::=
  category_relation
  | category_relation "," category_relations

category_relation ::=
  implements_protocols
  | extends_categories
  | complements_objects
```

2.1.4 Protocol definition

```
protocol ::=
  begin_protocol_directive ( protocol_directive )* end_protocol_directive

begin_protocol_directive ::=
  ":- protocol(" protocol_identifier ( "," extends_protocols )? ")."

end_protocol_directive ::=
  ":- end_protocol."
```

2.1.5 Entity relations

```

extends_protocols ::=
    "extends(" extended_protocols ")"

extends_objects ::=
    "extends(" extended_objects ")"

extends_categories ::=
    "extends(" extended_categories ")"

implements_protocols ::=
    "implements(" implemented_protocols ")"

imports_categories ::=
    "imports(" imported_categories ")"

instantiates_classes ::=
    "instantiates(" instantiated_objects ")"

specializes_classes ::=
    "specializes(" specialized_objects ")"

complements_objects ::=
    "complements(" complemented_objects ")"

```

Implemented protocols

```

implemented_protocols ::=
    implemented_protocol
    | implemented_protocol_sequence
    | implemented_protocol_list

implemented_protocol ::=
    protocol_identifier
    | scope "::" protocol_identifier

implemented_protocol_sequence ::=
    "(" implemented_protocol ("," implemented_protocol)* ")"

implemented_protocol_list ::=
    "[" implemented_protocol ("," implemented_protocol)* "]"

```

Extended protocols

```
extended_protocols ::=
  extended_protocol
  | extended_protocol_sequence
  | extended_protocol_list

extended_protocol ::=
  protocol_identifier
  | scope "::" protocol_identifier

extended_protocol_sequence ::=
  "(" extended_protocol ("," extended_protocol)* ")"

extended_protocol_list ::=
  "[" extended_protocol ("," extended_protocol)* "]"
```

Imported categories

```
imported_categories ::=
  imported_category
  | imported_category_sequence
  | imported_category_list

imported_category ::=
  category_identifier
  | scope "::" category_identifier

imported_category_sequence ::=
  "(" imported_category ("," imported_category)* ")"

imported_category_list ::=
  "[" imported_category ("," imported_category)* "]"
```

Extended objects

```
extended_objects ::=
  extended_object
  | extended_object_sequence
  | extended_object_list

extended_object ::=
  object_identifier
  | scope "::" object_identifier

extended_object_sequence ::=
  "(" extended_object ("," extended_object)* ")"

extended_object_list ::=
  "[" extended_object ("," extended_object)* "]"
```

Extended categories

```

extended_categories ::=
    extended_category
    | extended_category_sequence
    | extended_category_list

extended_category ::=
    category_identifier
    | scope "::" category_identifier

extended_category_sequence ::=
    "(" extended_category ("," extended_category)* ")"

extended_category_list ::=
    "[" extended_category ("," extended_category)* "]"

```

Instantiated objects

```

instantiated_objects ::=
    instantiated_object
    | instantiated_object_sequence
    | instantiated_object_list

instantiated_object ::=
    object_identifier
    | scope "::" object_identifier

instantiated_object_sequence ::=
    "(" instantiated_object ("," instantiated_object)* ")"

instantiated_object_list ::=
    "[" instantiated_object ("," instantiated_object)* "]"

```

Specialized objects

```

specialized_objects ::=
    specialized_object
    | specialized_object_sequence
    | specialized_object_list

specialized_object ::=
    object_identifier
    | scope "::" object_identifier

specialized_object_sequence ::=
    "(" specialized_object ("," specialized_object)* ")"

specialized_object_list ::=
    "[" specialized_object ("," specialized_object)* "]"

```

Complemented objects

```
complemented_objects ::=
  object_identifier
  | complemented_object_sequence
  | complemented_object_list

complemented_object_sequence ::=
  "(" object_identifier ("," object_identifier)* ")"

complemented_object_list ::=
  "[" object_identifier ("," object_identifier)* "]"
```

Entity and predicate scope

```
scope ::=
  "public"
  | "protected"
  | "private"
```

2.1.6 Entity identifiers

```
entity_identifier ::=
  object_identifier
  | protocol_identifier
  | category_identifier
```

Object identifiers

```
object_identifier ::=
  atom
  | compound
```

Category identifiers

```
category_identifier ::=
  atom
  | compound
```

Protocol identifiers

```
protocol_identifier ::=  
    atom
```

Module identifiers

```
module_identifier ::=  
    atom
```

2.1.7 Source files

```
source_file ::=  
    ( source_file_content )*
```

```
source_file_content ::=  
    source_file_directive  
    | clause  
    | grammar_rule  
    | entity
```

2.1.8 Source file names

```
source_file_name ::=  
    atom  
    | library_source_file_name
```

```
library_source_file_name ::=  
    library_name "(" atom ")"
```

```
library_name ::=  
    atom
```

2.1.9 Terms

Object terms

```
object_term ::=  
    object_directive  
    | clause  
    | grammar_rule
```

Category terms

```
category_term ::=
    category_directive
  | clause
  | grammar_rule
```

2.1.10 Directives

Source file directives

```
source_file_directive ::=
    ":- encoding(" atom ")."
  | ":- set_logtalk_flag(" atom "," nonvar ")."
  | ":- include(" source_file_name ")."
  | prolog_directive
```

Conditional compilation directives

```
conditional_compilation_directive ::=
    ":- if(" callable ")."
  | ":- elif(" callable ")."
  | ":- else."
  | ":- endif."
```

Object directives

```
object_directive ::=
    ":- initialization(" callable ")."
  | ":- built_in."
  | ":- threaded."
  | ":- dynamic."
  | ":- info(" entity_info_list ")."
  | ":- set_logtalk_flag(" atom "," nonvar ")."
  | ":- include(" source_file_name ")."
  | ":- uses(" object_alias_list ")."
  | ":- use_module(" module_alias_list ")."
  | conditional_compilation_directive
  | predicate_directive
```

Category directives

```
category_directive ::=
  "- built_in."
  | "- dynamic."
  | "- info(" entity_info_list ")."
  | "- set_logtalk_flag(" atom "," nonvar ")."
  | "- include(" source_file_name ")."
  | "- uses(" object_alias_list ")."
  | "- use_module(" module_alias_list ")."
  | conditional_compilation_directive
  | predicate_directive
```

Protocol directives

```
protocol_directive ::=
  "- built_in."
  | "- dynamic."
  | "- info(" entity_info_list ")."
  | "- set_logtalk_flag(" atom "," nonvar ")."
  | "- include(" source_file_name ")."
  | conditional_compilation_directive
  | predicate_directive
```

Predicate directives

```
predicate_directive ::=
  alias_directive
  | synchronized_directive
  | uses_directive
  | use_module_directive
  | scope_directive
  | mode_directive
  | mode_non_terminal_directive
  | meta_predicate_directive
  | meta_non_terminal_directive
  | info_directive
  | dynamic_directive
  | discontinuous_directive
  | multifile_directive
  | coinductive_directive
  | operator_directive

alias_directive ::=
  "- alias(" entity_identifier "," alias_directive_resource_list ")."

synchronized_directive ::=
  "- synchronized(" synchronized_directive_resource_term ")."

uses_directive ::=
```

(continues on next page)

(continued from previous page)

```

    "- uses(" ( object_identifier | parameter_variable ) "," uses_directive_resource_list ").
    ↪"

use_module_directive ::=
    "- use_module(" ( module_identifier | parameter_variable ) "," use_module_directive_
    ↪resource_list ")."

scope_directive ::=
    "- public(" scope_directive_resource_term ")."
    | "- protected(" scope_directive_resource_term ")."
    | "- private(" scope_directive_resource_term ")."

mode_directive ::=
    "- mode(" predicate_mode_term "," number_of_proofs ")."

mode_non_terminal_directive ::=
    "- mode_non_terminal(" non_terminal_mode_term "," number_of_proofs ")."

meta_predicate_directive ::=
    "- meta_predicate(" meta_predicate_template_term ")."

meta_non_terminal_directive ::=
    "- meta_non_terminal(" meta_non_terminal_template_term ")."

info_directive ::=
    "- info(" ( predicate_indicator | non_terminal_indicator ) "," predicate_info_list ")."

dynamic_directive ::=
    "- dynamic(" qualified_directive_resource_term ")."

discontiguous_directive ::=
    "- discontiguous(" qualified_directive_resource_term ")."

multifile_directive ::=
    "- multifile(" qualified_directive_resource_term ")."

coinductive_directive ::=
    "- coinductive(" ( predicate_indicator_term | coinductive_predicate_template_term ) ")."

parameter_variable ::=
    _variable_

scope_directive_resource_term ::=
    scope_directive_resource
    | scope_directive_resource_sequence
    | scope_directive_resource_list

scope_directive_resource ::=
    predicate_indicator
    | non_terminal_indicator
    | operator

```

(continues on next page)

(continued from previous page)

```

scope_directive_resource_sequence ::=
    "(" scope_directive_resource ("," scope_directive_resource)* ")"

scope_directive_resource_list ::=
    "[" scope_directive_resource ("," scope_directive_resource)* "]"

entity_resources_list ::=
    predicate_indicator_list
    | operator_list

predicate_indicator_term ::=
    predicate_indicator
    | predicate_indicator_sequence
    | predicate_indicator_list

predicate_indicator_sequence ::=
    "(" predicate_indicator ("," predicate_indicator)* ")"

predicate_indicator_list ::=
    "[" predicate_indicator ("," predicate_indicator)* "]"

alias_directive_resource ::=
    predicate_indicator_alias
    | non_terminal_indicator_alias

alias_directive_resource_sequence ::=
    "(" alias_directive_resource ("," alias_directive_resource)* ")"

alias_directive_resource_list ::=
    "[" alias_directive_resource ("," alias_directive_resource)* "]"

synchronized_directive_resource_term ::=
    synchronized_directive_resource
    | synchronized_directive_resource_sequence
    | synchronized_directive_resource_list

synchronized_directive_resource ::=
    predicate_indicator
    | non_terminal_indicator

synchronized_directive_resource_sequence ::=
    "(" synchronized_directive_resource ("," synchronized_directive_resource)* ")"

synchronized_directive_resource_list ::=
    "[" synchronized_directive_resource ("," synchronized_directive_resource)* "]"

uses_directive_resource ::=
    predicate_indicator_alias
    | non_terminal_indicator_alias
    | predicate_indicator
    | non_terminal_indicator
    | predicate_template_alias

```

(continues on next page)

(continued from previous page)

```

| operator

uses_directive_resource_sequence ::=
  "(" uses_directive_resource ("," uses_directive_resource)* ")"

uses_directive_resource_list ::=
  "[" uses_directive_resource ("," uses_directive_resource)* "]"

use_module_directive_resource ::=
  predicate_indicator_alias
  | non_terminal_indicator_alias
  | predicate_indicator
  | non_terminal_indicator
  | predicate_template_alias
  | operator

use_module_directive_resource_sequence ::=
  "(" use_module_directive_resource ("," use_module_directive_resource)* ")"

use_module_directive_resource_list ::=
  "[" use_module_directive_resource ("," use_module_directive_resource)* "]"

qualified_directive_resource_term ::=
  qualified_directive_resource
  | qualified_directive_resource_sequence
  | qualified_directive_resource_list

qualified_directive_resource_sequence ::=
  "(" qualified_directive_resource_term ("," qualified_directive_resource_term)* ")"

qualified_directive_resource_list ::=
  "[" qualified_directive_resource_term ("," qualified_directive_resource_term)* "]"

qualified_directive_resource ::=
  predicate_indicator
  | non_terminal_indicator
  | object_identifier ":" ( predicate_indicator | non_terminal_indicator)
  | category_identifier ":" ( predicate_indicator | non_terminal_indicator)
  | module_identifier ":" ( predicate_indicator | non_terminal_indicator)

predicate_indicator_alias ::=
  predicate_indicator "as" predicate_indicator

predicate_template_alias ::=
  callable "as" callable

non_terminal_indicator ::=
  functor "/" arity

non_terminal_indicator_alias ::=
  non_terminal_indicator "as" non_terminal_indicator

```

(continues on next page)

(continued from previous page)

```

operator_sequence ::=
  operator specification
  | operator specification "," operator_sequence

operator_list ::=
  "[" operator_sequence "]"

coinductive_predicate_template_term ::=
  coinductive_predicate_template
  | coinductive_predicate_template_sequence
  | coinductive_predicate_template_list

coinductive_predicate_template_sequence ::=
  "(" coinductive_predicate_template "," coinductive_predicate_template)* ")"

coinductive_predicate_template_list ::=
  "[" coinductive_predicate_template "," coinductive_predicate_template)* "]"

coinductive_predicate_template ::=
  atom "(" coinductive_mode_terms ")"

coinductive_mode_terms ::=
  coinductive_mode_term
  | coinductive_mode_terms "," coinductive_mode_terms

coinductive_mode_term ::=
  "+"
  | "-"

predicate_mode_term ::=
  atom "(" mode_terms ")"

non_terminal_mode_term ::=
  atom "(" mode_terms ")"

mode_terms ::=
  mode_term
  | mode_term "," mode_terms

mode_term ::=
  "@" type?
  | "+" type?
  | "-" type?
  | "?" type?
  | "++" type?
  | "--" type?

type ::=
  prolog_type | logtalk_type | user_defined_type

prolog_type ::=
  "term"

```

(continues on next page)

(continued from previous page)

```

    | "nonvar"
    | "var"
    | "compound"
    | "ground"
    | "callable"
    | "list"
    | "atomic"
    | "atom"
    | "number"
    | "integer"
    | "float"

logtalk_type ::=
    "object"
    | "category"
    | "protocol"
    | "event"

user_defined_type ::=
    atom
    | compound

number_of_proofs ::=
    "zero"
    | "zero_or_one"
    | "zero_or_more"
    | "one"
    | "one_or_more"
    | "zero_or_error"
    | "one_or_error"
    | "zero_or_one_or_error"
    | "zero_or_more_or_error"
    | "one_or_more_or_error"
    | "error"

meta_predicate_template_term ::=
    meta_predicate_template
    | meta_predicate_template_sequence
    | meta_predicate_template_list

meta_predicate_template_sequence ::=
    "(" meta_predicate_template ("," meta_predicate_template)* ")"

meta_predicate_template_list ::=
    "[" meta_predicate_template ("," meta_predicate_template)* "]"

meta_predicate_template ::=
    object_identifier ":" atom "(" meta_predicate_specifiers ")"
    | category_identifier ":" atom "(" meta_predicate_specifiers ")"
    | module_identifier ":" atom "(" meta_predicate_specifiers ")"
    | atom "(" meta_predicate_specifiers ")"

```

(continues on next page)

(continued from previous page)

```

meta_predicate_specifiers ::=
  meta_predicate_specifier
  | meta_predicate_specifier "," meta_predicate_specifiers

meta_predicate_specifier ::=
  non_negative_integer
  | "::"
  | "^"
  | "*"

meta_non_terminal_template_term ::=
  meta_predicate_template_term

entity_info_pair ::=
  entity_info_item "is" nonvar

entity_info_list ::=
  "[" entity_info_pair ("," entity_info_pair)* "]"

entity_info_item ::=
  "comment"
  | "remarks"
  | "author"
  | "version"
  | "date"
  | "copyright"
  | "license"
  | "parameters"
  | "parnames"
  | "see_also"
  | atom

predicate_info_pair ::=
  predicate_info_item "is" nonvar

predicate_info_list ::=
  "[" predicate_info_pair ("," predicate_info_pair)* "]"

predicate_info_item ::=
  "comment"
  | "remarks"
  | "arguments"
  | "argnames"
  | "redefinition"
  | "allocation"
  | "examples"
  | "exceptions"
  | "see_also"
  | atom

object_alias ::=
  object_identifier "as" object_identifier

```

(continues on next page)

(continued from previous page)

```

object_alias_list ::=
  "[" object_alias ("," object_alias)* "]"

module_alias ::=
  module_identifier "as" module_identifier

module_alias_list ::=
  "[" module_alias ("," module_alias)* "]"

```

2.1.11 Clauses and goals

```

clause ::=
  object_identifier ":" head ":" body
  | module_identifier ":" head ":" body
  | head ":" body
  | object_identifier ":" fact
  | module_identifier ":" fact
  | fact

goal ::=
  message_sending
  | super_call
  | external_call
  | context_switching_call
  | callable

message_sending ::=
  message_to_object
  | message_delegation
  | message_to_self

message_to_object ::=
  receiver ":" messages

message_delegation ::=
  "[" message_to_object "]"

message_to_self ::=
  ":" messages

super_call ::=
  "^^" message

messages ::=
  message
  | "(" message "," messages ")"
  | "(" message ";" messages ")"
  | "(" message "->" messages ")"

```

(continues on next page)

(continued from previous page)

```

message ::=
  callable
  | variable

receiver ::=
  "{" callable "}"
  | object_identifier
  | variable

external_call ::=
  "{" callable "}"

context_switching_call ::=
  object_identifier "<<" callable

```

2.1.12 Lambda expressions

```

lambda_expression ::=
  lambda_free_variables "/" lambda_parameters ">>" callable
  | lambda_free_variables "/" callable
  | lambda_parameters ">>" callable

lambda_free_variables ::=
  "{" variables? "}"

lambda_parameters ::=
  "[" terms? "]"

variables ::=
  variable
  | variable "," variables

terms ::=
  term
  | term "," terms

```

2.1.13 Entity properties

```

category_property ::=
  "static"
  | "dynamic"
  | "built_in"
  | "file(" atom ")"
  | "file(" atom "," atom ")"
  | "lines(" integer "," integer ")"
  | "directive(" integer "," integer ")"
  | "events"
  | "source_data"
  | "public(" entity_resources_list ")"

```

(continues on next page)

(continued from previous page)

```

| "protected(" entity_resources_list ")")
| "private(" entity_resources_list ")")
| "declares(" predicate_indicator "," predicate_declaration_property_list ")")
| "defines(" predicate_indicator "," predicate_definition_property_list ")")
| "includes(" predicate_indicator "," ( object_identifier | category_identifier ) ","_
↪predicate_definition_property_list ")")
| "provides(" predicate_indicator "," ( object_identifier | category_identifier ) ","_
↪predicate_definition_property_list ")")
| "references(" reference "," reference_property_list ")")
| "alias(" ( object_identifier | module_identifier ) "," entity_alias_property_list ")")
| "alias(" predicate_indicator "," predicate_alias_property_list ")")
| "calls(" predicate "," predicate_call_update_property_list ")")
| "updates(" predicate "," predicate_call_update_property_list ")")
| "number_of_clauses(" integer ")")
| "number_of_rules(" integer ")")
| "number_of_user_clauses(" integer ")")
| "number_of_user_rules(" integer ")")
| "debugging"

object_property ::=
  "static"
  | "dynamic"
  | "built_in"
  | "threaded"
  | "file(" atom ")")
  | "file(" atom "," atom ")")
  | "lines(" integer "," integer ")")
  | "directive(" integer "," integer ")")
  | "context_switching_calls"
  | "dynamic_declarations"
  | "events"
  | "source_data"
  | "complements(" ( "allow" | "restrict" ) ")")
  | "complements"
  | "public(" entity_resources_list ")")
  | "protected(" entity_resources_list ")")
  | "private(" entity_resources_list ")")
  | "declares(" predicate_indicator "," predicate_declaration_property_list ")")
  | "defines(" predicate_indicator "," predicate_definition_property_list ")")
  | "includes(" predicate_indicator "," ( object_identifier | category_identifier ) ","_
↪predicate_definition_property_list ")")
  | "provides(" predicate_indicator "," ( object_identifier | category_identifier ) ","_
↪predicate_definition_property_list ")")
  | "references(" reference "," reference_property_list ")")
  | "alias(" ( object_identifier | module_identifier ) "," entity_alias_property_list ")")
  | "alias(" predicate_indicator "," predicate_alias_property_list ")")
  | "calls(" predicate "," predicate_call_update_property_list ")")
  | "updates(" predicate "," predicate_call_update_property_list ")")
  | "number_of_clauses(" integer ")")
  | "number_of_rules(" integer ")")
  | "number_of_user_clauses(" integer ")")
  | "number_of_user_rules(" integer ")")

```

(continues on next page)

(continued from previous page)

```

    | "module"
    | "debugging"

protocol_property ::=
    "static"
    | "dynamic"
    | "built_in"
    | "source_data"
    | "file(" atom ") "
    | "file(" atom "," atom ") "
    | "lines(" integer "," integer ") "
    | "directive(" integer "," integer ") "
    | "public(" entity_resources_list ") "
    | "protected(" entity_resources_list ") "
    | "private(" entity_resources_list ") "
    | "declares(" predicate_indicator "," predicate_declaration_property_list ") "
    | "alias(" predicate_indicator "," predicate_alias_property_list ") "
    | "debugging"

predicate_declaration_property_list ::=
    "[" predicate_declaration_property ("," predicate_declaration_property)* "]"

predicate_declaration_property ::=
    "static"
    | "dynamic"
    | "scope(" scope ") "
    | "private"
    | "protected"
    | "public"
    | "coinductive"
    | "multifile"
    | "synchronized"
    | "meta_predicate(" meta_predicate_template ") "
    | "coinductive(" coinductive_predicate_template ") "
    | "non_terminal(" non_terminal_indicator ") "
    | "include(" atom ") "
    | "lines(" integer "," integer ") "
    | "line_count(" integer ") "
    | "mode(" predicate_mode_term "," number_of_proofs ") "
    | "info(" list ") "

predicate_definition_property_list ::=
    "[" predicate_definition_property ("," predicate_definition_property)* "]"

predicate_definition_property ::=
    "inline"
    | "auxiliary"
    | "non_terminal(" non_terminal_indicator ") "
    | "include(" atom ") "
    | "lines(" integer "," integer ") "
    | "line_count(" integer ") "
    | "number_of_clauses(" integer ") "

```

(continues on next page)

(continued from previous page)

```

    | "number_of_rules(" integer ")")

entity_alias_property_list ::=
    "[" entity_alias_property ("," entity_alias_property)* "]"

entity_alias_property ::=
    "object"
    | "module"
    | "for(" ( object_identifier | module_identifier ) ")")
    | "include(" atom ")")
    | "lines(" integer "," integer ")")
    | "line_count(" integer ")")

predicate_alias_property_list ::=
    "[" predicate_alias_property ("," predicate_alias_property)* "]"

predicate_alias_property ::=
    "predicate"
    | "for(" predicate_indicator ")")
    | "from(" entity_identifier ")")
    | "non_terminal(" non_terminal_indicator ")")
    | "include(" atom ")")
    | "lines(" integer "," integer ")")
    | "line_count(" integer ")")

predicate ::=
    predicate_indicator
    | "^^" predicate_indicator
    | "::" predicate_indicator
    | ( variable | object_identifier ) "::" predicate_indicator
    | ( variable | module_identifier ) ":" predicate_indicator

predicate_call_update_property_list ::=
    "[" predicate_call_update_property ("," predicate_call_update_property)* "]"

predicate_call_update_property ::=
    "caller(" predicate_indicator ")")
    | "include(" atom ")")
    | "lines(" integer "," integer ")")
    | "line_count(" integer ")")
    | "alias(" predicate_indicator ")")
    | "non_terminal(" non_terminal_indicator ")")

reference ::=
    object_identifier | category_identifier
    | ( object_identifier | category_identifier ) "::" predicate_indicator

reference_property ::=
    "in(" ( "multifile" | "dynamic" | "discontiguous" | "meta_predicate" | "meta_non_terminal
↪ | "clause" ) ")")
    | "include(" atom ")")
    | "lines(" integer "," integer ")")

```

(continues on next page)

(continued from previous page)

```

| "line_count(" integer ")")
| "non_terminal(" non_terminal_indicator ")")

reference_property_list ::=
  "[" reference_property ("," reference_property)* "]"

```

2.1.14 Predicate properties

```

predicate_property ::=
  "static"
  | "dynamic"
  | "scope(" scope ")")
  | "private"
  | "protected"
  | "public"
  | "logtalk"
  | "prolog"
  | "foreign"
  | "coinductive(" coinductive_predicate_template ")")
  | "multifile"
  | "synchronized"
  | "built_in"
  | "inline"
  | "recursive"
  | "declared_in(" entity_identifier ")")
  | "defined_in(" ( object_identifier | category_identifier ) ")")
  | "redefined_from(" ( object_identifier | category_identifier ) ")")
  | "meta_predicate(" meta_predicate_template ")")
  | "alias_of(" callable ")")
  | "alias_declared_in(" entity_identifier ")")
  | "non_terminal(" non_terminal_indicator ")")
  | "mode(" predicate_mode_term "," number_of_proofs ")")
  | "info(" list ")")
  | "number_of_clauses(" integer ")")
  | "number_of_rules(" integer ")")
  | "declared_in(" entity_identifier "," line_count ")")
  | "defined_in(" ( object_identifier | category_identifier ) "," line_count ")")
  | "redefined_from(" ( object_identifier | category_identifier ) "," line_count ")")
  | "alias_declared_in(" entity_identifier "," line_count ")")

line_count ::=
  integer

```

2.1.15 Compiler flags

```
compiler_flag ::=  
  flag "(" flag_value ")"
```

2.2 Control constructs

2.2.1 Message sending

control construct

`(::)/2`

Description

```
Object::Message  
{Proxy}::Message
```

Sends a message to an object. The message argument must match a *public* predicate of the receiver object. When the message corresponds to a *protected* or *private* predicate, the call is only valid if the *sender* matches the *predicate scope container*. When the predicate is declared but not defined, the message simply fails (as per the *closed-world assumption*).

When the predicate used to answer the message is a *meta-predicate*, the *calling context* for the predicate meta-arguments is the object sending the message.

The `{Proxy}::Message` syntax allows simplified access to *parametric object proxies*. Its operational semantics is equivalent to the conjunction `call(Proxy), Proxy::Message`. I.e. `Proxy` is proved within the context of the pseudo-object *user* and, if successful, the `Proxy` term is used as an object identifier. Exceptions thrown when proving `Proxy` are handled by the `(::)/2` control construct. This control construct supports backtracking over the `{Proxy}` goal.

The lookups for the message declaration and the corresponding method are performed using a depth-first strategy. Depending on the value of the *optimize* flag, these lookups are performed at compile-time whenever sufficient information is available. When the lookups are performed at runtime, a caching mechanism is used to improve performance for subsequent messages. See the User Manual section on *performance* for details.

Modes and number of proofs

```
+object_identifier::+callable - zero_or_more  
{+object_identifier}::+callable - zero_or_more
```

Errors

Either Object or Message is a variable:

```
instantiation_error
```

Object is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Object)
```

Message is neither a variable nor a callable term:

```
type_error(callable, Message)
```

Message, with predicate indicator Name/Arity, is declared private:

```
permission_error(access, private_predicate, Name/Arity)
```

Message, with predicate indicator Name/Arity, is declared protected:

```
permission_error(access, protected_predicate, Name/Arity)
```

Message, with predicate indicator Name/Arity, is not declared:

```
existence_error(predicate_declaration, Name/Arity)
```

Object does not exist:

```
existence_error(object, Object)
```

Proxy is a variable:

```
instantiation_error
```

Proxy is neither a variable nor a callable term:

```
type_error(callable, Proxy)
```

Proxy, with predicate indicator Name/Arity, does not exist in the user pseudo-object:

```
existence_error(procedure, Name/Arity)
```

Examples

```
| ?- list::member(X, [1, 2, 3]).
```

```
X = 1 ;
```

```
X = 2 ;
```

```
X = 3
```

```
yes
```

See also

```
(::)/1, (^ ^)/1, []/1
```

control construct

`(::)/1`

Description

`::Message`

Sends a message to *self*. Can only be used in the body of a predicate definition. The argument should match a *public* or *protected* predicate of *self*. It may also match a *private* predicate if the predicate is within the scope of the object where the method making the call is defined, if imported from a category, if used from within a category, or when using private inheritance. When the predicate is declared but not defined, the message simply fails (as per the *closed-world assumption*).

When the predicate used to answer the message is a *meta-predicate*, the *calling context* for the predicate meta-arguments is the object sending the message.

The lookups for the message declaration and the corresponding method are performed using a depth-first strategy. A message to *self* necessarily implies the use of *dynamic binding* but a caching mechanism is used to improve performance in subsequent messages. See the User Manual section on *performance* for details.

Modes and number of proofs

`::+callable - zero_or_more`

Errors

Message is a variable:

`instantiation_error`

Message is neither a variable nor a callable term:

`type_error(callable, Message)`

Message, with predicate indicator Name/Arity, is declared private:

`permission_error(access, private_predicate, Name/Arity)`

Message, with predicate indicator Name/Arity, is not declared:

`existence_error(predicate_declaration, Name/Arity)`

Examples

```
area(Area) :-
  ::width(Width),
  ::height(Height),
  Area is Width * Height.
```

See also

`(::)/2`, `(^ ^)/1`, `[]/1`

2.2.2 Message delegation

control construct

`[]/1`

Description

```
[Object::Message]
[{Proxy}::Message]
```

This control construct allows the programmer to send a message to an object while preserving the original sender and meta-call context. It is mainly used in the definition of object handlers for unknown messages. This functionality is usually known as *delegation* but be aware that this is an overloaded word that can mean different things in different object-oriented programming languages.

To prevent the use of this control construct to break object encapsulation, an attempt to delegate a message to the original sender results in an error. The remaining error conditions are the same as the `(::)/2` control construct.

Note that, despite the correct functor for this control construct being (traditionally) `'.'/2`, we refer to it as `[]/1` simply to emphasize that the syntax is a list with a single element.

Modes and number of proofs

```
[+object_identifier::+callable] - zero_or_more
[+{+object_identifier}::+callable] - zero_or_more
```

Errors

Object is a variable:

```
instantiation_error
```

Object is neither a variable nor an object identifier:

```
type_error(object_identifier, Object)
```

Object does not exist:

```
existence_error(object, Object)
```

Object and the original *sender* are the same object:

```
permission_error(access, object, Sender)
```

Proxy is a variable:

```
instantiation_error
```

Proxy is neither a variable nor an object identifier:

```
type_error(object_identifier, Proxy)
```

Proxy, with predicate indicator Name/Arity, does not exist in the user pseudo-object:

```
existence_error(procedure, Name/Arity)
```

Message is a variable:

```
instantiation_error
```

Message is neither a variable nor a callable term:

```
type_error(callable, Message)
```

Message, with predicate indicator Name/Arity, is declared private:

```
permission_error(access, private_predicate, Name/Arity)
```

Message, with predicate indicator Name/Arity, is declared protected:

```
permission_error(access, protected_predicate, Name/Arity)
```

Message, with predicate indicator Name/Arity, is not declared:

```
existence_error(predicate_declaration, Name/Arity)
```

Examples

```
% delegate unknown messages to the "backup" object:
forward(Message) :-
    [backup::Message].
```

➡ See also

```
(::)/2, (::)/1, (^ ^)/1, forward/1
```

2.2.3 Calling imported and inherited predicates

control construct

```
(^^)/1
```

Description

```
^^Predicate
```

Calls an imported or inherited predicate definition. The call fails if the predicate is declared but there is no imported or inherited predicate definition (as per the *closed-world assumption*). This control construct may be used within objects or categories in the body of a predicate definition.

This control construct preserves the implicit execution context *self* and *sender* arguments (plus the meta-call context and coinduction stack when applicable) when calling the inherited (or imported) predicate definition.

The lookups for the predicate declaration and the predicate definition are performed using a depth-first strategy. Depending on the value of the *optimize* flag, these lookups are performed at compile-time when the predicate is static and sufficient information is available. When the lookups are performed at runtime, a caching mechanism is used to improve performance in subsequent calls. See the User Manual section on *performance* for details.

When the call is made from within an object, the lookup for the predicate definition starts at the imported categories, if any. If an imported predicate definition is not found, the lookup proceeds to the ancestor objects. Calls from predicates defined in complementing categories lookup inherited definitions as if the calls were made from the complemented object, thus allowing more comprehensive object patching. For other categories, the predicate definition lookup is restricted to the extended categories.

The called predicate should be declared *public* or *protected*. It may also be declared *private* if within the scope of the entity where the method making the call is defined.

This control construct is a generalization of the Smalltalk *super* keyword to take into account Logtalk support for prototypes and categories besides classes.

Modes and number of proofs

```
^^+callable - zero_or_more
```

Errors

Predicate is a variable:

```
instantiation_error
```

Predicate is neither a variable nor a callable term:

```
type_error(callable, Predicate)
```

Predicate, with predicate indicator Name/Arity, is declared private:

```
permission_error(access, private_predicate, Name/Arity)
```

Predicate, with predicate indicator Name/Arity, is not declared:

```
existence_error(predicate_declaration, Name/Arity)
```

Examples

```
% specialize the inherited definition
% of the init/0 predicate:
init :-
    assertz(counter(0)),
    ^^init.
```

See also

```
(::)/2, (::)/1, []/1
```

2.2.4 Calling predicates in *this*

control construct

(@)/1

Description

@Predicate

Calls a predicate definition in *this*. The argument must be a callable term at compile-time. The predicate must be declared (by a scope directive). This control construct provides access to predicate definitions in *this* from categories. For example, it allows overriding a predicate definition from a complementing category with a new definition that calls goals before and after calling the overridden definition (the overriding definition is sometimes described in other programming languages as an *around method*). When used within an object, it's the same as calling its argument.

Modes and number of proofs

@ +callable - zero_or_more

Errors

Predicate, with predicate indicator Name/Arity, is not declared:
existence_error(predicate_declaration, Name/Arity)

Examples

Assuming an object declaring a `make_sound/0` predicate, define an *around method* in a complementing category:

```
make_sound :-  
    write('Started making sound...'), nl,  
    @make_sound,  
    write('... finished making sound.'), nl.
```

See also

(::)/2, (::)/1, []/1

2.2.5 Calling external predicates

control construct

`{}/1`

Description

```
{Goal}
{Closure}
{Term}
```

This control construct allows the programmer to bypass the Logtalk compiler (including its linter but not its optimizer) in multiple contexts:

- Calling a goal as-is (from within an object or category) in the context of the *user* pseudo-object.
- Extending a closure as-is with the remaining arguments of a *call/2-N* call in order to construct a goal that will be called within the context of the user pseudo-object.
- Wrapping a source file term (either a clause or a directive) or a source file goal to bypass the *term-expansion mechanism*.
- Using it in place of an object identifier when sending a message. In this case, its argument is proved as a goal within the context of the user pseudo-object with the resulting term being used as an object identifier in the message-sending goal. This feature is mainly used with *parametric objects* when their identifiers correspond to predicates defined in user.
- Using it as a message to an object. This is mainly useful when the message is a *conjunction of messages* where some of goals are calls to Prolog built-in predicates.

Note

This control construct is opaque to cuts when used to wrap a goal (thus ensuring the same semantics independently of the argument being bound at compile-time or at runtime).

Modes and number of proofs

```
{+callable} - zero_or_more
```

Errors

Goal is a variable:

```
instantiation_error
```

Goal is neither a variable nor a callable term:

```
type_error(callable, Goal)
```

Closure is a variable:

```
instantiation_error
```

Closure is neither a variable nor a callable term:

```
type_error(callable, Closure)
```

Term is a variable:

```
instantiation_error
```

Term is neither a variable nor a callable term:

```
type_error(callable, Term)
```

Examples

```
% overload the standard (<)/2 operator by
% calling its standard built-in definition:
N1/D1 < N2/D2 :-
    {N1*D2 < N2*D1}.

% call a closure in the context of "user":
call_in_user(F, X, Y, Z) :-
    call({F}, X, Y, Z).

% bypass the compiler for a proprietary backend directive:
{: load_foreign_resource(file)}.

% use parametric object proxies:
| ?- {circle(Id, Radius, Color)}::area(Area).
...

% use Prolog built-in predicates as messages:
| ?- logtalk::write('hello world!'), nl.
hello world!
yes
```

2.2.6 Context switching calls

control construct

`(<<)/2`

Description

```
Object<<Goal
{Proxy}<<Goal
```

Debugging control construct. Calls a goal within the context of the specified object. The goal is called with the following execution context:

- *sender*, *this*, and *self* values set to the object
- empty meta-call context
- empty coinduction stack

The goal may need to be written between parenthesis to avoid parsing errors due to operator conflicts. This control construct should only be used for debugging or for writing unit tests. This control construct can only be used for objects compiled with the *context_switching_calls* compiler flag set to allow. Set this compiler flag to deny to disable this control construct and thus prevent using it to break encapsulation when deploying applications.

The {Proxy}<<Goal syntax allows simplified access to *parametric object proxies*. Its operational semantics is equivalent to the goal conjunction (call(Proxy), Proxy<<Goal). I.e. Proxy is proved within the context of the pseudo-object *user* and, if successful, the goal term is used as a parametric object identifier. Exceptions thrown when proving Proxy are handled by the (<<)/2 control construct. This syntax construct supports backtracking over the {Proxy} goal.

Caveat: although the goal argument is fully compiled before calling, some necessary information for the second compiler pass may not be available at runtime.

Modes and number of proofs

```
+object_identifier<<+callable - zero_or_more
{+object_identifier}<<+callable - zero_or_more
```

Errors

Object is a variable:

```
instantiation_error
```

Object is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Object)
```

Object does not contain a local definition for the Goal predicate:

```
existence_error(procedure, Goal)
```

Object does not exist:

```
existence_error(object, Object)
```

Object was created/compiled with support for context-switching calls turned off:

```
permission_error(access, database, Goal)
```

Proxy is a variable:

```
instantiation_error
```

Proxy is neither a variable nor an object identifier:

```
type_error(object_identifier, Proxy)
```

The predicate Proxy does not exist in the user pseudo-object:

```
existence_error(procedure, ProxyFunctor/ProxyArity)
```

Goal is a variable:

```
instantiation_error
```

Goal is neither a variable nor a callable term:

```
type_error(callable, Goal)
```

Examples

```
% call the member/2 predicate in the
% context of the "list" object:
test(member) :-
    list << member(1, [1]).
```

2.3 Directives

2.3.1 Source file directives

directive

`encoding/1`

Description

```
encoding(Encoding)
```

Declares the source file text encoding. Requires a *backend Prolog compiler* supporting the chosen encoding. When used, this directive must be the first term in the source file in the first line. This directive is also supported in files included in a main file or in a dynamically created entity using *include/1* directives.

The encoding used in a source file (and, in the case of a Unicode encoding, any BOM present) will be used for the intermediate Prolog file generated by the compiler. Logtalk uses the encoding names specified by *IANA*. In those cases where a preferred MIME name alias is specified, the alias is used instead. Examples includes 'US-ASCII', 'ISO-8859-1', 'ISO-8859-2', 'ISO-8859-15', 'UCS-2', 'UCS-2LE', 'UCS-2BE', 'UTF-8', 'UTF-16', 'UTF-16LE', 'UTF-16BE', 'UTF-32', 'UTF-32LE', 'UTF-32BE', 'Shift_JIS', and 'EUC-JP'. When writing portable code that cannot be expressed using ASCII, 'UTF-8' is the most commonly supported Unicode encoding.

The backend Prolog compiler adapter files define a table that translates between the Logtalk and Prolog specific atoms that represent each supported encoding. The *encoding_directive* read-only flag can be used to find if a backend supports this directive and how.

Template and modes

```
encoding(+atom)
```

Examples

```
:- encoding('UTF-8').
```

directive

include/1

Description

```
include(File)
```

Includes a file contents, which must be valid terms, at the place of occurrence of the directive. The file can be specified as a relative path, an absolute path, or using *library notation* and is expanded as a source file name. Relative paths are interpreted as relative to the path of the file containing the directive. The file extension is optional (the recognized Logtalk and Prolog file extensions are defined in the *backend adapter files*).

When using the *reflection API*, predicates from an included file can be distinguished from predicates from the main file by looking for the `include/1` predicate declaration property or the `include/1` predicate definition property. For the included predicates, the `line_count/1` property stores the term line number in the included file.

This directive can be used as either a source file directive or an entity directive. As an entity directive, it can be used both in entities defined in source files and with the entity creation built-in predicates. In the latter case, the file should be specified using an absolute path or using library notation (which expands to a full path) to avoid a fragile dependency on the current working directory.

Included files may contain an *encoding/1* directive, which may specify the same encoding of the main file or a different encoding.

Warning

The use of nested included files (i.e., included files including other files) is only partially supported. Notably, the reflection API currently only keeps track of the included files dependency on the main file. Although this allows predicates such as `logtalk_make/0-1` to work correctly, it prevents full code navigation features in supported IDEs for predicates defined in the nested included files.

When using this directive as an argument in calls to the *create_object/4*, *create_category/4*, and *create_protocol/3* built-in predicates, the objects, categories, and protocols will not be recreated or redefined when the included file(s) are modified and the *logtalk_make/0* predicate or the *logtalk_make/1* (with target all) predicates are called.

Template and modes

```
include(@source_file_name)
```

Examples

```
% include the "raw_1.txt" text file found
% on the "data" library directory:
:- include(data('raw_1.txt')).

% include a "factbase.pl" file in the same directory
% of the source file containing the directive:
:- include('factbase.pl').

% include a file given its absolute path:
:- include('/home/me/databases/countries.pl').

% create a wrapper object for a Prolog file using
% library notation to define the file path:
| ?- create_object(cities, [], [public(city/4), include(geo('cities.pl'))], []).
```

directive

initialization/1

Description

```
initialization(Goal)
```

When used within an object, this directive defines a goal to be called after the object has been successfully loaded into memory. When used at a global level within a source file, this directive defines a goal to be called after the source file is successfully compiled and loaded into memory.

The loading context can be accessed from this directive by calling the *logtalk_load_context/2* predicate. Note that the usable loading context keys depend on the directive scope (file or object).

Multiple initialization directives can be used in a source file or in an object. Their goals will be called in the same order as the directives at loading time.

Note

Arbitrary goals cannot be used as directives in source files. Any goal that should be automatically called when a source file is loaded must be wrapped using this directive.

Categories and protocols cannot contain initialization/1 directives as the initialization goals would lack a complete *execution context* that is only available for objects.

Although technically a global initialization/1 directive in a source file is a Prolog directive, calls to Logtalk built-in predicates from it are usually compiled to improve portability, improve performance, and provide better support for embedded applications.

Warning

Some backend Prolog compilers declare the atom initialization as an operator for a lighter syntax. But this makes the code non-portable and is therefore a practice best avoided.

Template and modes

```
initialization(@callable)
```

Examples

```
% call the init/0 predicate after loading the
% source file containing the directive
```

```
:- initialization(init).
```

```
% print a debug message after loading a
% source file defining an object
```

```
:- object(log).
```

```
    :- initialization(start_date).
```

```
start_date :-
```

```
    os::date_time(Year, Month, Day, _, _, _, _),
```

```
    logtalk::print_message(debug, my_app, 'Starting date: ~d~d~d~n'+[Year,Month,Day]).
```

```
:- end_object.
```

See also

[*logtalk_load_context/2*](#)

directive

op/3

Description

```
op(Precedence, Associativity, Operator)
op(Precedence, Associativity, [Operator, ...])
```

Declares operators. Global operators can be declared inside a source file by writing the respective directives before the entity opening directives. Operators declared inside entities have local scope. Calls to the standard term input and output predicates take into account any locally defined operators.

Template and modes

```
op(+integer, +associativity, +atom_or_atom_list)
```

Examples

Some of the predicate argument instantiation mode operators used by Logtalk:

```
:- op(200, fy, +).
:- op(200, fy, ?).
:- op(200, fy, @).
:- op(200, fy, -).
```

An example of using entity local operators. Consider the following ops.lgt file:

```
:- initialization((write(<=>(1,2)), nl)).

:- object(ops).

    :- op(700, xfx, <=>).

    :- public(w/1).
    w(Term) :-
        write(Term), nl.

    :- public(r/1).
    r(Term) :-
        read(Term).

:- end_object.
```

Loading the file automatically calls the initialization goal. Compare its output with the output of the ops::w/1 predicate. Compare also reading a term from within the ops object versus reading from user.

```
| ?- {ops}.
<=>(1,2)
true.

| ?- ops::w(<=>(1,2)).
1<=>2
true.

| ?- ops::r(T).
|: 3<=>4.

T = <=>(3, 4).

| ?- read(T).
|: 5<=>6.

SYNTAX ERROR: operator expected
```

 See also[current_op/3](#)**directive**`set_logtalk_flag/2`**Description**`set_logtalk_flag(Flag, Value)`

Sets local flag values. The scope of this directive is the entity or the source file containing it. For global scope, use the corresponding [set_logtalk_flag/2](#) built-in predicate called from an [initialization/1](#) directive. For a description of the predefined compiler flags, consult the [Compiler flags](#) section in the User Manual. The directive affects the compilation of all terms that follow it within the scope of the directive.

Template and modes`set_logtalk_flag(+atom, +nonvar)`**Errors**

Flag is a variable:

`instantiation_error`

Value is a variable:

`instantiation_error`

Flag is not an atom:

`type_error(atom, Flag)`

Flag is neither a variable nor a valid flag:

`domain_error(flag, Flag)`

Value is not a valid value for flag Flag:

`domain_error(flag_value, Flag + Value)`

Flag is a read-only flag:

`permission_error(modify, flag, Flag)`**Examples**

```
% turn off the compiler unknown entity warnings
% during the compilation of this source file:
:- set_logtalk_flag(unknown_entities, silent).

:- object(...).
```

(continues on next page)

(continued from previous page)

```
% generate events for messages sent from this object:
:- set_logtalk_flag(events, allow).
...

% turn off suspicious call lint checks for the next predicate:
:- set_logtalk_flag(suspicious_calls, silent).
foo :-
    ...
:- set_logtalk_flag(suspicious_calls, warning).
...
```

2.3.2 Conditional compilation directives

directive

if/1

Description

if(*Goal*)

Starts an *if-then* branch when performing conditional compilation. The code following the directive is compiled iff *Goal* is true. If *Goal* throws an error instead of either succeeding or failing, the error is reported by the compiler and compilation of the enclosing source file or entity is aborted. The goal is *expanded* when the directive occurs in a source file. Conditional compilation directives can be nested.

Warning

Conditional compilation goals **cannot** depend on predicate definitions contained in the same source file that contains the conditional compilation directives (as those predicates only become available after the file is successfully compiled and loaded).

Conditional compilation directives are currently not supported in included files when the conditional code includes `op/3` directives or `uses/2`/`use_module/2` directives with `op/3` arguments.

Template and modes

if(@callable)

Examples

A common example is checking if a built-in predicate exists and providing a definition when the predicate is absent:

```
:- if(\+ predicate_property(length(_, _), built_in)).

    length(List, Length) :-
        ...

:- endif.
```

Another common example is conditionally including code for a specific backend Prolog compiler:

```
:- if(current_logtalk_flag(prolog_dialect, swi)).

    % SWI-Prolog specific code
    :- set_prolog_flag(double_quotes, codes).

:- endif.
```

If necessary, test goal errors can be converted into failures using the standard catch/3 control construct. For example:

```
:- if(catch(\+ log(7, _), _, fail)).

    % define the legacy log/2 predicate
    log(X, Y) :- Y is log(X).

:- endif.
```

See also

elif/1, else/0, endif/0

directive

elif/1

Description

elif(*Goal*)

Supports embedded conditionals when performing conditional compilation. The code following the directive is compiled iff *Goal* is true. If *Goal* throws an error instead of either succeeding or failing, the error is reported by the compiler and compilation of the enclosing source file or entity is aborted. The goal is *expanded* when the directive occurs in a source file.

Warning

Conditional compilation goals **cannot** depend on predicate definitions contained in the same source file that contains the conditional compilation directives (as those predicates only become available after the file is successfully compiled and loaded).

Conditional compilation directives are currently not supported in included files when the conditional code includes `op/3` directives or `uses/2/use_module/2` directives with `op/3` arguments.

Template and modes

```
elif(@callable)
```

Examples

```
:- if(current_prolog_flag(double_quotes, codes)).  
    ...  
:- elif(current_prolog_flag(double_quotes, chars)).  
    ...  
:- elif(current_prolog_flag(double_quotes, atom)).  
    ...  
:- endif.
```

➡ See also

else/0, endif/0, if/1

directive

else/0

Description

```
else
```

Starts an *else* branch when performing conditional compilation. The code following this directive is compiled iff the goal in the matching *if/1* or *elif/1* directive is false.

Template and modes

```
else
```

Examples

An example where a hypothetical application would have some limitations that the user should be made aware of when running on a backend Prolog compiler with bounded arithmetic:

```
:- if(current_prolog_flag(bounded, true)).

    :- initialization(
        logtalk::print_message(warning, app, bounded_arithmetic)
    ).

:- else.

    :- initialization(
        logtalk::print_message(comment, app, unbounded_arithmetic)
    ).

:- endif.
```

See also

elif/1, endif/0, if/1

directive

endif/0

Description

```
endif
```

Ends conditional compilation for the matching *if/1* directive.

Template and modes

```
endif
```

Examples

```
:- if(date::today(_,5,25)).  
    :- initialization(write('Happy Towel Day!\n')).  
:- endif.
```

See also

elif/1, else/0, if/1

2.3.3 Entity directives

directive

`built_in/0`

Description

`built_in`

Declares an entity as built-in. Built-in entities must be static and cannot be redefined once loaded. This directive is used in the pre-defined protocols, categories, and objects that are automatically loaded at startup.

Template and modes

`built_in`

Examples

```
:- built_in.
```

directive

`category/1-4`

Description

```
category(Category)  
category(Category,  
    implements(Protocols))
```

(continues on next page)

(continued from previous page)

```

category(Category,
  extends(Categories))

category(Category,
  complements(Objects))

category(Category,
  implements(Protocols),
  extends(Categories))

category(Category,
  implements(Protocols),
  complements(Objects))

category(Category,
  extends(Categories),
  complements(Objects))

category(Category,
  implements(Protocols),
  extends(Categories),
  complements(Objects))

```

Starting category directive.

Template and modes

```

category(+category_identifier)

category(+category_identifier,
  implements(+implemented_protocols))

category(+category_identifier,
  extends(+extended_categories))

category(+category_identifier,
  complements(+complemented_objects))

category(+category_identifier,
  implements(+implemented_protocols),
  extends(+extended_categories))

category(+category_identifier,
  implements(+implemented_protocols),
  complements(+complemented_objects))

category(+category_identifier,
  extends(+extended_categories),
  complements(+complemented_objects))

category(+category_identifier,

```

(continues on next page)

(continued from previous page)

```
implements(+implemented_protocols),
extends(+extended_categories),
complements(+complemented_objects))
```

Examples

```
:- category(monitoring).

:- category(monitoring,
  implements(monitoringp)).

:- category(attributes,
  implements(protected::variables)).

:- category(extended,
  extends(minimal)).

:- category(logging,
  implements(monitoring),
  complements(employee)).
```

➡ See also

[*end_category/0*](#)

directive

dynamic/0

Description

dynamic

Declares an entity and its contents as dynamic. Dynamic entities can be abolished at runtime.

As entities created at runtime (using the [*create_object/4*](#), [*create_protocol/3*](#), and [*create_category/4*](#) built-in predicates) are always dynamic, this directive is only necessary in the rare cases where we want to define dynamic entities in source files.

⚠ Warning

Some backend Prolog compilers declare the atom `dynamic` as an operator for a lighter syntax, forcing writing this atom between parenthesis when using this directive.

Template and modes

```
dynamic
```

Examples

```
:- dynamic.
```

See also

dynamic/1, object_property/2, protocol_property/2, category_property/2

directive

`end_category/0`

Description

```
end_category
```

Ending category directive.

Template and modes

```
end_category
```

Examples

```
:- end_category.
```

See also

category/1-4

directive

`end_object/0`

Description

```
end_object
```


Ending object directive.

Template and modes

```
end_object
```

Examples

```
:- end_object.
```

 **See also**

[object/1-5](#)

directive

`end_protocol/0`

Description

```
end_protocol
```


Ending protocol directive.

Template and modes

```
end_protocol
```

Examples

```
:- end_protocol.
```

 **See also**

[protocol/1-2](#)

directive

info/1**Description**

```
info([Key is Value, ...])
```

Documentation directive for objects, protocols, and categories. The directive argument is a list of pairs using the format *Key is Value*. See the [Entity documenting directives](#) section for a description of the default keys.

Template and modes

```
info(+entity_info_list)
```

Examples

```
:- info([
    version is 1:0:0,
    author is 'Paulo Moura',
    date is 2000-11-20,
    comment is 'List protocol.'
]).
```

 **See also**

[info/2](#), [object_property/2](#), [protocol_property/2](#), [category_property/2](#)

directive**object/1-5****Description**

Stand-alone objects (prototypes)

```
object(Object)

object(Object,
    implements(Protocols))

object(Object,
    imports(Categories))

object(Object,
    implements(Protocols),
    imports(Categories))
```

Prototype extensions

```
object(Object,  
  extends(Objects))  
  
object(Object,  
  implements(Protocols),  
  extends(Objects))  
  
object(Object,  
  imports(Categories),  
  extends(Objects))  
  
object(Object,  
  implements(Protocols),  
  imports(Categories),  
  extends(Objects))
```

Class instances

```
object(Object,  
  instantiates(Classes))  
  
object(Object,  
  implements(Protocols),  
  instantiates(Classes))  
  
object(Object,  
  imports(Categories),  
  instantiates(Classes))  
  
object(Object,  
  implements(Protocols),  
  imports(Categories),  
  instantiates(Classes))
```

Classes

```
object(Object,  
  specializes(Classes))  
  
object(Object,  
  implements(Protocols),  
  specializes(Classes))  
  
object(Object,  
  imports(Categories),  
  specializes(Classes))  
  
object(Object,  
  implements(Protocols),  
  imports(Categories),  
  specializes(Classes))
```

Classes with metaclasses

```

object(Object,
  instantiates(Classes),
  specializes(Classes))

object(Object,
  implements(Protocols),
  instantiates(Classes),
  specializes(Classes))

object(Object,
  imports(Categories),
  instantiates(Classes),
  specializes(Classes))

object(Object,
  implements(Protocols),
  imports(Categories),
  instantiates(Classes),
  specializes(Classes))

```

Starting object directive.

Template and modes

Stand-alone objects (prototypes)

```

object(+object_identifier)

object(+object_identifier,
  implements(+implemented_protocols))

object(+object_identifier,
  imports(+imported_categories))

object(+object_identifier,
  implements(+implemented_protocols),
  imports(+imported_categories))

```

Prototype extensions

```

object(+object_identifier,
  extends(+extended_objects))

object(+object_identifier,
  implements(+implemented_protocols),
  extends(+extended_objects))

object(+object_identifier,
  imports(+imported_categories),
  extends(+extended_objects))

object(+object_identifier,

```

(continues on next page)

(continued from previous page)

```
implements(+implemented_protocols),
imports(+imported_categories),
extends(+extended_objects))
```

Class instances

```
object(+object_identifier,
      instantiates(+instantiated_objects))

object(+object_identifier,
      implements(+implemented_protocols),
      instantiates(+instantiated_objects))

object(+object_identifier,
      imports(+imported_categories),
      instantiates(+instantiated_objects))

object(+object_identifier,
      implements(+implemented_protocols),
      imports(+imported_categories),
      instantiates(+instantiated_objects))
```

Classes

```
object(+object_identifier,
      specializes(+specialized_objects))

object(+object_identifier,
      implements(+implemented_protocols),
      specializes(+specialized_objects))

object(+object_identifier,
      imports(+imported_categories),
      specializes(+specialized_objects))

object(+object_identifier,
      implements(+implemented_protocols),
      imports(+imported_categories),
      specializes(+specialized_objects))
```

Class with metaclasses

```
object(+object_identifier,
      instantiates(+instantiated_objects),
      specializes(+specialized_objects))

object(+object_identifier,
      implements(+implemented_protocols),
      instantiates(+instantiated_objects),
      specializes(+specialized_objects))

object(+object_identifier,
      imports(+imported_categories),
```

(continues on next page)

(continued from previous page)

```

    instantiates(+instantiated_objects),
    specializes(+specialized_objects))

object(+object_identifier,
    implements(+implemented_protocols),
    imports(+imported_categories),
    instantiates(+instantiated_objects),
    specializes(+specialized_objects))

```

Examples

```

:- object(list).

:- object(list,
    implements(listp)).

:- object(list,
    extends(compound)).

:- object(list,
    implements(listp),
    extends(compound)).

:- object(object,
    imports(initialization),
    instantiates(class)).

:- object(abstract_class,
    instantiates(class),
    specializes(object)).

:- object(agent,
    imports(private::attributes)).

```

See also

[*end_object/0*](#)

directive

protocol/1-2

Description

```
protocol(Protocol)  
  
protocol(Protocol,  
        extends(Protocols))
```

Starting protocol directive.

Template and modes

```
protocol(+protocol_identifier)  
  
protocol(+protocol_identifier,  
        extends(+extended_protocols))
```

Examples

```
:- protocol(listp).  
  
:- protocol(listp,  
    extends(compoundp)).  
  
:- protocol(queuep,  
    extends(protected::listp)).
```

See also

[*end_protocol/0*](#)

directive

threaded/0

Description

```
threaded
```

Declares that an object supports threaded engines, concurrent calls, and asynchronous messages. Any object containing calls to the built-in multi-threading predicates (or importing a category that contains such calls) must include this directive.

This directive results in the automatic creation and set up of an object message queue when the object is loaded or created at runtime. Object message queues are used for exchanging thread notifications and for storing concurrent goal solutions and replies to the multi-threading calls made within the object. The

message queue for the *user* pseudo-object is automatically created at Logtalk startup (provided that multi-threading programming is supported and enabled for the chosen *backend Prolog compiler*).

Note

This directive requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *threads* flag is set to supported when that is the case.

Template and modes

```
threaded
```

Examples

```
:- threaded.
```

See also

synchronized/1, *object_property/2*

directive

uses/1

Description

```
uses([Object as Alias, ...])
```

Declares object aliases. Typically used to shorten long object names, to simplify and consistently send messages to parameterized objects, and to simplify using or experimenting with different object implementations of the same protocol when using explicit message-sending. Object aliases are local to the object (or category) where they are defined.

The objects being aliased can be *parameter variables* or parametric objects where one or more parameters are parameter variables when using the directive in a parametric object or a parametric category defined in a source file (the common case).

Declaring multiple aliases for the same object is allowed. But repeated declarations of the same alias, declaring an alias for an object alias, and redefining an alias to reference a different object are reported as compilation errors.

To enable the use of *static binding*, and thus optimal message-sending performance, the objects should be loaded before compiling the entities that call their predicates.

Template and modes

```
uses(+object_alias_list)
```

Examples

```
:- object(foo(_HeapType_, _OptionsObject_)).

    :- uses([
        fast_random as rnd,
        time(utc) as time,
        heap(_HeapType_) as heap,
        _OptionsObject_ as options
    ]).

    bar :-
        ...,
        % the same as fast_random::permutation(L, P)
        rnd::permutation(L, P),
        % the same as heap(_HeapType_)::as_heap(L, H)
        heap::as_heap(L, H),
        % the same as _OptionsObject_::get(foo, X)
        options::get(foo, X),
        % the same as time(utc)::now(T)
        time::now(T),
        ...
```

See also

[uses/2](#), [use_module/1](#), [use_module/2](#)

directive

[use_module/1](#)

Description

```
use_module([Module as Alias, ...])
```

Declares module aliases. Typically used to shorten long module names and to simplify using or experimenting with different module implementations of the same predicates when using explicitly-qualified calls. Module aliases are local to the object (or category) where they are defined.

The modules being aliased can be *parameter variables* when using the directive in a parametric object or a parametric category defined in a source file (the common case).

Declaring multiple aliases for the same module is allowed. But repeated declarations of the same alias, declaring an alias for a module alias, and redefining an alias to reference a different module are reported as compilation errors.

To enable the use of *static binding*, and thus optimal predicate call performance, the modules should be loaded before compiling the entities that call their predicates.

Note that this directive semantics differs from the directive with the same name found on some Prolog implementations where it is used to load a module and import all its exported predicates.

Template and modes

```
use_module(+module_alias_list)
```

Examples

```
:- object(foo(_DataModule_)).

    :- use_module([
        data_noise_handler as cleaner,
        _DataModule_ as data
    ]).

    bar :-
        ...,
        % the same as _DataModule_:xy(X, Y)
        data:xy(X, Y),
        % the same as data_noise_handler:filter(X, Y)
        cleaner:filter(X, Y, Z),
        ...
```

See also

[uses/1](#), [uses/2](#), [use_module/2](#)

2.3.4 Predicate directives

directive

alias/2

Description

```
alias(Entity, [Name/Arity as Alias/Arity, ...])
alias(Entity, [Name//Arity as Alias//Arity, ...])
```

Declares predicate and grammar rule non-terminal *aliases*. A predicate (non-terminal) alias is an alternative name for a predicate (non-terminal) declared or defined in an extended protocol, an implemented protocol, an extended category, an imported category, an extended prototype, an instantiated class, or a specialized class. Predicate aliases may be used to solve conflicts between imported or inherited predicates. It may also be used to give a predicate (non-terminal) a name more appropriate in its usage context. This directive may be used in objects, protocols, and categories.

Predicate (and non-terminal) aliases are specified using (preferably) the notation `Name/Arity` as `Alias/Arity` or, in alternative, the notation `Name/Arity::Alias/Arity`.

It is also possible to declare predicate and grammar rule non-terminal aliases in implicit qualification directives for sending messages to objects and calling module predicates.

Template and modes

```
alias(@entity_identifier, +list(predicate_indicator_alias))
alias(@entity_identifier, +list(non_terminal_indicator_alias))
```

Examples

```
% resolve a predicate name conflict:
:- alias(list, [member/2 as list_member/2]).
:- alias(set, [member/2 as set_member/2]).

% define an alternative name for a non-terminal:
:- alias(words, [singular//0 as peculiar//0]).
```

See also

[uses/2](#), [use_module/2](#), [uses/1](#)

directive

coinductive/1

Description

```
coinductive(Name/Arity)
coinductive((Name/Arity, ...))
coinductive([Name/Arity, ...])

coinductive(Name//Arity)
coinductive((Name//Arity, ...))
coinductive([Name//Arity, ...])

coinductive(Template)
coinductive((Template1, ...))
coinductive([Template1, ...])
```

This is an **experimental** directive, used for declaring *coinductive* predicates. Requires a *backend Prolog compiler* with minimal support for cyclic terms. The current implementation of coinduction allows the generation of only the *basic cycles* but all valid solutions should be recognized. Use a predicate indicator or a non-terminal indicator as argument when all the coinductive predicate arguments are relevant for coinductive success. Use a template when only some coinductive predicate arguments (represented by a “+”) should be considered when testing for coinductive success (represent the arguments that should be disregarded by

a “-“). It’s possible to define local *coinductive_success_hook/1-2* predicates that are automatically called with the coinductive predicate term resulting from a successful unification with an ancestor goal as first argument. The second argument, when present, is the coinductive hypothesis (i.e., the ancestor goal) used. These hook predicates can provide an alternative to the use of tabling when defining some coinductive predicates. There is no overhead when these hook predicates are not defined.

This directive must precede any calls to the declared coinductive predicates.

Template and modes

```
coinductive(+predicate_indicator_term)
coinductive(+non_terminal_indicator_term)
coinductive(+coinductive_predicate_template_term)
```

Examples

```
:- coinductive(comember/2).
:- coinductive(ones_and_zeros//0).
:- coinductive(controller(+,+,+,-,-)).
```

See also

coinductive_success_hook/1-2, *predicate_property/2*

directive

discontiguous/1

Description

```
discontiguous(Name/Arity)
discontiguous((Name/Arity, ...))
discontiguous([Name/Arity, ...])

discontiguous(Entity::Name/Arity)
discontiguous((Entity::Name/Arity, ...))
discontiguous([Entity::Name/Arity, ...])

discontiguous(Module:Name/Arity)
discontiguous((Module:Name/Arity, ...))
discontiguous([Module:Name/Arity, ...])

discontiguous(Name//Arity)
discontiguous((Name//Arity, ...))
discontiguous([Name//Arity, ...])

discontiguous(Entity::Name//Arity)
discontiguous((Entity::Name//Arity, ...))
```

(continues on next page)

(continued from previous page)

```
discontiguous([Entity::Name//Arity, ...])

discontiguous(Module:Name//Arity)
discontiguous((Module:Name//Arity, ...))
discontiguous([Module:Name//Arity, ...])
```

Declares *discontiguous* predicates and discontiguous grammar rule non-terminals. The use of this directive should be avoided as not all *backend Prolog compilers* support discontiguous predicates.

Warning

Some backend Prolog compilers declare the atom *discontiguous* as an operator for a lighter syntax. But this makes the code non-portable and is therefore a practice best avoided.

Template and modes

```
discontiguous(+predicate_indicator_term)
discontiguous(+non_terminal_indicator_term)
```

Examples

```
:- discontiguous(counter/1).

:- discontiguous((lives/2, works/2)).

:- discontiguous([db/4, key/2, file/3]).
```

directive

dynamic/1

Description

```
dynamic(Name/Arity)
dynamic((Name/Arity, ...))
dynamic([Name/Arity, ...])

dynamic(Entity::Name/Arity)
dynamic((Entity::Name/Arity, ...))
dynamic([Entity::Name/Arity, ...])

dynamic(Module:Name/Arity)
dynamic((Module:Name/Arity, ...))
dynamic([Module:Name/Arity, ...])

dynamic(Name//Arity)
```

(continues on next page)

(continued from previous page)

```
dynamic((Name//Arity, ...))
dynamic([Name//Arity, ...])

dynamic(Entity::Name//Arity)
dynamic((Entity::Name//Arity, ...))
dynamic([Entity::Name//Arity, ...])

dynamic(Module:Name//Arity)
dynamic((Module:Name//Arity, ...))
dynamic([Module:Name//Arity, ...])
```

Declares *dynamic* predicates and dynamic grammar rule non-terminals. Note that an object can be static and have both static and dynamic predicates/non-terminals. When the dynamic predicates are local to an object, declaring them also as *private predicates* allows the Logtalk compiler to generate optimized code for asserting and retracting predicate clauses. Categories can also contain dynamic predicate directives but cannot contain clauses for dynamic predicates.

The predicate indicators (or non-terminal indicators) can be explicitly qualified with an object, category, or module identifier when the predicates (or non-terminals) are also declared *multifile*.

Note that dynamic predicates cannot be declared synchronized (when necessary, declare the predicates updating the dynamic predicates as *synchronized*).

Warning

Some backend Prolog compilers declare the atom *dynamic* as an operator for a lighter syntax. But this makes the code non-portable and is therefore a practice best avoided.

Template and modes

```
dynamic(+qualified_predicate_indicator_term)
dynamic(+qualified_non_terminal_indicator_term)
```

Examples

```
:- dynamic(counter/1).

:- dynamic((lives/2, works/2)).

:- dynamic([db/4, key/2, file/3]).
```

See also

dynamic/0, *predicate_property/2*

directive

info/2

Description

```
info(Name/Arity, [Key is Value, ...])
info(Name//Arity, [Key is Value, ...])
```

Documentation directive for predicates and grammar rule non-terminals. The first argument is either a predicate indicator or a grammar rule non-terminal indicator. The second argument is a list of pairs using the format *Key is Value*. See the *Predicate documenting directives* section for a description of the default keys.

Template and modes

```
info(+predicate_indicator, +predicate_info_list)
info(+non_terminal_indicator, +predicate_info_list)
```

Examples

```
:- info(empty/1, [
    comment is 'True if the argument is an empty list.',
    argnames is ['List']
]).

:- info(sentence//0, [
    comment is 'Rewrites a sentence into a noun phrase and a verb phrase.'
]).
```

See also

info/1, mode/2, predicate_property/2

directive

meta_predicate/1

Description

```
meta_predicate(Template)
meta_predicate((Template, ...))
meta_predicate([Template, ...])

meta_predicate(Entity::Template)
meta_predicate((Entity::Template, ...))
meta_predicate([Entity::Template, ...])

meta_predicate(Module:Template)
```

(continues on next page)

(continued from previous page)

```
meta_predicate((Module:Template, ...))
meta_predicate([Module:Template, ...])
```

Declares *meta-predicates*, i.e., predicates that have arguments interpreted as goals, arguments interpreted as *closures* used to construct goals, or arguments with sub-terms that will be interpreted as goals or closures. Meta-arguments are always called in the meta-predicate *calling context*, not in the meta-predicate *definition context*.

Meta-arguments which are goals are represented by the integer 0. Meta-arguments which are closures are represented by a positive integer, N, representing the number of additional arguments that will be appended to the closure in order to construct the corresponding goal (typically by calling the *call/1-N* built-in method). Meta-arguments with sub-terms that will be interpreted as goals or closures are represented by `::`. Meta-arguments that will be called using the *bagof/3* or *setof/3* predicates and that can thus be existentially-qualified are represented by the atom `^`. Normal arguments are represented by the atom `*`.

Logtalk allows the use of this directive to override the original meta-predicate directive. This is sometimes necessary when calling Prolog built-in meta-predicates or Prolog module meta-predicates due to the lack of standardization of the syntax of the meta-predicate templates. Another case is when a meta-predicate directive is missing. The compiler requires access to correct and non-ambiguous meta-predicate templates to correctly compile calls to Prolog meta-predicates.

Warning

Some backend Prolog compilers declare the atom `meta_predicate` as an operator for a lighter syntax. But this makes the code non-portable and is therefore a practice best avoided.

Template and modes

```
meta_predicate(+meta_predicate_template_term)

meta_predicate(+object_identifier::+meta_predicate_template_term)
meta_predicate(+category_identifier::+meta_predicate_template_term)

meta_predicate(+module_identifier::+meta_predicate_template_term)
```

Examples

```
% findall/3 second argument is interpreted as a goal:
:- meta_predicate(findall(*, 0, *)).

% both forall/2 arguments are interpreted as goals:
:- meta_predicate(forall(0, 0)).

% maplist/3 first argument is interpreted as a closure
% that will be expanded to a goal by appending two
% arguments:
:- meta_predicate(maplist(2, *, *)).
```

 See also[*meta_non_terminal/1*](#), [*predicate_property/2*](#)**directive****meta_non_terminal/1****Description**

```

meta_non_terminal(Template)
meta_non_terminal((Template, ...))
meta_non_terminal([Template, ...])

meta_non_terminal(Entity::Template)
meta_non_terminal((Entity::Template, ...))
meta_non_terminal([Entity::Template, ...])

meta_non_terminal(Module:Template)
meta_non_terminal((Module:Template, ...))
meta_non_terminal([Module:Template, ...])

```

Declares meta-non-terminals, i.e., non-terminals that have arguments that will be called as non-terminals (or grammar rule bodies). An argument may also be a *closure* instead of a goal if the non-terminal uses the *call//1-N* built-in methods to construct and call the actual non-terminal from the closure and the additional arguments or the *phrase//1* built-in method.

Meta-arguments which are non-terminals are represented by the integer 0. Meta-arguments which are closures are represented by a positive integer, N, representing the number of additional arguments that will be appended to the closure in order to construct the corresponding meta-call. Normal arguments are represented by the atom *. Meta-arguments are always called in the meta-non-terminal calling context, not in the meta-non-terminal definition context.

Template and modes

```

meta_non_terminal(+meta_non_terminal_template_term)

meta_non_terminal(+object_identifier::+meta_non_terminal_template_term)
meta_non_terminal(+category_identifier::+meta_non_terminal_template_term)

meta_non_terminal(+module_identifier::+meta_non_terminal_template_term)

```

Examples

```
:- meta_non_terminal(phrase(1, *)).
phrase(X, T) --> call(X, T).
```

➡ See also

[meta_predicate/1](#), [predicate_property/2](#)

directive

`mode/2`

Description

```
mode(Mode, NumberOfProofs)
```

Most predicates can be used with several instantiations modes. This directive enables the specification of each *instantiation mode* and the corresponding *number of proofs* (but not necessarily distinct solutions). The instantiation mode of each argument can include type information.

Template and modes

```
mode(+predicate_mode_term, +number_of_proofs)
mode(+non_terminal_mode_term, +number_of_proofs)
```

Examples

```
:- mode(atom_concat(-atom, -atom, +atom), one_or_more).
:- mode(atom_concat(+atom, +atom, -atom), one).

:- mode(var(@term), zero_or_one).

:- mode(solve(+callable, -list(atom)), zero_or_one).
```

➡ See also

[mode_non_terminal/2](#), [info/2](#), [predicate_property/2](#)

directive

mode_non_terminal/2**Description**

```
mode_non_terminal(Mode, NumberOfProofs)
```

Most non-terminals can be used with several instantiations modes. This directive enables the specification of each *instantiation mode* and the corresponding *number of proofs* (but not necessarily distinct solutions). The instantiation mode of each argument can include type information.

Template and modes

```
mode_non_terminal(+predicate_mode_term, +number_of_proofs)
mode_non_terminal(+non_terminal_mode_term, +number_of_proofs)
```

Examples

```
:- mode_non_terminal(zero_or_more(-list(atomic)), one).
```

 **See also**

mode/2, *info/2*, *predicate_property/2*

directive**multifile/1****Description**

```
multifile(Name/Arity)
multifile((Name/Arity, ...))
multifile([Name/Arity, ...])

multifile(Entity::Name/Arity)
multifile((Entity::Name/Arity, ...))
multifile([Entity::Name/Arity, ...])

multifile(Module:Name/Arity)
multifile((Module:Name/Arity, ...))
multifile([Module:Name/Arity, ...])

multifile(Name//Arity)
multifile((Name//Arity, ...))
multifile([Name//Arity, ...])

multifile(Entity::Name//Arity)
multifile((Entity::Name//Arity, ...))
```

(continues on next page)

(continued from previous page)

```
multifile([Entity::Name//Arity, ...])

multifile(Module:Name//Arity)
multifile((Module:Name//Arity, ...))
multifile([Module:Name//Arity, ...])
```

Declares *multifile* predicates and multifile grammar rule non-terminals. In the case of object or category multifile predicates, the predicate (or non-terminal) must also have a *scope directive* in the object or category holding its *primary declaration* (i.e., the declaration without the `Entity::` prefix). Entities holding multifile predicate primary declarations must be compiled and loaded prior to any entities contributing with clauses for the multifile predicates (to prevent using multifile predicates to break entity encapsulation).

Protocols cannot declare or define multifile predicates as protocols cannot contain predicate definitions.

Warning

Some backend Prolog compilers declare the atom `multifile` as an operator for a lighter syntax. But this makes the code non-portable and is therefore a practice best avoided.

Template and modes

```
multifile(+qualified_predicate_indicator_term)
multifile(+qualified_non_terminal_indicator_term)
```

Examples

```
:- multifile(table/3).
:- multifile(user::hook/2).
```

See also

public/1, protected/1, private/1, predicate_property/2

directive

private/1

Description

```
private(Name/Arity)
private((Name/Arity, ...))
private([Name/Arity, ...])

private(Name//Arity)
private((Name//Arity, ...))
```

(continues on next page)

(continued from previous page)

```
private([Name//Arity, ...])

private(op(Precedence,Associativity,Operator))
private((op(Precedence,Associativity,Operator), ...))
private([op(Precedence,Associativity,Operator), ...])
```

Declares *private* predicates, private grammar rule non-terminals, and private operators. A private predicate can only be called from the object containing the private directive. A private non-terminal can only be used in a call of the *phrase/2* and *phrase/3* methods from the object containing the private directive.

Template and modes

```
private(+predicate_indicator_term)
private(+non_terminal_indicator_term)
private(+operator_declaration)
```

Examples

```
:- private(counter/1).

:- private((init/1, free/1)).

:- private([data/3, key/1, keys/1]).
```

See also

protected/1, *public/1*, *predicate_property/2*

directive

protected/1

Description

```
protected(Name/Arity)
protected((Name/Arity, ...))
protected([Name/Arity, ...])

protected(Name//Arity)
protected((Name//Arity, ...))
protected([Name//Arity, ...])

protected(op(Precedence,Associativity,Operator))
protected((op(Precedence,Associativity,Operator), ...))
protected([op(Precedence,Associativity,Operator), ...])
```

Declares *protected* predicates, protected grammar rule non-terminals, and protected operators. A protected predicate can only be called from the object containing the directive or from an object that inherits the directive. A protected non-terminal can only be used as an argument in a *phrase/2* and *phrase/3* calls from the object containing the directive or from an object that inherits the directive.

Note

Protected operators are not inherited but declaring them provides a reusable specification for using them in descendant objects (or categories).

Template and modes

```
protected(+predicate_indicator_term)
protected(+non_terminal_indicator_term)
protected(+operator_declaration)
```

Examples

```
:- protected(init/1).
:- protected((print/2, convert/4)).
:- protected([load/1, save/3]).
```

See also

private/1, *public/1*, *predicate_property/2*

directive

public/1

Description

```
public(Name/Arity)
public((Name/Arity, ...))
public([Name/Arity, ...])

public(Name//Arity)
public((Name//Arity, ...))
public([Name//Arity, ...])

public(op(Precedence,Associativity,Operator))
public((op(Precedence,Associativity,Operator), ...))
public([op(Precedence,Associativity,Operator), ...])
```

Declares *public* predicates, public grammar rule non-terminals, and public operators. A public predicate can be called from any object. A public non-terminal can be used as an argument in *phrase/2* and *phrase/3* calls from any object.

Note

Declaring a public operator does not make it global when the entity holding the scope directive is compiled and loaded. But declaring public operators provides a reusable specification for using them in the entity clients (e.g., in *uses/2* directives).

Template and modes

```
public(+predicate_indicator_term)
public(+non_terminal_indicator_term)
public(+operator_declaration)
```

Examples

```
:- public(ancestor/1).
:- public((instance/1, instances/1)).
:- public([leaf/1, leaves/1]).
```

See also

private/1, *protected/1*, *predicate_property/2*

directive

synchronized/1

Description

```
synchronized(Name/Arity)
synchronized((Name/Arity, ...))
synchronized([Name/Arity, ...])

synchronized(Name//Arity)
synchronized((Name//Arity, ...))
synchronized([Name//Arity, ...])
```

Declares *synchronized* predicates and synchronized grammar rule non-terminals. The most common use is for predicates that have side effects (e.g., asserting or retracting clauses for a dynamic predicate) in multi-threaded applications. A synchronized predicate (or synchronized non-terminal) is protected by a mutex in order to allow for thread synchronization when proving a call to the predicate (or non-terminal).

All predicates (and non-terminals) declared in the same synchronized directive share the same mutex. In order to use a separate mutex for each predicate (non-terminal) so that they are independently synchronized, a per-predicate synchronized directive must be used.

Warning

Declaring a predicate synchronized implicitly makes it **deterministic**. When using a single-threaded *backend Prolog compiler*, calls to synchronized predicates behave as wrapped by the standard *once/1* meta-predicate.

Note that synchronized predicates cannot be declared *dynamic* (when necessary, declare the predicates updating the dynamic predicates as synchronized).

Template and modes

```
synchronized(+predicate_indicator_term)
synchronized(+non_terminal_indicator_term)
```

Examples

```
:- synchronized(db_update/1).

:- synchronized((write_stream/2, read_stream/2)).

:- synchronized([add_to_queue/2, remove_from_queue/2]).
```

See also

predicate_property/2

directive

uses/2

Description

```
uses(Object, [Name/Arity, ...])
uses(Object, [Name/Arity as Alias/Arity, ...])

uses(Object, [Predicate as Alias, ...])

uses(Object, [Name//Arity, ...])
uses(Object, [Name//Arity as Alias//Arity, ...])

uses(Object, [op(Precedence, Associativity, Operator), ...])
```

Declares that all calls made from predicates (or non-terminals) defined in the category or object containing the directive to the specified predicates (or non-terminals) are to be interpreted as messages to the specified object. Thus, this directive may be used to simplify writing of predicate definitions by allowing the programmer to omit the `Object::` prefix when using the predicates listed in the directive (as long as the calls do not occur as arguments for non-standard Prolog meta-predicates not declared in the adapter files). It is also possible to include operator declarations in the second argument.

This directive is also taken into account when compiling calls to the *database* and *reflection* built-in methods by looking into these methods predicate arguments if bound at compile-time.

It is possible to specify a predicate alias using the notation `Name/Arity` as `Alias/Arity` or, in alternative, the notation `Name/Arity::Alias/Arity`. Aliases may be used to avoid conflicts between predicates specified in `use_module/2` and `uses/2` directives, for giving more meaningful names considering the calling context of the predicates, and to change the order of the predicate arguments when calling the predicate. For predicates, it is also possible to define alias shorthands using the notation `Predicate` as `Alias` or, in alternative, the notation `Predicate::Alias`, where `Predicate` and `Alias` are callable terms where some or all arguments may be instantiated.

To enable the use of *static binding*, and thus optimal message-sending performance, the objects should be loaded before compiling the entities that call their predicates.

The object identifier argument can also be a *parameter variable* when using the directive in a parametric object or a parametric category defined in a source file (the common case). In this case, *dynamic binding* will be used for all listed predicates (and non-terminals). The parameter variable must be instantiated at runtime when the messages are sent.

Template and modes

```
uses(+object_identifier, +predicate_indicator_list)
uses(+object_identifier, +predicate_indicator_alias_list)

uses(+object_identifier, +predicate_template_alias_list)

uses(+object_identifier, +non_terminal_indicator_list)
uses(+object_identifier, +non_terminal_indicator_alias_list)

uses(+object_identifier, +operator_list)
```

Examples

```
:- uses(list, [append/3, member/2]).
:- uses(store, [data/2]).
:- uses(user, [table/4]).

foo :-
    ...,
    % the same as findall(X, list::member(X, L), A)
    findall(X, member(X, L), A),
    % the same as list::append(A, B, C)
    append(A, B, C),
    % the same as store::assertz(data(X, C))
    assertz(data(X, C)),
```

(continues on next page)

(continued from previous page)

```
% call the table/4 predicate in "user"
table(X, Y, Z, T),
...
```

Another example, using the extended notation that allows us to define predicate aliases:

```
:- uses(btrees, [new/1 as new_btree/1]).
:- uses(queues, [new/1 as new_queue/1]).

btree_to_queue :-
    ...,
    % the same as btrees::new(Tree)
    new_btree(Tree),
    % the same as queues::new(Queue)
    new_queue(Queue),
    ...
```

An example of defining a predicate alias that is also a shorthand:

```
:- uses(logtalk, [
    print_message(debug, my_app, Message) as dbg(Message)
]).
```

Predicate aliases can also be used to change argument order:

```
:- uses(meta, [
    fold_left(Closure, Accumulator, List, Result) as foldl(Closure, List, Accumulator, Result)
]).
```

An example of using a *parameter variable* in place of the object identifier to allow using the same test set for checking multiple implementations of the same protocol:

```
:- object(tests(_HeapObject_),
    extends(lgtunit)).

:- uses(_HeapObject_, [
    as_heap/2, as_list/2, valid/1, new/1,
    insert/4, insert_all/3, delete/4, merge/3,
    empty/1, size/2, top/3, top_next/5
]).
```

➡ See also

[uses/1](#), [use_module/1](#), [use_module/2](#), [alias/2](#)

directive

use_module/2

Description

```

use_module(Module, [Name/Arity, ...])
use_module(Module, [Name/Arity as Alias/Arity, ...])

use_module(Module, [Predicate as Alias, ...])

use_module(Module, [Name//Arity, ...])
use_module(Module, [Name//Arity as Alias//Arity, ...])

use_module(Module, [op(Precedence, Associativity, Operator), ...])

```

This directive declares that all calls (made from predicates defined in the category or object containing the directive) to the specified predicates (or non-terminals) are to be interpreted as calls to explicitly-qualified module predicates (or non-terminals). Thus, this directive may be used to simplify the writing of predicate definitions by allowing the programmer to omit the `Module:` prefix when using the predicates listed in the directive (as long as the predicate calls do not occur as arguments for non-standard Prolog meta-predicates not declared on the adapter files). It is also possible to include operator declarations in the second argument.

This directive is also taken into account when compiling calls to the *database* and *reflection* built-in methods by looking into these methods predicate arguments if bound at compile-time.

It is possible to specify a predicate alias using the notation `Name/Arity as Alias/Arity` or, in alternative, the notation `Name/Arity:Alias/Arity`. Aliases may be used either for avoiding conflicts between predicates specified in `use_module/2` and *uses/2* directives or for giving more meaningful names considering the calling context of the predicates. For predicates, it is also possible to define alias shorthands using the notation `Predicate as Alias` or, in alternative, the notation `Predicate::Alias`, where `Predicate` and `Alias` are callable terms where some or all arguments may be instantiated.

Note that this directive differs from the directive with the same name found on some Prolog implementations by requiring the first argument to be a module name (an atom) instead of a file specification. In Logtalk, there's no mixing between *loading* a resource and (declaring the) *using* (of) a resource. As a consequence, this directive doesn't automatically load the module. Loading the module file is dependent on the used *back-end Prolog compiler* and must be done separately (usually, using a source file directive such as `use_module/1` or `use_module/2` in the entity file or preferably in the application *loader file* file). Also, note that the name of the module may differ from the name of the module file.

Warning

The modules **must** be loaded prior to the compilation of entities that call the module predicates. This is required in general to allow the compiler to check if the called module predicate is a meta-predicate and retrieve its meta-predicate template to ensure proper call compilation.

The module identifier argument can also be a *parameter variable* when using the directive in a parametric object or a parametric category defined in a source file (the common case). In this case, *dynamic binding* will be used for all listed predicates (and non-terminals). The parameter variable must be instantiated at runtime when the calls are made.

Template and modes

```
use_module(+module_identifier, +predicate_indicator_list)
use_module(+module_identifier, +module_predicate_indicator_alias_list)

use_module(+module_identifier, +predicate_template_alias_list)

use_module(+module_identifier, +non_terminal_indicator_list)
use_module(+module_identifier, +module_non_terminal_indicator_alias_list)

use_module(+module_identifier, +operator_list)
```

Examples

```
:- use_module(lists, [append/3, member/2]).
:- use_module(store, [data/2]).
:- use_module(user, [foo/1 as bar/1]).

foo :-
    ...,
    % same as findall(X, lists:member(X, L), A)
    findall(X, member(X, L), A),
    % same as lists:append(A, B, C)
    append(A, B, C),
    % same as assertz(store:data(X, C))
    assertz(data(X, C)),
    % same as retractall(user:foo(_))
    retractall(bar(_)),
    ...
```

Another example, using the extended notation that allows us to define predicate aliases:

```
:- use_module(ugraphs, [transpose_ugraph/2 as transpose/2]).

convert_graph :-
    ...,
    % the same as ugraphs:transpose_ugraph(Graph0, Graph)
    transpose(Graph0, Graph),
    ...
```

An example of defining a predicate alias that is also a shorthand:

```
:- use_module(pairs, [
    map_list_to_pairs(length, Lists, Pairs) as length_pairs(Lists, Pairs)
]).
```

An example of using a *parameter variable* in place of the module identifier to delay to runtime the actual module to use:

```
:- object(bar(_OptionsModule)).
```

(continues on next page)

(continued from previous page)

```
:- use_module(_OptionsModule_, [
    set/2, get/2, reset/0
]).
```

➡ See also

use_module/1, uses/2, uses/1, alias/2

2.4 Built-in predicates

2.4.1 Enumerating objects, categories and protocols

built-in predicate

`current_category/1`

Description

`current_category(Category)`

Enumerates, by backtracking, all currently defined categories. All categories are found, either static, dynamic, or built-in.

Modes and number of proofs

`current_category(?category_identifier) - zero_or_more`

Errors

Category is neither a variable nor a valid category identifier:
`type_error(category_identifier, Category)`

Examples

```
% enumerate the defined categories:
| ?- current_category(Category).

Category = core_messages ;
...
```

 See also

abolish_category/1, category_property/2, create_category/4, complements_object/2, extends_category/2-3, imports_category/2-3

built-in predicate

`current_object/1`

Description

`current_object(Object)`

Enumerates, by backtracking, all currently defined objects. All objects are found, either static, dynamic or built-in.

Modes and number of proofs

`current_object(?object_identifier) - zero_or_more`

Errors

Object is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Object)
```

Examples

```
% enumerate the defined objects:
| ?- current_object(Object).
```

```
Object = user ;
Object = logtalk ;
...
```

 See also

abolish_object/1, create_object/4, object_property/2, extends_object/2-3, instantiates_class/2-3, specializes_class/2-3, complements_object/2

built-in predicate

`current_protocol/1`

Description

`current_protocol(Protocol)`

Enumerates, by backtracking, all currently defined protocols. All protocols are found, either static, dynamic, or built-in.

Modes and number of proofs

`current_protocol(?protocol_identifier) - zero_or_more`

Errors

Protocol is neither a variable nor a valid protocol identifier:

```
type_error(protocol_identifier, Protocol)
```

Examples

```
% enumerate the defined protocols:  
| ?- current_protocol(Protocol).  
  
Protocol = expanding ;  
Protocol = monitoring ;  
Protocol = forwarding ;  
...
```

See also

abolish_protocol/1, create_protocol/3, protocol_property/2, conforms_to_protocol/2-3, extends_protocol/2-3, implements_protocol/2-3

2.4.2 Enumerating objects, categories and protocols properties

built-in predicate

category_property/2

Description

```
category_property(Category, Property)
```

Enumerates, by backtracking, the properties associated with the defined categories. The valid properties are listed in the language grammar section on *entity properties* and described in the User Manual section on *category properties*.

Modes and number of proofs

```
category_property(?category_identifier, ?category_property) - zero_or_more
```

Errors

Category is neither a variable nor a valid category identifier:

```
type_error(category_identifier, Category)
```

Property is neither a variable nor a callable term:

```
type_error(callable, Property)
```

Property is a callable term but not a valid category property:

```
domain_error(category_property, Property)
```

Examples

```
% enumerate the properties of the core_messages built-in category:
| ?- category_property(core_messages, Property).

Property = source_data ;
Property = static ;
Property = built_in ;
...
```

➡ See also

abolish_category/1, *create_category/4*, *current_category/1*, *complements_object/2*, *extends_category/2-3*, *imports_category/2-3*

built-in predicate

object_property/2

Description

```
object_property(Object, Property)
```

Enumerates, by backtracking, the properties associated with the defined objects. The valid properties are listed in the language grammar section on *entity properties* and described in the User Manual section on *object properties*.

Modes and number of proofs

```
object_property(?object_identifier, ?object_property) - zero_or_more
```

Errors

Object is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Object)
```

Property is neither a variable nor a callable term:

```
type_error(callable, Property)
```

Property is a callable term but not a valid object property:

```
domain_error(object_property, Property)
```

Examples

```
% enumerate the properties of the logtalk built-in object:
| ?- object_property(logtalk, Property).

Property = context_switching_calls ;
Property = source_data ;
Property = threaded ;
Property = static ;
Property = built_in ;
...
```

See also

abolish_object/1, *create_object/4*, *current_object/1*, *extends_object/2-3*, *instantiates_class/2-3*, *specializes_class/2-3*, *complements_object/2*

built-in predicate

protocol_property/2

Description

```
protocol_property(Protocol, Property)
```

Enumerates, by backtracking, the properties associated with the currently defined protocols. The valid properties are listed in the language grammar section on *entity properties* and described in the User Manual section on *protocol properties*.

Modes and number of proofs

```
protocol_property(?protocol_identifier, ?protocol_property) - zero_or_more
```

Errors

Protocol is neither a variable nor a valid protocol identifier:

```
type_error(protocol_identifier, Protocol)
```

Property is neither a variable nor a callable term:

```
type_error(callable, Property)
```

Property is a callable term but not a valid protocol property:

```
domain_error(protocol_property, Property)
```

Examples

```
% enumerate the properties of the monitoring built-in protocol:
| ?- protocol_property(monitoring, Property).

Property = source_data ;
Property = static ;
Property = built_in ;
...
```

See also

abolish_protocol/1, create_protocol/3, current_protocol/1, conforms_to_protocol/2-3, extends_protocol/2-3, implements_protocol/2-3

2.4.3 Creating new objects, categories and protocols

built-in predicate

`create_category/4`

Description

```
create_category(Identifier, Relations, Directives, Clauses)
```

Creates a new, dynamic category. This predicate is often used as a primitive to implement high-level category creation methods.

Note that, when opting for runtime generated category identifiers, it's possible to run out of identifiers when using a *backend Prolog compiler* with bounded integer support. The portable solution, when creating a large number of dynamic categories in long-running applications, is to recycle, whenever possible, the identifiers.

When creating a new dynamic parametric category, access to the object parameters must use the *parameter/2* built-in execution context method.

When using Logtalk multi-threading features, predicates calling this built-in predicate may need to be declared synchronized in order to avoid race conditions.

Modes and number of proofs

```
create_category(?category_identifier, @list(category_relation), @list(category_directive), _  
↪ @list(clause)) - one
```

Errors

Relations, Directives, or Clauses is a variable:

```
instantiation_error
```

Identifier is neither a variable nor a valid category identifier:

```
type_error(category_identifier, Identifier)
```

Identifier is a valid object identifier but one of its arguments, Parameter, is not a variable:

```
type_error(variable, Parameter)
```

Identifier is already in use:

```
permission_error(modify, category, Identifier)
```

```
permission_error(modify, object, Identifier)
```

```
permission_error(modify, protocol, Identifier)
```

Relations is neither a variable nor a proper list:

```
type_error(list, Relations)
```

Repeated entity relation clause:

```
permission_error(repeat, entity_relation, implements/1)
```

```
permission_error(repeat, entity_relation, extends/1)
```

```
permission_error(repeat, entity_relation, complements/1)
```

Directives is neither a variable nor a proper list:

```
type_error(list, Directives)
```

Clauses is neither a variable nor a proper list:

```
type_error(list, Clauses)
```

Examples

```
| ?- create_category(
    tolerances,
    [implements(comparing)],
    [],
    [epsilon(1e-15), (equal(X, Y) :- epsilon(E), abs(X-Y) =< E)]
).
```

➡ See also

[abolish_category/1](#), [category_property/2](#), [current_category/1](#), [complements_object/2](#), [extends_category/2-3](#), [imports_category/2-3](#)

built-in predicate

create_object/4

Description

```
create_object(Identifier, Relations, Directives, Clauses)
```

Creates a new, dynamic object. The word *object* is used here as a generic term. This predicate can be used to create new prototypes, instances, and classes. This predicate is often used as a primitive to implement high-level object creation methods.

Note that, when opting for runtime generated object identifiers, it's possible to run out of identifiers when using a *backend Prolog compiler* with bounded integer support. The portable solution, when creating a large number of dynamic objects in long-running applications, is to recycle, whenever possible, the identifiers.

When creating a new dynamic parametric object, access to the object parameters must use the *parameter/2* built-in execution context method.

Declared predicates (using scope directives in the *Directives* argument) are implicitly declared also as dynamic.

When using Logtalk multi-threading features, predicates calling this built-in predicate may need to be declared synchronized in order to avoid race conditions.

Modes and number of proofs

```
create_object(?object_identifier, @list(object_relation), @list(object_directive),
↳ @list(clause)) - one
```

Errors

Relations, Directives, or Clauses is a variable:

```
instantiation_error
```

Identifier is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Identifier)
```

Identifier is a valid object identifier but one of its arguments, Parameter, is not a variable:

```
type_error(variable, Parameter)
```

Identifier is already in use:

```
permission_error(modify, category, Identifier)
```

```
permission_error(modify, object, Identifier)
```

```
permission_error(modify, protocol, Identifier)
```

Relations is neither a variable nor a proper list:

```
type_error(list, Relations)
```

Repeated entity relation clause:

```
permission_error(repeat, entity_relation, implements/1)
```

```
permission_error(repeat, entity_relation, imports/1)
```

```
permission_error(repeat, entity_relation, extends/1)
```

```
permission_error(repeat, entity_relation, instantiates/1)
```

```
permission_error(repeat, entity_relation, specializes/1)
```

Directives is neither a variable nor a proper list:

```
type_error(list, Directives)
```

Clauses is neither a variable nor a proper list:

```
type_error(list, Clauses)
```

Examples

```
% create a stand-alone object (a prototype):
| ?- create_object(
    translator,
    [],
    [public(int/2)],
    [int(0, zero)]
).

% create a prototype derived from a parent prototype:
| ?- create_object(
    mickey,
    [extends(mouse)],
    [public(alias/1)],
    [alias(mortimer)]
```

(continues on next page)

(continued from previous page)

```

    ).

% create a class instance:
| ?- create_object(
    p1,
    [instantiates(person)],
    [],
    [name('Paulo Moura'), age(42)]
    ).

% create a subclass:
| ?- create_object(
    hovercraft,
    [specializes(vehicle)],
    [public([propeller/2, fan/2])],
    []
    ).

% create an object with an initialization goal:
| ?- create_object(
    runner,
    [instantiates(runners)],
    [initialization::start],
    [length(22), time(60)]
    ).

% create an object supporting dynamic predicate declarations:
| ?- create_object(
    database,
    [],
    [set_logtalk_flag(dynamic_declarations, allow)],
    []
    ).

```

See also

abolish_object/1, current_object/1, object_property/2, extends_object/2-3, instantiates_class/2-3, specializes_class/2-3, complements_object/2

built-in predicate

create_protocol/3

Description

```
create_protocol(Identifier, Relations, Directives)
```

Creates a new, dynamic, protocol. This predicate is often used as a primitive to implement high-level protocol creation methods.

Note that, when opting for runtime generated protocol identifiers, it's possible to run out of identifiers when using a *backend Prolog compiler* with bounded integer support. The portable solution, when creating a large number of dynamic protocols in long-running applications, is to recycle, whenever possible, the identifiers.

When using Logtalk multi-threading features, predicates calling this built-in predicate may need to be declared synchronized in order to avoid race conditions.

Modes and number of proofs

```
create_protocol(?protocol_identifier, @list(protocol_relation), @list(protocol_directive)) -  
↳one
```

Errors

Either Relations or Directives is a variable:

```
instantiation_error
```

Identifier is neither a variable nor a valid protocol identifier:

```
type_error(protocol_identifier, Identifier)
```

Identifier is already in use:

```
permission_error(modify, category, Identifier)
```

```
permission_error(modify, object, Identifier)
```

```
permission_error(modify, protocol, Identifier)
```

Relations is neither a variable nor a proper list:

```
type_error(list, Relations)
```

Repeated entity relation clause:

```
permission_error(repeat, entity_relation, extends/1)
```

Directives is neither a variable nor a proper list:

```
type_error(list, Directives)
```

Examples

```
| ?- create_protocol(  
    logging,  
    [extends(monitored)],  
    [public([log_file/1, log_on/0, log_off/0])]  
).  
|
```

 See also

abolish_protocol/1, current_protocol/1, protocol_property/2, conforms_to_protocol/2-3, extends_protocol/2-3, implements_protocol/2-3

2.4.4 Abolishing objects, categories and protocols

built-in predicate

`abolish_category/1`

Description

`abolish_category(Category)`

Abolishes a dynamic category. The category identifier can then be reused when creating a new category.

Modes and number of proofs

`abolish_category(+category_identifier) - one`

Errors

Category is a variable:

`instantiation_error`

Category is neither a variable nor a valid category identifier:

`type_error(category_identifier, Category)`

Category is an identifier of a static category:

`permission_error(modify, static_category, Category)`

Category does not exist:

`existence_error(category, Category)`

Examples

| ?- `abolish_category(monitored).`

 See also

category_property/2, create_category/4, current_category/1 complements_object/2, extends_category/2-3, imports_category/2-3

built-in predicate

`abolish_object/1`

Description

```
abolish_object(Object)
```

Abolishes a dynamic object. The object identifier can then be reused when creating a new object.

Modes and number of proofs

```
abolish_object(+object_identifier) - one
```

Errors

Object is a variable:

`instantiation_error`

Object is neither a variable nor a valid object identifier:

`type_error(object_identifier, Object)`

Object is an identifier of a static object:

`permission_error(modify, static_object, Object)`

Object does not exist:

`existence_error(object, Object)`

Examples

```
| ?- abolish_object(list).
```

See also

[create_object/4](#), [current_object/1](#), [object_property/2](#), [extends_object/2-3](#), [instantiates_class/2-3](#), [specializes_class/2-3](#), [complements_object/2](#)

built-in predicate

`abolish_protocol/1`

Description

```
abolish_protocol(Protocol)
```

Abolishes a dynamic protocol. The protocol identifier can then be reused when creating a new protocol.

Modes and number of proofs

```
abolish_protocol(@protocol_identifier) - one
```

Errors

Protocol is a variable:

```
instantiation_error
```

Protocol is neither a variable nor a valid protocol identifier:

```
type_error(protocol_identifier, Protocol)
```

Protocol is an identifier of a static protocol:

```
permission_error(modify, static_protocol, Protocol)
```

Protocol does not exist:

```
existence_error(protocol, Protocol)
```

Examples

```
| ?- abolish_protocol(listp).
```

➡ See also

create_protocol/3, current_protocol/1, protocol_property/2, conforms_to_protocol/2-3, extends_protocol/2-3, implements_protocol/2-3

2.4.5 Objects, categories, and protocols relations

built-in predicate

`extends_object/2-3`

Description

```
extends_object(Prototype, Parent)
extends_object(Prototype, Parent, Scope)
```

Enumerates, by backtracking, all pairs of objects such that the first one extends the second. The relation scope is represented by the atoms `public`, `protected`, and `private`.

Modes and number of proofs

```
extends_object(?object_identifier, ?object_identifier) - zero_or_more  
extends_object(?object_identifier, ?object_identifier, ?scope) - zero_or_more
```

Errors

Prototype is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Prototype)
```

Parent is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Parent)
```

Scope is neither a variable nor an atom:

```
type_error(atom, Scope)
```

Scope is an atom but an invalid entity scope:

```
domain_error(scope, Scope)
```

Examples

```
% enumerate objects derived from the state_space prototype:  
| ?- extends_object(Object, state_space).  
  
% enumerate objects publicly derived from the list prototype:  
| ?- extends_object(Object, list, public).
```

See also

current_object/1, instantiates_class/2-3, specializes_class/2-3

built-in predicate

extends_protocol/2-3

Description

```
extends_protocol(Protocol, ParentProtocol)  
extends_protocol(Protocol, ParentProtocol, Scope)
```

Enumerates, by backtracking, all pairs of protocols such that the first one extends the second. The relation scope is represented by the atoms public, protected, and private.

Modes and number of proofs

```
extends_protocol(?protocol_identifier, ?protocol_identifier) - zero_or_more
extends_protocol(?protocol_identifier, ?protocol_identifier, ?scope) - zero_or_more
```

Errors

Protocol is neither a variable nor a valid protocol identifier:

```
type_error(protocol_identifier, Protocol)
```

ParentProtocol is neither a variable nor a valid protocol identifier:

```
type_error(protocol_identifier, ParentProtocol)
```

Scope is neither a variable nor an atom:

```
type_error(atom, Scope)
```

Scope is an atom but an invalid entity scope:

```
domain_error(scope, Scope)
```

Examples

```
% enumerate the protocols extended by the listp protocol:
| ?- extends_protocol(listp, Protocol).

% enumerate protocols that privately extend the term protocol:
| ?- extends_protocol(Protocol, term, private).
```

See also

current_protocol/1, implements_protocol/2-3, conforms_to_protocol/2-3

built-in predicate

extends_category/2-3

Description

```
extends_category(Category, ParentCategory)
extends_category(Category, ParentCategory, Scope)
```

Enumerates, by backtracking, all pairs of categories such that the first one extends the second. The relation scope is represented by the atoms public, protected, and private.

Modes and number of proofs

```
extends_category(?category_identifier, ?category_identifier) - zero_or_more  
extends_category(?category_identifier, ?category_identifier, ?scope) - zero_or_more
```

Errors

Category is neither a variable nor a valid protocol identifier:

```
type_error(category_identifier, Category)
```

ParentCategory is neither a variable nor a valid protocol identifier:

```
type_error(category_identifier, ParentCategory)
```

Scope is neither a variable nor an atom:

```
type_error(atom, Scope)
```

Scope is an atom but an invalid entity scope:

```
domain_error(scope, Scope)
```

Examples

```
% enumerate the categories extended by the derailleur category:  
| ?- extends_category(derailleur, Category).  
  
% enumerate categories that privately extend the basics category:  
| ?- extends_category(Category, basics, private).
```

See also

current_category/1, complements_object/2, imports_category/2-3

built-in predicate

implements_protocol/2-3

Description

```
implements_protocol(Object, Protocol)  
implements_protocol(Category, Protocol)  
  
implements_protocol(Object, Protocol, Scope)  
implements_protocol(Category, Protocol, Scope)
```

Enumerates, by backtracking, all pairs of entities such that an object or a category implements a protocol. The relation scope is represented by the atoms *public*, *protected*, and *private*. This predicate only returns direct implementation relations. For a transitive closure, see the *conforms_to_protocol/2-3* predicate.

Modes and number of proofs

```
implements_protocol(?object_identifier, ?protocol_identifier) - zero_or_more
implements_protocol(?category_identifier, ?protocol_identifier) - zero_or_more

implements_protocol(?object_identifier, ?protocol_identifier, ?scope) - zero_or_more
implements_protocol(?category_identifier, ?protocol_identifier, ?scope) - zero_or_more
```

Errors

Object is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Object)
```

Category is neither a variable nor a valid category identifier:

```
type_error(category_identifier, Category)
```

Protocol is neither a variable nor a valid protocol identifier:

```
type_error(protocol_identifier, Protocol)
```

Scope is neither a variable nor an atom:

```
type_error(atom, Scope)
```

Scope is an atom but an invalid entity scope:

```
domain_error(scope, Scope)
```

Examples

```
% check that the list object implements the listp protocol:
| ?- implements_protocol(list, listp).

% check that the list object publicly implements the listp protocol:
| ?- implements_protocol(list, listp, public).

% enumerate only objects that implement the listp protocol:
| ?- current_object(Object), implements_protocol(Object, listp).

% enumerate only categories that implement the serialization protocol:
| ?- current_category(Category), implements_protocol(Category, serialization).
```

See also

[current_object/1](#), [current_protocol/1](#), [current_category/1](#), [conforms_to_protocol/2-3](#)

built-in predicate

`conforms_to_protocol/2-3`

Description

```
conforms_to_protocol(Object, Protocol)
conforms_to_protocol(Category, Protocol)

conforms_to_protocol(Object, Protocol, Scope)
conforms_to_protocol(Category, Protocol, Scope)
```

Enumerates, by backtracking, all pairs of entities such that an object or a category conforms to a protocol. The relation scope is represented by the atoms public, protected, and private. This predicate implements a transitive closure for the protocol implementation relation.

Modes and number of proofs

```
conforms_to_protocol(?object_identifier, ?protocol_identifier) - zero_or_more
conforms_to_protocol(?category_identifier, ?protocol_identifier) - zero_or_more

conforms_to_protocol(?object_identifier, ?protocol_identifier, ?scope) - zero_or_more
conforms_to_protocol(?category_identifier, ?protocol_identifier, ?scope) - zero_or_more
```

Errors

Object is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Object)
```

Category is neither a variable nor a valid category identifier:

```
type_error(category_identifier, Category)
```

Protocol is neither a variable nor a valid protocol identifier:

```
type_error(protocol_identifier, Protocol)
```

Scope is neither a variable nor an atom:

```
type_error(atom, Scope)
```

Scope is an atom but an invalid entity scope:

```
domain_error(scope, Scope)
```

Examples

```
% enumerate objects and categories that conform to the listp protocol:
| ?- conforms_to_protocol(Object, listp).

% enumerate objects and categories that privately conform to the listp protocol:
| ?- conforms_to_protocol(Object, listp, private).

% enumerate only objects that conform to the listp protocol:
| ?- current_object(Object), conforms_to_protocol(Object, listp).
```

(continues on next page)

(continued from previous page)

```
% enumerate only categories that conform to the serialization protocol:
| ?- current_category(Category), conforms_to_protocol(Category, serialization).
```

➡ See also

current_object/1, current_protocol/1, current_category/1, implements_protocol/2-3

built-in predicate

complements_object/2

Description

complements_object(Category, Object)

Enumerates, by backtracking, all category–object pairs such that the category explicitly complements the object.

Modes and number of proofs

complements_object(?category_identifier, ?object_identifier) – zero_or_more

Errors

Category is neither a variable nor a valid category identifier:

```
type_error(category_identifier, Prototype)
```

Object is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Parent)
```

Examples

```
% check that the logging category complements the employee object:
| ?- complements_object(logging, employee).
```

➡ See also

current_category/1, imports_category/2-3

built-in predicate

`imports_category/2-3`

Description

```
imports_category(Object, Category)
imports_category(Object, Category, Scope)
```

Enumerates, by backtracking, importation relations between objects and categories. The relation scope is represented by the atoms `public`, `protected`, and `private`.

Modes and number of proofs

```
imports_category(?object_identifier, ?category_identifier) - zero_or_more
imports_category(?object_identifier, ?category_identifier, ?scope) - zero_or_more
```

Errors

Object is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Object)
```

Category is neither a variable nor a valid category identifier:

```
type_error(category_identifier, Category)
```

Scope is neither a variable nor an atom:

```
type_error(atom, Scope)
```

Scope is an atom but an invalid entity scope:

```
domain_error(scope, Scope)
```

Examples

```
% check that the xref_diagram object imports the diagram category:
| ?- imports_category(xref_diagram, diagram).

% enumerate the objects that privately import the diagram category:
| ?- imports_category(Object, diagram, private).
```

See also

[current_category/1](#), [complements_object/2](#)

built-in predicate

instantiates_class/2-3

Description

```
instantiates_class(Instance, Class)
instantiates_class(Instance, Class, Scope)
```

Enumerates, by backtracking, all pairs of objects such that the first one instantiates the second. The relation scope is represented by the atoms public, protected, and private.

Modes and number of proofs

```
instantiates_class(?object_identifier, ?object_identifier) - zero_or_more
instantiates_class(?object_identifier, ?object_identifier, ?scope) - zero_or_more
```

Errors

Instance is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Instance)
```

Class is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Class)
```

Scope is neither a variable nor an atom:

```
type_error(atom, Scope)
```

Scope is an atom but an invalid entity scope:

```
domain_error(scope, Scope)
```

Examples

```
% check that the water_jug is an instance of state_space:
| ?- instantiates_class(water_jug, state_space).

% enumerate the state_space instances where the
% instantiation relation is public:
| ?- instantiates_class(Space, state_space, public).
```

See also

current_object/1, extends_object/2-3, specializes_class/2-3

built-in predicate

specializes_class/2-3

Description

```
specializes_class(Class, Superclass)
specializes_class(Class, Superclass, Scope)
```

Enumerates, by backtracking, all pairs of objects such that the first one specializes the second. The relation scope is represented by the atoms public, protected, and private.

Modes and number of proofs

```
specializes_class(?object_identifier, ?object_identifier) - zero_or_more
specializes_class(?object_identifier, ?object_identifier, ?scope) - zero_or_more
```

Errors

Class is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Class)
```

Superclass is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Superclass)
```

Scope is neither a variable nor an atom:

```
type_error(atom, Scope)
```

Scope is an atom but an invalid entity scope:

```
domain_error(scope, Scope)
```

Examples

```
% enumerate the state_space subclasses:
| ?- specializes_class(Subclass, state_space).

% enumerate the state_space subclasses where the
% specialization relation is public:
| ?- specializes_class(Subclass, state_space, public).
```

See also

current_object/1, extends_object/2-3, instantiates_class/2-3

2.4.6 Event handling

built-in predicate

`abolish_events/5`

Description

```
abolish_events(Event, Object, Message, Sender, Monitor)
```

Abolishes all matching events. The two types of events are represented by the atoms before and after. When the predicate is called with the first argument unbound, both types of events are abolished.

Modes and number of proofs

```
abolish_events(@term, @term, @term, @term, @term) - one
```

Errors

Event is neither a variable nor a valid event identifier:

```
type_error(event, Event)
```

Object is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Object)
```

Message is neither a variable nor a callable term:

```
type_error(callable, Message)
```

Sender is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Sender)
```

Monitor is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Monitor)
```

Examples

```
% abolish all events for messages sent to the "list"
% object being monitored by the "debugger" object:
| ?- abolish_events(_, list, _, _, debugger).
```

See also

current_event/5, define_events/5, before/3, after/3

built-in predicate

`current_event/5`

Description

```
current_event(Event, Object, Message, Sender, Monitor)
```

Enumerates, by backtracking, all defined events. The two types of events are represented by the atoms `before` and `after`.

Modes and number of proofs

```
current_event(?event, ?term, ?term, ?term, ?object_identifier) - zero_or_more
```

Errors

Event is neither a variable nor a valid event identifier:

```
type_error(event, Event)
```

Object is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Object)
```

Message is neither a variable nor a callable term:

```
type_error(callable, Message)
```

Sender is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Sender)
```

Monitor is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Monitor)
```

Examples

```
% enumerate all events monitored by the "debugger" object:  
| ?- current_event(Event, Object, Message, Sender, debugger).
```

See also

[abolish_events/5](#), [define_events/5](#), [before/3](#), [after/3](#)

built-in predicate

define_events/5

Description

```
define_events(Event, Object, Message, Sender, Monitor)
```

Defines a new set of events. The two types of events are represented by the atoms `before` and `after`. When the predicate is called with the first argument unbound, both types of events are defined. The object `Monitor` must define the event handler methods required by the `Event` argument.

Modes and number of proofs

```
define_events(@term, @term, @term, @term, +object_identifier) - one
```

Errors

Event is neither a variable nor a valid event identifier:

```
type_error(event, Event)
```

Object is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Object)
```

Message is neither a variable nor a callable term:

```
type_error(callable, Message)
```

Sender is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Sender)
```

Monitor is a variable:

```
instantiation_error
```

Monitor is neither a variable nor a valid object identifier:

```
type_error(object_identifier, Monitor)
```

Monitor does not define the required `before/3` method:

```
existence_error(procedure, before/3)
```

Monitor does not define the required `after/3` method:

```
existence_error(procedure, after/3)
```

Examples

```
% define "debugger" as a monitor for member/2 messages
% sent to the "list" object:
| ?- define_events(_, list, member(_, _), _ , debugger).
```

See also

abolish_events/5, current_event/5, before/3, after/3

2.4.7 Multi-threading

built-in predicate

`threaded/1`

Description

```
threaded(Conjunction)
threaded(Disjunction)
```

Proves each goal in a conjunction or a disjunction of goals in its own thread. This meta-predicate is deterministic and opaque to cuts. The predicate argument is **not** flattened.

When the argument is a conjunction of goals, a call to this predicate blocks until either all goals succeed, one of the goals fail, or one of the goals generate an exception; the failure of one of the goals or an exception on the execution of one of the goals results in the termination of the remaining threads. The predicate call is true *iff* all goals are true. The predicate call fails if all goals fail. When one of the goals throws an exception, the predicate call re-throws that exception.

When the argument is a disjunction of goals, a call to this predicate blocks until either one of the goals succeeds or all the goals fail or throw exceptions; the success of one of the goals results in the termination of the remaining threads. The predicate call is true *iff* one of the goals is true. The predicate call fails if all goals fails. When no goal succeeds and one of the goals throws an exception, the predicate call re-throws that exception.

When the predicate argument is neither a conjunction nor a disjunction of goals, no threads are used. In this case, the predicate call is equivalent to a `once/1` predicate call.

A dedicated message queue is used per call of this predicate to collect the individual goal results.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *threads* flag is set to supported when that is the case.

Meta-predicate template

```
threaded(θ)
```

Modes and number of proofs

```
threaded(+callable) - zero_or_one
```

Errors

Goals is a variable:

```
instantiation_error
```

A goal in Goals is a variable:

```
instantiation_error
```

Goals is neither a variable nor a callable term:

```
type_error(callable, Goals)
```

A goal Goal in Goals is neither a variable nor a callable term:

```
type_error(callable, Goal)
```

Examples

Prove a conjunction of goals, each one in its own thread:

```
threaded((Goal, Goals))
```

Prove a disjunction of goals, each one in its own thread:

```
threaded((Goal; Goals))
```

See also

[threaded_call/1-2](#), [threaded_once/1-2](#), [threaded_ignore/1](#), [synchronized/1](#)

built-in predicate

threaded_call/1-2

Description

```
threaded_call(Goal)
threaded_call(Goal, Tag)
```

Proves Goal asynchronously using a new thread. The argument can be a message-sending goal. Calls to this predicate always succeed and return immediately. The results (success, failure, or exception) are sent back to the message queue of the object containing the call (*this*) and can be retrieved by calling the [threaded_exit/1](#) predicate.

The `threaded_call/2` variant returns a threaded call identifier tag that can be used with the [threaded_exit/2](#) and [threaded_cancel/1](#) predicates. Tags shall be regarded as opaque terms; users shall not rely on their type.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only `threads` flag is set to supported when that is the case.

Meta-predicate template

```
threaded_call(0)  
threaded_call(0, *)
```

Modes and number of proofs

```
threaded_call(@callable) - one  
threaded_call(@callable, --nonvar) - one
```

Errors

Goal is a variable:

instantiation_error

Goal is neither a variable nor a callable term:

type_error(callable, Goal)

Tag is not a variable:

uninstantiation_error(Tag)

Examples

Prove Goal asynchronously in a new thread:

threaded_call(Goal)

Prove ::Message asynchronously in a new thread:

threaded_call(::Message)

Prove Object::Message asynchronously in a new thread:

threaded_call(Object::Message)

See also

threaded_exit/1-2, *threaded_ignore/1*, *threaded_once/1-2*, *threaded_peek/1-2*, *threaded_cancel/1*,
threaded/1, *synchronized/1*

built-in predicate

threaded_once/1-2

Description

```
threaded_once(Goal)
threaded_once(Goal, Tag)
```

Proves Goal asynchronously using a new thread. Only the first goal solution is found. The argument can be a message-sending goal. This call always succeeds. The result (success, failure, or exception) is sent back to the message queue of the object containing the call (*this*).

The threaded_once/2 variant returns a threaded call identifier tag that can be used with the *threaded_exit/2* and *threaded_cancel/1* predicates. Tags shall be regarded as opaque terms; users shall not rely on their type.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *threads* flag is set to supported when that is the case.

Meta-predicate template

```
threaded_once(0)
threaded_once(0, *)
```

Modes and number of proofs

```
threaded_once(@callable) - one
threaded_once(@callable, --nonvar) - one
```

Errors

Goal is a variable:

instantiation_error

Goal is neither a variable nor a callable term:

type_error(callable, Goal)

Tag is not a variable:

uninstantiation_error(Tag)

Examples

Prove Goal asynchronously in a new thread:

```
threaded_once(Goal)
```

Prove ::Message asynchronously in a new thread:

```
threaded_once(::Message)
```

Prove Object::Message asynchronously in a new thread:

```
threaded_once(Object::Message)
```

See also

[threaded_call/1-2](#), [threaded_exit/1-2](#), [threaded_ignore/1](#), [threaded_peek/1-2](#), [threaded_cancel/1](#),
[threaded/1](#), [synchronized/1](#)

built-in predicate

threaded_ignore/1

Description

threaded_ignore(Goal)

Proves Goal asynchronously using a new thread. Only the first goal solution is found. The argument can be a message-sending goal. This call always succeeds, independently of the result (success, failure, or exception), which is simply discarded instead of being sent back to the message queue of the object containing the call (*this*).

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *threads* flag is set to supported when that is the case.

Meta-predicate template

```
threaded_ignore(0)
```

Modes and number of proofs

```
threaded_ignore(@callable) - one
```

Errors

Goal is a variable:

```
instantiation_error
```

Goal is neither a variable nor a callable term:

```
type_error(callable, Goal)
```

Examples

Prove Goal asynchronously in a new thread:

```
threaded_ignore(Goal)
```

Prove ::Message asynchronously in a new thread:

```
threaded_ignore(::Message)
```

Prove Object::Message asynchronously in a new thread:

```
threaded_ignore(Object::Message)
```

See also

[threaded_call/1-2](#), [threaded_exit/1-2](#), [threaded_once/1-2](#), [threaded_peek/1-2](#), [threaded/1](#), [synchronized/1](#)

built-in predicate

threaded_exit/1-2

Description

```
threaded_exit(Goal)
threaded_exit(Goal, Tag)
```

Retrieves the result of proving Goal in a new thread. This predicate blocks execution until the reply is sent to the *this* message queue by the thread executing the goal. When there is no thread proving the goal, the predicate generates an exception. This predicate is non-deterministic, providing access to any alternative solutions of its argument.

The argument of this predicate should be a *variant* of the argument of the corresponding *threaded_call/1* or *threaded_once/1* call. When the predicate argument is subsumed by the *threaded_call/1* or *threaded_once/1* call argument, the *threaded_exit/1* call will succeed iff its argument is a solution of the (more general) goal.

The *threaded_exit/2* variant accepts a threaded call identifier tag generated by the calls to the *threaded_call/2* and *threaded_once/2* predicates. Tags shall be regarded as an opaque term; users shall not rely on their type.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *threads* flag is set to supported when that is the case.

Modes and number of proofs

```
threaded_exit(+callable) - zero_or_more
threaded_exit(+callable, +nonvar) - zero_or_more
```

Errors

Goal is a variable:

instantiation_error

Goal is neither a variable nor a callable term:

type_error(callable, Goal)

No thread is running for proving Goal in the Object calling context:

existence_error(thread, Object)

Tag is a variable:

instantiation_error

Examples

To retrieve an asynchronous goal proof result:

```
threaded_exit(Goal)
```

To retrieve an asynchronous message to *self* result:

```
threaded_exit(::Goal)
```

To retrieve an asynchronous message result:

```
threaded_exit(Object::Goal)
```

See also

threaded_call/1-2, *threaded_ignore/1*, *threaded_once/1-2*, *threaded_peek/1-2*, *threaded_cancel/1*, *threaded/1*

built-in predicate

threaded_peek/1-2**Description**

```
threaded_peek(Goal)
threaded_peek(Goal, Tag)
```

Checks if the result of proving *Goal* in a new thread is already available. This call succeeds or fails without blocking execution waiting for a reply to be available. When there is no thread proving the goal, the predicate generates an exception.

The argument of this predicate should be a *variant* of the argument of the corresponding *threaded_call/1* or *threaded_once/1* call. When the predicate argument is subsumed by the *threaded_call/1* or *threaded_once/1* call argument, the *threaded_peek/1* call will succeed iff its argument unifies with an already available solution of the (more general) goal.

The *threaded_peek/2* variant accepts a threaded call identifier tag generated by the calls to the *threaded_call/2* and *threaded_once/2* predicates. Tags shall be regarded as an opaque term; users shall not rely on their type.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *threads* flag is set to supported when that is the case.

Modes and number of proofs

```
threaded_peek(+callable) - zero_or_one
threaded_peek(+callable, +nonvar) - zero_or_one
```

Errors

Goal is a variable:

```
instantiation_error
```

Goal is neither a variable nor a callable term:

```
type_error(callable, Goal)
```

No thread is running for proving *Goal* in the Object calling context:

```
existence_error(thread, Object)
```

Tag is a variable:

```
instantiation_error
```

Examples

To check for an asynchronous goal proof result:

```
threaded_peek(Goal)
```

To check for an asynchronous message to *self* result:

```
threaded_peek(::Goal)
```

To check for an asynchronous message result:

```
threaded_peek(Object::Goal)
```

See also

[threaded_call/1-2](#), [threaded_exit/1-2](#), [threaded_ignore/1](#), [threaded_once/1-2](#), [threaded_cancel/1](#), [threaded/1](#)

built-in predicate

threaded_cancel/1

Description

threaded_cancel(Tag)

Cancels a tagged threaded call. When there is no asynchronous call with the given tag, calling this predicate succeeds assuming the asynchronous call has already terminated or canceled. The threaded call identifier tag is generated by calls to the [threaded_call/2](#) and [threaded_once/2](#) predicates. Tags shall be regarded as an opaque term; users shall not rely on their type.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *threads* flag is set to supported when that is the case.

Modes and number of proofs

```
threaded_cancel(+nonvar) - one
```

Errors

Tag is a variable:

```
instantiation_error
```

Examples

(none)

➡ See also

threaded_call/1-2, threaded_exit/1-2, threaded_ignore/1, threaded_once/1-2, threaded/1

built-in predicate

threaded_wait/1

Description

```
threaded_wait(Term)
threaded_wait([Term| Terms])
```

Suspends the thread making the call until a notification is received that unifies with Term. The call must be made within the same object (*this*) containing the calls to the *threaded_notify/1* predicate that will eventually send the notification. The argument may also be a list of notifications, [Term| Terms]. In this case, the thread making the call will suspend until all notifications in the list are received.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *threads* flag is set to supported when that is the case.

Modes and number of proofs

```
threaded_wait(?term) - one
threaded_wait(+list(term)) - one
```

Errors

(none)

Examples

```
% wait until the "data_available" notification is received:  
..., threaded_wait(data_available), ...
```

See also

[*threaded_notify/1*](#)

built-in predicate

threaded_notify/1

Description

```
threaded_notify(Term)  
threaded_notify([Term| Terms])
```

Sends Term as a notification to any thread suspended waiting for it in order to proceed. The call must be made within the same object (*this*) containing the calls to the [*threaded_wait/1*](#) predicate waiting for the notification. The argument may also be a list of notifications, [Term| Terms]. In this case, all notifications in the list will be sent to any threads suspended waiting for them in order to proceed.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *threads* flag is set to supported when that is the case.

Modes and number of proofs

```
threaded_notify(@term) - one  
threaded_notify(@list(term)) - one
```

Errors

(none)

Examples

```
% send a "data_available" notification:
..., threaded_notify(data_available), ...
```

➡ See also

threaded_wait/1

2.4.8 Multi-threading engines

built-in predicate

`threaded_engine_create/3`

Description

```
threaded_engine_create(AnswerTemplate, Goal, Engine)
```

Creates a new engine for proving the given goal and defines an answer template for retrieving the goal solution bindings. A message queue for passing arbitrary terms to the engine is also created. If the name for the engine is not given, a unique name is generated and returned. Engine names shall be regarded as opaque terms; users shall not rely on its type.

i Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *engines* flag is set to supported when that is the case.

Meta-predicate template

```
threaded_engine_create(*, 0, *)
```

Modes and number of proofs

```
threaded_engine_create(@term, @callable, @nonvar) - one
threaded_engine_create(@term, @callable, --nonvar) - one
```

Errors

Goal is a variable:

instantiation_error

Goal is neither a variable nor a callable term:

type_error(callable, Goal)

Engine is the name of an existing engine:

permission_error(create, engine, Engine)

Examples

```
% create a new engine to enumerate list elements:
| ?- threaded_engine_create(X, member(X, [1,2,3]), worker_1).
```

➡ See also

threaded_engine_destroy/1, threaded_engine_self/1, threaded_engine/1, threaded_engine_next/2, threaded_engine_next_reified/2

built-in predicate

threaded_engine_destroy/1

Description

```
threaded_engine_destroy(Engine)
```

Stops and destroys an engine.

📘 Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *engines* flag is set to supported when that is the case.

Modes and number of proofs

```
threaded_engine_destroy(@nonvar) - one
```

Errors

Engine is a variable:

```
instantiation_error
```

Engine is neither a variable nor the name of an existing engine:

```
existence_error(engine, Engine)
```

Examples

```
% stop the worker_1 engine:
| ?- threaded_engine_destroy(worker_1).

% stop all engines:
| ?- forall(
    threaded_engine(Engine),
    threaded_engine_destroy(Engine)
).
```

➡ See also

[threaded_engine_create/3](#), [threaded_engine_self/1](#), [threaded_engine/1](#)

built-in predicate

`threaded_engine/1`

Description

```
threaded_engine(Engine)
```

Enumerates, by backtracking, all existing engines. Engine names shall be regarded as opaque terms; users shall not rely on their type.

i Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *engines* flag is set to supported when that is the case.

Modes and number of proofs

```
threaded_engine(?nonvar) - zero_or_more
```

Errors

(none)

Examples

```
% check that the worker_1 engine exists:  
| ?- threaded_engine(worker_1).  
  
% write the names of all existing engines:  
| ?- forall(  
    threaded_engine(Engine),  
    (writeq(Engine), nl)  
    ).
```

See also

[threaded_engine_create/3](#), [threaded_engine_self/1](#), [threaded_engine_destroy/1](#)

built-in predicate

`threaded_engine_self/1`

Description

```
threaded_engine_self(Engine)
```

Queries the name of engine calling the predicate. Fails if not called from within an engine or if the argument doesn't unify with the engine name.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *engines* flag is set to supported when that is the case.

Modes and number of proofs

```
threaded_engine_self(?nonvar) - zero_or_one
```

Errors

(none)

Examples

```
% find the name of the engine making the query:
..., threaded_engine_self(Engine), ...

% check if the the engine making the query is worker_1:
..., threaded_engine_self(worker_1), ...
```

➔ See also

[threaded_engine_create/3](#), [threaded_engine_destroy/1](#), [threaded_engine/1](#)

built-in predicate

`threaded_engine_next/2`

Description

```
threaded_engine_next(Engine, Answer)
```

Retrieves an answer from an engine and signals it to start computing the next answer. This predicate blocks until an answer becomes available. The predicate fails when there are no more solutions to the engine goal. If the engine goal throws an exception, calling this predicate will re-throw the exception and subsequent calls will fail.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *engines* flag is set to supported when that is the case.

Modes and number of proofs

```
threaded_engine_next(@nonvar, ?term) - zero_or_one
```

Errors

Engine is a variable:

instantiation_error

Engine is neither a variable nor the name of an existing engine:

existence_error(engine, Engine)

Examples

```
% get the next answer from the worker_1 engine:  
| ?- threaded_engine_next(worker_1, Answer).
```

See also

[threaded_engine_create/3](#), [threaded_engine_next_reified/2](#), [threaded_engine_yield/1](#)

built-in predicate

`threaded_engine_next_reified/2`

Description

```
threaded_engine_next_reified(Engine, Answer)
```

Retrieves an answer from an engine and signals it to start computing the next answer. This predicate always succeeds and blocks until an answer becomes available. Answers are returned using the terms `the(Answer)`, `no`, and `exception(Error)`.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *engines* flag is set to supported when that is the case.

Modes and number of proofs

```
threaded_engine_next_reified(@nonvar, ?nonvar) - one
```

Errors

Engine is a variable:

```
instantiation_error
```

Engine is neither a variable nor the name of an existing engine:

```
existence_error(engine, Engine)
```

Examples

```
% get the next reified answer from the worker_1 engine:
| ?- threaded_engine_next_reified(worker_1, Answer).
```

➡ See also

```
threaded_engine_create/3, threaded_engine_next/2, threaded_engine_yield/1
```

built-in predicate

`threaded_engine_yield/1`

Description

```
threaded_engine_yield(Answer)
```

Returns an answer independent of the solutions of the engine goal. Fails if not called from within an engine. This predicate is usually used when the engine goal is a call to a recursive predicate processing terms from the engine term queue.

This predicate blocks until the returned answer is consumed.

Note that this predicate should not be called as the last element of a conjunction resulting in an engine goal solution as, in this case, an answer will always be returned. For example, instead of `(threaded_engine_yield(ready); member(X,[1,2,3]))` use `(X=ready; member(X,[1,2,3]))`.

i Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *engines* flag is set to supported when that is the case.

Modes and number of proofs

```
threaded_engine_yield(@term) - zero_or_one
```

Errors

(none)

Examples

```
% returns the atom "ready" as an engine answer:  
..., threaded_engine_yield(ready), ...
```

See also

[threaded_engine_create/3](#), [threaded_engine_next/2](#), [threaded_engine_next_reified/2](#)

built-in predicate

`threaded_engine_post/2`

Description

```
threaded_engine_post(Engine, Term)
```

Posts a term to the engine term queue.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *engines* flag is set to supported when that is the case.

Modes and number of proofs

```
threaded_engine_post(@nonvar, @term) - one
```

Errors

Engine is a variable:

`instantiation_error`

Engine is neither a variable nor the name of an existing engine:

`existence_error(engine, Engine)`

Examples

```
% post the atom "ready" to the worker_1 engine queue:
| ?- threaded_engine_post(worker_1, ready).
```

➡ See also

[*threaded_engine_fetch/1*](#)

built-in predicate

`threaded_engine_fetch/1`

Description

`threaded_engine_fetch(Term)`

Fetches a term from the engine term queue. Blocks until a term is available. Fails if not called from within an engine.

Note

This predicate requires a *backend Prolog compiler* providing compatible multi-threading primitives. The value of the read-only *engines* flag is set to supported when that is the case.

Modes and number of proofs

`threaded_engine_fetch(?term)` - zero_or_one

Errors

(none)

Examples

```
% fetch a term from the engine term queue:  
..., threaded_engine_fetch(Term), ...
```

See also

[threaded_engine_post/2](#)

2.4.9 Compiling and loading source files

built-in predicate

`logtalk_compile/1`

Description

```
logtalk_compile(File)  
logtalk_compile(Files)
```

Compiles to disk a *source file* or a list of source files using the default compiler flag values. The Logtalk source file name extension (by default, `.lgt`) can be omitted. Source file paths can be absolute, relative to the current directory, or use *library notation*. This predicate can also be used to compile Prolog source files as Logtalk source code. When no recognized Logtalk or Prolog extension is specified, the compiler tries first to append a Logtalk source file extension and then a Prolog source file extension. If that fails, the compiler tries to use the file name as-is. The recognized Logtalk and Prolog file extensions are defined in the *backend adapter files*.

Note

This predicate does not load into memory the compiled source file. If you want to both compile and load a source file, use instead the [logtalk_load/1](#) built-in predicate.

When this predicate is called from the top-level interpreter, relative source file paths are resolved using the current working directory. When the calls are made from a source file, relative source file paths are resolved using the source file directory.

Note that only the errors related to problems in the predicate argument are listed below. This predicate fails on the first error found during compilation of a source file. In this case, no file with the compiled code is written to disk.

Modes and number of proofs

```
logtalk_compile(@source_file_name) - zero_or_one
logtalk_compile(@list(source_file_name)) - zero_or_one
```

Errors

File is a variable:

```
instantiation_error
```

Files is a variable or a list with an element which is a variable:

```
instantiation_error
```

File, or an element File of the Files list, is neither a variable nor a source file name:

```
type_error(source_file_name, File)
```

File, or an element File of the Files list, uses library notation but the library does not exist:

```
existence_error(library, Library)
```

File or an element File of the Files list does not exist:

```
existence_error(file, File)
```

Examples

```
% compile to disk the "set" source file in the
% current directory:
| ?- logtalk_compile(set).

% compile to disk the "tree" source file in the
% "types" library directory:
| ?- logtalk_load(types(tree)).

% compile to disk the "listp" and "list" source
% files in the current directory:
| ?- logtalk_compile([listp, list]).
```

See also

logtalk_compile/2, logtalk_load/1, logtalk_load/2, logtalk_make/0, logtalk_make/1, logtalk_library_path/2

built-in predicate

logtalk_compile/2

Description

```
logtalk_compile(File, Flags)
logtalk_compile(Files, Flags)
```

Compiles to disk a *source file* or a list of source files using a list of compiler flags. The Logtalk source file name extension (by default, `.lgt`) can be omitted. Source file paths can be absolute, relative to the current directory, or use *library notation*. This predicate can also be used to compile Prolog source files as Logtalk source code. When no recognized Logtalk or Prolog extension is specified, the compiler tries first to append a Logtalk source file extension and then a Prolog source file extension. If that fails, the compiler tries to use the file name as-is. Compiler flags are represented as *flag(value)*. For a description of the available compiler flags, please see the *Compiler flags* section in the User Manual. The recognized Logtalk and Prolog file extensions are defined in the *backend adapter files*.

Note

This predicate does not load into memory the compiled source file. If you want to both compile and load a source file, use instead the *logtalk_load/2* built-in predicate.

When this predicate is called from the top-level interpreter, relative source file paths are resolved using the current working directory. When the calls are made from a source file, relative source file paths are resolved by default using the source file directory (unless a *relative_to* flag is passed).

Note that only the errors related to problems in the predicate argument are listed below. This predicate fails on the first error found during compilation of a source file. In this case, no file with the compiled code is written to disk.

Warning

The compiler flags specified in the second argument only apply to the files listed in the first argument. Notably, if you are compiling a *loader file*, the flags only apply to the loader file itself.

Modes and number of proofs

```
logtalk_compile(@source_file_name, @list(compiler_flag)) - zero_or_one
logtalk_compile(@list(source_file_name), @list(compiler_flag)) - zero_or_one
```

Errors

File is a variable:

instantiation_error

Files is a variable or a list with an element which is a variable:

instantiation_error

File, or an element File of the Files list, is neither a variable nor a source file name:

type_error(source_file_name, File)

File, or an element File of the Files list, uses library notation but the library does not exist:

```
existence_error(library, Library)
```

File or an element File of the Files list, does not exist:

```
existence_error(file, File)
```

Flags is a variable or a list with an element which is a variable:

```
instantiation_error
```

Flags is neither a variable nor a proper list:

```
type_error(list, Flags)
```

An element Flag of the Flags list is not a valid compiler flag:

```
type_error(compiler_flag, Flag)
```

An element Flag of the Flags list defines a value for a read-only compiler flag:

```
permission_error(modify, flag, Flag)
```

An element Flag of the Flags list defines an invalid value for a flag:

```
domain_error(flag_value, Flag+Value)
```

Examples

```
% compile to disk the "list" source file in the
% current directory using default compiler flags:
| ?- logtalk_compile(list, []).

% compile to disk the "tree" source file in the "types"
% library directory with the source_data flag turned on:
| ?- logtalk_compile(types(tree), [source_data(on)]).

% compile to disk the "file_system" source file in the
% current directory with portability warnings suppressed:
| ?- logtalk_compile(file_system, [portability(silent)]).
```

➡ See also

[logtalk_compile/1](#), [logtalk_load/1](#), [logtalk_load/2](#), [logtalk_make/0](#), [logtalk_make/1](#),
[logtalk_library_path/2](#)

built-in predicate

logtalk_load/1

Description

```
logtalk_load(File)
logtalk_load(Files)
```

Compiles to disk and then loads to memory a *source file* or a list of source files using the default compiler flag values. The Logtalk source file name extension (by default, .lgt) can be omitted. Source file paths can be absolute, relative to the current directory, or use *library notation*. This predicate can also be used to compile Prolog source files as Logtalk source code. When no recognized Logtalk or Prolog extension is specified, the

compiler tries first to append a Logtalk source file extension and then a Prolog source file extension. If that fails, the compiler tries to use the file name as-is. The recognized Logtalk and Prolog file extensions are defined in the *backend adapter files*.

When this predicate is called from the top-level interpreter, relative source file paths are resolved using the current working directory. When the calls are made from a source file, relative source file paths are resolved using the source file directory.

Note that only the errors related to problems in the predicate argument are listed below. This predicate fails on the first error found during compilation of a source file. In this case, the source file contents is not loaded.

Depending on the *backend Prolog compiler*, the shortcuts {File} or {File1, File2, ...} may be used as an alternative. Check the adapter files for the availability of these shortcuts as they are not part of the language (and thus should only be used at the top-level interpreter).

Modes and number of proofs

```
logtalk_load(@source_file_name) - zero_or_one
logtalk_load(@list(source_file_name)) - zero_or_one
```

Errors

File is a variable:

instantiation_error

Files is a variable or a list with an element which is a variable:

instantiation_error

File, or an element File of the Files list, is neither a variable nor a source file name:

type_error(source_file_name, File)

File, or an element File of the Files list, uses library notation but the library does not exist:

existence_error(library, Library)

File or an element File of the Files list, does not exist:

existence_error(file, File)

Examples

```
% compile and load the "set" source file in the
% current directory:
| ?- logtalk_load(set).

% compile and load the "tree" source file in the
% "types" library directory:
| ?- logtalk_load(types(tree)).

% compile and load the "listp" and "list" source
% files in the current directory:
| ?- logtalk_load([listp, list]).
```

 See also

[logtalk_compile/1](#), [logtalk_compile/2](#), [logtalk_load/2](#), [logtalk_make/0](#), [logtalk_make/1](#),
[logtalk_library_path/2](#)

built-in predicate**logtalk_load/2****Description**

```
logtalk_load(File, Flags)
logtalk_load(Files, Flags)
```

Compiles to disk and then loads to memory a *source file* or a list of source files using a list of compiler flags. The Logtalk source file name extension (by default, .lgt) can be omitted. Source file paths can be absolute, relative to the current directory, or use *library notation*. Compiler flags are represented as *flag(value)*. This predicate can also be used to compile Prolog source files as Logtalk source code. When no recognized Logtalk or Prolog extension is specified, the compiler tries first to append a Logtalk source file extension and then a Prolog source file extension. If that fails, the compiler tries to use the file name as-is. For a description of the available compiler flags, please see the *Compiler flags* section in the User Manual. The recognized Logtalk and Prolog file extensions are defined in the *backend adapter files*. The recognized Logtalk and Prolog file extensions are defined in the *backend adapter files*.

When this predicate is called from the top-level interpreter, relative source file paths are resolved using the current working directory. When the calls are made from a source file, relative source file paths are resolved by default using the source file directory (unless a *relative_to* flag is passed).

Note that only the errors related to problems in the predicate argument are listed below. This predicate fails on the first error found during compilation of a source file. In this case, the source file contents is not loaded.

 **Warning**

The compiler flags specified in the second argument only apply to the files listed in the first argument and not to any files that those files may load or compile. Notably, if you are loading a *loader file*, the flags only apply to the loader file itself and not to the files loaded by it.

Modes and number of proofs

```
logtalk_load(@source_file_name, @list(compiler_flag)) - zero_or_one
logtalk_load(@list(source_file_name), @list(compiler_flag)) - zero_or_one
```

Errors

File is a variable:

`instantiation_error`

Files is a variable or a list with an element which is a variable:

`instantiation_error`

File, or an element File of the Files list, is neither a variable nor a source file name:

`type_error(source_file_name, File)`

File, or an element File of the Files list, uses library notation but the library does not exist:

`existence_error(library, Library)`

File or an element File of the Files list, does not exist:

`existence_error(file, File)`

Flags is a variable or a list with an element which is a variable:

`instantiation_error`

Flags is neither a variable nor a proper list:

`type_error(list, Flags)`

An element Flag of the Flags list is not a valid compiler flag:

`type_error(compiler_flag, Flag)`

An element Flag of the Flags list defines a value for a read-only compiler flag:

`permission_error(modify, flag, Flag)`

An element Flag of the Flags list defines an invalid value for a flag:

`domain_error(flag_value, Flag+Value)`

Examples

```
% compile and load the "list" source file in the
% current directory using default compiler flags:
| ?- logtalk_load(list, []).

% compile and load the "tree" source file in the "types"
% library directory with the source_data flag turned on:
| ?- logtalk_load(types(tree)).

% compile and load the "file_system" source file in the
% current directory with portability warnings suppressed:
| ?- logtalk_load(file_system, [portability(silent)]).
```

See also

logtalk_compile/1, logtalk_compile/2, logtalk_load/1, logtalk_make/0, logtalk_make/1, logtalk_library_path/2

built-in predicate

`logtalk_make/0`

Description

`logtalk_make`

Reloads all Logtalk source files that have been modified since the time they were last loaded. Only source files loaded using the `logtalk_load/1` and `logtalk_load/2` predicates are reloaded. Non-modified files will also be reloaded when a previous attempt to load them failed or when there is a change to the compilation mode (i.e., when the files were loaded without explicit `debug` or `optimize` flags and the default values of these flags changed after loading; no check is made, however, for other implicit compiler flags that may have changed since loading). When an included file is modified, this predicate reloads its main file (i.e., the file that contains the `include/1` directive).

Depending on the *backend Prolog compiler*, the shortcut `{*}` may be used as an alternative. Check the *adapter files* for the availability of the shortcut as it is not part of the language.

Warning

Only use the `{*}` shortcut at the top-level interpreter and never in source files.

This predicate can be extended by the user by defining clauses for the `logtalk_make_target_action/1` multifile and dynamic hook predicate using the argument `all`. The additional user defined actions are run after the default one.

Modes and number of proofs

`logtalk_make` - one

Errors

(none)

Examples

```
% reload all files modified since last loaded:
| ?- logtalk_make.
```

See also

`logtalk_compile/1`, `logtalk_compile/2`, `logtalk_load/1`, `logtalk_load/2`, `logtalk_make/1`,
`logtalk_make_target_action/1`

built-in predicate

logtalk_make/1

Description

`logtalk_make(Target)`

Runs a make target. Prints a warning message and fails when the target is not valid.

Allows reloading all Logtalk source files that have been modified since last loaded when called with the target `all`, deleting all intermediate files generated by the compilation of Logtalk source files when called with the target `clean`, checking for code issues when called with the target `check`, listing of circular dependencies between pairs or trios of objects when called with the target `circular`, generating documentation when called with the target `documentation`, and deleting the *dynamic binding* caches with the target `caches`.

There are also four variants of the `all` target: `debug`, `normal`, `optimal`, and `force`. The first three targets change the compilation mode (by changing the default value of the *debug* and *optimize* flags) and reload all affected files (i.e., all files loaded without an explicit `debug/1` or `optimize/1` compiler option). The `force` target forces reloading of all files that were loaded without an explicit `reload/1` compiler option. This target is typically used after changing some compiler linter flag (e.g., the portability flag).

When using the `all` target, only source files loaded using the *logtalk_load/1* and *logtalk_load/2* predicates are reloaded. Non-modified files will also be reloaded when a previous attempt to load them failed or when there is a change to the compilation mode (i.e., when the files were loaded without explicit *debug* or *optimize* flags and the default values of these flags changed after loading; no check is made, however, for other implicit compiler flags that may have changed since loading). When an included file is modified, this target reloads its main file (i.e., the file that contains the *include/1* directive).

When using the `check` or `circular` targets, be sure to compile your source files with the *source_data* flag turned on for complete and detailed reports.

The `check` target scans for missing entities (objects, protocols, categories, and modules), missing entity predicates, and duplicated library aliases. Predicates for messages sent to objects that implement the *forwarding* built-in protocol are not reported. While this usually avoids only false positives, it may also result in failure to report true missing predicates in some cases.

When using the `circular` target, be prepared for a lengthy computation time for applications with a large combined number of objects and message calls. Only mutual and triangular dependencies are checked due to the computational cost. Circular dependencies occur when an object sends a message to a second object that, in turn, sends a message to the first object. These circular dependencies are often a consequence of lack of separation of concerns. But, when they cannot be fixed, the only practical consequence is a small performance cost as some of the messages would be forced to use dynamic binding.

The `documentation` target requires the *doclet* tool and a single *doclet object* to be loaded. See the *doclet* tool documentation for more details.

Depending on the *backend Prolog compiler*, the following top-level shortcuts are usually defined:

- `{*}` - `logtalk_make(all)`
- `{!}` - `logtalk_make(clean)`
- `{?}` - `logtalk_make(check)`
- `{@}` - `logtalk_make(circular)`
- `{#}` - `logtalk_make(documentation)`
- `{$}` - `logtalk_make(caches)`
- `{+d}` - `logtalk_make(debug)`

- {+n} - `logtalk_make(normal)`
- {+o} - `logtalk_make(optimal)`
- {+f} - `logtalk_make(force)`

Check the *adapter files* for the availability of these shortcuts as they are not part of the language.

Warning

Only use the shortcuts at the top-level interpreter and never in source files.

The target actions can be extended by defining clauses for the multifile and dynamic hook predicate `logtalk_make_target_action(Target)` where `Target` is one of the targets listed above. The additional user-defined actions are run after the default ones.

Modes and number of proofs

```
logtalk_make(+atom) - zero_or_one
```

Errors

(none)

Examples

```
% reload loaded source files in debug mode:
| ?- logtalk_make(debug).

% check for code issues in the loaded source files:
| ?- logtalk_make(check).

% delete all intermediate files generated by
% the compilation of Logtalk source files:
| ?- logtalk_make(clean).
```

See also

`logtalk_compile/1`, `logtalk_compile/2`, `logtalk_load/1`, `logtalk_load/2`, `logtalk_make/0`,
`logtalk_make_target_action/1`

built-in predicate

`logtalk_make_target_action/1`

Description

`logtalk_make_target_action(Target)`

Multifile and dynamic hook predicate that allows defining user actions for the *logtalk_make/1* targets. The user defined actions are run after the default ones using a failure driven loop. This loop does not catch any exceptions thrown when calling the user-defined actions.

Modes and number of proofs

`logtalk_make_target_action(+atom) - zero_or_more`

Errors

(none)

Examples

```
% integrate the dead_code_scanner tool with logtalk_make/1

:- multifile(logtalk_make_target_action/1).
:- dynamic(logtalk_make_target_action/1).

logtalk_make_target_action(check) :-
    dead_code_scanner::all.
```

See also

logtalk_make/1, *logtalk_make/0*

built-in predicate

`logtalk_library_path/2`

Description

`logtalk_library_path(Library, Path)`

Dynamic and multifile user-defined predicate, allowing the declaration of aliases to *library* paths. Library aliases may also be used in the second argument (using the notation *alias(path)*). Paths must always end with the path directory separator character ('/').

Warning

Library aliases should be unique. The `logtalk make/1` built-in predicate can be used to detect and report duplicated library aliases using the check target.

Clauses for this predicate should preferably be facts. Defining rules to dynamically compute at runtime both library alias names and their paths, although sometimes handy, can inadvertently result in endless loops when those rules also attempt to expand library paths. When asserting facts for this predicate, use preferably `asserta/1` instead of `assertz/1` to help ensure the facts will be used before any rules.

Relative paths (e.g., `'../'` or `'./'`) should only be used within the `alias(path)` notation so that library paths can always be expanded to absolute paths independently of the (usually unpredictable) current directory at the time the `logtalk_library_path/2` predicate is called.

When working with a relocatable application, the actual application installation directory can be retrieved by calling the `logtalk_load_context/2` predicate with the `directory` key and using the returned value to define the `logtalk_library_path/2` predicate. On a *settings file* or a *loader file*, simply use an *initialization/1* directive to wrap the call to the `logtalk_load_context/2` predicate and the assert of the `logtalk_library_path/2` fact.

This predicate may be used to override the default *scratch directory* by defining the library alias `scratch_directory` in a backend Prolog initialization file (assumed to be loaded prior to Logtalk loading). This allows e.g. Logtalk to be installed in a read-only directory by setting this alias to the operating-system directory for temporary files. It also allows several Logtalk instances to run concurrently without conflict by using a unique scratch directory per instance (e.g., using a process ID or a UUID generator).

This predicate may be used to override the default location used by the *packs* tool to store registries and packs by defining the `logtalk_packs` library alias in settings file or in a backend Prolog initialization file (assumed to be loaded prior to Logtalk loading).

The *logtalk* built-in object provides an `expand_library_path/2` predicate that can be used to expand library aliases and files expressed using library notation.

Modes and number of proofs

```
logtalk_library_path(?atom, -atom) - zero_or_more
logtalk_library_path(?atom, -compound) - zero_or_more
```

Errors

(none)

Examples

```
:- initialization((
    logtalk_load_context(directory, Directory),
    asserta(logtalk_library_path(my_application_root, Directory))
)).
```

```
| ?- logtalk_library_path(viewpoints, Path).
```

```
Path = examples('viewpoints/')
```

```
yes
```

```
| ?- logtalk_library_path(Library, Path).
```

```
Library = home,
```

```
Path = '$HOME/' ;
```

```
Library = logtalk_home,
```

```
Path = '$LOGTALKHOME/' ;
```

```
Library = logtalk_user
```

```
Path = '$LOGTALKUSER/' ;
```

```
Library = examples
```

```
Path = logtalk_user('examples/') ;
```

```
Library = library
```

```
Path = logtalk_user('library/') ;
```

```
Library = viewpoints
```

```
Path = examples('viewpoints/')
```

```
yes
```

```
| ?- logtalk::expand_library_path(viewpoints, Path).
```

```
Path = '/Users/pmoura/logtalk/examples/viewpoints/'.
```

```
yes
```

```
| ?- logtalk::expand_library_path(viewpoints('loader.lgt'), Path).
```

```
Path = '/Users/pmoura/logtalk/examples/viewpoints/loader.lgt'.
```

```
yes
```

See also

logtalk_compile/1, logtalk_compile/2, logtalk_load/1, logtalk_load/2

built-in predicate

logtalk_load_context/2

Description

```
logtalk_load_context(Key, Value)
```

Provides access to the Logtalk compilation/loading context. The following keys are currently supported:

- `entity_identifier` - identifier of the entity being compiled if any
- `entity_prefix` - internal prefix for the entity compiled code
- `entity_type` - returns the value module when compiling a module as an object
- `entity_relation` - returns the entity relations as declared in the entity opening directive
- `source` - full path of the source file being compiled
- `file` - the actual file being compiled, different from `source` only when processing an `include/1` directive
- `basename` - source file basename
- `directory` - source file directory
- `stream` - input stream being used to read source file terms
- `target` - the full path of the intermediate Prolog file
- `flags` - the list of the explicit flags used for the compilation of the source file
- `term` - the source file term being compiled
- `term_position` - the position of the term being compiled (StartLine-EndLine)
- `term_position(Term)` - the position of the term being compiled (StartLine-EndLine)
- `variables` - the variables of the term being compiled ([Variable1, ...])
- `variable_names` - the variable names of the term being compiled ([Name1=Variable1, ...])
- `variable_names(Term)` - the variable names of the term being compiled ([Name1=Variable1, ...])
- `singletons` - the singleton variables of the term being compiled ([Name1=Variable1, ...])
- `singletons(Term)` - the singleton variables of the term being compiled ([Name1=Variable1, ...])
- `parameter_variables` - list of parameter variable names and positions ([Name1-Position1, ...])

For the `entity_relation` key, the possible values are:

- `extends_protocol(Protocol, ParentProtocol, Scope)`
- `implements_protocol(ObjectOrCategory, Protocol, Scope)`
- `extends_category(Category, ParentCategory, Scope)`
- `imports_category(Object, Category, Scope)`
- `extends_object(Prototype, Parent, Scope)`
- `instantiates_class(Instance, Class, Scope)`
- `specializes_class(Class, Superclass, Scope)`
- `complements_object(Category, Object)`

Calling this predicate with the `parameter_variables` key only succeeds when compiling a parametric entity containing parameter variables.

This predicate is usually called by the `term_expansion/2` and `goal_expansion/2` methods. It can also be called directly from `initialization/1` directives in a source file. Note that the entity keys are only available when compiling an entity term or from an object initialization/1 directive.

Warning

The `term_position` key is only supported in *backend Prolog compilers* that provide access to the start and end lines of a read term. When such support is not available, the value `-1` is returned for both the start and the end lines.

Variables in the values of the `term`, `variables`, `variable_names`, and `singletons` keys are not shared with, respectively, the `term` and `goal` arguments of the `term_expansion/2` and `goal_expansion/2` methods. Use instead the `variable_names(Term)` and `singletons(Term)` keys when possible.

Modes and number of proofs

```
logtalk_load_context(?callable, -nonvar) - zero_or_more
```

Errors

Key is neither a variable nor a callable term:

```
type_error(callable, Key)
```

Key is a callable term but not a valid key:

```
domain_error(logtalk_load_context_key, Key)
```

Examples

```
% expand source file terms only if they are entity terms
term_expansion(Term, ExpandedTerms) :-
    logtalk_load_context(entity_identifier, _),
    ....

% expand source file term while accessing its variable names
term_expansion(Term, ExpandedTerms) :-
    logtalk_load_context(variable_names(Term), VariableNames),
    ....

% define a library alias based on the source directory
:- initialization((
    logtalk_load_context(directory, Directory),
    assertz(logtalk_library_path(my_app, Directory))
)).
```

 See also*term_expansion/2, goal_expansion/2, initialization/1*

2.4.10 Flags

built-in predicate

`current_logtalk_flag/2`

Description

`current_logtalk_flag(Flag, Value)`

Enumerates, by backtracking, the current Logtalk flag values. For a description of the predefined compiler flags, please see the *Compiler flags* section in the User Manual.

Modes and number of proofs

`current_logtalk_flag(?atom, ?atom) - zero_or_more`

Errors

Flag is neither a variable nor an atom:

`type_error(atom, Flag)`

Flag is an atom but an invalid flag:

`domain_error(flag, Value)`

Examples

```
% get the current value of the source_data flag:
| ?- current_logtalk_flag(source_data, Value).
```

 See also*create_logtalk_flag/3, set_logtalk_flag/2*

built-in predicate

`set_logtalk_flag/2`

Description

```
set_logtalk_flag(Flag, Value)
```

Sets global, default, flag values. For local flag scope, use the corresponding *set_logtalk_flag/2* directive. To set a global flag value when compiling and loading a source file, wrap the calls to this built-in predicate with an *initialization/1* directive. For a description of the predefined compiler flags, please see the *Compiler flags* section in the User Manual.

Modes and number of proofs

```
set_logtalk_flag(+atom, +nonvar) - one
```

Errors

Flag is a variable:

instantiation_error

Value is a variable:

instantiation_error

Flag is neither a variable nor an atom:

type_error(atom, Flag)

Flag is an atom but an invalid flag:

domain_error(flag, Flag)

Value is not a valid value for flag Flag:

domain_error(flag_value, Flag + Value)

Flag is a read-only flag:

permission_error(modify, flag, Flag)

Examples

```
% turn off globally and by default the compiler
% unknown entities warnings:
| ?- set_logtalk_flag(unknown_entities, silent).
```

See also

create_logtalk_flag/3, *current_logtalk_flag/2*

built-in predicate

create_logtalk_flag/3

Description

```
create_logtalk_flag(Flag, Value, Options)
```

Creates a new Logtalk flag and sets its default value. User-defined flags can be queried and set in the same way as predefined flags by using, respectively, the [current_logtalk_flag/2](#) and [set_logtalk_flag/2](#) built-in predicates. For a description of the predefined compiler flags, please see the [Compiler flags](#) section in the User Manual.

This predicate is based on the specification of the SWI-Prolog `create_prolog_flag/3` built-in predicate and supports the same options: `access(Access)`, where `Access` can be either `read_write` (the default) or `read_only`; `keep(Keep)`, where `Keep` can be either `false` (the default) or `true`, for deciding if an existing definition of the flag should be kept or replaced by the new one; and `type(Type)` for specifying the type of the flag, which can be `boolean`, `atom`, `integer`, `float`, or `term` (which only restricts the flag value to ground terms). When the `type/1` option is not specified, the type of the flag is inferred from its initial value.

Modes and number of proofs

```
create_logtalk_flag(+atom, +ground, +list(ground)) - one
```

Errors

Flag is a variable:

```
instantiation_error
```

Value is not a ground term:

```
instantiation_error
```

Options is not a ground term:

```
instantiation_error
```

Flag is neither a variable nor an atom:

```
type_error(atom, Flag)
```

Options is neither a variable nor a list:

```
type_error(atom, Flag)
```

Value is not a valid value for flag Flag:

```
domain_error(flag_value, Flag + Value)
```

Flag is a system-defined flag:

```
permission_error(modify, flag, Flag)
```

An element Option of the list Options is not a valid option

```
domain_error(flag_option, Option)
```

The list Options contains a type(Type) option and Value is not of type Type

```
type_error(Type, Value)
```

Examples

```
% create a new boolean flag with default value set to false:
| ?- create_logtalk_flag(pretty_print_blobs, false, []).
```

➡ See also

current_logtalk_flag/2, set_logtalk_flag/2

2.4.11 Linter

built-in predicate

logtalk_linter_hook/7

Description

```
logtalk_linter_hook(Goal, Flag, File, Lines, Type, Entity, Warning)
```

Multifile user-defined predicate, supporting the definition of custom linter warnings. Experimental. The Goal argument can be a message-sending goal, `Object::Message`, a call to a Prolog built-in predicate, or a call to a module predicate, `Module:Predicate`. The Flag argument must be a supported linter flag. The Warning argument must be a valid core message term. For a given Goal, only the first successful call to this predicate is considered.

Modes and number of proofs

```
logtalk_linter_hook(@callable, +atom, +atom, +pair(integer), +atom, @object_identifier, --
↳callable) - zero_or_one
```

Errors

(none)

Examples

```
:- multifile(user::logtalk_linter_hook/7).
% warn about using list::append/3 to construct a list from an head and a tail
user::logtalk_linter_hook(
    list::append(L1,L2,L), suspicious_calls,
    File, Lines, Type, Entity,
    suspicious_call(File, Lines, Type, Entity, list::append(L1,L2,L), [L=[Head|L2]])
) :-
    nonvar(L1),
    L1 = [Head].
```

2.5 Built-in methods

2.5.1 Logic and control

built-in method

`!/0`

Description

`!`

Always succeeds with the side-effect of discarding choice-points. See also the ISO Prolog standard definition. This built-in method is declared as a public method and can be used as a message to an object.

Modes and number of proofs

`! - one`

Errors

(none)

Examples

(none)

➡ See also

true/0, fail/0, false/0, repeat/0

built-in method

`true/0`

Description

`true`

Always succeeds. See also the ISO Prolog standard definition. This built-in method is declared as a public method and can be used as a message to an object.

Modes and number of proofs

true – one

Errors

(none)

Examples

(none)

See also

!/0, fail/0, false/0, repeat/0

built-in method

fail/0

Description

fail

Always fails. See also the ISO Prolog standard definition. This built-in method is declared as a public method and can be used as a message to an object.

Modes and number of proofs

fail – one

Errors

(none)

Examples

(none)

See also

!/0, true/0, false/0, repeat/0

built-in method

`false/0`

Description

`false`

Always fails. See also the ISO Prolog standard definition. This built-in method is declared as a public method and can be used as a message to an object.

Modes and number of proofs

`false` - one

Errors

(none)

Examples

(none)

See also

!/0, true/0, fail/0, repeat/0

built-in method

`repeat/0`

Description

`repeat`

Always succeeds when called and when backtracking into its call with an infinite number of choice-points. See also the ISO Prolog standard definition. This built-in method is declared as a public method and can be used as a message to an object.

Modes and number of proofs

repeat - one_or_more

Errors

(none)

Examples

(none)

➡ See also

!/0, true/0, fail/0, false/0

2.5.2 Execution context

built-in method

`context/1`

Description

context(Context)

Returns the execution context for a predicate clause using the term `logtalk(Head,ExecutionContext)` where `Head` is the head of the clause containing the call. This private predicate is mainly used for providing an error context when type-checking predicate arguments. The `ExecutionContext` term should be regarded as an opaque term, which can be decoded using the `logtalk::execution_context/7` predicate. Calls to this predicate are inlined at compilation time.

⚠ Warning

As the execution context term includes the clause head, the head of the clause calling the `context(Context)` method cannot contain the `Context` variable as that would result in the creation of a cyclic term. The compiler detects and reports any offending clauses by throwing a `representation_error(acyclic_term)` error.

Modes and number of proofs

```
context(--callable) - one
```

Errors

(none)

Examples

```
foo(A, N) :-
    % type-check arguments
    context(Context),
    type::check(atom, A, Context),
    type::check(integer, N, Context),
    % arguments are fine; go ahead
    ...
```

See also

[parameter/2](#), [self/1](#), [sender/1](#), [this/1](#)

built-in method

parameter/2

Description

```
parameter(Number, Term)
```

Used in *parametric objects* (and parametric categories), this private method provides runtime access to the parameter values of the entity that contains the predicate clause whose body is being executed by using the argument number in the entity identifier. This predicate is implemented as a unification between its second argument and the corresponding implicit execution-context argument in the predicate clause making the call. This unification occurs at the clause head when the second argument is not instantiated (the most common case). When the second argument is instantiated, the unification must be delayed to runtime and thus occurs at the clause body.

Entity parameters can also be accessed using *parameter variables*, which use the syntax `_VariableName_`. The compiler recognizes occurrences of these variables in directives and clauses. Parameter variables allow us to abstract parameter positions, thus simplifying code maintenance.

Modes and number of proofs

```
parameter(+integer, ?term) - zero_or_one
```

Errors

Number is a variable:

instantiation_error

Number is neither a variable nor an integer value:

type_error(integer, Number)

Number is smaller than one or greater than the parametric entity identifier arity:

domain_error(out_of_range, Number)

Entity identifier is not a compound term:

type_error(compound, Entity)

Examples

```
:- object(box(_Color, _Weight)).

...

% this clause is translated into
% a fact upon compilation
color(Color) :-
    parameter(1, Color).

% upon compilation, the >/2 call will be
% the single goal in the clause body
heavy :-
    parameter(2, Weight),
    Weight > 10.

...
```

The same example using *parameter variables*:

```
:- object(box(_Color_, _Weight_)).

...

color(_Color_).

heavy :-
    _Weight_ > 10.

...
```

 See also

context/1, self/1, sender/1, this/1

built-in method

self/1

Description

self(Self)

Unifies its argument with the object that received the message under processing.

This private method is compiled into a unification between its argument and the corresponding implicit context argument in the predicate clause making the call. This unification occurs at the clause head when the argument is not bound at compile-time (the most common case).

Modes and number of proofs

self(?object_idenfier) - zero_or_one

Errors

(none)

Examples

```
% upon compilation, the write/1 call will be
% the first goal in the clause body
test :-
    self(Self),
    write('executing a method in behalf of '),
    writeq(Self), nl.
```

➡ See also

context/1, parameter/2, sender/1, this/1

built-in method

sender/1

Description

```
sender(Sender)
```

Unifies its argument with the object that sent the message under processing.

This private method is translated into a unification between its argument and the corresponding implicit context argument in the predicate clause making the call. This unification occurs at the clause head when the argument is not bound at compile-time (the most common case).

Modes and number of proofs

```
sender(?object_identifier) - zero_or_one
```

Errors

(none)

Examples

```
% after compilation, the write/1 call will
% be the first goal in the clause body
test :-
    sender(Sender),
    write('executing a method to answer a message sent by '),
    writeq(Sender), nl.
```

See also

context/1, parameter/2, self/1, this/1

built-in method

this/1

Description

```
this(This)
```

Unifies its argument with the identifier of the object calling this method. When this method is called from a category, the argument is unified with the object importing the category on whose behalf the clause containing the call is being used to prove the current goal.

This private method is implemented as a unification between its argument and the corresponding implicit execution-context argument in the predicate clause making the call. This unification occurs at the clause head when the argument is not bound at compile-time (the most common case).

This method is useful for avoiding hard-coding references to an object identifier or for retrieving all object parameters with a single call when using parametric objects.

Modes and number of proofs

```
this(?object_identifier) - zero_or_one
```

Errors

(none)

Examples

```
% after compilation, the write/1 call will
% be the first goal in the clause body
test :-
    this(This),
    write('Using a predicate clause contained in '),
    writeq(This), nl.
```

➡ See also

context/1, parameter/2, self/1, sender/1

2.5.3 Reflection

built-in method

current_op/3

Description

```
current_op(Priority, Specifier, Operator)
```

Enumerates, by backtracking, the visible operators declared for an object. Operators not declared using a scope directive are not enumerated.

Modes and number of proofs

```
current_op(?operator_priority, ?operator_specifier, ?atom) - zero_or_more
```

Errors

Priority is neither a variable nor an integer:

```
type_error(integer, Priority)
```

Priority is an integer but not a valid operator priority:

```
domain_error(operator_priority, Priority)
```

Specifier is neither a variable nor an atom:

```
type_error(atom, Specifier)
```

Specifier is an atom but not a valid operator specifier:

```
domain_error(operator_specifier, Specifier)
```

Operator is neither a variable nor an atom:

```
type_error(atom, Operator)
```

Examples

To enumerate, by backtracking, the local operators or the operators visible in *this*:

```
current_op(Priority, Specifier, Operator)
```

To enumerate, by backtracking, the public and protected operators visible in *self*:

```
::current_op(Priority, Specifier, Operator)
```

To enumerate, by backtracking, the public operators visible for an explicit object:

```
Object::current_op(Priority, Specifier, Operator)
```

See also

[current_predicate/1](#), [predicate_property/2](#), [op/3](#)

built-in method

current_predicate/1

Description

```
current_predicate(Predicate)
```

Enumerates, by backtracking, visible, user-defined, object predicates. Built-in predicates and predicates not declared using a scope directive are not enumerated.

This predicate also succeeds for any predicates listed in [uses/2](#) and [use_module/2](#) directives.

When Predicate is bound at compile-time to a `(:)/2` term, this predicate enumerates module predicates (assuming that the *backend Prolog compiler* supports modules).

Modes and number of proofs

```
current_predicate(?predicate_indicator) - zero_or_more
```

Errors

Predicate is neither a variable nor a valid predicate indicator:

```
type_error(predicate_indicator, Predicate)
```

Predicate is a Name/Arity term but Functor is neither a variable nor an atom:

```
type_error(atom, Name)
```

Predicate is a Name/Arity term but Arity is neither a variable nor an integer:

```
type_error(integer, Arity)
```

Predicate is a Name/Arity term but Arity is a negative integer:

```
domain_error(not_less_than_zero, Arity)
```

Examples

To enumerate, by backtracking, the locally visible user predicates or the user predicates visible in *this*:

```
current_predicate(Predicate)
```

To enumerate, by backtracking, the public and protected user predicates visible in *self*:

```
::current_predicate(Predicate)
```

To enumerate, by backtracking, the public user predicates visible for an explicit object:

```
Object::current_predicate(Predicate)
```

An example of enumerating locally visible object predicates. These include predicates listed using *uses/2* and *use_module/2* directives:

```
:- object(foo).

   :- uses(bar, [
       baz/1, quux/2
     ]).

   :- public(pred/1).
   pred(X) :-
       current_predicate(X).

:- end_object.
```

```
| ?- foo::pred(X).
X = pred/1 ;
X = baz/1 ;
X = quux/2 ;
no
```

 **See also**

[current_op/3](#), [predicate_property/2](#), [uses/2](#), [use_module/2](#)

built-in method

`predicate_property/2`

Description

```
predicate_property(Predicate, Property)
```

Enumerates, by backtracking, the properties of a visible object predicate. Properties for predicates not declared using a scope directive are not enumerated. The valid predicate properties are listed in the language grammar section on *predicate properties* and described in the User Manual section on *predicate properties*.

When Predicate is listed in a [uses/2](#) or [use_module/2](#) directive, properties are enumerated for the referenced object or module predicate.

When Predicate is bound at compile-time to a `(:)/2` term, this predicate enumerates properties for module predicates (assuming that the *backend Prolog compiler* supports modules).

Modes and number of proofs

```
predicate_property(+callable, ?predicate_property) - zero_or_more
```

Errors

Predicate is a variable:

`instantiation_error`

Predicate is neither a variable nor a callable term:

`type_error(callable, Predicate)`

Property is neither a variable nor a valid predicate property:

`domain_error(predicate_property, Property)`

Examples

To enumerate, by backtracking, the properties of a locally visible user predicate or a user predicate visible in *this*:

```
predicate_property(Predicate, Property)
```

To enumerate, by backtracking, the properties of a public or protected predicate visible in *self*:

```
:::predicate_property(Predicate, Property)
```

To enumerate, by backtracking, the properties of a public predicate visible in an explicit object:

```
Object:::predicate_property(Predicate, Property)
```

An example of enumerating properties for locally visible object predicates. These include predicates listed using *uses/2* and *use_module/2* directives:

```
:- object(foo).

    :- uses(bar, [
        baz/1, quux/2
    ]).

    :- public(pred/1).
    pred_prop(Pred, Prop) :-
        predicate_property(Pred, Prop).

:- end_object.
```

```
| ?- foo::pred(baz(_), Prop).
Prop = logtalk ;
Prop = scope(public) ;
Prop = public ;
Prop = declared_in(bar) ;
...
```

➡ See also

current_op/3, *current_predicate/1*, *uses/2*, *use_module/2*

2.5.4 Database

built-in method

abolish/1

Description

abolish(Predicate)

Abolishes a runtime declared object dynamic predicate or an object local dynamic predicate. Only predicates that are dynamically declared at runtime (using a call to the *asserta/1* or *assertz/1* built-in methods) can be abolished.

When the predicate indicator is declared in a *uses/2* or *use_module/2* directive, the predicate is abolished in the referenced object or module. When the backend Prolog compiler supports a module system, the predicate argument can also be module qualified.

Modes and number of proofs

```
abolish(@predicate_indicator) - one
```

Errors

Predicate is a variable:

```
instantiation_error
```

Functor is a variable:

```
instantiation_error
```

Arity is a variable:

```
instantiation_error
```

Predicate is neither a variable nor a valid predicate indicator:

```
type_error(predicate_indicator, Predicate)
```

Functor is neither a variable nor an atom:

```
type_error(atom, Functor)
```

Arity is neither a variable nor an integer:

```
type_error(integer, Arity)
```

Predicate is statically declared:

```
permission_error(modify, predicate_declaration, Name/Arity)
```

Predicate is a private predicate:

```
permission_error(modify, private_predicate, Name/Arity)
```

Predicate is a protected predicate:

```
permission_error(modify, protected_predicate, Name/Arity)
```

Predicate is a static predicate:

```
permission_error(modify, static_predicate, Name/Arity)
```

Predicate is not declared for the object receiving the message:

```
existence_error(predicate_declaration, Name/Arity)
```

Examples

To abolish a local dynamic predicate or a dynamic predicate in *this*:

```
abolish(Predicate)
```

To abolish a public or protected dynamic predicate in *self*:

```
::abolish(Predicate)
```

To abolish a public dynamic predicate in an explicit object:

```
Object::abolish(Predicate)
```

See also

asserta/1, assertz/1, clause/2, retract/1, retractall/1 dynamic/0, dynamic/1, uses/2, use_module/2

built-in method

asserta/1**Description**

```
asserta(Head)
asserta((Head:-Body))
```

Asserts a clause as the first one for an object dynamic predicate.

When the predicate was not previously declared (using a scope directive), a dynamic predicate declaration is added to the object. In this case, the predicate scope depends on how this method is called:

asserta(Clause)

The predicate is dynamically declared as *private*.

::asserta(Clause)

The predicate is dynamically declared as *protected*.

Object::asserta(Clause)

The predicate is dynamically declared as *public*.

Note, however, that dynamically declaring a new predicate requires either a local assert or the *dynamic_declarations* compiler flag set to allow when the object was created or compiled.

When the predicate indicator for Head is declared in a *uses/2* or *use_module/2* directive, the clause is asserted in the referenced object or module. When the backend Prolog compiler supports a module system, the predicate argument can also be module qualified.

This method may be used to assert clauses for predicates that are not declared dynamic for dynamic objects provided that the predicates are declared in *this*. This allows easy initialization of dynamically created objects when writing constructors.

Modes and number of proofs

```
asserta(+clause) - one
```

Errors

Head is a variable:

```
instantiation_error
```

Head is neither a variable nor a callable term:

```
type_error(callable, Head)
```

Body cannot be converted to a goal:

```
type_error(callable, Body)
```

The predicate indicator of Head, Name/Arity, is that of a private predicate:

```
permission_error(modify, private_predicate, Name/Arity)
```

The predicate indicator of Head, Name/Arity, is that of a protected predicate:

```
permission_error(modify, protected_predicate, Name/Arity)
```

The predicate indicator of Head, Name/Arity, is that of a static predicate:

```
permission_error(modify, static_predicate, Name/Arity)
```

The predicate indicator of Head, Name/Arity, does not match a declared predicate and the target object was created or compiled with support for dynamic declaration of predicates turned off:

```
permission_error(create, predicate_declaration, Name/Arity)
```

Examples

To assert a clause as the first one for a local dynamic predicate or a dynamic predicate in *this*:

```
asserta(Clause)
```

To assert a clause as the first one for any public or protected dynamic predicate in *self*:

```
::asserta(Clause)
```

To assert a clause as the first one for any public dynamic predicate in an explicit object:

```
Object::asserta(Clause)
```

An example of asserting clauses in *this* and in *self* from a category:

```
:- category(attributes,
    implements(attributes_protocol)).

:- private(attr_/1).
:- dynamic(attr_/1).

set_in_this(A, X) :-
    asserta(attr_(A, X)).

set_in_self(A, X) :-
    ::asserta(attr_(A, X)).

...
```

An example of asserting clauses into another object with the predicates listed using a *uses/2* directive (similar when using a *use_module/2* directive):


```
:- object(reasoner(_KnowledgeBase_)).

:- uses(_KnowledgeBase_, [
    foo/1, bar/1
]).

baz(X) :-
    % compiled as _KnowledgeBase_::assertz(foo(X))
    asserta(foo(X)).

foobar(Name, Argument) :-
    Clause =.. [Name, Argument],
    % runtime resolved to _KnowledgeBase_::assertz(Clause)
    % when Name is either foo or bar
    asserta(Clause).

...
```

 See also

```
abolish/1, assertz/1, clause/2, retract/1, retractall/1 dynamic/0, dynamic/1, uses/2, use_module/2
```

built-in method

assertz/1

Description

```
assertz(Head)
assertz((Head:-Body))
```

Asserts a clause as the last one for a dynamic predicate.

When the predicate was not previously declared (using a scope directive), a dynamic predicate declaration is added to the object. In this case, the predicate scope depends on how this method is called:

assertz(Clause)

The predicate is dynamically declared as *private*.

::assertz(Clause)

The predicate is dynamically declared as *protected*.

Object::assertz(Clause)

The predicate is dynamically declared as *public*.

Note, however, that dynamically declaring a new predicate requires either a local assert or the *dynamic_declarations* compiler flag set to allow when the object was created or compiled.

When the predicate indicator for Head is declared in a *uses/2* or *use_module/2* directive, the clause is asserted in the referenced object or module. When the backend Prolog compiler supports a module system, the predicate argument can also be module qualified.

This method may be used to assert clauses for predicates that are not declared dynamic for dynamic objects provided that the predicates are declared in *this*. This allows easy initialization of dynamically created objects when writing constructors.

Modes and number of proofs

```
assertz(+clause) - one
```

Errors

Head is a variable:

```
instantiation_error
```

Head is neither a variable nor a callable term:

```
type_error(callable, Head)
```

Body cannot be converted to a goal:

```
type_error(callable, Body)
```

The predicate indicator of Head, Name/Arity, is that of a private predicate:

```
permission_error(modify, private_predicate, Name/Arity)
```

The predicate indicator of Head, Name/Arity, is that of a protected predicate:

```
permission_error(modify, protected_predicate, Name/Arity)
```

The predicate indicator of Head, Name/Arity, is that of a static predicate:

```
permission_error(modify, static_predicate, Name/Arity)
```

The predicate indicator of Head, Name/Arity, does not match a declared predicate and the target object was created/compiled with support for dynamic declaration of predicates turned off:

```
permission_error(create, predicate_declaration, Name/Arity)
```

Examples

To assert a clause as the last one for a local dynamic predicate or a dynamic predicate in *this*:

```
assertz(Clause)
```

To assert a clause as the last one for any public or protected dynamic predicate in *self*:

```
::assertz(Clause)
```

To assert a clause as the last one for any public dynamic predicate in an explicit object:

```
Object::assertz(Clause)
```

An example of asserting clauses in *this* and in *self* from a category:

```
:- category(attributes,
    implements(attributes_protocol)).

:- private(attr_/1).
:- dynamic(attr_/1).

set_in_this(A, X) :-
    assertz(attr_(A, X)).

set_in_self(A, X) :-
    ::assertz(attr_(A, X)).

...
```

An example of asserting clauses into another object with the predicates listed using a *uses/2* directive (similar when using a *use_module/2* directive):

```
:- object(reasoner(_KnowledgeBase_)).

:- uses(_KnowledgeBase_, [
    foo/1, bar/1
]).

baz(X) :-
    % compiled as _KnowledgeBase_::assertz(foo(X))
    assertz(foo(X)).

foobar(Name, Argument) :-
    Clause =.. [Name, Argument],
    % runtime resolved to _KnowledgeBase_::assertz(Clause)
    % when Name is either foo or bar
    assertz(Clause).
```

(continues on next page)

(continued from previous page)

...

 **See also***abolish/1, asserta/1, clause/2, retract/1, retractall/1 dynamic/0, dynamic/1, uses/2, use_module/2***built-in method****clause/2****Description****clause**(Head, Body)

Enumerates, by backtracking, the clauses of a dynamic predicate.

When the predicate indicator for Head is declared in a *uses/2* or *use_module/2* directive, the predicate enumerates the clauses in the referenced object or module. When the backend Prolog compiler supports a module system, the head argument can also be module qualified.

This method may be used to enumerate clauses for predicates that are not declared dynamic for dynamic objects provided that the predicates are declared in *this*.

Modes and number of proofs**clause**(+callable, ?body) - zero_or_more**Errors**

Head is a variable:

instantiation_error

Head is neither a variable nor a callable term:

type_error(callable, Head)

Body is neither a variable nor a callable term:

type_error(callable, Body)

The predicate indicator of Head, Name/Arity, is that of a private predicate:

permission_error(access, private_predicate, Name/Arity)

The predicate indicator of Head, Name/Arity, is that of a protected predicate:

permission_error(access, protected_predicate, Name/Arity)

The predicate indicator of Head, Name/Arity, is that of a static predicate:

permission_error(access, static_predicate, Name/Arity)

Head is not a declared predicate:

existence_error(predicate_declaration, Name/Arity)

Examples

To retrieve a matching clause of a local dynamic predicate or a dynamic predicate in *this*:

```
clause(Head, Body)
```

To retrieve a matching clause of a public or protected dynamic predicate in *self*:

```
:::clause(Head, Body)
```

To retrieve a matching clause of a public dynamic predicate in an explicit object:

```
Object:::clause(Head, Body)
```

See also

[abolish/1](#), [asserta/1](#), [assertz/1](#), [retract/1](#), [retractall/1](#) [dynamic/0](#), [dynamic/1](#), [uses/2](#), [use_module/2](#)

built-in method

retract/1

Description

```
retract(Head)
retract((Head:-Body))
```

Retracts a clause for an object dynamic predicate. On backtracking, the predicate retracts the next matching clause.

When the predicate indicator for Head is declared in a [uses/2](#) or [use_module/2](#) directive, the clause is retracted in the referenced object or module. When the backend Prolog compiler supports a module system, the predicate argument can also be module qualified.

This method may be used to retract clauses for predicates that are not declared dynamic for dynamic objects provided that the predicates are declared in *this*.

Modes and number of proofs

```
retract(+clause) - zero_or_more
```

Errors

Head is a variable:

```
instantiation_error
```

Head is neither a variable nor a callable term:

```
type_error(callable, Head)
```

The predicate indicator of Head, Name/Arity, is that of a private predicate:

```
permission_error(modify, private_predicate, Name/Arity)
```

The predicate indicator of Head, Name/Arity, is that of a protected predicate:

```
permission_error(modify, protected_predicate, Name/Arity)
```

The predicate indicator of `Head`, `Name/Arity`, is that of a static predicate:

```
permission_error(modify, static_predicate, Name/Arity)
```

The predicate indicator of `Head`, `Name/Arity`, is not declared:

```
existence_error(predicate_declaration, Name/Arity)
```

Examples

To retract a matching clause of a dynamic predicate in *this*:

```
retract(Clause)
```

To retract a matching clause of a public or protected dynamic predicate in *self*:

```
::retract(Clause)
```

To retract a matching clause of a public dynamic predicate in an explicit object:

```
Object::retract(Clause)
```

➡ See also

[*abolish/1*](#), [*asserta/1*](#), [*assertz/1*](#), [*clause/2*](#), [*retractall/1*](#), [*dynamic/0*](#), [*dynamic/1*](#), [*uses/2*](#), [*use_module/2*](#)

built-in method

`retractall/1`

Description

`retractall`(*Head*)

Retracts all clauses with a matching head for an object dynamic predicate.

When the predicate indicator for `Head` is declared in a [*uses/2*](#) or [*use_module/2*](#) directive, the clauses are retracted in the referenced object or module. When the backend Prolog compiler supports a module system, the predicate argument can also be module qualified.

This method may be used to retract clauses for predicates that are not declared dynamic for dynamic objects provided that the predicates are declared in [*this*](#).

Modes and number of proofs

`retractall`(@callable) - one

Errors

Head is a variable:

```
instantiation_error
```

Head is neither a variable nor a callable term:

```
type_error(callable, Head)
```

The predicate indicator of Head, Name/Arity, is that of a private predicate:

```
permission_error(modify, private_predicate, Name/Arity)
```

The predicate indicator of Head, Name/Arity, is that of a protected predicate:

```
permission_error(modify, protected_predicate, Name/Arity)
```

The predicate indicator of Head, Name/Arity, is that of a static predicate:

```
permission_error(modify, static_predicate, Name/Arity)
```

The predicate indicator of Head, Name/Arity, is not declared:

```
existence_error(predicate_declaration, Name/Arity)
```

Examples

To retract all clauses with a matching head of a dynamic predicate in *this*:

```
retractall(Head)
```

To retract all clauses with a matching head of a public or protected dynamic predicate in *self*:

```
::retractall(Head)
```

To retract all clauses with a matching head of a public dynamic predicate in an explicit object:

```
Object::retractall(Head)
```

See also

[abolish/1](#), [asserta/1](#), [assertz/1](#), [clause/2](#), [retract/1](#), [dynamic/0](#), [dynamic/1](#), [uses/2](#), [use_module/2](#)

2.5.5 Meta-calls

built-in method

call/1-N

Description

```
call(Goal)
call(Closure, Arg1, ...)
```

Calls a goal constructed by appending additional arguments to a *closure*. The upper limit for N depends on the upper limit for the arity of a compound term of the *backend Prolog compiler*. This built-in meta-predicate is declared as a private method and thus cannot be used as a message to an object. The Closure argument can also be a lambda expression or a Logtalk control construct. When using a backend Prolog compiler supporting a module system, calls in the format `call(Module:Closure, Arg1, ...)` may also be used.

This meta-predicate is opaque to cuts in its arguments.

Meta-predicate template

```
call(0)
call(1, *)
call(2, *, *)
...
```

Modes and number of proofs

```
call(+callable) - zero_or_more
call(+callable, ?term) - zero_or_more
call(+callable, ?term, ?term) - zero_or_more
...
```

Errors

Goal is a variable:

```
instantiation_error
```

Goal is neither a variable nor a callable term:

```
type_error(callable, Goal)
```

Closure is a variable:

```
instantiation_error
```

Closure is neither a variable nor a callable term:

```
type_error(callable, Closure)
```

Examples

Call a goal, constructed by appending additional arguments to a closure, in the context of the object or category containing the call:

```
call(Closure, Arg1, Arg2, ...)
```

To send a goal, constructed by appending additional arguments to a closure, as a message to *self*:

```
call(:Closure, Arg1, Arg2, ...)
```

To send a goal, constructed by appending additional arguments to a closure, as a message to an explicit object:

```
call(Object::Closure, Arg1, Arg2, ...)
```

See also

ignore/1, once/1, (\+)/1

built-in method

ignore/1

Description

```
ignore(Goal)
```

This predicate succeeds whether its argument succeeds or fails and it is not re-executable. This built-in meta-predicate is declared as a private method and thus cannot be used as a message to an object.

This meta-predicate is opaque to cuts in its argument.

Meta-predicate template

```
ignore(0)
```

Modes and number of proofs

```
ignore(+callable) - one
```

Errors

Goal is a variable:

instantiation_error

Goal is neither a variable nor a callable term:

type_error(callable, Goal)

Examples

Call a goal and succeeding even if it fails:

```
ignore(Goal)
```

To send a message succeeding even if it fails to *self*:

```
ignore(:Goal)
```

To send a message succeeding even if it fails to an explicit object:

```
ignore(Object::Goal)
```

See also

call/1-N, *once/1*, *(\+)/1*

built-in method

`once/1`

Description

`once(Goal)`

This predicate behaves as `call(Goal)` but it is not re-executable. This built-in meta-predicate is declared as a private method and thus cannot be used as a message to an object.

This meta-predicate is opaque to cuts in its argument.

Meta-predicate template

`once(0)`

Modes and number of proofs

`once(+callable) - zero_or_one`

Errors

Goal is a variable:

`instantiation_error`

Goal is neither a variable nor a callable term:

`type_error(callable, Goal)`

Examples

Call a goal deterministically in the context of the object or category containing the call:

`once(Goal)`

To send a goal as a non-backtracable message to *self*:

`once(:Goal)`

To send a goal as a non-backtracable message to an explicit object:

`once(Object::Goal)`

See also

call/1-N, ignore/1, (\+)/1

built-in method

`(\+)/1`

Description

`\+ Goal`

Not-provable meta-predicate. True iff `call(Goal)` is false. This built-in meta-predicate is declared as a private method and thus cannot be used as a message to an object.

Warning

The argument is always compiled (for improved performance). As a consequence, when the argument is a control construct (e.g., a conjunction), any meta-variables will be wrapped with the equivalent to the `call/1` control construct. Note that these semantics differ from the ISO Prolog Core standard specification for the `(\+)/1` built-in predicate. For example, assuming a conforming system:

```
| ?- X = !, \+ (member(Y,[1,2,3]), X, write(Y), fail).
1
X = !
```

But in Logtalk `X` is compiled into a meta-call, which is not cut-transparent:

```
yes
| ?- logtalk << (X = !, \+ (member(Y,[1,2,3]), X, write(Y), fail)).
123
X = !
```

Note that the ISO Prolog Core standard doesn't specify a cut-transparent alternative to the `call/1` control construct.

Meta-predicate template

`\+ 0`

Modes and number of proofs

```
\+ +callable - zero_or_one
```

Errors

Goal is a variable:

```
instantiation_error
```

Goal is neither a variable nor a callable term:

```
type_error(callable, Goal)
```

Examples

Not-provable goal in the context of the object or category containing the call:

```
\+ Goal
```

Not-provable goal sent as a message to *self*:

```
\+ ::Goal
```

Not-provable goal sent as a message to an explicit object:

```
\+ Object::Goal
```

See also

call/1-N, ignore/1, once/1

2.5.6 Error handling

built-in method

catch/3

Description

```
catch(Goal, Catcher, Recovery)
```

Catches exceptions thrown by a goal. See also the ISO Prolog standard definition. This built-in meta-predicate is declared as a private method and thus cannot be used as a message to an object.

This method can also be used in grammar rules. In this case, the meta-arguments are interpreted as grammar rule bodies.

Modes and number of proofs

```
catch(?callable, ?term, ?callable) - zero_or_more
```

Errors

(none)

Examples

(none)

See also

throw/1, context/1, instantiation_error/0, un instantiation_error/1, type_error/2, domain_error/2, consistency_error/3, existence_error/2, permission_error/3, evaluation_error/1, representation_error/1 resource_error/1, syntax_error/1, system_error/0

built-in method

throw/1

Description

```
throw(Exception)
```

Throws an exception. See also the ISO Prolog standard definition. This built-in method is declared private and thus cannot be used as a message to an object.

Modes and number of proofs

```
throw(+nonvar) - error
```

Errors

Exception is a variable:

instantiation_error

Exception does not unify with the second argument of any call of *catch/3*:

system_error

Examples

(none)

See also

catch/3, context/1, instantiation_error/0, un instantiation_error/1, type_error/2, domain_error/2, consistency_error/3, existence_error/2, permission_error/3, evaluation_error/1, representation_error/1 resource_error/1, syntax_error/1, system_error/0

built-in method

instantiation_error/0

Description

instantiation_error

Throws an instantiation error. Used when an argument or one of its sub-arguments is a variable but a non-variable term is required. For example, trying to open a file with a variable for the input/output mode.

This built-in method is declared private and thus cannot be used as a message to an object. Calling this predicate is equivalent to the following sequence of calls:

```
...,
context(Context),
throw(error(instantiation_error, Context)).
```

This allows the user to generate errors in the same format used by the runtime.

Modes and number of proofs

instantiation_error - error

Errors

When called:

instantiation_error

Examples

```
...,  
var(Handler),  
in instantiation_error.
```

See also

catch/3, throw/1, context/1, uninstantiation_error/1, type_error/2, domain_error/2, consistency_error/3, existence_error/2, permission_error/3, representation_error/1, evaluation_error/1, resource_error/1, syntax_error/1, system_error/0

built-in method

`uninstantiation_error/1`

Description

```
uninstantiation_error(Culprit)
```

Throws an uninstantiation error. Used when an argument or one of its sub-arguments is bound but a variable is required. For example, trying to open a file with a stream argument bound.

This built-in method is declared private and thus cannot be used as a message to an object. Calling this predicate is equivalent to the following sequence of calls:

```
...,  
context(Context),  
throw(error(uninstantiation_error(Culprit), Context)).
```

This allows the user to generate errors in the same format used by the runtime.

Modes and number of proofs

```
uninstantiation_error(@nonvar) - error
```

Errors

When called:

`uninstantiation_error(Culprit)`

Examples

```
...,
var(Handler),
uninstantiation_error(my_stream).
```

➡ See also

catch/3, throw/1, context/1, instantiation_error/0, type_error/2, domain_error/2, consistency_error/3, existence_error/2, permission_error/3, representation_error/1, evaluation_error/1, resource_error/1, syntax_error/1, system_error/0

built-in method

`type_error/2`

Description

```
type_error(Type, Culprit)
```

Throws a type error. Used when the type of an argument is incorrect. For example, trying to use a non-callable term as a message. This built-in method is declared private and thus cannot be used as a message to an object. Calling this predicate is equivalent to the following sequence of goals:

```
...,
context(Context),
throw(error(type_error(Type, Culprit), Context)).
```

This allows the user to generate errors in the same format used by the runtime.

Possible values for Type include all the types defined by the type library object and by other libraries such as `os`, `expecteds`, and `optionals`. The value of Culprit is the argument or one of its sub-terms that caused the error.

Modes and number of proofs

```
type_error(@nonvar, @term) - error
```

Errors

When called:

```
type_error(Type, Culprit)
```

Examples

```
...,
\+ atom(Name),
type_error(atom, Name).
```

➡ See also

catch/3, throw/1, context/1, instantiation_error/0, uninstantiation_error/1, domain_error/2, consistency_error/3, existence_error/2, permission_error/3, representation_error/1, evaluation_error/1, resource_error/1, syntax_error/1, system_error/0,

built-in method

domain_error/2

Description

```
domain_error(Domain, Culprit)
```

Throws a domain error. Used when an argument is of the correct type but outside the valid domain. For example, trying to use an atom as an operator specifier that is not a valid specifier. This built-in method is declared private and thus cannot be used as a message to an object. Calling this predicate is equivalent to the following sequence of goals:

```
...,
context(Context),
throw(error(domain_error(Domain,Culprit), Context)).
```

This allows the user to generate errors in the same format used by the runtime.

Possible values for Domain include:

- character_code_list
- close_option
- flag_option
- flag_value
- compiler_flag
- flag
- prolog_flag
- io_mode
- non_empty_list
- not_less_than_zero
- operator_priority
- operator_specifier
- read_option

- `source_sink`
- `stream`
- `stream_option`
- `stream_or_alias`
- `stream_position`
- `stream_property`
- `write_option`
- `character_code_list`
- `text_encoding`
- `directive`
- `object_directive`
- `protocol_directive`
- `category_directive`
- `object_relation`
- `protocol_relation`
- `category_relation`
- `object_property`
- `protocol_property`
- `category_property`
- `predicate_property`
- `meta_argument_specifier`
- `meta_directive_template`
- `closure`
- `allocation`
- `redefinition`
- `message_sending_goal`
- `class`
- `prototype`
- `scope`
- `boolean`

The value of `Culprit` is the argument or one of its sub-terms that caused the error.

Modes and number of proofs

```
domain_error(+atom, @nonvar) - error
```

Errors

When called:

```
domain_error(Domain, Culprit)
```

Examples

```
...,
atom(Color),
\+ color(Color),
domain_error(color, Color).
```

See also

catch/3, throw/1, context/1, instantiation_error/0, un instantiation_error/1, type_error/2, consistency_error/3, existence_error/2, permission_error/3, representation_error/1, evaluation_error/1, resource_error/1, syntax_error/1, system_error/0

built-in method

consistency_error/3

Description

```
consistency_error(Consistency, Argument1, Argument2)
```

Throws a consistency error. Used when two directives or predicate arguments are individually correct but together are not consistent. For example, a predicate and its alias having different arity in a `uses/2` directive. This built-in method is declared private and thus cannot be used as a message to an object. Calling this predicate is equivalent to the following sequence of goals:

```
...,
context(Context),
throw(error(consistency_error(Consistency,Argument1,Argument2), Context)).
```

This allows the user to generate errors in the same format used by the runtime.

Possible values representing Consistency checks include:

- `same_arity`
- `same_number_of_parameters`
- `same_number_of_arguments`
- `same_closure_specification`

Modes and number of proofs

```
consistency_error(+atom, @nonvar, @nonvar) - error
```

Errors

When called:

```
consistency_error(Consistency, Argument1, Argument2)
```

Examples

```
% code that will trigger consistency errors when compiled:

% predicates (and non-terminals) aliases must have the same
% arity as the original predicates (and non-terminals)
:- uses(list, [
    member/2 as in/1
]).

% meta-predicate templates should be consistent with how closures
% are used regarding the number of additional arguments
:- public(p/2).
:- meta_predicate(p(1, *)).

p(G, A) :-
    call(G, A, 2).
```

➡ See also

catch/3, throw/1, context/1, instantiation_error/0, un instantiation_error/1, type_error/2, domain_error/2, existence_error/2, permission_error/3, representation_error/1, evaluation_error/1, resource_error/1, syntax_error/1, system_error/0

built-in method

existence_error/2

Description

```
existence_error(Thing, Culprit)
```

Throws an existence error. Used when the subject of an operation does not exist. This built-in method is declared private and thus cannot be used as a message to an object. Calling this predicate is equivalent to the following sequence of goals:

```
...,  
context(Context),  
throw(error(existence_error(Thing,Culprit), Context)).
```

This allows the user to generate errors in the same format used by the runtime.

Possible values for Thing include:

- predicate
- non_terminal
- predicate_declaration
- procedure
- source_sink
- stream
- object
- protocol
- category
- module
- ancestor
- library
- file
- directive
- engine
- thread
- goal_thread

The value of Culprit is the argument or one of its sub-terms that caused the error.

Modes and number of proofs

```
existence_error(@nonvar, @nonvar) - error
```

Errors

When called:

```
existence_error(Thing, Culprit)
```

Examples

```
...,
\+ current_object(payload),
existence_error(object, payroll).
```

➡ See also

catch/3, throw/1, context/1, instantiation_error/0, un instantiation_error/1, type_error/2, domain_error/2, consistency_error/3, evaluation_error/1, permission_error/3, representation_error/1, resource_error/1, syntax_error/1, system_error/0

built-in method

permission_error/3

Description

```
permission_error(Operation, PermissionType, Culprit)
```

Throws a permission error. Used when an operation is not allowed. For example, sending a message for a predicate that is not within the scope of the sender. This built-in method is declared private and thus cannot be used as a message to an object. Calling this predicate is equivalent to the following sequence of goals:

```
...,
context(Context),
throw(error(permission_error(Operation, PermissionType, Culprit), Context)).
```

This allows the user to generate errors in the same format used by the runtime.

Possible values for Operation include:

- access
- create
- declare
- define
- modify
- open
- input
- output
- reposition
- repeat

Possible values for PermissionType include:

- predicate_declaration
- protected_predicate

- `private_predicate`
- `static_predicate`
- `dynamic_predicate`
- `predicate`
- `non_terminal`
- `database`
- `object`
- `static_object`
- `static_protocol`
- `static_category`
- `entity_relation`
- `operator`
- `flag`
- `engine`
- `binary_stream`
- `text_stream`
- `source_sink`
- `stream`
- `past_end_of_stream`

The value of `Culprit` is the argument or one of its sub-terms that caused the error.

Modes and number of proofs

```
permission_error(@nonvar, @nonvar, @nonvar) - error
```

Errors

When called:

```
permission_error(Operation, PermissionType, Culprit)
```

Examples

```
...,
\+ writable(File),
permission_error(modify, file, File).
```

➡ See also

catch/3, throw/1, context/1, instantiation_error/0, un instantiation_error/1, type_error/2, domain_error/2, consistency_error/3, existence_error/2, representation_error/1, evaluation_error/1, resource_error/1, syntax_error/1, system_error/0

built-in method**representation_error/1****Description****representation_error(Flag)**

Throws a representation error. Used when some representation limit is exceeded. For example, trying to construct a compound term that exceeds the maximum arity supported by the backend Prolog system. This built-in method is declared private and thus cannot be used as a message to an object. Calling this predicate is equivalent to the following sequence of goals:

```
...,
context(Context),
throw(error(representation_error(Flag), Context)).
```

This allows the user to generate errors in the same format used by the runtime.

Possible values for Flag include:

- character
- character_code
- in_character_code
- max_arity
- max_integer
- min_integer
- acyclic_term
- lambda_parameters
- entity_prefix

Modes and number of proofs**representation_error(+atom) - error**

Errors

When called:

```
representation_error(Flag)
```

Examples

```
...,  
Code > 127,  
representation_error(character_code).
```

See also

catch/3, throw/1, context/1, instantiation_error/0, un instantiation_error/1, type_error/2, domain_error/2, consistency_error/3, existence_error/2, permission_error/3, evaluation_error/1, resource_error/1, syntax_error/1, system_error/0

built-in method

`evaluation_error/1`

Description

```
evaluation_error(Error)
```

Throws an evaluation error. Used when evaluating an arithmetic expression generates an exception. This built-in method is declared private and thus cannot be used as a message to an object. Calling this predicate is equivalent to the following sequence of goals:

```
...,  
context(Context),  
throw(error(evaluation_error(Error), Context)).
```

This allows the user to generate errors in the same format used by the runtime.

Possible values for Error include:

- float_overflow
- int_overflow
- undefined
- underflow
- zero_divisor

Modes and number of proofs

```
evaluation_error(@nonvar) - error
```

Errors

When called:

```
evaluation_error(Exception)
```

Examples

```
...,
Divisor ::= 0,
evaluation_error(zero_divisor).
```

➡ See also

catch/3, throw/1, context/1, instantiation_error/0, un instantiation_error/1, type_error/2, domain_error/2, consistency_error/3, existence_error/2, permission_error/3, representation_error/1, resource_error/1, syntax_error/1, system_error/0

built-in method

resource_error/1

Description

```
resource_error(Resource)
```

Throws a resource error. Used when a required resource (e.g., memory or disk space) to complete execution is not available. This built-in method is declared private and thus cannot be used as a message to an object. Calling this predicate is equivalent to the following sequence of goals:

```
...,
context(Context),
throw(error(resource_error(Resource), Context)).
```

This allows the user to generate errors in the same format used by the runtime.

Possible values for Resource include:

- engines
- threads
- coinduction
- soft_cut_support
- text_encoding_support

Modes and number of proofs

```
resource_error(@nonvar) - error
```

Errors

When called:

```
resource_error(Resource)
```

Examples

```
...,  
empty(Tank),  
resource_error(gas).
```

See also

catch/3, throw/1, context/1, instantiation_error/0, un instantiation_error/1, type_error/2, domain_error/2, consistency_error/3, existence_error/2, permission_error/3, representation_error/1, syntax_error/1, system_error/0

built-in method

`syntax_error/1`

Description

```
syntax_error(Description)
```

Throws a syntax error. Used when the sequence of characters being read are not syntactically valid. This built-in method is declared private and thus cannot be used as a message to an object. Calling this predicate is equivalent to the following sequence of goals:

```
...,  
context(Context),  
throw(error(syntax_error(Description), Context)).
```

This allows the user to generate errors in the same format used by the runtime.

Modes and number of proofs

```
syntax_error(@nonvar) - error
```

Errors

When called:

```
syntax_error(Description)
```

Examples

(none)

See also

catch/3, throw/1, context/1, instantiation_error/0, un instantiation_error/1, type_error/2, domain_error/2, consistency_error/3, existence_error/2, permission_error/3, representation_error/1, system_error/0 resource_error/1

built-in method

`system_error/0`

Description

```
system_error
```

Throws a system error. Used when runtime execution can no longer proceed. For example, an exception is thrown without an active catcher. This built-in method is declared private and thus cannot be used as a message to an object. Calling this predicate is equivalent to the following sequence of goals:

```
...,
context(Context),
throw(error(system_error, Context)).
```

This allows the user to generate errors in the same format used by the runtime.

Modes and number of proofs

```
system_error - error
```

Errors

When called:

`system_error`

Examples

(none)

See also

catch/3, throw/1, context/1, instantiation_error/0, un instantiation_error/1, type_error/2, domain_error/2, consistency_error/3, existence_error/2, permission_error/3, representation_error/1 evaluation_error/1, resource_error/1, syntax_error/1,

2.5.7 All solutions

built-in method

bagof/3

Description

bagof(Template, Goal, List)

Collects a bag of solutions for the goal for each set of instantiations of the free variables in the goal. The order of the elements in the bag follows the order of the goal solutions. The free variables in the goal are the variables that occur in the goal but not in the template. Free variables can be ignored, however, by using the $\wedge/2$ existential qualifier. For example, if T is term containing all the free variables that we want to ignore, we can write $T \wedge \text{Goal}$. Note that the term T can be written as $V1 \wedge V2 \wedge \dots$

When there are free variables, this method is re-executable on backtracking. This method fails when there are no solutions, never returning an empty list.

This built-in meta-predicate is declared as a private method and thus cannot be used as a message to an object.

Meta-predicate template

bagof(*, \wedge , *)

Modes and number of proofs

```
bagof(@term, +callable, -list) - zero_or_more
```

Errors

Goal is a variable:

```
instantiation_error
```

Goal is neither a variable nor a callable term:

```
type_error(callable, Goal)
```

Goal is a call to a non-existing predicate:

```
existence_error(procedure, Predicate)
```

Examples

To find a bag of solutions in the context of the object or category containing the call:

```
bagof(Template, Goal, List)
```

To find a bag of solutions of sending a message to *self*:

```
bagof(Template, ::Message, List)
```

To find a bag of solutions of sending a message to an explicit object:

```
bagof(Template, Object::Message, List)
```

See also

findall/3, findall/4, forall/2, setof/3

built-in method

findall/3

Description

```
findall(Template, Goal, List)
```

Collects a list of solutions for the goal. The order of the elements in the list follows the order of the goal solutions. It succeeds returning an empty list when the goal has no solutions.

This built-in meta-predicate is declared as a private method and thus cannot be used as a message to an object.

Meta-predicate template

```
findall(*, 0, *)
```

Modes and number of proofs

```
findall(?term, +callable, ?list) - zero_or_one
```

Errors

Goal is a variable:

instantiation_error

Goal is neither a variable nor a callable term:

type_error(callable, Goal)

Goal is a call to a non-existing predicate:

existence_error(procedure, Predicate)

Examples

To find all solutions in the context of the object or category containing the call:

```
findall(Template, Goal, List)
```

To find all solutions of sending a message to *self*:

```
findall(Template, ::Message, List)
```

To find all solutions of sending a message to an explicit object:

```
findall(Template, Object::Message, List)
```

See also

bagof/3, *findall/4*, *forall/2*, *setof/3*

built-in method

findall/4

Description

```
findall(Template, Goal, List, Tail)
```

Variant of the *findall/3* method that allows passing the tail of the results list. It succeeds returning the tail argument when the goal has no solutions.

This built-in meta-predicate is declared as a private method and thus cannot be used as a message to an object.

Meta-predicate template

```
findall(*, 0, *, *)
```

Modes and number of proofs

```
findall(?term, +callable, ?list, ?term) - zero_or_one
```

Errors

Goal is a variable:

instantiation_error

Goal is neither a variable nor a callable term:

type_error(callable, Goal)

Goal is a call to a non-existing predicate:

existence_error(procedure, Predicate)

Examples

To find all solutions in the context of the object or category containing the call:

```
findall(Template, Goal, List, Tail)
```

To find all solutions of sending a message to *self*:

```
findall(Template, ::Message, List, Tail)
```

To find all solutions of sending a message to an explicit object:

```
findall(Template, Object::Message, List, Tail)
```

See also

bagof/3, findall/3, forall/2, setof/3

built-in method

forall/2

Description

```
forall(Generator, Test)
```

For all solutions of Generator, Test is true. This meta-predicate implements a *generate-and-test* loop using a definition equivalent to `\+ (Generator, \+ Test)`. As a consequence, no variables in the arguments are bound by a call to this predicate. This predicate often provides a better alternative to a *failure-driven loop* as an unexpected Test failure will not be ignored as it will make the forall/2 call fail.

This built-in meta-predicate is declared as a private method and thus cannot be used as a message to an object.

Meta-predicate template

```
forall(0, 0)
```

Modes and number of proofs

```
forall(@callable, @callable) - zero_or_one
```

Errors

Either Generator or Test is a variable:

```
instantiation_error
```

Generator is neither a variable nor a callable term:

```
type_error(callable, Generator)
```

Test is neither a variable nor a callable term:

```
type_error(callable, Test)
```

Examples

To call both goals in the context of the object or category containing the call:

```
forall(Generator, Test)
```

To send both goals as messages to *self*:

```
forall(::Generator, ::Test)
```

To send both goals as messages to explicit objects:

```
forall(Object1::Generator, Object2::Test)
```

See also

bagof/3, findall/3, findall/4, setof/3

built-in method

setof/3**Description**

```
setof(Template, Goal, List)
```

Collects a set of solutions for the goal for each set of instantiations of the free variables in the goal. The solutions are sorted using standard term order. The free variables in the goal are the variables that occur in the goal but not in the template. Free variables can be ignored, however, by using the \exists existential qualifier. For example, if T is term containing all the free variables that we want to ignore, we can write T^{\exists} Goal. Note that the term T can be written as $V1^{\exists}V2^{\exists} \dots$

When there are free variables, this method is re-executable on backtracking. This method fails when there are no solutions, never returning an empty list.

This built-in meta-predicate is declared as a private method and thus cannot be used as a message to an object.

Meta-predicate template

```
setof(*, ^, *)
```

Modes and number of proofs

```
setof(@term, +callable, -list) - zero_or_more
```

Errors

Goal is a variable:

instantiation_error

Goal is neither a variable nor a callable term:

type_error(callable, Goal)

Goal is a call to a non-existing predicate:

existence_error(procedure, Predicate)

Examples

To find a set of solutions in the context of the object or category containing the call:

```
setof(Template, Goal, List)
```

To find a set of solutions of sending a message to *self*:

```
setof(Template, ::Message, List)
```

To find a set of solutions of sending a message to an explicit object:

```
setof(Template, Object::Message, List)
```

 **See also**

bagof/3, findall/3, findall/4, forall/2

2.5.8 Event handling

built-in method

`before/3`

Description

```
before(Object, Message, Sender)
```

User-defined method for handling *before events*. This method is declared in the `monitoring` built-in protocol as a public predicate and is automatically called by the runtime for messages sent using the `(:)/2` control construct from within objects compiled with the *events* flag set to allow.

Note that you can make this predicate scope protected or private by using, respectively, *protected or private implementation* of the monitoring protocol.

Modes and number of proofs

```
before(?object_identifier, ?callable, ?object_identifier) - zero_or_more
```

Errors

(none)

Examples

```
:- object(...,  
    implements(monitoring),  
    ...).  
  
% write a log message when a message is sent:  
before(Object, Message, Sender) :-  
    writeln(Object), write('::'), writeln(Message),  
    write(' from '), writeln(Sender), nl.
```

 **See also**

after/3, abolish_events/5, current_event/5, define_events/5

built-in method

after/3**Description**

```
after(Object, Message, Sender)
```

User-defined method for handling *after events*. This method is declared in the *monitoring* built-in protocol as a public predicate and is automatically called by the runtime for messages sent using the *(::)/2* control construct from within objects compiled with the *events* flag set to allow.

Note that you can make this predicate scope protected or private by using, respectively, *protected* or *private implementation* of the *monitoring* protocol.

Modes and number of proofs

```
after(?object_identifier, ?callable, ?object_identifier) - zero_or_more
```

Errors

(none)

Examples

```
:- object(...,
    implements(monitoring),
    ...).

% write a log message when a message is successful:
after(Object, Message, Sender) :-
    writeq(Object), write('::'), writeq(Message),
    write(' from '), writeq(Sender), nl.
```

 **See also**

before/3, *abolish_events/5*, *current_event/5*, *define_events/5*

2.5.9 Message forwarding**built-in method**

forward/1

Description

forward(*Message*)

User-defined method for forwarding unknown messages sent to an object (using the `(::)/2` control construct), automatically called by the runtime when defined. This method is declared in the `forwarding` built-in protocol as a *public* predicate. Note that you can make its scope protected or private by using, respectively, *protected* or *private implementation* of the forwarding protocol.

Modes and number of proofs

forward(+callable) - zero_or_more

Errors

(none)

Examples

```
:- object(proxy,
    implements(forwarding),
    ...).

forward(Message) :-
    % delegate unknown messages to the "real" object
    [real::Message].
```

See also

[\[\]/1](#)

2.5.10 Definite clause grammar rules

built-in method

`call//1-N`

Description

```

call(Closure)
call(Closure, Arg1, ...)
call(Object::Closure, Arg1, ...)
call(::Closure, Arg1, ...)
call(^^Closure, Arg1, ...)
...

```

This non-terminal takes a *closure* and is processed by appending the two implicit grammar rule arguments to the arguments of the closure. This built-in non-terminal is interpreted as a private non-terminal and thus cannot be used as a message to an object.

Using this non-terminal is recommended when calling a predicate whose last two arguments are the two implicit grammar rule arguments to avoid hard-coding assumptions about how grammar rules are compiled into clauses. Note that the compiler ensures zero overhead when using this non-terminal with a bound argument at compile-time. To call a predicate with a different argument order, use a *lambda expression* or define a *predicate alias*. For example:

```

square -->
  call([Number, Double]>>(Double is Number*Number)).

```

When using a *backend Prolog compiler* supporting a module system, calls in the format `call(Module:Closure)` may also be used.

Meta-non-terminal template

```

call(0)
call(1, *)
call(2, *, *)
...

```

Modes and number of proofs

```
call(+callable) - zero_or_more
call(+callable, ?term) - zero_or_more
call(+callable, ?term, ?term) - zero_or_more
...
```

Errors

Closure is a variable:

instantiation_error

Closure is neither a variable nor a callable term:

type_error(callable, Closure)

Examples

Calls a goal, constructed by appending the two implicit grammar rule arguments to the closure, in in the context of the object or category containing the call:

call(Closure)

To make a *super* call, constructed by appending the two implicit grammar rule arguments to the closure:

call(^ ^ Closure)

To send a goal, constructed by appending the two implicit grammar rule arguments to the closure, as a message to *self*:

call(: : Closure)

To send a goal, constructed by appending the two implicit grammar rule arguments to the closure, as a message to an explicit object:

call(Object : : Closure)

See also

eos//0, phrase//1, phrase/2, phrase/3

built-in method

eos//0

Description

eos

This non-terminal matches the end-of-input. It is implemented by checking that the implicit difference list unifies with []-[].

Modes and number of proofs

```
eos - zero_or_one
```

Errors

(none)

Examples

```
abc --> a, b, c, eos.
```

See also

call//1-N, phrase//1, phrase/2, phrase/3

built-in method

`phrase//1`

Description

```
phrase(GrammarRuleBody)
```

This non-terminal takes a grammar rule body and parses it using the two implicit grammar rule arguments. A common use is to wrap what otherwise would be a *naked meta-variable* in a grammar rule body when defining a meta non-terminal.

Meta-non-terminal template

```
phrase(0)
```

Modes and number of proofs

```
phrase(+callable) - zero_or_more
```

Errors

GrammarRuleBody is a variable:

```
instantiation_error
```

GrammarRuleBody is neither a variable nor a callable term:

```
type_error(callable, GrammarRuleBody)
```

Examples

(none)

See also

call//1-N, phrase/2, phrase/3

built-in method

phrase/2

Description

```
phrase(GrammarRuleBody, Input)
phrase(::GrammarRuleBody, Input)
phrase(Object::GrammarRuleBody, Input)
```

True when the GrammarRuleBody grammar rule body can be applied to the Input list of tokens. In the most common case, GrammarRuleBody is a non-terminal defined by a grammar rule. This built-in method is declared private and thus cannot be used as a message to an object. When using a *backend Prolog compiler* supporting a module system, calls in the format `phrase(Module:GrammarRuleBody, Input)` may also be used.

This method is opaque to cuts in the first argument. When the first argument is sufficiently instantiated at compile-time, the method call is compiled in order to eliminate the implicit overheads of converting the grammar rule body into a goal and meta-calling it. For performance reasons, the second argument is only type-checked at compile-time.

Meta-predicate template

```
phrase(2, *)
```

Modes and number of proofs

```
phrase(+callable, ?list) - zero_or_more
```

Errors

GrammarRuleBody is a variable:

```
instantiation_error
```

GrammarRuleBody is neither a variable nor a callable term:

```
type_error(callable, GrammarRuleBody)
```

Examples

To parse a list of tokens using a local non-terminal:

```
phrase(NonTerminal, Input)
```

To parse a list of tokens using a non-terminal within the scope of *self*:

```
phrase(::NonTerminal, Input)
```

To parse a list of tokens using a public non-terminal of an explicit object:

```
phrase(Object::NonTerminal, Input)
```

See also

call//1-N, phrase//1, phrase/3

built-in method

phrase/3

Description

```
phrase(GrammarRuleBody, Input, Rest)
phrase(::GrammarRuleBody, Input, Rest)
phrase(Object::GrammarRuleBody, Input, Rest)
```

True when the GrammarRuleBody grammar rule body can be applied to the Input-Rest difference list of tokens. In the most common case, GrammarRuleBody is a non-terminal defined by a grammar rule. This built-in method is declared private and thus cannot be used as a message to an object. When using a *backend Prolog compiler* supporting a module system, calls in the format phrase(Module:GrammarRuleBody, Input, Rest) may also be used.

This method is opaque to cuts in the first argument. When the first argument is sufficiently instantiated at compile-time, the method call is compiled in order to eliminate the implicit overheads of converting the grammar rule body into a goal and meta-calling it. For performance reasons, the second and third arguments are only type-checked at compile time.

Meta-predicate template

```
phrase(2, *, *)
```

Modes and number of proofs

```
phrase(+callable, ?list, ?list) - zero_or_more
```

Errors

GrammarRuleBody is a variable:

```
instantiation_error
```

GrammarRuleBody is neither a variable nor a callable term:

```
type_error(callable, GrammarRuleBody)
```

Examples

To parse a list of tokens using a local non-terminal:

```
phrase(NonTerminal, Input, Rest)
```

To parse a list of tokens using a non-terminal within the scope of *self*:

```
phrase(::NonTerminal, Input, Rest)
```

To parse a list of tokens using a public non-terminal of an explicit object:

```
phrase(Object::NonTerminal, Input, Rest)
```

See also

[call//1-N](#), [phrase/2](#), [phrase/3](#)

2.5.11 Term and goal expansion

built-in method

[expand_term/2](#)

Description

```
expand_term(Term, Expansion)
```

Expands a term. The most common use is to expand a grammar rule into a clause. Users may override the default Logtalk grammar rule translator by defining clauses for the [term_expansion/2](#) hook predicate.

The expansion works as follows: if the first argument is a variable, then it is unified with the second argument; if the first argument is not a variable and there are local or inherited clauses for the `term_expansion/2`

hook predicate within scope, then this predicate is called to provide an expansion that is then unified with the second argument; if the `term_expansion/2` predicate is not used and the first argument is a compound term with functor (`-->`)/2 then the default Logtalk grammar rule translator is used, with the resulting clause being unified with the second argument; when the translator is not used, the two arguments are unified. The `expand_term/2` predicate may return a single term or a list of terms.

This built-in method may be used to expand a grammar rule into a clause for use with the built-in database methods.

Automatic term expansion is only performed at compile-time (to expand terms read from a source file) when using a *hook object*. This predicate can be used by the user to manually perform term expansion at runtime (for example, to convert a grammar rule into a clause).

Modes and number of proofs

```
expand_term(?term, ?term) - one
```

Errors

(none)

Examples

(none)

See also

`expand_goal/2`, `goal_expansion/2`, `term_expansion/2`

built-in method

`term_expansion/2`

Description

```
term_expansion(Term, Expansion)
```

Defines an expansion for a term. This predicate, when defined and within scope, is automatically called by the `expand_term/2` method. When that is not the case, the `expand_term/2` method only uses the default expansions. Use of this predicate by the `expand_term/2` method may be restricted by changing its default public scope.

The `term_expansion/2` predicate may return a list of terms. Returning an empty list effectively suppresses the term.

Term expansion may also be applied when compiling source files by defining the object providing access to the `term_expansion/2` clauses as a *hook object*. Clauses for the `term_expansion/2` predicate defined within an object or a category are **never** used in the compilation of the object or the category itself. Moreover, in this context, terms wrapped using the `{}/1` compiler bypass control construct are not expanded and any expanded term wrapped in this control construct will not be further expanded.

Objects and categories implementing this predicate should declare that they implement the [expanding](#) protocol if no ancestor already declares it. This protocol implementation relation can be declared as either *protected* or *private* to restrict the scope of this predicate.

Modes and number of proofs

```
term_expansion(+nonvar, -nonvar) - zero_or_one
term_expansion(+nonvar, -list(nonvar)) - zero_or_one
```

Errors

(none)

Examples

```
term_expansion((:- license(default)), (:- license(gplv3))).
term_expansion(data(Millimeters), data(Meters)) :- Meters is Millimeters / 1000.
```

See also

[expand_goal/2](#), [expand_term/2](#), [goal_expansion/2](#), [logtalk_load_context/2](#)

built-in method

[expand_goal/2](#)

Description

```
expand_goal(Goal, ExpandedGoal)
```

Expands a goal. The expansion works as follows: if the first argument is a variable, then it is unified with the second argument; if the first argument is not a variable and there are local or inherited clauses for the [goal_expansion/2](#) hook predicate within scope, then this predicate is recursively called until a fixed-point is reached to provide an expansion that is then unified with the second argument; if the [goal_expansion/2](#) predicate is not within scope, the two arguments are unified.

Automatic goal expansion is only performed at compile-time (to expand the body of clauses and meta-directives read from a source file) when using [hook objects](#). This predicate can be used by the user to manually perform goal expansion at runtime (for example, before asserting a clause).

Modes and number of proofs

```
expand_goal(?term, ?term) - one
```

Errors

(none)

Examples

(none)

See also

[expand_term/2](#), [goal_expansion/2](#), [term_expansion/2](#)

built-in method

`goal_expansion/2`

Description

```
goal_expansion(Goal, ExpandedGoal)
```

Defines an expansion for a goal. The first argument is the goal to be expanded. The expanded goal is returned in the second argument. This predicate is called recursively on the expanded goal until a fixed point is reached. Thus, care must be taken to avoid compilation loops. This predicate, when defined and within scope, is automatically called by the [expand_goal/2](#) method. Use of this predicate by the `expand_goal/2` method may be restricted by changing its default public scope.

Goal expansion may also be applied when compiling source files by defining the object providing access to the `goal_expansion/2` clauses as a *hook object*. Clauses for the `goal_expansion/2` predicate defined within an object or a category are **never** used in the compilation of the object or the category itself. Moreover, in this context, goals wrapped using the `{}/1` compiler bypass control construct are not expanded and any expanded goal wrapped in this control construct will not be further expanded.

Objects and categories implementing this predicate should declare that they implement the [expanding](#) built-in protocol if no ancestor already declares it. This protocol implementation relation can be declared as either *protected* or *private* to restrict the scope of this predicate.

Modes and number of proofs

```
goal_expansion(+callable, -callable) - zero_or_one
```

Errors

(none)

Examples

```
goal_expansion(write(Term), (write_term(Term, []), nl)).  
goal_expansion(read(Term), (write('Input: '), {read(Term)})).
```

See also

expand_goal/2, expand_term/2, term_expansion/2, logtalk_load_context/2

2.5.12 Coinduction hooks

built-in method

`coinductive_success_hook/1-2`

Description

```
coinductive_success_hook(Head, Hypothesis)  
coinductive_success_hook(Head)
```

User-defined hook predicates that are automatically called in case of coinductive success when proving a query for a coinductive predicates. The hook predicates are called with the head of the coinductive predicate on coinductive success and, optionally, with the hypothesis that has used to reach coinductive success.

When both hook predicates are defined, the `coinductive_success_hook/1` clauses are only used if no `coinductive_success_hook/2` clause applies. The compiler ensures zero performance penalties when defining coinductive predicates without a corresponding definition for the coinductive success hook predicates.

The compiler assumes that these hook predicates are defined as static predicates in order to optimize their use.

Modes and number of proofs

```
coinductive_success_hook(+callable, +callable) - zero_or_one
coinductive_success_hook(+callable) - zero_or_one
```

Errors

(none)

Examples

```
% Are there "occurrences" of arg1 in arg2?
:- public(member/2).
:- coinductive(member/2).

member(X, [X| _]).
member(X, [_| T]) :-
    member(X, T).

% Are there infinitely many "occurrences" of arg1 in arg2?
:- public(comember/2).
:- coinductive(comember/2).
comember(X, [_| T]) :-
    comember(X, T).

coinductive_success_hook(member(_, _)) :-
    fail.
coinductive_success_hook(comember(X, L)) :-
    member(X, L).
```

See also

[coinductive/1](#)

2.5.13 Message printing

built-in method

`print_message/3`

Description

```
print_message(Kind, Component, Term)
```

Built-in method for printing a message represented by a term, which is converted to the message text using the `logtalk::message_tokens(Term, Component)` hook non-terminal. This method is declared in the `logtalk` built-in object as a public predicate. The line prefix and the output stream used for each Kind-Component pair can be found using the `logtalk::message_prefix_stream(Kind, Component, Prefix, Stream)` hook predicate.

This predicate starts by converting the message term to a list of tokens and by calling the `logtalk::message_hook(Message, Kind, Component, Tokens)` hook predicate. If this predicate succeeds, the `print_message/3` predicate assumes that the message has been successfully printed.

By default: messages of kind `debug` or `debug(_)` are only printed when the debug flag is turned on; messages of kind `banner`, `comment`, or `comment(_)` are only printed when the report flag is set to on; messages of kind `warning` and `warning(_)` are not printed when the report flag is set to off; messages of kind `silent` and `silent()` are not printed (but can be intercepted).

Modes and number of proofs

```
print_message(+nonvar, +nonvar, +nonvar) - one
```

Errors

(none)

Examples

```
..., logtalk::print_message(information, core, redefining_entity(object, foo)), ...
```

See also

`message_hook/4`, `message_prefix_stream/4`, `message_tokens//2`, `print_message_tokens/3`,
`print_message_token/4`, `ask_question/5`, `question_hook/6`, `question_prompt_stream/4`

built-in method

`message_tokens//2`

Description

`message_tokens(Message, Component)`

User-defined non-terminal hook used to rewrite a message term into a list of tokens and declared in the `logtalk` built-in object as a public, multifile, and dynamic non-terminal. The list of tokens can be printed by calling the `print_message_tokens/3` method. This non-terminal hook is automatically called by the `print_message/3` method.

Modes and number of proofs

`message_tokens(+nonvar, +nonvar) - zero_or_more`

Errors

(none)

Examples

```
:- multifile(logtalk::message_tokens//2).
:- dynamic(logtalk::message_tokens//2).

logtalk::message_tokens(redefining_entity(Type, Entity), core) -->
    ['Redefining ~w ~q'-[Type, Entity], nl].
```

See also

`message_hook/4`, `message_prefix_stream/4`, `print_message/3`, `print_message_tokens/3`,
`print_message_token/4`, `ask_question/5`, `question_hook/6`, `question_prompt_stream/4`

built-in method

`message_hook/4`

Description

`message_hook(Message, Kind, Component, Tokens)`

User-defined hook method for intercepting printing of a message, declared in the `logtalk` built-in object as a public, multifile, and dynamic predicate. This hook method is automatically called by the `print_message/3` method. When the call succeeds, the `print_message/3` method assumes that the message has been successfully printed.

Modes and number of proofs

```
message_hook(@nonvar, @nonvar, @nonvar, @list(nonvar)) - zero_or_one
```

Errors

(none)

Examples

```
:- multifile(logtalk::message_hook/4).
:- dynamic(logtalk::message_hook/4).

% print silent messages instead of discarding them as default
logtalk::message_hook(_, silent, core, Tokens) :-
    logtalk::message_prefix_stream(silent, core, Prefix, Stream),
    logtalk::print_message_tokens(Stream, Prefix, Tokens).
```

See also

[message_prefix_stream/4](#), [message_tokens//2](#), [print_message/3](#), [print_message_tokens/3](#),
[print_message_token/4](#), [ask_question/5](#), [question_hook/6](#), [question_prompt_stream/4](#)

built-in method

`message_prefix_stream/4`

Description

```
message_prefix_stream(Kind, Component, Prefix, Stream)
```

User-defined hook method for specifying the default prefix and stream for printing a message for a given kind and *component*. This method is declared in the `logtalk` built-in object as a public, multifile, and dynamic predicate.

The prefix is printed for each line in the list of tokens generated by the message tokenization (new lines are specified using the `nl` token; see the documentation of the `logtalk` built-in object for details).

Modes and number of proofs

```
message_prefix_stream(?nonvar, ?nonvar, ?atom, ?stream_or_alias) - zero_or_more
```

Errors

(none)

Examples

```
:- multifile(logtalk::message_prefix_stream/4).
:- dynamic(logtalk::message_prefix_stream/4).

logtalk::message_prefix_stream(information, core, '% ', user_output).
```

See also

message_prefix_file/6, message_hook/4, message_tokens//2, print_message/3, print_message_tokens/3, print_message_token/4, ask_question/5, question_hook/6, question_prompt_stream/4

built-in method

`message_prefix_file/6`

Description

```
message_prefix_file(Kind, Component, Prefix, File, Mode, Options)
```

Experimental user-defined hook method for specifying the prefix and file for copying messages for a given kind and *component*. This method is declared in the `logtalk` built-in object as a public, multifile, and dynamic predicate. The valid values for *Mode* are `write` and `append`. The valid options are the same as the standard `open/4` predicate.

The prefix is printed for each line in the list of tokens generated by the message tokenization (new lines are specified using the `nl` token; see the documentation of the `logtalk` built-in object for details).

Modes and number of proofs

```
message_prefix_file(?nonvar, ?nonvar, ?atom, ?atom, ?atom, ?list(compound)) - zero_or_more
```

Errors

(none)

Examples

```
:- multifile(logtalk::message_prefix_file/6).
:- dynamic(logtalk::message_prefix_file/6).

logtalk::message_prefix_file(comment, app, '% ', 'comments.txt', append, []).
```

See also

message_prefix_stream/4, message_hook/4, message_tokens//2, print_message/3, print_message_tokens/3, print_message_token/4, ask_question/5, question_hook/6, question_prompt_stream/4

built-in method

`print_message_tokens/3`

Description

```
print_message_tokens(Stream, Prefix, Tokens)
```

Built-in method for printing a list of message tokens, declared in the `logtalk` built-in object as a public predicate. This method is automatically called by the `print_message/3` method (assuming that the message was not intercepted by a `message_hook/4` definition) and calls the user-defined hook predicate `print_message_token/4` for each token. When a call to this hook predicate succeeds, the `print_message_tokens/3` predicate assumes that the token has been printed. When the call fails, the `print_message_tokens/3` predicate uses a default printing procedure for the token.

Modes and number of proofs

```
print_message_tokens(@stream_or_alias, +atom, @list(nonvar)) - zero_or_one
```

Errors

(none)

Examples

```
...,
logtalk::print_message_tokens(user_error, '% ', ['Redefining ~w ~q'-[object,foo], nl]),
...
```

➡ See also

message_hook/4, message_prefix_stream/4, message_tokens//2, print_message/3, print_message_token/4, ask_question/5, question_hook/6, question_prompt_stream/4

built-in method

`print_message_token/4`

Description

```
print_message_token(Stream, Prefix, Token, Tokens)
```

User-defined hook method for printing a message token, declared in the `logtalk` built-in object as a public, multifile, and dynamic predicate. It allows the user to intercept the printing of a message token. This hook method is automatically called by the `print_message_tokens/3` built-in method for each token.

Modes and number of proofs

```
print_message_token(@stream_or_alias, @atom, @nonvar, @list(nonvar)) - zero_or_one
```

Errors

(none)

Examples

```
:- multifile(logtalk::print_message_token/4).
:- dynamic(logtalk::print_message_token/4).

% ignore all flush tokens
logtalk::print_message_token(_Stream, _Prefix, flush, _Tokens).
```

➡ See also

message_hook/4, message_prefix_stream/4, message_tokens//2, print_message/3, print_message_tokens/3, ask_question/5, question_hook/6, question_prompt_stream/4

2.5.14 Question asking

built-in method

`ask_question/5`

Description

```
ask_question(Kind, Component, Question, Check, Answer)
```

Built-in method for asking a question represented by a term, *Question*, which is converted to the question text using the `logtalk::message_tokens(Question, Component)` hook predicate. This method is declared in the logtalk built-in object as a public predicate. The default question prompt and the input stream used for each Kind-Component pair can be found using the `logtalk::question_prompt_stream(Kind, Component, Prompt, Stream)` hook predicate. The *Check* argument is a *closure* that is converted into a checking goal by extending it with the user supplied answer. This predicate implements a read-loop that terminates when the check goal succeeds.

This predicate starts by calling the `logtalk::question_hook(Question, Kind, Component, Tokens, Check, Answer)` hook predicate. If this predicate succeeds, the `ask_question/5` predicate assumes that the question have been successfully asked and replied.

Modes and number of proofs

```
ask_question(+nonvar, +nonvar, +nonvar, +callable, -term) - one
```

Meta-predicate template

```
ask_question(*, *, *, 1, *)
```

Errors

(none)

Examples

```
...,
logtalk::ask_question(enter_age, question, my_app, integer, Age),
...
```

See also

`question_hook/6`, `question_prompt_stream/4`, `message_hook/4`, `message_prefix_stream/4`, `message_tokens//2`, `print_message/3`, `print_message_tokens/3`, `print_message_token/4`

built-in method

question_hook/6

Description

```
question_hook(Question, Kind, Component, Tokens, Check, Answer)
```

User-defined hook method for intercepting asking a question, declared in the `logtalk` built-in object as a public, multifile, and dynamic predicate. This hook method is automatically called by the `ask_question/5` method. When the call succeeds, the `ask_question/5` method assumes that the question have been successfully asked and replied.

Modes and number of proofs

```
question_hook(+nonvar, +nonvar, +nonvar, +list(nonvar), +callable, -term) - zero_or_one
```

Meta-predicate template

```
question_hook(*, *, *, *, 1, *)
```

Errors

(none)

Examples

```
:- multifile(logtalk::question_hook/6).
:- dynamic(logtalk::question_hook/6).

% use a pre-defined answer instead of asking the user
logtalk::question_hook(upper_limit, question, my_app, _, _, 3.7).
```

See also

`ask_question/5`, `question_prompt_stream/4`, `message_hook/4`, `message_prefix_stream/4`, `message_tokens//2`, `print_message/3`, `print_message_tokens/3`, `print_message_token/4`,

built-in method

`question_prompt_stream/4`

Description

`question_prompt_stream(Kind, Component, Prompt, Stream)`

User-defined hook method for specifying the default prompt and input stream for asking a question for a given kind and *component*. This method is declared in the `logtalk` built-in object as a public, multifile, and dynamic predicate.

Modes and number of proofs

`question_prompt_stream(?nonvar, ?nonvar, ?atom, ?stream_or_alias) - zero_or_more`

Errors

(none)

Examples

```
:- multifile(logtalk::question_prompt_stream/4).
:- dynamic(logtalk::question_prompt_stream/4).

logtalk::question_prompt_stream(question, debugger, '    > ', user_input).
```

See also

ask_question/5, *question_hook/6*, *message_hook/4*, *message_prefix_stream/4*, *message_tokens//2*,
print_message/3, *print_message_tokens/3*, *print_message_token/4*

3.1 List predicates

In this example, we will illustrate the use of:

- objects
- protocols

by using common list utility predicates.

3.1.1 Defining a list object

We will start by defining an object, `list`, containing predicate definitions for some common list predicates like `append/3`, `length/2`, and `member/2`:

```
:- object(list).

  :- public([
    append/3, length/2, member/2
  ]).

  append([], List, List).
  append([Head| Tail], List, [Head| Tail2]) :-
    append(Tail, List, Tail2).

  length(List, Length) :-
    length(List, 0, Length).

  length([], Length, Length).
  length([_| Tail], Acc, Length) :-
    Acc2 is Acc + 1,
    length(Tail, Acc2, Length).

  member(Element, [Element| _]).
  member(Element, [_| List]) :-
    member(Element, List).

:- end_object.
```

What is different here from a regular Prolog program? The definitions of the list predicates are the usual ones. We have two new directives, `object/1-5` and `end_object/0`, that encapsulate the object's code. In Logtalk,

by default, all object predicates are private; therefore, we have to explicitly declare all predicates that we want to be public, that is, that we want to call from outside the object. This is done using the [public/1](#) scope directive.

After we copy the object code to a text file and saved it under the name `list.lgt`, we need to change the Prolog working directory to the one used to save our file (consult your Prolog compiler reference manual). Then, after starting Logtalk (see the [Installing and running Logtalk](#) section on the User Manual), we can compile and load the object using the [logtalk_load/1](#) Logtalk built-in predicate:

```
| ?- logtalk_load(list).  
  
object list loaded  
yes
```

We can now try goals like:

```
| ?- list::member(X, [1, 2, 3]).  
  
X = 1;  
X = 2;  
X = 3;  
no
```

or:

```
| ?- list::length([1, 2, 3], L).  
  
L = 3  
yes
```

The infix operator [\(::\)/2](#) is used in Logtalk to send a message to an object. The message must match a public object predicate. If we try to call a non-public predicate such as the `length/3` auxiliary predicate an exception will be generated:

```
| ?- list::length([1, 2, 3], 0, L).  
  
uncaught exception:  
  error(  
    existence_error(predicate_declaration, length/3),  
    logtalk(list::length([1,2,3],0,_), ...)  
  )
```

The exception term describes the type of error and the context where the error occurred.

3.1.2 Defining a list protocol

As we saw in the above example, a Logtalk object may contain predicate directives and predicate definitions (clauses). The set of predicate directives defines what we call the object's *protocol* or interface. An interface may have several implementations. For instance, we may want to define a new object that implements the list predicates using difference lists. However, we do not want to repeat the predicate directives in the new object. Therefore, what we need is to split the object's protocol from the object's implementation by defining a new Logtalk entity known as a protocol. Logtalk protocols are compilation units, at the same level as objects and categories. That said, let us define a `listp` protocol:

```
:- protocol(listp).

    :- public([
        append/3, length/2, member/2
    ]).

:- end_protocol.
```

Similar to what we have done for objects, we use the *protocol/1-2* and *end_protocol/0* directives to encapsulate the predicate directives. We can improve this protocol by documenting the call/return modes and the number of proofs of each predicate using the *mode/2* directive:

```
:- protocol(listp).

    :- public(append/3).
    :- mode(append(?list, ?list, ?list), zero_or_more).

    :- public(length/2).
    :- mode(length(?list, ?integer), zero_or_more).

    :- public(member/2).
    :- mode(member(?term, ?list), zero_or_more).

:- end_protocol.
```

We now need to change our definition of the list object by removing the predicate directives and by declaring that the object implements the listp protocol:

```
:- object(list,
    implements(listp)).

    append([], List, List).
    append([Head| Tail], List, [Head| Tail2]) :-
        append(Tail, List, Tail2).
    ...

:- end_object.
```

The protocol declared in listp may now be alternatively implemented using difference lists by defining a new object, difflist:

```
:- object(difflist,
    implements(listp)).

    append(L1-X, X-L2, L1-L2).
    ...

:- end_object.
```

3.1.3 Summary

- It is easy to define a simple object: just write your Prolog code inside starting and ending object directives and add the necessary scope directives. The object will be self-defining and ready to use.
- Define a protocol when you may want to provide or enable several alternative definitions to a given set of predicates. This way we avoid needless repetition of predicate directives.

3.2 Dynamic object attributes

In this example, we will illustrate the use of:

- categories
- category predicates
- dynamic predicates

by defining a category that implements a set of predicates for handling dynamic object attributes.

3.2.1 Defining a category

We want to define a set of predicates to handle dynamic object attributes. We need public predicates to set, get, and delete attributes, and a private dynamic predicate to store the attribute values. Let us name these predicates `set_attribute/2` and `get_attribute/2`, for getting and setting an attribute value, `del_attribute/2` and `del_attributes/2`, for deleting attributes, and `attribute_/2`, for storing the attributes values.

But we do not want to encapsulate these predicates in an object. Why? Because they are a set of useful, closely related, predicates that may be used by several, unrelated, objects. If defined at an object level, we would be constrained to use inheritance in order to have the predicates available to other objects. Furthermore, this could force us to use multi-inheritance or to have some kind of generic root object containing all kinds of possible useful predicates.

For this kind of situation, Logtalk enables the programmer to encapsulate the predicates in a *category*, so that they can be used in any object. A category is a Logtalk entity, at the same level as objects and protocols. It can contain predicate directives and predicate definitions. Category predicates can be imported by any object, without code duplication and without resorting to inheritance.

When defining category predicates, we need to remember that a category can be imported by more than one object. Thus, calls to the built-in methods that handle the private dynamic predicate (such as [assertz/1](#) or [retract/1](#)) must be made either in the context of *self*, using the *message to self* control structure, `(::)/1`, or in the context of *this* (i.e., in the context of the object importing the category). This way, we ensure that when we call one of the attribute predicates on an object, the intended object own definition of `attribute_/2` will be used. The predicate definitions are straightforward. For example, if opting to store the attributes in *self*:

```
:- category(attributes).

:- public(set_attribute/2).
:- mode(set_attribute(+nonvar, +nonvar), one).

:- public(get_attribute/2).
:- mode(get_attribute(?nonvar, ?nonvar), zero_or_more).

:- public(del_attribute/2).
```

(continues on next page)

(continued from previous page)

```

:- mode(del_attribute(?nonvar, ?nonvar), zero_or_more).

:- public(del_attributes/2).
:- mode(del_attributes(@term, @term), one).

:- private(attribute_/2).
:- mode(attribute_(?nonvar, ?nonvar), zero_or_more).
:- dynamic(attribute_/2).

set_attribute(Attribute, Value):-
    ::retractall(attribute_(Attribute, _)),
    ::assertz(attribute_(Attribute, Value)).

get_attribute(Attribute, Value):-
    ::attribute_(Attribute, Value).

del_attribute(Attribute, Value):-
    ::retract(attribute_(Attribute, Value)).

del_attributes(Attribute, Value):-
    ::retractall(attribute_(Attribute, Value)).

:- end_category.

```

The alternative, opting to store the attributes on *this*, is similar: just delete the `(::)/1` operator from the code above.

We have two new directives, `category/1-4` and `end_category/0`, that encapsulate the category code. If needed, we can put the predicate directives inside a protocol that will be implemented by the category:

```

:- category(attributes,
    implements(attributes_protocol)).

    ...

:- end_category.

```

Any protocol can be implemented by either an object, a category, or both.

3.2.2 Importing the category

We reuse a category's predicates by importing them into an object:

```

:- object(person,
    imports(attributes)).

    ...

:- end_object.

```

After compiling and loading this object and our category, we can now try queries like:

```
| ?- person::set_attribute(name, paulo).  
  
yes  
  
| ?- person::set_attribute(gender, male).  
  
yes  
  
| ?- person::get_attribute(Attribute, Value).  
  
Attribute = name, Value = paulo ;  
Attribute = gender, Value = male ;  
no
```

3.2.3 Summary

- Categories are similar to objects: we just write our predicate directives and definitions bracketed by opening and ending category directives.
- An object reuses a category by importing it. The imported predicates behave as if they have been defined in the object itself.
- When do we use a category instead of an object? Whenever we have a set of closely related predicates that we want to reuse in several, unrelated, objects without being constrained by inheritance relations. Thus, categories can be interpreted as object building components.

3.3 A reflective class-based system

When compiling an object, Logtalk distinguishes prototypes from instances or classes by examining the object relations. If an object instantiates and/or specializes another object, then it is compiled as an instance or class, otherwise it is compiled as a prototype. A consequence of this is that, in order to work with instance or classes, we always have to define root objects for the instantiation and specialization hierarchies (however, we are not restricted to a single hierarchy). The best solution is often to define a reflective class-based system [Maes87], where every class is also an object and, as such, an instance of some class.

In this example, we are going to define the basis for a reflective class-based system, based on an extension of the ideas presented in [Cointe87]. This extension provides, along with root objects for the instantiation and specialization hierarchies, explicit support for abstract classes [Moura94].

3.3.1 Defining the base classes

We will start by defining three classes: `object`, `abstract_class`, and `class`. The class `object` will contain all predicates common to all objects. It will be the root of the inheritance graph:

```
:- object(object,  
    instantiates(class)).  
  
    % predicates common to all objects  
  
:- end_object.
```

The class `abstract_class` specializes `object` by adding predicates common to all classes. It will be the default meta-class for abstract classes:

```
:- object(abstract_class,
    instantiates(class),
    specializes(object)).

    % predicates common to all classes

:- end_object.
```

The class `class` specializes `abstract_class` by adding predicates common to all instantiable classes. It will be the root of the instantiation graph and the default meta-class for instantiable classes:

```
:- object(class,
    instantiates(class),
    specializes(abstract_class)).

    % predicates common to all instantiable classes

:- end_object.
```

Note that all three objects are instances of class `class`. The instantiation and specialization relationships are chosen so that each object may use the predicates defined in itself and in the other two objects, with no danger of *message lookup* endless loops.

3.3.2 Summary

- An object that does not instantiate or specialize other objects is always compiled as a prototype.
- An instance must instantiate at least one object (its class). Similarly, a class must at least specialize or instantiate other object.
- The distinction between abstract classes and instantiable classes is an operational one, depending on the class inherited methods. A class is instantiable if inherits methods for creating instances. Conversely, a class is abstract if does not inherit any instance creation method.

3.4 Profiling programs

In this example, we will illustrate the use of:

- events
- monitors

by defining a simple profiler that prints the starting and ending time for processing a message sent to an object.

3.4.1 Messages as events

In a pure object-oriented system, all computations start by sending messages to objects. We can thus define an *event* as the sending of a message to an object. An event can then be specified by the tuple (Object, Message, Sender). This definition can be refined by interpreting the sending of a message and the return of the control to the object that has sent the message as two distinct events. We call these events respectively *before* and *after*. Therefore, we end up by representing an event by the tuple (Event, Object, Message, Sender). For instance, if we send the message:

```
| ?- foo::bar(X).  
  
X = 1  
yes
```

the two corresponding events will be:

```
(before, foo, bar(X), user)  
(after, foo, bar(1), user)
```

Note that the second event is only generated if the message succeeds. If the message as a goal has multiple solutions, then an after event will be generated for each solution.

Events are automatically generated by the message-sending mechanisms for each public message sent using the `::/2` operator.

3.4.2 Profilers as monitors

A monitor is an object that reacts whenever a spied event occurs. The monitor actions are defined by two event handlers: *before/3* for before events and *after/3* for after events. These predicates are automatically called by the message-sending mechanisms when an event registered for the monitor occurs. These event handlers are declared as public predicates in the monitoring built-in protocol.

In our example, we need a way to get the current time before and after we process a message. We will assume that we have a `time` object implementing a `cpu_time/1` predicate that returns the current CPU time for the Prolog session:

```
:- object(time).  
  
    :- public(cpu_time/1).  
    :- mode(cpu_time(-number), one).  
    ...  
  
:- end_object.
```

Our profiler will be named `stop_watch`. It must define event handlers for the before and after events that will print the event description (object, message, and sender) and the current time:

```
:- object(stop_watch,  
    % event handler predicates protocol  
    implements(monitored)).  
  
    :- uses(time, [cpu_time/1]).  
  
    before(Object, Message, Sender) :-
```

(continues on next page)

(continued from previous page)

```

write(Object), write(' <-- '), writeq(Message),
write(' from '), write(Sender), nl, write('STARTING at '),
cpu_time(Seconds), write(Seconds), write(' seconds'), nl.

after(Object, Message, Sender) :-
  write(Object), write(' <-- '), writeq(Message),
  write(' from '), write(Sender), nl, write('ENDING at '),
  cpu_time(Seconds), write(Seconds), write(' seconds'), nl.

:- end_object.

```

After compiling and loading the `stop_watch` object (and the objects that we want to profile), we can use the `define_events/5` built-in predicate to set up our profiler. For example, to profile all messages that are sent to the object `foo`, we need to call the goal:

```

| ?- define_events(_, foo, _, _, stop_watch).

yes

```

This call will register `stop_watch` as a monitor to all messages sent to object `foo`, for both before and after events. Note that we say “as a monitor”, not “the monitor”: we can have any number of monitors over the same events.

From now on, every time we sent a message to `foo`, the `stop_watch` monitor will print the starting and ending times for the message execution. For instance:

```

| ?- foo::bar(X).

foo <-- bar(X) from user
STARTING at 12.87415 seconds
foo <-- bar(1) from user
ENDING at 12.87419 seconds

X = 1
yes

```

To stop profiling the messages sent to `foo` we use the `abolish_events/5` built-in predicate:

```

| ?- abolish_events(_, foo, _, _, stop_watch).

yes

```

This call will abolish all events defined over the object `foo` assigned to the `stop_watch` monitor.

3.4.3 Summary

- An event is defined as the sending of a (public) message to an object.
- There are two kinds of events: before events, generated before a message is processed, and after events, generated after the message processing has completed successfully.
- Any object can be declared as a monitor to any event. A monitor shall reference the `monitoring` built-in protocol in the object opening directive.
- A monitor defines event handlers, the predicates `before/3` and `after/3`, that are automatically called by the runtime engine when a spied event occurs.
- Three built-in predicates, `define_events/5`, `current_event/5`, and `abolish_events/5`, enables us to define, query, and abolish both events and monitors.

4.1 General

- *Why are all versions of Logtalk numbered 2.x or 3.x?*
- *Why do I need a Prolog compiler to use Logtalk?*
- *Is the Logtalk implementation based on Prolog modules?*
- *Does the Logtalk implementation use term-expansion?*

4.1.1 Why are all versions of Logtalk numbered 2.x or 3.x?

The numbers “2” and “3” in the Logtalk version string refer to, respectively, the second and the third generations of the Logtalk language. Development of Logtalk 2 started in January 1998, with the first alpha release for registered users in July and the first public beta on October. The first stable version of Logtalk 2 was released in February 9, 1999. Development of Logtalk 3 started in April 2012, with the first public alpha released in August 21, 2012. The first stable version of Logtalk 3 was released in January 7, 2015.

4.1.2 Why do I need a Prolog compiler to use Logtalk?

Currently, the Logtalk language is implemented as a *trans-compiler* to Prolog instead of as a standalone compiler. Compilation of Logtalk source files is performed in two steps. First, the Logtalk compiler converts a source file to a Prolog file. Second, the chosen Prolog compiler is called by Logtalk to compile the intermediate Prolog file generated on the first step. The implementation of Logtalk as a trans-compiler allows users to use Logtalk together with features only available on specific Prolog compilers.

4.1.3 Is the Logtalk implementation based on Prolog modules?

No. Logtalk is (currently) implemented in plain Prolog code. Only a few Prolog compilers include a module system, with several compatibility problems between them. Moreover, the current ISO Prolog standard for modules is next to worthless and is ignored by most of the Prolog community. Nevertheless, the Logtalk compiler is able to compile simple modules (using a common subset of module directives) as objects for backward-compatibility with existing code (see the *Prolog integration and migration* for details).

4.1.4 Does the Logtalk implementation use term-expansion?

No. Term-expansion mechanisms are not standard and are not available in all supported Prolog compilers.

4.2 Compatibility

- *What are the backend Prolog compiler requirements to run Logtalk?*
- *Can I use constraint-based packages with Logtalk?*
- *Can I use Logtalk objects and Prolog modules at the same time?*

4.2.1 What are the backend Prolog compiler requirements to run Logtalk?

See the [backend Prolog compiler requirements guide](#).

4.2.2 Can I use constraint-based packages with Logtalk?

Usually, yes. Some constraint-based packages may define operators which clash with the ones defined by Logtalk. In these cases, compatibility with Logtalk depends on the constraint-based packages providing an alternative for accessing the functionality provided by those operators. When the constraint solver is encapsulated using a Prolog module, a possible workaround is to use either explicit module qualification or encapsulate the call using the `{}/1` control construct (thus bypassing the Logtalk compiler).

4.2.3 Can I use Logtalk objects and Prolog modules at the same time?

Yes. In order to call a module predicate from within an object (or category) you may use an `use_module/2` directive or use explicit module qualification (possibly wrapping the call using the Logtalk control construct `{}/1` that allows bypassing of the Logtalk compiler when compiling a predicate call). Logtalk also allows modules to be compiled as objects (see the [Prolog integration and migration](#) for details).

4.3 Installation

- *The integration scripts/shortcuts are not working!*
- *I get errors when starting up Logtalk after upgrading to the latest version!*

4.3.1 The integration scripts/shortcuts are not working!

Check that the `LOGTALKHOME` and `LOGTALKUSER` environment variables are defined, that the Logtalk user folder is available in the location pointed by `LOGTALKUSER` (you can create this folder by running the `logtalk_user_setup` shell script), and that the Prolog compilers that you want to use are supported and available from the system path. If the problem persists, run the shell script that creates the integration script or shortcut manually and check for any error message or additional instructions. For some Prolog compilers such as XSB and Ciao, the first call of the integration script or shortcut must be made by an administrator user. If you are using Windows, make sure that any anti-virus or other security software that you might have installed is not silently blocking some of the installer tasks.

4.3.2 I get errors when starting up Logtalk after upgrading to the latest version!

Changes in the Logtalk compiler between releases may render Prolog adapter files from older versions incompatible with new ones. You may need to update your local Logtalk user files by running e.g. the `logtalk_user_setup` shell script. Check the `UPGRADING.md` file on the root of the Logtalk installation directory and the release notes for any incompatible changes to the adapter files.

4.4 Portability

- *Are my Logtalk applications portable across Prolog compilers?*
- *Are my Logtalk applications portable across operating systems?*

4.4.1 Are my Logtalk applications portable across Prolog compilers?

Yes, as long you don't use built-in predicates or special features only available on some Prolog compilers. There is a *portability* compiler flag that you can set to instruct Logtalk to print a warning for each occurrence of non-ISO Prolog standard features such as proprietary built-in predicates. In addition, it is advisable that you constrain, if possible, the use of platform or compiler dependent code to a small number of objects with clearly defined protocols. You may also use Logtalk support for conditional compilation to compile different entity or predicate definitions depending on the backend Prolog compiler being used.

4.4.2 Are my Logtalk applications portable across operating systems?

Yes, as long you don't use built-in predicates or special features that your chosen backend Prolog compiler only supports on some operating-systems. You may need to change the end-of-line characters of your source files to match the ones on the target operating system and the expectations of your Prolog compiler. Some Prolog compilers silently fail to compile source files with the wrong end-of-line characters.

4.5 Programming

- *Should I use prototypes or classes in my application?*
- *Can I use both classes and prototypes in the same application?*
- *Can I mix classes and prototypes in the same hierarchy?*
- *Can I use a protocol or a category with both prototypes and classes?*
- *What support is provided in Logtalk for defining and using components?*
- *What support is provided in Logtalk for reflective programming?*

4.5.1 Should I use prototypes or classes in my application?

Prototypes and classes provide different patterns of code reuse. A prototype encapsulates code that can be used by itself and by its descendent prototypes. A class encapsulates code to be used by its descendent instances. Prototypes provide the best replacement to the use of modules as encapsulation units, avoiding the need to instantiate a class in order to access its code.

4.5.2 Can I use both classes and prototypes in the same application?

Yes. In addition, you may freely exchange messages between prototypes, classes, and instances.

4.5.3 Can I mix classes and prototypes in the same hierarchy?

No. However, you may define as many prototype hierarchies and class hierarchies and classes as needed by your application.

4.5.4 Can I use a protocol or a category with both prototypes and classes?

Yes. A protocol may be implemented by both prototypes and classes in the same application. Likewise, a category may be imported by both prototypes and classes in the same application.

4.5.5 What support is provided in Logtalk for defining and using components?

Logtalk supports component-based programming (since its inception on January 1998), by using *categories* (which are first-class entities like objects and protocols). Logtalk categories can be used with both classes and prototypes and are inspired by the Smalltalk-80 (documentation-only) concept of method categories and by Objective-C categories, hence the name. For more information, please consult the [Categories](#) section and the examples provided with the distribution.

4.5.6 What support is provided in Logtalk for reflective programming?

Logtalk supports meta-classes, behavioral reflection through the use of event-driven programming, and structural reflection through the use of a set of built-in predicates and built-in methods which allow us to query the system about existing entities, entity relations, and entity predicates.

4.6 Troubleshooting

- *Using compiler options on calls to the Logtalk compiling and loading predicates does not work!*
- *Gecko-based browsers (e.g., Firefox) show non-rendered HTML entities when browsing XML documenting files!*
- *Compiling a source file results in errors or warnings but the Logtalk compiler reports a successful compilation with zero errors and zero warnings!*

4.6.1 Using compiler options on calls to the Logtalk *logtalk_compile/2* and *logtalk_load/2* built-in predicates does not work!

Using compiler options on calls to the Logtalk *logtalk_compile/2* and *logtalk_load/2* built-in predicates only apply to the file being compiled. If the first argument is a *loader file*, the compiler options will only be used in the compilation of the loader file itself, not in the compilation of the files loaded by the loader file. The solution is to edit the loader file and add the compiler options to the calls that compile/load the individual files.

4.6.2 Gecko-based browsers (e.g., Firefox) show non-rendered HTML entities when browsing XML documenting files!

Using Gecko-based browsers (e.g., Firefox) show non-rendered HTML entities (e.g., –) when browsing XML documenting files after running the `1gt2xml` shell script in the directory containing the XML documenting files. This is a consequence of the lack of support for the `disable-output-escaping` attribute in the browser XSLT processor. The workaround is to use another browser (e.g., Safari or Opera) or to use instead the `1gt2html` shell script in the directory containing the XML documenting files to convert them to (X)HTML files for browsing.

4.6.3 Compiling a source file results in errors or warnings but the Logtalk compiler reports a successful compilation with zero errors and zero warnings!

This may happen when your Prolog compiler implementation of the ISO Prolog standard `write_canonical/2` built-in predicate is buggy and writes terms that cannot be read back when consulting the intermediate Prolog files generated by the Logtalk compiler. Often, syntax errors found when consulting result in error messages but not in exceptions as the Prolog compiler tries to continue the compilation despite the problems found. As the Logtalk compiler relies on the exception mechanisms to catch compilation problems, it may report zero errors and zero warnings despite the error messages. Send a bug report to the Prolog compiler developers asking them to fix the `write_canonical/2` buggy implementation.

4.7 Usability

- *Is there a shortcut for compiling and loading source files?*
- *Is there an equivalent directive to the `ensure_loaded/1` Prolog directive?*
- *Are there shortcuts for the `make` functionality?*

4.7.1 Is there a shortcut for compiling and loading source files?

Yes. With most backend Prolog compilers, you can use `{File}` as a shortcut for `logtalk_load(File)`. For compiling and loading multiple files simply use `{File1, File2, ...}`. See the documentation of the *logtalk_load/1* predicate for details.

4.7.2 Is there an equivalent directive to the `ensure_loaded/1` Prolog directive?

You can use the goal `logtalk_load(File, [reload(skip)])` to ensure that `File` it is only loaded once. See the documentation of the [logtalk_load/2](#) predicate for details.

4.7.3 Are there shortcuts for the `make` functionality?

Yes. With most backend Prolog compilers, you can use `{*}` as a shortcut for `logtalk_make(all)` to reload all files modified since last compiled and loaded, `{!}` as a shortcut for `logtalk_make(clean)` to delete all intermediate Prolog files generated by the compilation of Logtalk source files, `{?}` as a shortcut for `logtalk_make(missing)` to list missing entities and predicates, and `{@}` as a shortcut for `logtalk_make(circular)` to list circular references. See the documentation of the [logtalk_make/1](#) predicate for details.

4.8 Deployment

- *Can I create standalone applications with Logtalk?*

4.8.1 Can I create standalone applications with Logtalk?

It depends on the Prolog compiler that you use to run Logtalk. Assuming that your Prolog compiler supports the creation of standalone executables, your application must include the adapter file for your compiler and the Logtalk compiler and runtime. The distribution includes embedding scripts for selected backend Prolog compilers and embedding examples.

For instructions on how to embed Logtalk and Logtalk applications, see the [embedding guide](#).

4.9 Performance

- *Is Logtalk implemented as a meta-interpreter?*
- *What kind of code Logtalk generates when compiling objects? Dynamic code? Static code?*
- *How about message-sending performance? Does Logtalk use static binding or dynamic binding?*
- *How does Logtalk performance compare with plain Prolog and with Prolog modules?*

4.9.1 Is Logtalk implemented as a meta-interpreter?

No. Objects and their encapsulated predicates are compiled, not meta-interpreted. In particular, inheritance relations are pre-compiled for improved performance. Moreover, no meta-interpreter is used even for objects compiled in debug mode.

4.9.2 What kind of code Logtalk generates when compiling objects? Dynamic code? Static code?

Static objects are compiled to static code. Static objects containing dynamic predicates are also compiled to static code, except, of course, for the dynamic predicates themselves. Dynamic objects are necessarily compiled to dynamic code. As in Prolog programming, for best performance, dynamic object predicates and dynamic objects should only be used when truly needed.

4.9.3 How about message-sending performance? Does Logtalk use static binding or dynamic binding?

Logtalk supports both static binding and dynamic binding. When static binding is not possible, Logtalk uses dynamic binding coupled with a caching mechanism that avoids repeated lookups of predicate declarations and predicate definitions. This is a solution common to other programming languages supporting dynamic binding. Message lookups are automatically cached the first time a message is sent. Cache entries are automatically removed when loading entities or using Logtalk dynamic features that invalidate the cached lookups. Whenever static binding is used, message-sending performance is essentially the same as a predicate call in plain Prolog. Performance of dynamic binding when lookups are cached is close to the performance that would be achieved with static binding. See the User Manual section on [performance](#) for more details.

4.9.4 Which Prolog-dependent factors are most crucial for good Logtalk performance?

Logtalk compiles objects assuming first-argument indexing for static code. First-argument indexing of dynamic code, when available, helps to improve performance due to the automatic caching of method lookups and the necessary use of book-keeping tables by the runtime engine (this is specially important when using event-driven programming). Dynamic objects and static objects containing dynamic predicates also benefit from first-argument indexing of dynamic predicates. The availability of multi-argument indexing, notably for dynamic predicates, also benefits dynamic binding performance.

4.9.5 How does Logtalk performance compare with plain Prolog and with Prolog modules?

Plain Prolog, Prolog modules, and Logtalk objects provide different trade-offs between performance and features. In general, for a given predicate definition, the best performance will be attained using plain Prolog, second will be Prolog modules (assuming no explicitly qualified calls are used), and finally Logtalk objects. Whenever static binding is used, the performance of Logtalk is equal or close to that of plain Prolog (depending on the Prolog virtual machine implementation and compiler optimizations). See the [simple benchmark test results](#) using some popular Prolog compilers.

4.10 Licensing

- *What's the Logtalk distribution license?*
- *Can Logtalk be used in commercial applications?*
- *What's the final license for a combination of Logtalk with a Prolog compiler?*

4.10.1 What's the Logtalk distribution license?

Logtalk follows the [Apache License 2.0](#).

4.10.2 Can Logtalk be used in commercial applications?

Yes, the Apache License 2.0 allows commercial use. See e.g. the [Apache License and Distribution FAQ](#).

4.10.3 What's the final license for a combination of Logtalk with a Prolog compiler?

See the [licensing guide](#) for details and relevant resources.

4.11 Support

- *Are there professional consulting, training and supporting services?*

4.11.1 Are there professional consulting, training and supporting services?

Yes. Please visit logtalk.pt for professional consulting, development, training, and other support services.

DEVELOPER TOOLS

The documentation of each developer tool can also be found in the tool directory in the `NOTES.md` file.

5.1 Overview

The following developer tools are available, each one with its own `loader.lgt` loader file (except for the built-in linter and make tools, which are integrated with the compiler/runtime) and `NOTES.md` documentation files:

- `asdf`
- `assertions`
- `code_metrics`
- `dead_code_scanner`
- `debug_messages`
- `debugger`
- `diagrams`
- `doclet`
- `help`
- `issue_creator`
- `lgtdoc`
- `lgtunit`
- `linter`
- `linter_reporter`
- `make`
- `mutation_testing`
- `packs`
- `ports_profiler`
- `profiler`
- `sarif`
- `sbom`

- `tool_diagnostics`
- `tutor`
- `wrapper`

5.1.1 Loading the developer tools

To load the main developer tools, use the following goal:

```
| ?- logtalk_load(tools(loader)).
```

The `linter_reporter` and `sarif` tools are not loaded by default as they are mostly useful for CI/CD pipelines.

The `ports_profiler` tool is not loaded by default; however, as it conflicts with the `debugger` tool as both provide a debug handler that must be unique in a running session.

The `profiler` tool is also not loaded by default, as it provides integration with selected backend Prolog compiler profilers that are not portable.

The `tutor` tool is also not loaded by default, given its usefulness mainly for new users that need help understanding compiler warning and error messages.

The `wrapper` tool is also not loaded by default, given its specialized purpose and beta status.

The `sbom` tool is also not loaded by default, given its specialized purpose.

To load a specific tool, either change your Prolog working directory to the tool folder and then compile and load the corresponding loader utility file or simply use library notation as argument for the compiling and loading predicates. For example:

```
| ?- logtalk_load(lgtunit(loader)).
```

5.1.2 Tools documentation

Specific notes about each tool can be found in the corresponding `NOTES.md` files. HTML documentation for each tool API can be found in the `docs` directory (open the `../docs/handbook/index.html` file with your web browser).

5.1.3 Tools common flags

The `lgtdoc`, `lgtunit`, and `issue_creator` tools share a `suppress_path_prefix` flag that can be used to suppress a prefix when printing file paths. For example (after loading the tools):

```
| ?- set_logtalk_flag(suppress_path_prefix, '/home/jdoe/').
```

This flag is mainly used to avoid user specific path prefixes appearing in documentation, test logs, and bug reports.

5.1.4 Tools requirements

Some of the developer tools have third-party dependencies. For example, the lgtdoc tool depends on XSLT processors to generate documentation final formats and uses Sphinx for the preferred HTML final format. Be sure to consult the tools documentation details on those requirements and possible alternatives. For convenience, follows a global list of the main tool requirements and suggestions for installing them per operating-system. If your operating-system or a dependency for it is not listed, see the dependency websites for installation instructions.

Tool dependencies for full functionality

- code_metric: cloc, sed
- diagrams: d2, Graphviz
- help: lynx info
- issue_creator: gh, glab
- lgtdoc: Sphinx, libxslt, fop, texlive, texinfo, pandoc, xsltproc
- lgtunit: Allure, coreutils, gsed
- packs: coreutils, libarchive, gnupg2, git, curl, wget, direnv

Python dependencies (all operating-systems)

```
$ pip install --upgrade pygments
$ pip install --upgrade sphinx
$ pip install --upgrade sphinx_rtd_theme
```

In alternative, if your want to use the same versions used to build release documentation:

```
$ cd manuals/sources
$ python3 -m pip install -r requirements.txt
```

macOS - MacPorts

```
$ sudo port install cloc
$ sudo port install d2 graphviz
$ sudo port install lynx texinfo pandoc
$ sudo port install libxslt fop texlive
$ sudo port install gsed
$ sudo port install coreutils gsed libarchive gnupg2 git curl wget direnv
$ sudo port install gh glab
```

macOS - Homebrew

```
$ brew install cloc
$ brew install d2 graphviz
$ brew install lynx texinfo pandoc
$ brew install libxslt fop texlive
$ brew install allure gsed
$ brew install coreutils libarchive gnupg2 git curl wget direnv
$ brew install gh glab
```

Ubuntu

```
$ sudo apt install cloc
$ sudo apt install graphviz
$ sudo apt install lynx texinfo pandoc
$ sudo apt install xsltproc fop texlive-latex-extra tex-gyre latexmk
$ sudo apt install coreutils libarchive-tools gnupg2 git curl wget direnv
$ sudo apt install gh
```

RedHat

```
$ sudo dnf install cloc
$ sudo dnf install graphviz
$ sudo dnf install lynx texinfo pandoc
$ sudo dnf install libxslt fop
$ sudo dnf install bsdtar gnupg2 git curl wget direnv
```

Windows - Chocolatey

```
> choco install cloc sed
> choco install pandoc xsltproc
> choco install d2 graphviz
> choco install xsltproc apache-fop texlive
> choco install gnupg git
> choco install gh glab
> choco install wget
```

Installers

- <https://d2lang.com/tour/install>
- <https://www.graphviz.org/download/>
- <https://allurereport.org/docs/>
- <https://www.gnupg.org/>
- <https://gitforwindows.org>
- <https://cli.github.com>

- <https://glab.readthedocs.io>
- <https://eternallybored.org/misc/wget/>

Windows - PowerShell add-ons

```
PS> Install-Module -Name Set-PsEnv
```

5.2 asdf

A Logtalk plugin for the `asdf` extendable version manager is available at:

<https://github.com/LogtalkDotOrg/asdf-logtalk>

This plugin provides an alternative to the `logtalk_version_select` script that can be useful to manage Logtalk versions when developing solutions that use other languages and tools that can also be handled by `asdf`.

5.3 assertions

The `assertions.lgt` file contains definitions for two meta-predicates, `assertion/1-2`, which allows the use of assertions on your source code to print warning and error messages (using the message printing mechanism). The `assertions_messages.lgt` file defines the default message translations generated when assertions succeed, fail, or throw an exception.

5.3.1 API documentation

This tool API documentation is available at:

../apis/library_index.html#assertions

5.3.2 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(assertions(loader)).
```

5.3.3 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(assertions(tester)).
```

5.3.4 Adding assertions to your source code

The `assertion/1` predicate takes a goal as argument. For example:

```
foo(L) :-
    assertions::assertion(non_empty_list(L)),
    ...
```

The `assertion/2` predicate takes as arguments a term for passing context information and a goal. Using again a unit test as an example:

```
foo(L) :-
    assertions::assertion(foo_list_alerts, non_empty_list(L)),
    ...
```

When using a large number of assertions, you can use a lighter syntax by adding a `uses/2` directive. For example:

```
:- uses(assertions, [assertion/1, assertion/2]).
```

5.3.5 Automatically adding file and line context information to assertions

The `assertions/1` parametric object can be used as a hook object to automatically add file and line context information, represented by the term `file_lines(File, BeginLine-EndLine)`, to calls to the `assertion/1` predicate by goal-expanding it to calls to the `assertion/2` predicate (the expansion assumes that a `uses/2` directive is being used in the code that will be expanded to direct `assertion/1` calls to the `assertions` object). For example, assuming the file using assertions is named `source`, it would be compiled and loaded using the call:

```
logtalk_load(source, [hook(assertions(debug))])
```

5.3.6 Suppressing assertion calls from source code

The `assertions/1` parametric object can be used as a hook object to suppress calls to the `assertion/1-2` predicates using goal-expansion (the expansion assumes `assertions::assertion/1-2` messages). For example, assuming the file using assertions is named `source`, it would be compiled and loaded using the call:

```
logtalk_load(source, [hook(assertions(production))])
```

5.3.7 Redirecting assertion failure messages

By default, assertion failures and errors are printed to the standard output stream. These messages, however, can be intercepted by defining the `logtalk::message_hook/4` multifile predicate. For example:

```
:- category(redirect_assertions_messages).

:- multifile(logtalk::message_hook/4).
:- dynamic(logtalk::message_hook/4).

logtalk::message_hook(Message, error, assertions, _) :-
```

(continues on next page)

(continued from previous page)

```

writeq(my_log_file, Message), write(my_log_file, '.\n').

:- end_category.

```

5.3.8 Converting assertion failures into errors

If you want an assertion failure to result in a failure or a runtime error, you can intercept the assertion failure messages, optionally still printing them, and throw an error. For example:

```

:- category(assertions_failures_to_errors).

:- multifile(logtalk::message_hook/4).
:- dynamic(logtalk::message_hook/4).

logtalk::message_hook(Message, error, assertions, Tokens) :-
    % uncomment the next two lines to also print the default message
    % logtalk::message_prefix_stream(error, assertions, Prefix, Stream),
    % logtalk::print_message_tokens(Stream, Prefix, Tokens),
    throw(error(Message, _)).

:- end_category.

```

In alternative, if you want assertions to always trigger an exception, use instead the `lgtunit` tool `assertions/1-2` public predicates.

5.4 code_metrics

The purpose of this tool is to assess qualities of source code that may predict negative aspects such as entity coupling, cohesion, complexity, error-proneness, and overall maintainability. It is meant to be extensible via the addition of objects implementing new metrics.

This tool provides predicates for computing metrics for source files, entities, libraries, files, and directories. The actual availability of a particular predicate depends on the specific metric. A set of predicates prints, by default, the computed metric values to the standard output. A second set of predicates computes and returns a score (usually a compound term with the computed metric values as arguments).

5.4.1 API documentation

This tool API documentation is available at:

../apis/library_index.html#code-metrics

5.4.2 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(code_metrics(loader)).
```

5.4.3 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(code_metrics(tester)).
```

5.4.4 Available metrics

Currently, the following metrics are provided:

- Number of Clauses (`noc_metric`)
- Number of Rules (`nor_metric`)
- Unique Predicate Nodes (`upn_metric`)
- Cyclomatic Complexity (`cc_metric`)
- Depth of Inheritance (`dit_metric`)
- Efferent coupling, afferent coupling, instability, and abstractness (`coupling_metric`)
- Lack of cohesion of methods (`lcom_metric`)
- Weighted Methods per Class (`wmc_metric`)
- Response For a Class (`rfc_metric`)
- Cognitive complexity (`cogc_metric`)
- Documentation (`doc_metric`)
- Lines of code, comments, and blanks (`lines_metric`)
- Source code size (`size_metric`)
- Maintainability index (`mi_metric` and `mi_metric(Stroud)`)
- Halstead complexity (`halstead_metric` and `halstead_metric(Stroud)`)

A helper object, `code_metrics`, is also provided, allowing running all loaded individual metrics. For code coverage metrics, see the `lgtunit` tool documentation.

5.4.5 Coupling metrics

- Efferent coupling (C_e): Number of entities that an entity depends on. These include objects receiving messages from the entity, plus the implemented protocols, imported categories, and extended/instantiated/specialized objects.
- Afferent coupling (C_a): Number of entities that depend on an entity. For a protocol, the number of protocols that extend it, plus the number of objects and categories that implement it. For a category, the number of objects that import it. For an object, the number of categories and objects that send messages to it plus the number of objects that extend/instantiate/specialize it.

- **Instability:** Computed as $C_e / (C_e + C_a)$. Measures an entity resilience to change. Ranging from 0.0 to 1.0, with 0.0 indicating a maximally stable entity and 1.0 indicating a maximally unstable entity. Ideally, an entity is either maximally stable or maximally unstable.
- **Abstractness:** Computed as the ratio between the number of static predicates with scope directives without a local definition and the number of static predicates with scope directives. Measures the rigidity of an entity. Ranging from 0.0 to 1.0, with 0.0 indicating a fully concrete entity and 1.0 indicating a fully abstract entity.
- **Distance from main sequence:** Computed as $\text{abs}(A + I - 1)$. Measures how far an entity is from the idealized line $A + I = 1$ (the “main sequence”). Ranging from 0.0 to 1.0, with 0.0 indicating the entity sits exactly on the main sequence and 1.0 indicating maximum deviation. Entities in the “zone of pain” ($A \approx 0, I \approx 0$: concrete and stable) are difficult to change; entities in the “zone of uselessness” ($A \approx 1, I \approx 1$: abstract and unstable) are not used. Ideally, entities should have a distance close to 0.0.

The dependencies count includes direct entity relations plus predicate calls or dynamic updates to predicates in external objects or categories.

For more information on the interpretation of the coupling metric scores, see, e.g., the original paper by Robert Martin:

```
@inproceedings{citeulike:1579528,
  author = "Martin, Robert",
  booktitle = "Workshop Pragmatic and Theoretical Directions in Object-Oriented Software_
↳Metrics",
  citeulike-article-id = 1579528,
  citeulike-linkout-0 = "http://www.objectmentor.com/resources/articles/oodmetrc.pdf",
  keywords = "diplomarbeit",
  organization = "OOPSLA'94",
  posted-at = "2007-08-21 11:08:44",
  priority = 0,
  title = "OO Design Quality Metrics - An Analysis of Dependencies",
  url = "http://www.objectmentor.com/resources/articles/oodmetrc.pdf",
  year = 1994
}
```

The coupling metric was also influenced by the metrics rating system in Microsoft Visual Studio and aims to eventually emulate the functionality of a maintainability index score.

5.4.6 Lack of cohesion metric

The lack of cohesion of methods (LCOM4) metric is computed as the number of connected components in the undirected graph whose nodes are the locally defined (non-auxiliary) predicates of an object or category and whose edges represent direct internal calls between those predicates. The metric is reported as the compound term `lcom(Components, Predicates)` where `Components` is the number of connected components and `Predicates` is the total number of locally defined predicates.

An `lcom(1, N)` score indicates a fully cohesive entity: all predicates are reachable from any other predicate via internal calls. Higher `Components` values suggest the entity may benefit from being split into separate, more focused entities. Entities with no predicates or a single predicate always score `lcom(1, N)`.

Protocols are not scored by this metric (they never define predicate clauses). Only direct bare-predicate calls (as reported by the Logtalk reflection API) are counted as internal edges; external message-sends (`Object::Predicate`) are excluded.

For more details on the LCOM4 variant, see:

```
@inproceedings{Hitz:1995,  
  author = "Hitz, M. and Montazeri, B.",  
  title = "Measuring Coupling and Cohesion in Object-Oriented Systems",  
  booktitle = "Proceedings of the International Symposium on Applied Corporate Computing",  
  year = 1995  
}
```

5.4.7 WMC metric

The Weighted Methods per Class (WMC) metric counts the number of locally defined (non-auxiliary) predicates in an object or category. Unit weights are used: each predicate contributes 1 to the score regardless of its complexity. Protocols are not scored as they cannot define predicates.

A cyclomatic-complexity-weighted variant is not provided. The Logtalk reflection API exposes inter-predicate call edges but not intra-clause branching structure (if-then-else, disjunction). A CC-weighted WMC would inherit the same approximation limitations already present in `cc_metric` and would also be numerically equivalent to the number of user clauses reported by `noc_metric`.

For more details on the WMC metric, see the original CK paper:

```
@article{Chidamber:1994,  
  author = "Chidamber, S. R. and Kemerer, C. F.",  
  title = "A Metrics Suite for Object Oriented Design",  
  journal = "IEEE Transactions on Software Engineering",  
  volume = 20,  
  number = 6,  
  pages = "476--493",  
  year = 1994  
}
```

5.4.8 RFC metric

The Response For a Class (RFC) metric is computed as the size of the response set for an object or category. The response set $RS(E)$ for an entity E is the union of its locally defined (non-auxiliary) predicates and all distinct predicates directly called or updated by those predicates. $RFC = |RS(E)|$.

Internal calls (to predicates within the entity) are included in the response set but are not double-counted: they are already part of the locally defined predicates. Only additional external callees increase the score beyond the WMC score of the same entity.

When the receiver of a message is only bound at runtime (e.g. `Obj::foo/1` where `Obj` is a variable), all such calls for the same predicate indicator are counted as one callee entry, regardless of how many distinct call sites exist.

Calls to built-in predicates are never returned by the Logtalk reflection API and therefore do not contribute to the score.

Protocols are not scored as they cannot define predicates.

Higher RFC scores indicate entities with larger potential execution footprints, which increases testing effort.

For more details on the RFC metric, see the original CK paper:

```
@article{Chidamber:1994,
  author = "Chidamber, S. R. and Kemerer, C. F.",
  title = "A Metrics Suite for Object Oriented Design",
  journal = "IEEE Transactions on Software Engineering",
  volume = 20,
  number = 6,
  pages = "476--493",
  year = 1994
}
```

5.4.9 Cognitive complexity metric

The cognitive complexity (`cogc_metric`) metric is an approximation of the Cognitive Complexity metric proposed by Campbell (2018). For each non-auxiliary predicate, the score contribution is the number of extra clauses (number of clauses minus one, representing multi-clause branching choices) plus one if the predicate is directly recursive (i.e. calls itself within the same entity). The entity score is the sum of all predicate contributions.

A score of 0 means all predicates are single-clause and non-recursive. Higher scores indicate entities with more branching and recursion, which typically require more effort to understand and test.

Approximation note: the Logtalk reflection API does not expose the use of control constructs such as cuts, if-then-else, disjunction, or catch/3 in the predicate clause bodies. Thus, these constructs cannot contribute to the score, which should be interpreted as a lower bound for the cognitive complexity of the entity.

Protocols are not scored as they cannot define predicates.

5.4.10 Lines metric

The lines metric (`lines_metric`) computes the number of code lines, comment lines, and blank lines, represented as the compound term `lines(Code, Comments, Blanks)`.

The implementation calls the external `cloc` tool and parses its CSV report. For entity scores, it first obtains the entity source line range from the reflection API using the `lines(BeginLine, EndLine)` property, writes only that range to a temporary file, and runs `cloc` on that temporary file.

For file, directory, and library scores, the metric computes and sums per-file `lines/3` scores.

This metric requires the `cloc` and `sed` commands to be available in the operating system path.

5.4.11 Maintainability index metric

The maintainability index metric (`mi_metric`) computes the original maintainability index formula:

$$MI = 171 - 5.2 * \ln(V) - 0.23 * C - 16.2 * \ln(L)$$

where V is the Halstead volume, C is cyclomatic complexity, and L is the number of code lines. The score is represented as `mi(Index)`.

The metric is available as `mi_metric` (default Stroud coefficient 18) and as `mi_metric(Stroud)` for a user-defined Stroud coefficient passed to the Halstead metric.

Protocols are not scored as they cannot define predicates.

This metric requires the lines metric and thus the `cloc` and `sed` commands to be available in the operating system path.

5.4.12 Halstead metric

Predicates declared, user-defined, and called are interpreted as *operators*. Built-in predicates and built-in control constructs are ignored. Predicate arguments are abstracted, assumed distinct, and interpreted as *operands*. Note that this definition of operands is a significant deviation from the original definition, which used syntactic literals. A computation closer to the original definition of the metric would require switching to use the parser to collect information on syntactic literals, which would imply a much larger computation cost. The number of predicate calls doesn't include calls to built-in predicates and can underestimate recursive calls.

The computation of this metric is parameterized by the *Stroud* coefficient for computing the time required to program (default is 18). The following individual measures are computed:

- Number of distinct predicates (declared, defined, called, or updated; P_n).
- Number of predicate arguments (assumed distinct; PAn).
- Number of predicate calls/updates + number of clauses (C_n).
- Number of predicate call/update arguments + number of clause head arguments (CAn).
- Entity vocabulary (EV). Computed as $EV = P_n + PAn$.
- Entity length (EL). Computed as $EL = C_n + CAn$.
- Volume (V). Computed as $V = EL * \log_2(EV)$.
- Difficulty (D). Computed as $D = (P_n/2) * (CAn/P_n)$.
- Effort (E). Computed as $E = D * V$.
- Time required to program (T). Computed as $T = E/k$ seconds (where k is the Stroud number; defaults to 18).
- Number of delivered bugs (B). Computed as $B = V/3000$.

5.4.13 UPN metric

The Unique Predicate Nodes (UPN) metric is described in the following paper:

```
@article{MOORES199845,
  title = "Applying Complexity Measures to Rule-Based Prolog Programs",
  journal = "Journal of Systems and Software",
  volume = "44",
  number = "1",
  pages = "45 - 52",
  year = "1998",
  issn = "0164-1212",
  doi = "https://doi.org/10.1016/S0164-1212(98)10042-0",
  url = "http://www.sciencedirect.com/science/article/pii/S0164121298100420",
  author = "Trevor T Moores"
}
```

The nodes include called and updated predicates independently of where they are defined. It also includes multifile predicates contributed to other entities.

5.4.14 Cyclomatic complexity metric

The cyclomatic complexity metric evaluates an entity code complexity by measuring the number of linearly independent paths through the code. In its current implementation, all defined predicates that are not called or updated are counted as graph-connected components (the reasoning being that these predicates can be considered entry points). The implementation uses the same predicate abstraction as the UPN metric. The defined predicates include multifile predicate definitions contributed by the entity to other entities.

For more details on this metric, see the original paper by Thomas J. McCabe:

```
@inproceedings{McCabe:1976:CM:800253.807712,
  author = "McCabe, Thomas J.",
  title = "A Complexity Measure",
  booktitle = "Proceedings of the 2Nd International Conference on Software Engineering",
  series = "ICSE '76",
  year = 1976,
  location = "San Francisco, California, USA",
  pages = "407--",
  url = "http://dl.acm.org/citation.cfm?id=800253.807712",
  acmid = 807712,
  publisher = "IEEE Computer Society Press",
  address = "Los Alamitos, CA, USA",
  keywords = "Basis, Complexity measure, Control flow, Decomposition, Graph theory,
  ↪Independence, Linear, Modularization, Programming, Reduction, Software, Testing",
}
```

5.4.15 Usage

All metrics require the source code to be analyzed to be loaded with the `source_data` flag turned on. For usage examples, see the `SCRIPT.txt` file in the tool directory.

Be sure to fully understand the metrics individual meanings and any implementation limitations before using them to support any evaluation or decision process.

5.4.16 Excluding code from analysis

A set of options is available to specify code that should be excluded when applying code metrics:

- `exclude_directories(Directories)`
list of directories to exclude (default is `[]`); all sub-directories of the excluded directories are also excluded; directories may be listed by full or relative path
- `exclude_files(Files)`
list of source files to exclude (default is `[]`); files may be listed by full path or basename, with or without extension
- `exclude_libraries(Libraries)`
list of libraries to exclude (default is `[startup, scratch_directory]`)
- `exclude_entities(Entities)`
list of entities to exclude (default is `[]`)

5.4.17 Defining new metrics

New metrics can be implemented by defining an object that imports the `code_metric` category and implements its score predicates. There is also a `code_metrics_utilities` category that defines useful predicates for the definition of metrics.

5.4.18 Third-party tools

The following open-source command-line programs can count blank lines, comment lines, and lines of source code in many programming languages, including Logtalk:

- `cloc` - <https://github.com/AlDanial/cloc>
- `ohcount` - <https://github.com/blackducksoftware/ohcount>
- `tokei` - <https://github.com/XAMPPRocky/tokei>

5.4.19 Applying metrics to Prolog modules

Some of the metrics can also be applied to Prolog modules that Logtalk is able to compile as objects. For example, if the Prolog module file is named `module.pl`, try:

```
| ?- logtalk_load(module, [source_data(on)]).
```

Due to the lack of standardization of module systems and the abundance of proprietary extensions, this solution is not expected to work for all cases.

5.4.20 Applying metrics to plain Prolog code

Some of the metrics can also be applied to plain Prolog code. For example, if the Prolog file is named `code.pl`, simply define an object including its code:

```
:- object(code).  
    :- include('code.pl').  
:- end_object.
```

Save the object to an e.g. `code.lgt` file in the same directory as the Prolog file and then load it in debug mode:

```
| ?- logtalk_load(code, [source_data(on)]).
```

In alternative, use the `object_wrapper_hook` provided by the `hook_objects` library:

```
| ?- logtalk_load(hook_objects(loader)).  
...  
  
| ?- logtalk_load(code, [hook(object_wrapper_hook), source_data(on)]).
```

With either wrapping solution, pay special attention to any compilation warnings that may signal issues that could prevent the plain Prolog code from working when wrapped by an object.

5.5 dead_code_scanner

This tool detects *likely* dead code in Logtalk entities and in Prolog modules compiled as objects. Predicates (and non-terminals) are classified as dead code when:

- There is no scope directive for them, and they are not called, directly or indirectly, by any predicate with a (local or inherited) scope directive.
- They are listed in `uses/2` and `use_module/2` directives but not called.

Besides dead code, this tool can also help detect other problems in the code that often result in reporting false positives. For example, typos in `alias/2` directives, missing scope directives, and missing `meta_non_terminal/1` and `meta_predicate/1` directives.

Given the possibility of false positives, care must be taken before deleting reported dead code to ensure that it's, in fact, code that is not used. A common cause of false positives is the use of conditional compilation directives to provide implementations for predicates missing in some systems or different predicate implementations per operating-system.

The `dead_code_scanner.lgt` source file implements the scanning predicates for finding dead code in entities, libraries, and directories. The source file `dead_code_scanner_messages.lgt` defines the default translations for the messages printed when scanning for dead code. These messages can be intercepted to customize the output, e.g., to make it less verbose or for integration with, e.g., GUI IDEs and continuous integration servers.

5.5.1 API documentation

This tool API documentation is available at:

../apis/library_index.html#dead-code-scanner

For sample queries, please see the `SCRIPT.txt` file in the tool directory.

5.5.2 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(dead_code_scanner(loader)).
```

5.5.3 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(dead_code_scanner(tester)).
```

5.5.4 Usage

This tool provides a set of predicates that allows scanning entities, libraries, files, and directories. See the tool API documentation for details. The source code to be analyzed should be loaded with the `source_data` and `optimize` flags turned on (possibly set from a loader file).

For machine-readable integrations, the `diagnostics/2-3`, `diagnostic/2-3`, `diagnostics_summary/2-3`, and `diagnostics_preflight/2-3` predicates can be used to collect dead-code results and post-filter counts without relying on printed messages. Diagnostics are returned as terms of the form:

- `diagnostic(RuleId, Severity, Confidence, Message, Context, File, Lines, Properties)`

where:

- `RuleId` distinguishes local dead code from unused `uses/2` and `use_module/2` resources
- `Severity` is the reported diagnostic severity (error, warning, or note)
- `Confidence` is a triage-oriented rating (high, medium, or low)
- `Context` identifies the entity kind and entity where the diagnostic was found
- `Properties` contains additional structured metadata, including the analysis settings seen for the source file and any class-specific details needed to explain the finding

For machine-readable export integrations, the `export/3-4` predicates serialize scan results in SARIF format using the `json` library.

As an example, assume that we want to scan an application with a library alias `my_app`. The following goals could be used:

```
| ?- set_logtalk_flag(source_data, on),  
    set_logtalk_flag(optimize, on).  
yes  
  
| ?- logtalk_load(my_app(loader)).  
...  
yes  
  
| ?- dead_code_scanner::library(my_app).  
...
```

For complex applications that make use of sub-libraries, there are also `rlibrary/1-2` predicates that perform a recursive scan of a library and all its sub-libraries. Conversely, we may be interested in scanning a single entity:

```
| ?- dead_code_scanner::entity(some_object).  
...
```

For other usage examples, see the `SCRIPT.txt` file in the tool directory.

5.5.5 Excluding code from analysis

A set of options is available to specify code that should be excluded when looking for unused predicates (and non-terminals):

- `exclude_directories(Directories)`
list of directories to exclude (default is `[]`); all sub-directories of the excluded directories are also excluded; directories may be listed by full or relative path
- `exclude_files(Files)`
list of source files to exclude (default is `[]`); files may be listed by full path or basename, with or without extension
- `exclude_libraries(Libraries)`
list of libraries to exclude (default is `[startup, scratch_directory]`)
- `exclude_entities(Entities)`
list of entities to exclude (default is `[]`)
- `exclude_predicates(Predicates)`
list of predicate and non-terminal indicators to exclude from findings (default is `[]`); supports local indicators such as `foo/1` and `bar//2` plus qualified indicators such as `object::baz/2` and `module:qux/3`
- `waive_findings(Findings)`
list of finding patterns to suppress from the results (default is `[]`); intended for reusable CI baselines and supports partially instantiated structured finding terms such as `dead_predicate(local_dead_code, medium, _, object, some_object, helper/2, _, _)`

5.5.6 Machine-readable summaries

The `diagnostics_summary/2-3` predicates return a term of the form:

- `diagnostics_summary(Target, TotalContexts, TotalDiagnostics, Breakdown, ContextSummaries)`

where `Breakdown` is a term of the form:

- `diagnostic_breakdown(RuleCounts, SeverityCounts, ConfidenceCounts)`

with `RuleCounts` containing `rule_count(RuleId, Count)` terms, `SeverityCounts` containing `severity_count(Severity, Count)` terms, and `ConfidenceCounts` containing `confidence_count(Confidence, Count)` terms.

`ContextSummaries` is a list of terms of the form:

- `context_summary(Context, DiagnosticsCount, Breakdown)`

The summary is computed after applying all exclusions and finding waivers, making it suitable for CI thresholds and regression tracking.

5.5.7 Machine-readable preflight warnings

The `diagnostics_preflight/2-3` predicates return an ordered set of warnings describing analysis prerequisites that affect result quality for a target. Warning terms currently use the form:

- `preflight_issue(missing_analysis_prerequisite, Severity, Message, Context, File, Lines, Properties)`

where `Prerequisite` is currently one of:

- `source_data`
- `optimize`

This API is intended for CI and editor integrations that need to surface analysis quality issues without scraping console messages.

5.5.8 Finding classes and confidence

The current finding classes are:

- `local_dead_code`
- `unused_uses_resource`
- `unused_use_module_resource`

Confidence is a triage-oriented rating attached to every finding. The JSON schema allows the values `high`, `medium`, and `low`, but the tool currently emits only `high` and `medium`:

- `high`
Used when the tool has stronger evidence that the reported code is really unused.
- `medium`
Used when the result is still useful, but can be affected by missing compilation information such as optimized call graph data.
- `low`
Reserved for future use. It is accepted by the export schema but is not currently generated by the scanner.

For `local_dead_code`, the confidence depends on the effective compilation mode of the analyzed file:

- If the file was loaded in normal mode, or otherwise not loaded with `optimize(on)`, the confidence is `medium`.
- If the file was loaded with `optimize(on)`, the confidence is upgraded to `high` because the compiler can provide more complete call graph information, reducing the risk of false positives caused by missed meta-calls.

The effective mode is inferred from the loaded file itself, not from the current scanner defaults. Therefore, results reflect how the analyzed code was actually compiled.

Unused `uses/2` and `use_module/2` resources are always reported with high confidence because they are derived from the generated linking clauses and the recorded call graph rather than from heuristic reachability alone.

When scanning using the text-based predicates such as `entity/1-2`, `file/1-2`, and `all/0-1`, the tool now emits preflight warnings whenever the analyzed files were not loaded with `source_data(on)` or `optimize(on)` so that potentially less reliable results can be triaged more carefully.

5.5.9 SARIF reports

This tool is a diagnostics producer. To generate SARIF reports from its diagnostics, load the standalone `sarif` tool and call `sarif::generate(dead_code_scanner, Target, Sink, Options)` where `Target` and `Options` are the same ones accepted by the diagnostics predicates.

The generated report uses the SARIF 2.1.0 JSON format and is suitable for code scanning integrations that consume SARIF reports. The shared generator includes deterministic result fingerprints derived from the canonical finding and preserves the rule descriptors and triage-oriented severity mapping described below.

SARIF findings are exported using one rule descriptor per finding class:

- `local_dead_code`
- `unused_uses_resource`
- `unused_use_module_resource`

Result severities are also triage-aware instead of using a single static SARIF level for all findings:

- `local_dead_code` with high confidence is exported as warning
- `local_dead_code` with medium or low confidence is exported as note
- `unused_uses_resource` with high confidence is exported as error
- `unused_use_module_resource` with high confidence is exported as error

This mapping allows CI consumers to distinguish advisory dead-code findings from stronger unused-resource findings without reinterpreting the custom class and confidence properties.

SARIF reports include the same preflight information as invocation-level `toolExecutionNotifications`, allowing consumers to distinguish prerequisite warnings from dead code findings. When the analyzed code is inside a git repository, the shared generator also includes the current branch and commit hash as optional run properties and a `versionControlProvenance` entry when the repository URI can be derived.

5.5.10 Integration with the make tool

The `loader.lgt` file sets a make target action that will call the `dead_code_scanner::all` goal whenever the `logtalk_make(check)` goal (or its top-level abbreviation, `{?}`) is called.

5.5.11 Caveats

Use of local meta-calls with goal arguments only known at runtime can result in false positives. When using library or user-defined meta-predicates, compilation of the source files with the `optimize` flag turned on may allow meta-calls to be resolved at compile-time and thus allow calling information for the meta-arguments to be recorded, avoiding false positives for predicates that are only meta-called.

5.5.12 Scanning Prolog modules

This tool can also be applied to Prolog modules that Logtalk is able to compile as objects. For example, if the Prolog module file is named `module.pl`, try:

```
| ?- logtalk_load(module, [source_data(on)]).
```

Due to the lack of standardization of module systems and the abundance of proprietary extensions, this solution is not expected to work for all cases.

5.5.13 Scanning plain Prolog files

This tool can also be applied to plain Prolog code. For example, if the Prolog file is named `code.pl`, simply define an object including its code:

```
:- object(code).  
    :- include('code.pl').  
:- end_object.
```

Save the object to an e.g. `code.lgt` file in the same directory as the Prolog file and then load it in debug mode:

```
| ?- logtalk_load(code, [source_data(on), optimize(on)]).
```

In alternative, use the `object_wrapper_hook` provided by the `hook_objects` library:

```
| ?- logtalk_load(hook_objects(loader)).  
...  
| ?- logtalk_load(code, [hook(object_wrapper_hook), source_data(on), optimize(on)]).
```

With either wrapping solution, pay special attention to any compilation warnings that may signal issues that could prevent the plain Prolog from being fully analyzed when wrapped by an object.

5.5.14 Known issues

Some tests fail when using the ECLIPSe and Trealla Prolog backends due to their implementation of the `multifile/1` directive, specifically when the `multifile` predicates are also dynamic with clauses defined in a source file: reconsulting the file adds the new clauses to the old ones instead of replacing them.

5.6 debug_messages

By default, `debug` and `debug(Group)` messages are only printed when the `debug` flag is turned on. These messages are also suppressed when compiling code with the `optimize` flag turned on. This tool supports selective enabling of `debug` and `debug(Group)` messages in normal and debug modes.

5.6.1 API documentation

This tool API documentation is available at:

../apis/library_index.html#debug-messages

For general information on debugging, open in a web browser the following link and consult the debugging section of the User Manual:

../handbook/userman/debugging.html

5.6.2 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(debug_messages(loader)).
```

5.6.3 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(debug_messages(tester)).
```

5.6.4 Usage

The tool provides two sets of predicates. The first set allows enabling and disabling of all debug and `debug(Group)` messages for a given component. The second set allows enabling and disabling of `debug(Group)` messages for a given group and component for fine-grained control.

Upon loading the tool, all debug messages are skipped. The user is then expected to use the tool API to selectively enable the messages that will be printed. As an example, consider the following object, part of a `xyz` component:

```
:- object(foo).

   :- public([bar/0, baz/0]).
   :- uses(logtalk, [print_message/3]).

   bar :-
       print_message(debug(bar), xyz, @'bar/0 called').

   baz :-
       print_message(debug(baz), xyz, @'baz/0 called').

:- end_object.
```

Assuming the object `foo` is compiled and loaded in normal or debug mode, after also loading this tool, `bar/0` and `baz/0` messages will not print any debug messages:

```
| ?- {debug_messages(loader), foo}.
...
yes
```

(continues on next page)

(continued from previous page)

```
| ?- foo::(bar, baz).  
yes
```

We can then enable all debug messages for the xyz component:

```
| ?- debug_messages::enable(xyx).  
yes  
  
| ?- foo::(bar, baz).  
bar/0 called  
baz/0 called  
yes
```

Or we can selectively enable only debug messages for a specific group:

```
| ?- debug_messages::disable(xyx).  
yes  
  
| ?- debug_messages::enable(xyx, bar).  
yes  
  
| ?- foo::(bar, baz).  
bar/0 called  
yes
```

5.7 debugger

This tool provides the default Logtalk command-line debugger. Unlike Prolog systems, the Logtalk debugger is a regular application, using a public API. As a consequence, it must be explicitly loaded by the programmer, either manually at the top-level interpreter or automatically from a settings file.

5.7.1 API documentation

This tool API documentation is available at:

[../apis/library_index.html#debugger](http://logtalk.org/..../apis/library_index.html#debugger)

5.7.2 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(debugger(loader)).
```

When the code to be debugged runs computationally expensive initializations, loading this tool after the code may have a noticeable impact on loading times.

5.7.3 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(debugger(tester)).
```

5.7.4 Usage

Debugging Logtalk source code (with this debugger) requires compiling source files using the `debug(on)` compiler flag. For example:

```
| ?- logtalk_load(my_buggy_code, [debug(on)]).
```

In alternative, you may also turn on the debug flag globally by typing:

```
| ?- set_logtalk_flag(debug, on).
```

But note that loader files may override this flag setting (e.g., by using `debug(off)` or `optimize(on)` options for loaded files). If that's the case, you will need to either edit the loader files or write customized loader files enabling debugging. For detailed information on using the debugger, consult the debugging section of the User Manual:

../handbook/userman/debugging.html

The `debugger_messages.lgt` source file defines the default debugger message translations.

The `dump_trace.lgt` provides a simple solution for dumping a goal trace to a file. For example:

```
| ?- dump_trace::start_redirect_to_file('trace.txt', some_goal),
    dump_trace::stop_redirect_to_file.
```

A full trace can also be obtained at the top-level by using the `S` (Skip) command at the call port for the top-level goal when tracing it.

5.7.5 Alternative debugger tools

Logtalk provides basic support for the SWI-Prolog graphical tracer. The **required** settings are described in the `samples/settings-sample.lgt` file. Logtalk queries can be traced using this tool by using the `gtrace/0-1` predicates. For example:

```
| ?- gtrace(foo::bar).
```

Or alternatively:

```
| ?- gtrace, foo::bar.
```

You can also use the `gspy/1` predicate to spy on a Logtalk predicate specified as `Entity::Functor/Arity` when using the graphical tracer. When using this tool, internal Logtalk compiler/runtime predicates and compiled predicates that resulted from the term-expansion mechanism may be exposed in some cases. This issue is shared with Prolog code and results from the non-availability of source code for the predicates being traced.

5.7.6 Known issues

Clause breakpoints require a Prolog backend compiler that supports accessing read term starting line but only some backends (B-Prolog, GNU Prolog, JIProlog, SICStus Prolog, SWI-Prolog, Trealla Prolog, XVM, and YAP) provide accurate line numbers. As a workaround, you can check the start line number for an entity predicate definition using a query such as:

```
| ?- object_property(Entity, defines(Functor/Arity, Properties)).
```

Check the returned `line_count/1` property to find if there's any offset to the source file number of the predicate clause that you want to trace. This issue, if present, usually only affects the first predicate clause.

Clause breakpoints are currently not available when using XSB as this backend doesn't provide line information.

Using the port command `p` (print) requires a backend supporting the user-defined `portray/1` hook predicate called via the `format/2-3` predicates `~p` control sequence.

5.8 diagrams

This tool generates *library*, *directory*, *file*, *entity*, and *predicate* diagrams for source files and for libraries of source files using the Logtalk reflection API to collect the relevant information and a graph language for representing the diagrams. Limited support is also available for generating diagrams for Prolog module applications. It's also possible in general to generate predicate cross-referencing diagrams for plain Prolog files.

Linking *library diagrams* to *entity diagrams* to *predicate cross-referencing diagrams* and linking *directory diagrams* to *file diagrams* is also supported when using SVG output. This feature allows using diagrams for understanding the architecture of applications by navigating complex code and zooming into details. SVG output can also easily link to both source code repositories and API documentation. This allows diagrams to be used for source code navigation.

Diagrams can also be used to uncover code issues. For example, comparing *loading diagrams* with *dependency diagrams* can reveal implicit dependencies. Loading diagrams can reveal circular dependencies that may warrant code refactoring. Entity diagrams can provide a good overview of code coupling. Predicate cross-referencing diagrams can be used to visually access entity code complexity, complementing the `code_metrics` tool.

All diagrams support a comprehensive set of options, discussed below, to customize the final contents and appearance.

Diagram generation can be easily automated using the `doclet` tool and the `logtalk_doclet` scripts. See the `doclet` tool examples and documentation for details. See also the diagrams tool own `lgt2svg` Bash and PowerShell scripts.

5.8.1 Requirements

Recent versions of d2 (0.6.9 or later) or Graphviz are required for generating diagrams in the final formats:

- <https://d2lang.com>
- <https://www.graphviz.org/>

These can be installed using their own installers or using operating-system package managers:

macOS - MacPorts

```
$ sudo port install d2 graphviz
```

macOS - Homebrew

```
$ brew install d2 graphviz
```

Windows - Chocolatey

```
> choco install d2 graphviz
```

Installers

- <https://d2lang.com/tour/install>
- <https://www.graphviz.org/download/>

5.8.2 API documentation

This tool API documentation is available at:

[../..../apis/library_index.html#diagrams](http://localhost:8080/..../apis/library_index.html#diagrams)

For sample queries, please see the SCRIPT.txt file in the tool directory.

5.8.3 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(diagrams(loader)).
```

5.8.4 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(diagrams(tester)).
```

5.8.5 Supported diagrams

The following entity diagrams are supported:

- *entity diagrams* showing entity public interfaces, entity inheritance relations, and entity predicate (or non-terminal) cross-reference relations
- *predicate/non-terminal cross-reference diagrams* (between entities or within an entity)
- *caller diagrams* showing direct and indirect callers of a given predicate (or non-terminal)
- *inheritance diagrams* showing entity inheritance relations
- *uses diagrams* showing which entities use resources from other entities

Entity nodes can optionally show a coupling metrics overlay.

The following library diagrams are supported:

- *library loading diagrams* showing which libraries load other libraries
- *library dependency diagrams* showing which libraries contain entities with references to entities defined in other libraries

The following file diagrams are supported:

- *file loading diagrams* showing which files load or include other files
- *file dependency diagrams* showing which files contain entities with references to entities defined in other files

File dependency diagrams are specially useful in revealing dependencies that are not represented in file loading diagrams due to files being loaded indirectly by files external to the libraries being documented.

The following directory diagrams are supported:

- *directory loading diagrams* showing which directories contain files that load files from other directories
- *directory dependency diagrams* showing which directories contain entities with references to entities defined in other directories

Comparing directory (or file) loading diagrams with directory (or file) dependency diagrams allows comparing what is explicitly loaded with the actual directory (or file) dependencies, which are inferred from the source code.

Library and directory dependency diagrams are specially useful for large applications where file diagrams would be too large and complex to be useful, notably when combined with the *zoom* option to link to, respectively, entity and file diagrams.

A utility object, `diagrams`, is provided for generating all supported diagrams in one step. This object provides an interface common to all diagrams, but note that some predicates that generate diagrams only make sense for some types of diagrams. For best results and fine-grained customization of each diagram, the individual diagram objects should be used with the intended set of options.

5.8.6 Graph elements

Limitations in both the graph languages and UML force the invention of a modeling language that can represent all kinds of Logtalk entities and entity relations. Currently we use the following Graphviz node shapes (libraries, entities, predicates, and files) and arrowheads (entity, predicate, and file relations):

- libraries
tab (lightsalmon)
- library loading and dependency relations
normal (arrow ending with a black triangle)
- objects (classes, instances, and prototypes)
box (rectangle, yellow for instances/classes and beige for prototypes)
- protocols
note (aqua marine rectangle with folded right-upper corners)
- categories
component (light cyan rectangle with two small rectangles intercepting the left side)
- modules
box (plum rectangle with small tab at top)
- public predicates
box (springgreen)
- public, multifile, predicates
box (skyblue)
- protected predicates
box (yellow)
- private predicates
box (indianred)
- external predicates
box (beige)
- exported module predicates
box (springgreen)
- directories
tab (lightsalmon)
- directory loading and dependency relations
normal (arrow ending with a black triangle)
- files
box (pale turquoise rectangle)
- file loading and dependency relations
normal (arrow ending with a black triangle)
- specialization relation
onormal (arrow ending with a white triangle)
- instantiation relation
normal (arrow ending with a black triangle)
- extends relation

vee (arrow ending with a “v”)

- implements relation
dot (arrow ending with a black circle)
- imports relation
box (arrow ending with a black square)
- complements relation
obox (arrow ending with a white square)
- uses and use module relations
rdiamond (arrow ending with a black half diamond)
- predicate calls
normal (arrow ending with a black triangle)
- dynamic predicate updates
diamond (arrow ending with a black diamond)

When using the d2 graph language, we use similar node shapes and arrowheads when available. As d2 evolves, we hope that these graph elements will converge further.

The library, directory, file, entity, and predicate nodes that are not part of the predicates, entities, files, or libraries for which we are generating a diagram use a dashed border, a darker color, and are described as external.

Note that all the elements above can have captions. See below the diagrams `node_type_captions/1` and `relation_labels/1` output options.

5.8.7 Supported graph languages

Currently, both DOT and d2 graph languages support all the features of the diagrams tool. There's also preliminary support for Mermaid (its current version lacks several required features for parity with d2 and Graphviz, notably support for links on edges).

JSON output format is also supported for use with Cytoscape. The Cytoscape backend generates `.cx2` files using the Cytoscape Exchange Format version 2 (CX2), including core aspects such as `attributeDeclarations`, `networkAttributes`, `nodes`, `edges`, `visualProperties`, and `status`. Node and edge styling is exported using CX2 visual property mappings based on node and edge attributes. The JSON backend includes optional metadata such as title, description, generation timestamp, and Logtalk/Prolog version information. A default HTML viewer, `cytoscapejs_viewer.html`, is provided for convenience. This viewer currently requires starting a local web server in the generated diagrams directory and manually entering the server address.

PlantUML output is also supported. The PlantUML backend generates `.puml` files using custom stereotypes with spotted characters to distinguish Logtalk entity kinds (O/blue for objects, P/green for protocols, T/orange for categories). All five Logtalk relationships (implements, imports, extends, instantiates, specializes, complements) are mapped to distinct labeled arrow styles. A shared `logtalk.iuml` include file defines all `skinparam` blocks, stereotype colors, and CSS styles for consistent rendering. PlantUML version 1.2024.0 or later is recommended. See <https://plantuml.com> for installation instructions.

The diagrams `.d2` and `.dot` files are created by default in the `'./dot_dias'` sub-directory of the current directory. These files can be easily converted into a printable format such as SVG, PDF, or Postscript. Sample helper scripts are provided for converting a directory of `.d2` or `.dot` files to `.svg` files:

- `lgt2svg.sh` for POSIX systems
- `lgt2svg.ps1` for Windows systems

The scripts assume that the d2 and Graphviz command-line executables are available from the system path. For Graphviz, the default is the dot executable, but the scripts accept a command-line option to select in alternative the circo, fdp, or neato executables). For d2, the default layout engine is elk, but the scripts accept a command-line option to select in alternative the dagre or tala layout engines.

The recommended output format is SVG, as it supports tooltips and URL links, which can be used for showing, e.g., entity types, relation types, file paths, and for navigating to files and directories of files (libraries) or to API documentation. See the relevant diagram options below in order to take advantage of these features (see the discussion below on “linking diagrams”).

To convert to formats other than SVG, you will need to use the d2 and Graphviz executables directly. For example, using the Graphviz dot executable, we can generate a PDF with the command:

```
dot -Tpdf diagram.dot > diagram.pdf
```

This usually works fine for entity and predicate call cross-referencing diagrams. For directory and file diagrams, the fdp and circo command-line executables may produce better results. For example:

```
fdp -Tsvg diagram.dot > diagram.svg
circo -Tsvg diagram.dot > diagram.svg
```

It’s also worth experimenting with different layouts to find the one that produces the best results (see the layout/1 option described below).

When generating diagrams for multiple libraries or directories, it’s possible to split a diagram with several disconnected library or directory graphs using the ccomps command-line executable. For example:

```
ccomps -x -o subdiagram.dot diagram.dot
```

For more information on the DOT language and related tools, see:

```
http://www.graphviz.org/
```

When using Windows, there are known issues with some Prolog compilers due to the internal representation of paths. If you encounter problems with a specific backend Prolog compiler, try, if possible, to use another supported backend Prolog compiler when generating diagrams.

For printing large diagrams, you will need to either use a tool to slice the diagram into page-sized pieces or, preferably, use software capable of tiled printing (e.g., Adobe Reader). You can also hand-edit the generated .dot files and play with settings such as aspect ratio for fine-tuning the diagrams layout.

5.8.8 Customization

This tool provides parametric objects implementing the different types of diagrams where the parameter is the export format: dot, d2, cx2, and puml. The non-parametric versions of the diagram objects default to dot.

A set of options is available to specify the details to include in the generated diagrams. For entity diagrams, the options are:

- layout(Layout)
diagram layout (one of the atoms {top_to_bottom,bottom_to_top,left_to_right,right_to_left}; default is bottom_to_top)
- title(Title)
diagram title (an atom; default is '')
- date(Boolean)

print current date and time (true or false; default is true)

- `versions(Boolean)`
print Logtalk and backend version data (true or false; default is false)
- `interface(Boolean)`
print public predicates (true or false; default is true)
- `file_labels(Boolean)`
print file labels (true or false; default is true)
- `file_extensions(Boolean)`
print file name extensions (true or false; default is true)
- `relation_labels(Boolean)`
print entity relation labels (true or false; default is true)
- `externals(Boolean)`
print external nodes (true or false; default is true)
- `node_type_captions(Boolean)`
print node type captions (true or false; default is true)
- `metrics_overlay(Boolean)`
print coupling metrics overlay in entity nodes showing afferent coupling (Ca), efferent coupling (Ce), instability (I), abstractness (A), and distance from main sequence (D) values (true or false; default is false)
- `cycle_detection(Boolean)`
highlight cyclic relations by coloring the edges red (true or false; default is false)
- `inheritance_relations(Boolean)`
print inheritance relations (true or false; default is true for entity inheritance diagrams and false for other entity diagrams)
- `provide_relations(Boolean)`
print provide relations (true or false; default is false)
- `xref_relations(Boolean)`
print predicate call cross-reference relations (true or false; default depends on the specific diagram)
- `xref_calls(Boolean)`
print predicate cross-reference calls (true or false; default depends on the specific diagram)
- `output_directory(Directory)`
directory for the .d2 and .dot files (an atom; default is `'./dot_dias'`)
- `exclude_directories(Directories)`
list of directories to exclude except as external nodes (default is `[]`); all sub-directories of the excluded directories are also excluded; directories may be listed by full or relative path
- `exclude_files(Files)`
list of source files to exclude except as external nodes (default is `[]`); files may be listed by full path or basename, with or without extension
- `exclude_libraries(Libraries)`
list of libraries to exclude except as external nodes (default is `[startup, scratch_directory]`)
- `exclude_entities(Entities)`
list of entities to exclude except as external nodes (default is `[]`)

- `path_url_prefixes(PathPrefix, CodeURLPrefix, DocURLPrefix)`
code and documenting URL prefixes for a path prefix used when generating cluster, library, directory, file, and entity links (atoms; no default; can be specified multiple times)
- `url_prefixes(CodeURLPrefix, DocURLPrefix)`
default URL code and documenting URL prefixes used when generating cluster, library, file, and entity links (atoms; no default)
- `entity_url_suffix_target(Suffix, Target)`
extension for entity documenting URLs (an atom; default is `'.html'`) and target separating symbols (an atom; default is `'#'`)
- `omit_path_prefixes(Prefixes)`
omit common path prefixes when printing directory paths and when constructing URLs (a list of atoms; default is a list with the user home directory)
- `zoom(Boolean)`
generate sub-diagrams and add links and zoom icons to library and entity nodes (true or false; default is false)
- `zoom_url_suffix(Suffix)`
extension for linked diagrams (an atom; default is `'.svg'`)

In the particular case of cross-referencing diagrams, there are also the options:

- `recursive_relations(Boolean)`
print recursive predicate relations (true or false; default is false)
- `url_line_references(Host)`
syntax for the URL source file line part (an atom; possible values are {github,gitlab,bitbucket}; default is github); when using this option, the CodeURLPrefix should be a permanent link (i.e., it should include the commit SHA1)
- `predicate_url_target_format(Generator)`
documentation final format generator (an atom; possible values are {sphinx,other}; default is sphinx)

For directory and file diagrams, the options are:

- `layout(Layout)`
diagram layout (one of the atoms {top_to_bottom,bottom_to_top,left_to_right,right_to_left}; default is top_to_bottom)
- `title(Title)`
diagram title (an atom; default is `''`)
- `date(Boolean)`
print current date and time (true or false; default is true)
- `versions(Boolean)`
print Logtalk and backend version data (true or false; default is false)
- `directory_paths(Boolean)`
print file directory paths (true or false; default is false)
- `file_extensions(Boolean)`
print file name extensions (true or false; default is true)
- `path_url_prefixes(PathPrefix, CodeURLPrefix, DocURLPrefix)`

code and documenting URL prefixes for a path prefix used when generating cluster, directory, file, and entity links (atoms; no default; can be specified multiple times)

- `url_prefixes(CodeURLPrefix, DocURLPrefix)`
default URL code and documenting URL prefixes used when generating cluster, library, file, and entity links (atoms; no default)
- `omit_path_prefixes(Prefixes)`
omit common path prefixes when printing directory paths and when constructing URLs (a list of atoms; default is a list with the user home directory)
- `relation_labels(Boolean)`
print entity relation labels (true or false; default is false)
- `externals(Boolean)`
print external nodes (true or false; default is true)
- `node_type_captions(Boolean)`
print node type captions (true or false; default is false)
- `output_directory(Directory)`
directory for the .d2 and .dot files (an atom; default is './dot_dias')
- `exclude_directories(Directories)`
list of directories to exclude except as external nodes (default is [])
- `exclude_files(Files)`
list of source files to exclude except as external nodes (default is [])
- `zoom(Boolean)`
generate sub-diagrams and add links and zoom icons to library and entity nodes (true or false; default is false)
- `zoom_url_suffix(Suffix)`
extension for linked diagrams (an atom; default is '.svg')

For directory dependency diagrams, there is also the option:

- `cycle_detection(Boolean)` highlight cyclic relations by coloring the edges red (true or false; default is false)

For library diagrams, the options are:

- `layout(Layout)`
diagram layout (one of the atoms {top_to_bottom, bottom_to_top, left_to_right, right_to_left}; default is top_to_bottom)
- `title(Title)`
diagram title (an atom; default is '')
- `date(Boolean)`
print current date and time (true or false; default is true)
- `versions(Boolean)`
print Logtalk and backend version data (true or false; default is false)
- `directory_paths(Boolean)`
print file directory paths (true or false; default is false)
- `path_url_prefixes(PathPrefix, CodeURLPrefix, DocURLPrefix)`

code and documenting URL prefixes for a path prefix used when generating cluster, library, file, and entity links (atoms; no default; can be specified multiple times)

- `url_prefixes(CodeURLPrefix, DocURLPrefix)`
default URL code and documenting URL prefixes used when generating cluster, library, file, and entity links (atoms; no default)
- `omit_path_prefixes(Prefixes)`
omit common path prefixes when printing directory paths and when constructing URLs (a list of atoms; default is a list with the user home directory)
- `relation_labels(Boolean)`
print entity relation labels (true or false; default is false)
- `externals(Boolean)`
print external nodes (true or false; default is true)
- `node_type_captions(Boolean)`
print node type captions (true or false; default is false)
- `output_directory(Directory)`
directory for the .d2 and .dot files (an atom; default is './dot_dias')
- `exclude_directories(Directories)`
list of directories to exclude except as external nodes (default is [])
- `exclude_files(Files)`
list of source files to exclude except as external nodes (default is [])
- `exclude_libraries(Libraries)`
list of libraries to exclude except as external nodes (default is [startup, scratch_directory])
- `zoom(Boolean)`
generate sub-diagrams and add links and zoom icons to library and entity nodes (true or false; default is false)
- `zoom_url_suffix(Suffix)`
extension for linked diagrams (an atom; default is '.svg')

For library dependency diagrams, there is also the option:

- `cycle_detection(Boolean)` highlight cyclic relations by coloring the edges red (true or false; default is false)

When using the `zoom(true)` option, the `layout(Layout)` option applies only to the top diagram; sub-diagrams will use their own layout default.

The option `omit_path_prefixes(Prefixes)` with a non-empty list of prefixes should preferably be used together with the option `directory_paths(true)` when generating library or file diagrams that reference external libraries or files. To confirm the exact default options used by each type of diagram, send the `default_options/1` message to the diagram object.

Be sure to set the `source_data` flag on before compiling the libraries or files for which you want to generate diagrams.

Support for displaying Prolog modules and Prolog module files in diagrams of Logtalk applications:

- ECLiPSe
file diagrams don't display module files
- SICStus Prolog

file diagrams don't display module files

- SWI-Prolog
full support (uses the SWI-Prolog `prolog_xref` library)
- YAP
full support (uses the YAP `prolog_xref` library)

5.8.9 Linking diagrams

When using SVG output, it's possible to generate diagrams that link to other diagrams, to API documentation, to local files and directories, and to source code repositories.

For generating links between diagrams, use the `zoom(true)` option. This option allows (1) linking library diagrams to entity diagrams to predicate cross-referencing diagrams and (2) linking directory diagrams to file diagrams to entity diagrams to predicate cross-referencing diagrams. The sub-diagrams are automatically generated. For example, using the predicates that generate library diagrams will also automatically generate the entity and predicate cross-referencing diagrams.

To generate local links for opening directories, files, and file locations in selected text editors, set the URL code prefix:

- VSCode: `url_prefixes('vscode://file/', DocPrefix)`
- VSCodium: `url_prefixes('vscodium://file/', DocPrefix)`
- Cursor: `url_prefixes('cursor://file/', DocPrefix)`
- PearAI: `url_prefixes('pearai://file/', DocPrefix)`
- Windsurf: `url_prefixes('windsurf://file/', DocPrefix)`
- Zed: `url_prefixes('zed://file/', DocPrefix)`
- BBEdit: `url_prefixes('x-bbedit://open?url=file://', DocPrefix)`
- MacVim: `url_prefixes('mvim://open?url=file://', DocPrefix)`
- TextMate: `url_prefixes('txmt://open?url=file://', DocPrefix)`
- IDEA: `url_prefixes('idea://open?file=', DocPrefix)`
- PyCharm: `url_prefixes('pycharm://open?file=', DocPrefix)`

In this case, the `DocPrefix` argument should be the path to the directory containing the HTML version of the application APIs.

As most of the text editor URL scheme handlers require local links to use absolute paths, the `omit_path_prefixes/1` option is ignored. Note that local links require text editor support for URL schemes that can handle both file and directory links.

To generate links to API documentation and source code repositories, use the options `path_url_prefixes/3` (or `url_prefixes/2` for simpler cases) and `omit_path_prefixes/1`. The idea is that the `omit_path_prefixes/1` option specifies local file prefixes that will be cut and replaced by the URL prefixes (which can be path prefix specific when addressing multiple code repositories). To generate local file system URLs, define the empty atom, `''`, as a prefix. As an example, consider the Logtalk library. Its source code is available from a GitHub repository, and its documentation is published on the Logtalk website. The relevant URLs in this case are:

- <https://github.com/LogtalkDotOrg/logtalk3/> (source code)
- <https://logtalk.org/library/> (API documentation)

Git source code URLs should include the commit SHA1 to ensure that entity and predicate file line information in the URLs remain valid if the code changes in later commits. Assuming a GitHub variable bound to the SHA1 commit URL we want to reference, an inheritance diagram can be generated using the goal:

```
| ?- GitHub = 'https://github.com/LogtalkDotOrg/logtalk3/commit/
↪eb156d46e135ac47ef23adcc5d20d49dd8b66abb',
   APIDocs = 'https://logtalk.org/library/',
   logtalk_load(diagrams(loader)),
   set_logtalk_flag(source_data, on),
   logtalk_load(library(all_loader)),
   inheritance_diagram::rlibrary(library, [
       title('Logtalk library'),
       node_type_captions(true),
       zoom(true),
       path_url_prefixes('$LOGTALKUSER/', GitHub, APIDocs),
       path_url_prefixes('$LOGTALKHOME/', GitHub, APIDocs),
       omit_path_prefixes(['$LOGTALKUSER/', '$LOGTALKHOME/', '$HOME/'])
   ]).
```

The two `path_url_prefixes/3` options take care of source code and API documentation for entities loaded either from the Logtalk installation directory (whose location is given by the `LOGTALKHOME` environment variable) or from the Logtalk user directory (whose location is given by the `LOGTALKUSER` environment variable). As we also don't want any local operating-system paths to be exposed in the diagram, we use the `omit_path_prefixes/1` option to suppress those path prefixes. Note that all the paths and URLs must end with a slash for proper handling. The `git` library may be useful to retrieve the commit SHA1 from a local repo directory.

For both `path_url_prefixes/3` and `omit_path_prefixes/1` options, when a path prefix is itself a prefix of another path, the shorter path must come last to ensure correct links.

See the `SCRIPT.txt` file in the tool directory for additional examples. To avoid retyping such complex goals when updating diagrams, use the `doclet` tool to save and reapply them easily (e.g., by using the `make` tool with the `documentation` target).

5.8.10 Creating diagrams for Prolog module applications

Currently limited to SWI-Prolog and YAP Prolog module applications due to the lack of a comprehensive reflection API in other Prolog systems.

Simply load your Prolog module application and its dependencies and then use `diagram` entity, `directory`, or `file` predicates. Library diagram predicates are not supported. See the `SCRIPT.txt` file in the tool directory for some usage examples. Note that support for diagrams with links to API documentation is quite limited, however, due to the lack of Prolog standards.

5.8.11 Creating diagrams for plain Prolog files

This tool can also be used to create predicate cross-referencing diagrams for plain Prolog files. For example, if the Prolog file is named `code.pl`, simply define an object including its code:

```
:- object(code).
   :- include('code.pl').
:- end_object.
```

Save the object to an e.g. `code.lgt` file in the same directory as the Prolog file and then load it and create the diagram:

```
| ?- logtalk_load(code),  
    xref_diagram::entity(code).
```

An alternative is to use the `object_wrapper_hook` provided by the `hook_objects` library:

```
| ?- logtalk_load(hook_objects(loader)).  
...  
  
| ?- logtalk_load(code, [hook(object_wrapper_hook)]),  
    xref_diagram::entity(code).
```

5.8.12 Other notes

Caller graph diagrams can be generated for a specific predicate using the `caller_diagram` object. For example, to generate a caller graph for the `list::member/2` predicate:

```
| ?- caller_diagram::predicate(list::member/2).
```

Non-terminal indicators can also be used. For example:

```
| ?- caller_diagram::predicate(number_grammars(chars)::digit//1).
```

Generating complete diagrams requires that all referenced entities are loaded. When that is not the case, notably when generating cross-referencing diagrams, missing entities can result in incomplete diagrams.

For complex applications, diagrams can often be made simpler and more readable by omitting external nodes (see the `externals/1` option) and/or using one of the alternatives to `dot` provided by Graphviz depending on the type of the diagram (see the section above on supported graph languages for more details).

When generating entity predicate call cross-reference diagrams, caller nodes are not created for auxiliary predicates. For example, if the `meta_compiler` library is used to optimize meta-predicate calls, the diagrams may show predicates that are not apparently called by any other predicate when the callers are from the optimized meta-predicate goals (which are called via library generated auxiliary predicates). A workaround in this case would be creating a dedicated loader file that doesn't load (and apply) the `meta_compiler` library when generating the diagrams.

When generating diagrams in SVG format, a copy of the `diagrams.css` file must exist in any directory used for publishing diagrams using it. The `lgt2svg` scripts also take care of copying this file.

The Graphviz command-line utilities, e.g., `dot`, are notorious for random crashes (segmentation faults usually), often requiring retrying conversions from `.dot` files to other formats. A possible workaround is to repeat the command until it completes without error. See, for example, the `lgt2svg.sh` script.

The conversion by the `d2` command-line executable of `.d2` files to `.svg` files can be quite slow (as of its 0.6.8 version) with the default `elk` layout engine. The `dagre` layout engine is much faster but doesn't support a node referencing itself (notably, a node representing a metaclass that instantiates itself).

Using the default `d2` layout engine (`elk`) works fine with graphs with a relatively small number of nodes and edges. When that's not the case, it's a good idea to experiment with other layout engines.

5.9 doclet

This folder provides a simple tool for (re)generating documentation for a project. The tool defines a doclet object that is expected to be extended by the user to specify a sequence of goals and a sequence of shell commands that load the application and (re)generate its documentation.

Doclet source files are preferably named `doclet.lgt` (or `doclet.logtalk`) and the doclet objects are usually named after the application or library to be documented with a `_doclet` suffix. By using an `initialization/1` directive to automatically send the `update/0` message that generates the documentation upon doclet loading, we can abstract the name of the doclet object. The usual query to load and run a doclet is therefore:

```
| ?- logtalk_load([doclet(loader), doclet]).
```

For usage examples, see the `sample_doclet.lgt`, `doclet1.lgt`, `zoom_doclet.lgt`, and `tools_doclet.lgt` source files.

5.9.1 API documentation

This tool API documentation is available at:

../apis/library_index.html#doclet

For sample queries, please see the `SCRIPT.txt` file in the tool directory.

5.9.2 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(doclet(loader)).
```

5.9.3 Automating running doclets

You can use the `scripts/logtalk_doclet.sh` Bash shell script for automating running doclets. The script expects the doclet source files to be named either `doclet.lgt` or `doclet.logtalk`. See the `scripts/NOTES.md` file or the script man page for details.

5.9.4 Integration with the make tool

Loading this tool adds a definition for the `logtalk_make_target_action/1` hook predicate for the target documentation. The hook definition sends an `update/0` message to each loaded doclet.

5.10 help

This tool provides help for Logtalk features and libraries when running in most operating-systems. For help on the Logtalk compiler error and warning messages, see the tutor tool.

5.10.1 Requirements

On Windows, the start command must be available. On Linux, the xdg-open command must be available. On macOS, the command open is used. On POSIX systems bsdtar is also required.

Browsing the Handbook and APIs documentation at the top-level requires a POSIX system and one of the following terminal-based browsers installed:

- <https://invisible-island.net/lynx/>
- <https://w3m.sourceforge.net/>
- <http://links.twibright.com/>
- <https://sr.ht/~bptato/chawan/>

On macOS, these browsers and bsdtar can be installed with either MacPorts:

```
$ sudo port install lynx
$ sudo port install w3m
$ sudo port install links
$ sudo port install libarchive
```

Or using Homebrew:

```
$ brew install lynx
$ brew install w3m
$ brew install links
$ brew install chawan
$ brew install libarchive
```

On Linux systems, use the distribution's own package manager to install the browsers and bsdtar. For example, in Ubuntu systems:

```
$ sudo apt install lynx
$ sudo apt install w3m
$ sudo apt install links2
$ sudo apt install libarchive-tools
```

For RedHat systems:

```
$ sudo dnf install lynx
$ sudo dnf install w3m
$ sudo dnf install links2
$ sudo dnf install bsdtar
```

If you're running Logtalk from a git clone of its repo, you will need to run the docs/apis/sources/build.sh or docs/apis/sources/build.ps1 scripts to generate APIs documentation HTML files and also run the docs/handbook/sources/build.sh or docs/handbook/sources/build.ps1 scripts to generate the Handbook HTML files. Alternatively, you can download the documentation for the latest stable release from the Logtalk website and save them to the docs directories.

5.10.2 Customization

The preferred browser can be set in a `settings.lgt` file by defining the `help_default_browser` user-defined flag as follows:

```
:- initialization(
    create_logtalk_flag(help_default_browser, Browser, [type(atom), keep(true)])
).
```

The valid values for `Browser` are: `lynx` (default value), `w3m`, `links`, `cha`, and `default` (default means use the operating-system default browser instead of one of the terminal-based browsers).

We can also select between browsing the HTML files generated by Sphinx for web browsing or the XHTML files generated by Sphinx for the ePub version of the documentation by defining the `help_default_files` user-defined flag as follows:

```
:- initialization(
    create_logtalk_flag(help_default_files, Format, [type(atom), keep(true)])
).
```

The valid values for `Format` are `html` (default value) or `xhtml`.

The `samples/settings-sample.lgt` file contains commented out code for defining these flags.

On POSIX systems, one of the supported terminal-based browsers must be installed unless you prefer using the default browser. By default, the tool checks first for `lynx`, second for `w3m`, third for `links`, and finally for `chawan`. When none of these terminal-based browsers is available in their default installation paths, the documentation is open in the operating-system default browser.

On Windows systems, the documentation is open in the operating-system default browser.

When using the terminal-based browsers, after finishing consulting the documentation and quitting the process, you will be back to the top-level prompt (if you find that the top-level have scrolled from its last position, try to set your terminal terminfo to `xterm-256colour`).

5.10.3 API documentation

This tool API documentation is available at:

../apis/library_index.html#help

5.10.4 Loading

```
| ?- logtalk_load(help(loader)).
```

5.10.5 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(help(tester)).
```

5.10.6 Supported operating-systems

Currently, support is limited to Linux, macOS, and Windows.

This tool relies on the library portable operating-system access abstraction.

5.10.7 Usage

After loading the tool, use the query `help::help` to get started:

```
| ?- logtalk_load(help(loader)).  
...  
  
| ?- help::help.  
...
```

When using ECLiPSe, you will need to write the object name, `help`, between parentheses to avoid a clash with the `help` built-in operator:

```
| ?- (help)::help.  
...
```

Same sample help queries:

```
% get et on-line help for the `logtalk_load/2` built-in predicate:  
  
| ?- help::logtalk_load/2.  
...  
  
% get on-line help for the `eos//0` built-in non-terminal:  
  
| ?- help::eos//0.  
...  
  
% consult the Handbook:  
  
| ?- help::handbook.  
...  
  
% consult the APIs documentation:  
  
| ?- help::apis.  
...  
  
% consult the APIs documentation about a library predicate:
```

(continues on next page)

(continued from previous page)

```

| ?- help::apis(member/2).
...

% consult the APIs documentation about a library non-terminal:

| ?- help::apis(one_or_more//0).
...

% consult the documentation of a library:

| ?- help::library(random).
...

% consult the documentation of an entity:

| ?- help::entity(logtalk).
...

% consult the developer tools documentation:

| ?- help::tools.
...

% consult the documentation of the lgtunit tool:

| ?- help::tool(lgtunit).
...

% consult the documentation of the logtalk_tester script:

| ?- help::man(logtalk_tester).
...

```

5.10.8 Known issues

When using the lynx terminal-based browser, the `help_default_files` flag setting is ignored as this browser doesn't support XHTML.

When using the terminal-based browsers, the Handbook and APIs search boxes are not usable as they require JavaScript support. Use instead the indexes.

Some CSS support is only provided by the Chawan terminal-based browser. This results in formatting issues when browsing the Logtalk documentation. Some of these issues can be avoided by setting the `help_default_files` flag to `xhtml`.

The open commands used to open documentation URLs in the default browser drop the fragment part, thus preventing navigation to the specified position on the documentation page.

ECLiPSe defines a help prefix operator that forces wrapping this atom between parentheses when sending messages to the tool. E.g. use `(help)::help` instead of `help::help`.

5.11 issue_creator

This is a complementary tool for the `lgtunit` tool for automatically creating bug report issues for failed tests in GitHub or GitLab servers.

5.11.1 Requirements

This tool requires that the GitHub `gh` and GitLab `glab` CLIs be installed. For the installation instructions see:

- GitHub: <https://cli.github.com>
- GitLab: <https://glab.readthedocs.io>

5.11.2 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(issue_creator(loader)).
```

But, in the most common usage scenario, this tool is automatically loaded by the `logtalk_tester` automation script.

5.11.3 Usage

The `logtalk_tester` automation script accepts a `-b` option for automatically using this tool (see the script man page for details). In the most simple case, this option possible values are `github` and `gitlab`. For example:

```
$ logtalk_tester \
  -p gnu \
  -b github \
  -s "/home/jdoe/foo/" \
  -u https://github.com/jdoe/foo/tree/55aa900775befa135e0d5b48ea63098df8b97f5c/
```

The `logtalk_tester` script **must** be called from a git repo directory or one of its sub-directories, which is a common setup in CI/CD pipelines. Moreover, prior to running the tests, the CLI must be used, if required, to authenticate and login to the server where the bug report issues will be created:

- GitHub: `gh auth login --hostname <string> --with-token < token.txt`
- GitLab: `glab auth login --hostname <string> --token <string>`

The access token must have the necessary scopes that allow bug reports to be created. See the CLIs documentation for details. Typically, the `auth` command is called from the CI/CD pipeline definition scripts. However, depending on the CI/CD workflow, the authentication may be done implicitly.

The bug reports are created using by default the label `bug` and assigned to the author of the latest commit of the git repo. The `-b` option can also be used to override the label with a comma separated set of labels. For example, to use both `bug` and `auto` labels:

```
$ logtalk_tester \
  -p gnu \
  -b github:bug,auto \
```

(continues on next page)

(continued from previous page)

```
-s "/home/jdoe/foo/" \
-u https://github.com/jdoe/foo/tree/55aa900775befa135e0d5b48ea63098df8b97f5c/
```

Note that the labels **must** be predefined in the issue tracker server for the bug report to be successfully created. The auto label can be used to simplify filtering of auto-generated bug reports when browsing the issue tracker.

The bug reports use Markdown formatting, which is the default in GitHub and GitLab issue trackers.

But reports are only created for non-flaky tests. The bug report title and labels are used to prevent creating duplicated bug reports. Therefore, the title should not be manually edited and the same labels should be used for multiple runs of the same tests and preserved when editing the bug reports.

There are cases where we may want to postpone or temporarily disable the automatic creation of bug reports. E.g. a WIP branch that's known to break multiple tests. A solution is to define a pull/merge request label, e.g. NO_AUTO_BUG_REPORTS, that can then be checked by the CI/CD workflow. For example, we can test the presence of that label to set a AUTO_BUG_REPORTS environment variable to either an empty string or a -b option and use:

```
logtalk_tester.sh $AUTO_BUG_REPORTS -p ...
```

5.11.4 Known issues

GitLab creates CI/CD pipelines in a detached HEAD state by default. As a consequence, the git branch would be reported as HEAD. To workaround this issue, this tool uses the value of the GitLab CI/CD pipeline variable CI_COMMIT_REF_NAME when defined as the branch name.

This tool is in a beta stage of development. Your feedback is most appreciated.

5.12 lgtdoc

This is the default Logtalk documenting tool for generating API documentation for libraries and applications. It uses the structural reflection API to extract and output in XML format relevant documentation about a source file, a library or directory of source files, or all loaded source files. The tool predicates accept several options for generating the XML files, including the output directory.

The lgtdoc/xml directory contains several ready-to-use Bash and PowerShell scripts for converting the XML documenting files into final formats, including XHTML 1.1, HTML 4.01, HTML 5, PDF, Markdown, and reStructuredText (for use with Sphinx), and plain text files. The scripts are described in their man pages and made available in the system path by default.

5.12.1 Requirements

This tool requirements for converting the XML files it generates to a final format are as follows:

1. Converting XML files to (X)HTML, reStructuredText, Markdown, or plain text files requires a XSLT processor. The supported XSLT processors for Bash conversion scripts are:
 - xsltproc: <https://gitlab.gnome.org/GNOME/libxslt/-/wikis/home>
 - Xalan: <https://xml.apache.org/xalan-c/index.html>
 - Saxon: <https://www.saxonica.com>

On Windows, the PowerShell scripts use the .Net XSLT classes.

The reStructuredText files output is usually used as an intermediate step to generate Sphinx HTML, PDF, ePub, and Texinfo files. The additional requirements are:

- Sphinx: <https://www.sphinx-doc.org/>
 - Pygments: <https://pygments.org/>
 - Read the Docs theme: https://github.com/readthedocs/sphinx_rtd_theme
2. Converting XML files to PDF files require a XSL-FO processor. The supported XSL-FO processors for Bash and PowerShell conversion scripts are:
- Apache FOP processor: <https://xmlgraphics.apache.org/fop/>
 - RenderX XEP processor: <https://www.renderx.com/>

For additional details, including compatible dependency versions and available conversion scripts, see the `xml/NOTES.md` file. See the `tools/NOTES.md` file for per operating-system installation instructions for the above dependencies.

5.12.2 API documentation

This tool API documentation is available at:

`../..apis/library_index.html#lgtdoc`

5.12.3 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(lgtdoc(loader)).
```

5.12.4 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(lgtdoc(tester)).
```

5.12.5 Documenting source code

For information on documenting your source code, notably on documenting directives, consult the documenting section of the User Manual:

`../..handbook/userman/documenting.html`

Extracting documenting information from your source code using this tool requires compiling the source files using the `source_data(on)` compiler flag. For example:

```
| ?- logtalk_load(source_file, [source_data(on)]).
```

Usually, this flag is set for all application source files in the corresponding loader file. In alternative, you may also turn on the `source_data` flag globally by typing:

```
| ?- set_logtalk_flag(source_data, on).
```

The tool API allows generating documentation for libraries, directories, and files, complemented with library, directory, entity, and predicate indexes. Note that the source files to be documented **must** be loaded prior to using this tool predicates to generate the documentation.

5.12.6 Generating documentation

For a simple application, assuming a library alias is defined for it (e.g. `my_app`), and at the top-level interpreter, we can generate the application documentation by typing:

```
| ?- {my_app(loader)}.
...

| ?- {lgt doc(loader)}.
...

| ?- lgt doc::library(my_app).
...
```

By default, the documenting XML files are created in a `xml_docs` directory in the current working directory. But usually all documenting files are collected for both the application and the libraries it uses in a common directory so that all documentation links resolved properly. The `lgt doc` predicates can take a list of options to customize the generated XML documenting files. See the remarks section in the [lgt doc](#) protocol documentation for details on the available options.

After generating the XML documenting files, these can be easily converted into final formats using the provided scripts. For example, assuming that we want to generate HTML documentation:

```
$ cd xml_docs
$ lgt2html -t "My app"
```

To generate the documentation in Sphinx format instead (as used by Logtalk itself for its APIs):

```
$ cd xml_docs
$ lgt2rst -s -- -q -p "Application name" -a "Author name" -v "Version X.YZ.P"
$ make html
```

In this case, the generated documentation will be in the `xml_docs/_build/html/` directory. See the scripts man pages or call them using the `-h` option to learn more about their supported options.

For more complex applications, you can use the `doclet` tool to define a *doclet* to automate all the steps required to generate documentation. The *doclet* message that triggers the process can also be sent automatically when the `make` tool is used with the documentation target.

5.12.7 Documentation linter checks

When the `lgtdoc_missing_directives` flag is set to warning (its usual default value), the `lgtdoc` tool prints warnings on missing entity `info/1` directives and missing predicate `info/2` and `mode/2` directives.

When the `lgtdoc_missing_info_key` flag is set to warning (its usual default value), the `lgtdoc` tool prints warnings on entity `info/1` directive and predicate `info/2` directive missing de facto required keys (e.g., `comment`, `parameters` or `parnames` for parametric entities, `arguments` or `argnames` for predicates/non-terminals with arguments).

When the `lgtdoc_invalid_dates` flag is set to warning (its usual default value), the `lgtdoc` tool prints warnings on invalid dates (including dates in the future) in `info/1` directives.

When the `lgtdoc_non_standard_exceptions` flag is set to warning (its usual default value), the `lgtdoc` tool prints warnings on non-standard exceptions. This linter check is particularly effective in detecting typos when specifying standard exceptions.

When the `lgtdoc_missing_punctuation` flag is set to warning (its usual default value), the `lgtdoc` tool prints warnings on missing ending periods (full stops), exclamation marks, or question marks in `info/1-2` directives (in comments, remarks, parameter descriptions, and argument descriptions).

Set a flag value to `silent` to turn off the corresponding linter warnings.

5.13 lgtunit

The `lgtunit` tool provides testing support for Logtalk. It can also be used for testing plain Prolog code and Prolog module code.

This tool is inspired by the xUnit frameworks architecture and by the works of Joachim Schimpf (ECLiPSe library `test_util`) and Jan Wielemaker (SWI-Prolog `plunit` package).

Tests are defined in objects, which represent a *test set* or *test suite*. In simple cases, we usually define a single object containing the tests. But it is also possible to use parametric test objects or multiple objects defining parametrizable tests or test subsets for testing more complex units and facilitating tests maintenance. Parametric test objects are specially useful to test multiple implementations of the same protocol using a single set of tests by passing the implementation object as a parameter value.

5.13.1 Main files

The `lgtunit.lgt` source file implements a framework for defining and running unit tests in Logtalk. The `lgtunit_messages.lgt` source file defines the default translations for the messages printed when running unit tests. These messages can be intercepted to customize the output, e.g. to make it less verbose or for integration with e.g. GUI IDEs and continuous integration servers.

Other files that are part of this tool provide support for alternative output formats of test results and are discussed below.

5.13.2 API documentation

This tool API documentation is available at:

../apis/library_index.html#lgtunit

5.13.3 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(lgtunit(loader)).
```

5.13.4 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(lgtunit(tester)).
```

5.13.5 Writing and running tests

In order to write your own unit tests, define objects extending the `lgtunit` object. You may start by copying the `samples/tests-sample.lgt` file (at the root of the Logtalk distribution) to a `tests.lgt` file in your project directory and editing it to add your tests:

```
:- object(tests,
    extends(lgtunit)).

    % test definitions
    ...

:- end_object.
```

The section on *test dialects* below describes in detail how to write tests. See the `tests` top directory for examples of actual unit tests. Other sources of examples are the `library` and `examples` directories.

The tests must be term-expanded by the `lgtunit` object by compiling the source files defining the test objects using the option `hook(lgtunit)`. For example:

```
| ?- logtalk_load(lgtunit(loader)),
    logtalk_load(tests, [hook(lgtunit)]).
```

As the term-expansion mechanism applies to all the contents of a source file, the source files defining the test objects should preferably not contain entities other than the test objects. Additional code necessary for the tests should go to separate files. In general, the tests themselves can be compiled in *optimized* mode. Assuming that's the case, also use the `optimize(on)` compiler option for faster execution.

The term-expansion performed by the `lgtunit` object sets the test object `source_data` flag to on and the `context_switching_calls` flag to allow for code coverage and debugging support. But these settings can always be overridden in the test objects.

The `samples/tester-sample.lgt` file (at the root of the Logtalk distribution) exemplifies how to compile and load `lgtunit` tool, the source code under testing, the unit tests, and how to automatically run all the tests after loading:

```
:- initialization((
    % minimize compilation reports to the essential ones (errors and warnings)
    set_logtalk_flag(report, warnings),
    % load any necessary library files for your application; for example
    logtalk_load(basic_types(loader)),
    % load the unit test tool
    logtalk_load(lgtunit(loader)),
    % load your application files (e.g., "source.lgt") enabling support for
    % code coverage, which requires compilation in debug mode and collecting
    % source data information; if code coverage is not required, remove the
    % "debug(on)" option for faster execution
    logtalk_load(source, [source_data(on), debug(on)]),
    % compile the unit tests file expanding it using "lgtunit" as the hook
    % object to preprocess the tests; if you have failing tests, add the
    % option debug(on) to debug them (see "tools/lgtunit/NOTES.md" for
    % debugging advice); tests should be loaded after the code being tested
    % is loaded to avoid warnings such as references to unknown entities
    logtalk_load(tests, [hook(lgtunit)]),
    % run all the unit tests; assuming your tests object is named "tests"
    tests::run
)).
```

You may copy this sample file to a `tester.lgt` file in your project directory and edit it to load your project and test files. The `logtalk_tester` testing automation script defaults to looking for test driver files named `tester.lgt` or `tester.logtalk` (if you have work-in-progress test sets that you don't want to run by default, simply use a different file name such as `tester_wip.lgt`; you can still run them automated by using `logtalk_tester -n tester_wip`).

Debugged test sets should preferably be compiled in optimal mode, specially when containing deterministic tests and when using the utility benchmarking predicates.

Assuming a `tester.lgt` driver file as exemplified above, the tests can be run by simply loading this file:

```
| ?- logtalk_load(tester).
```

Assuming your test object is named `tests`, you can re-run the tests by typing:

```
| ?- tests::run.
```

You can also re-run a single test (or a list of tests) using the `run/1` predicate:

```
| ?- tests::run(test_identifier).
```

When testing complex *units*, it is often desirable to split the tests between several test objects or use parametric test objects to be able to run the same tests using different parameters (e.g., different data sets or alternative implementations of the same protocol). In this case, you can run all test subsets using the goal:

```
| ?- lgtunit::run_test_sets([test_set_1, test_set_2, ...]).
```

where the `run_test_sets/1` predicate argument is a list of two or more test object identifiers. This predicate makes it possible to get a single code coverage report that takes into account all the tests.

It's also possible to automatically run loaded tests when using the `make` tool by calling the goal that runs the tests from a definition of the hook predicate `logtalk_make_target_action/1`. For example, by adding to the `tester.lgt` driver file the following code:

```
% integrate the tests with logtalk_make/1
:- multifile(logtalk_make_target_action/1).
:- dynamic(logtalk_make_target_action/1).

logtalk_make_target_action(check) :-
    tests::run.
```

Alternatively, you can define the predicate `make/1` inside the test set object. For example:

```
:- object(tests, extends(lgtunit)).

    make(check).
    ...

:- end_object.
```

This clause will cause all tests to be run when calling the `logtalk_make/1` predicate with the target `check` (or its top-level shortcut, `{?}`). The other possible target is `all` (with top-level shortcut `{*}`).

Note that you can have multiple test driver files. For example, one driver file that runs the tests collecting code coverage data and a quicker driver file that skips code coverage and compiles the code to be tested in optimized mode.

5.13.6 Automating running tests

You can use the `scripts/logtalk_tester.sh` Bash shell script or the `scripts/logtalk_tester.ps1` PowerShell script for automating running unit tests (e.g., from a CI/CD pipeline). When using one of the Logtalk installers, the `.sh` extension can usually be omitted. For example, assuming your current directory (or sub-directories) contains one or more `tester.lgt` files:

```
$ logtalk_tester -p gnu
```

The only required argument is the identifier of the backend Prolog system. For other options, see the `scripts/NOTES.md` file or type:

```
$ logtalk_tester -h
```

On POSIX systems, you can also access extended documentation by consulting the script man page:

```
$ man logtalk_tester
```

The scripts support the same set of options. But the option for passing additional arguments to the tests uses different syntax. For example:

```
$ logtalk_tester -p gnu -- foo bar baz
PS> logtalk_tester -p gnu -a foo,bar,baz
```

On POSIX systems, assuming Logtalk was installed using one of the provided installers or installation scripts, there is also a man page for the script:

```
$ man logtalk_tester
```

Alternatively, an HTML version of this man page can be found at:

https://logtalk.org/man/logtalk_tester.html

On POSIX systems, the `logtalk_tester.sh` Bash script timeout option requires either a `timeout` or a `gtimeout` command (provided by the GNU coreutils package). The `logtalk_tester.ps1` PowerShell script timeout option requires that Git for Windows is also installed, as it requires the GNU timeout command bundled with it.

In addition to using the `logtalk_tester.ps1` PowerShell script, the Bash shell version of the automation script can also be used in Windows operating-systems with selected backends by using the Bash shell included in the Git for Windows installer. That requires defining a `.profile` file setting the paths to the Logtalk scripts and the Prolog backend executables. For example:

```
$ cat ~/.profile
# YAP
export PATH="/C/Program Files/Yap64/bin":$PATH
# GNU Prolog
export PATH="/C/GNU-Prolog/bin":$PATH
# SWI/Prolog
export PATH="/C/Program Files/swipl/bin":$PATH
# ECLiPse
export PATH="/C/Program Files/ECLiPse 7.0/lib/x86_64_nt":$PATH
# SICStus Prolog
export PATH="/C/Program Files/SICStus Prolog VC16 4.6.0/bin":$PATH
# Logtalk
export PATH="$LOGTALKHOME/scripts":"$LOGTALKHOME/integration":$PATH
```

The Git for Windows installer also includes GNU coreutils and its `timeout` command, which is used by the `logtalk_tester` script `-t` option.

Note that some tests may give different results when run from within the Bash shell compared with running the tests manually using a Windows GUI version of the Prolog backend. Some backends may also not be usable for automated testing due to the way they are made available as Windows applications.

Additional advice on testing and on automating testing using continuous integration servers can be found at:

<https://logtalk.org/testing.html>

5.13.7 Parametric test objects

Parameterized unit tests can be easily defined by using parametric test objects. A typical example is testing multiple implementations of the same protocol. In this case, we can use a parameter to pass the specific implementation being tested. For example, assume that we want to run the same set of tests for the library `random_protocol` protocol. We can write:

```
:- object(tests(_RandomObject_),
    extends(lgtunit)).

    :- uses(_RandomObject_, [
        random/1, between/3, member/2,
        ...
    ]).

    test(between_3_in_interval) :-
        between(1, 10, Random),
```

(continues on next page)

(continued from previous page)

```

1 =< Random, Random =< 10.

...

:- end_object.

```

We can then test a specific implementation by instantiating the parameter. For example:

```
| ?- tests(fast_random)::run.
```

Or use the `lgtunit::run_test_sets/1` predicate to test all the implementations:

```
| ?- lgtunit::run_test_sets([
    tests(backend_random),
    tests(fast_random),
    tests(random)
]).
```

5.13.8 Test dialects

Multiple test *dialects* are supported by default. See the next section on how to define your own test dialects. In all dialects, a **ground callable term**, usually an atom, is used to uniquely identify a test. This simplifies reporting failed tests and running tests selectively. An error message is printed if invalid or duplicated test identifiers are found. These errors must be corrected; otherwise the reported test results can be misleading. Ideally, tests should have descriptive names that clearly state the purpose of the test and what is being tested.

Unit tests can be written using any of the following predefined dialects:

```
test(Test) :- Goal.
```

This is the most simple dialect, allowing the specification of tests that are expected to succeed. The argument of the `test/1` predicate is the test identifier, which must be unique. A more versatile dialect is:

```

succeeds(Test) :- Goal.
deterministic(Test) :- Goal.
fails(Test) :- Goal.
throws(Test, Ball) :- Goal.
throws(Test, Balls) :- Goal.

```

This is a straightforward dialect. For `succeeds/1` tests, `Goal` is expected to succeed. For `deterministic/1` tests, `Goal` is expected to succeed once without leaving a choice-point. For `fails/1` tests, `Goal` is expected to fail. For `throws/2` tests, `Goal` is expected to throw the exception term `Ball` or one of the exception terms in the list `Balls`. The specified exception must subsume the actual exception for the test to succeed.

An alternative test dialect that can be used with more expressive power is:

```
test(Test, Outcome) :- Goal.
```

The possible values of the outcome argument are:

- `true`
The test is expected to succeed.
- `true(Assertion)`

The test is expected to succeed and satisfy the Assertion goal.

- `deterministic`

The test is expected to succeed once without leaving a choice-point.

- `deterministic(Assertion)`

The test is expected to succeed once without leaving a choice-point and satisfy the Assertion goal.

- `subsumes(Expected, Result)`

The test is expected to succeed by binding `Result` to a term that is subsumed by the `Expected` term.

- `variant(Term1, Term2)`

The test is expected to succeed by binding `Term1` to a term that is a variant of the `Term2` term.

- `exists(Assertion)`

A solution exists for the test goal that satisfies the Assertion goal.

- `all(Assertion)`

All test goal solutions satisfy the Assertion goal.

- `fail`

The test is expected to fail.

- `false`

The test is expected to fail.

- `error(Error)`

The test is expected to throw the exception term `error(ActualError, _)` where `ActualError` is subsumed `Error`.

- `errors(Errors)`

The test is expected to throw an exception term `error(ActualError, _)` where `ActualError` is subsumed by an element of the list `Errors`.

- `ball(Ball)`

The test is expected to throw the exception term `ActualBall` where `ActualBall` is subsumed `Ball`.

- `balls(Balls)`

The test is expected to throw an exception term `ActualBall` where `ActualBall` is subsumed by an element of the list `Balls`.

In the case of the `true(Assertion)`, `deterministic(Assertion)`, and `all(Assertion)` outcomes, a message that includes the assertion goal is printed for assertion failures and errors to help to debug failed unit tests. Same for the `subsumes(Expected, Result)` and `variant(Term1, Term2)` assertions. Note that this message is only printed when the test goal succeeds, as its failure will prevent the assertion goal from being called. This allows distinguishing between test goal failure and assertion failure.

Note that the `all(Assertion)` outcome simplifies pinpointing which test goal solution failed the assertion. See also the section below on testing non-deterministic predicates.

The `fail` and `false` outcomes are better reserved for cases where there is a single test goal. With multiple test goals, the test will succeed when *any* of those goals fail.

Some tests may require individual condition, setup, or cleanup goals. In this case, the following alternative test dialect can be used:

```
test(Test, Outcome, Options) :- Goal.
```

The currently supported options are (non-recognized options are ignored):

- `condition(Goal)`

Condition for deciding if the test should be run or skipped (default goal is true).

- `setup(Goal)`
Setup goal for the test (default goal is true).
- `cleanup(Goal)`
Cleanup goal for the test (default goal is true).
- `flaky`
Declare the test as a flaky test.
- `note(Term)`
Annotation to print (between parentheses by default) after the test result (default is ''); the annotation term can share variables with the test goal, which can be used to pass additional information about the test result.

Also supported is QuickCheck testing, where random tests are automatically generated and run given a predicate mode template with type information for each argument (see the section below for more details):

```
quick_check(Test, Template, Options).
quick_check(Test, Template).
```

The valid options are the same as for the `test/3` dialect plus all the supported QuickCheck specific options (see the QuickCheck section below for details).

For examples of how to write unit tests, check the `tests` folder or the `testing` example in the `examples` folder in the Logtalk distribution. Most of the provided examples also include unit tests, some of them with code coverage.

5.13.9 User-defined test dialects

Additional test dialects can be easily defined by extending the `lgtunit` object and by term-expanding the new dialect into one of the default dialects. As an example, suppose that you want a dialect where you can simply write a file with tests defined by clauses using the format:

```
test_identifier :-
    test_goal.
```

First, we define an expansion for this file into a test object:

```
:- object(simple_dialect,
    implements(expanding)).

    term_expansion(begin_of_file, [(:- object(tests,extends(lgtunit))]).
    term_expansion((Head :- Body), [test(Head) :- Body]).
    term_expansion(end_of_file, [(:- end_object)]).

:- end_object.
```

Then we can use this hook object to expand and run tests written in this dialect by using a `tester.lgt` driver file with contents such as:

```
:- initialization((
    set_logtalk_flag(report, warnings),
    logtalk_load(lgtunit(loader)),
```

(continues on next page)

(continued from previous page)

```

logtalk_load(hook_flows(loader)),
logtalk_load(simple_dialect),
logtalk_load(tests, [hook(hook_pipeline([simple_dialect,lgtunit]))]),
tests::run
)).

```

The hook pipeline first applies our `simple_dialect` expansion, followed by the default `lgtunit` expansion. This solution allows other hook objects (e.g., required by the code being tested) to also be used by updating the pipeline.

5.13.10 QuickCheck

QuickCheck was originally developed for Haskell. Implementations for several other programming languages soon followed. QuickCheck provides support for *property-based testing*. The idea is to express properties that predicates must comply with and automatically generate tests for those properties. The `lgtunit` tool supports both `quick_check/2-3` test dialects, as described above, and `quick_check/1-3` public predicates for interactive use:

```

quick_check(Template, Result, Options).
quick_check(Template, Options).
quick_check(Template).

```

The following options are supported:

- `n/1`: number of random tests that will be generated and run (default is 100).
- `s/1`: maximum number of shrink operations when a counter-example is found (default is 64).
- `ec/1`: boolean option deciding if type edge cases are tested before generating random tests (default is true).
- `rs/1`: starting seed to be used when generating the random tests (no default).
- `pc/1`: pre-condition closure for generated tests (extended with the test arguments; no default).
- `l/1`: label closure for classifying the generated tests (extended with the test arguments plus the label argument; no default).
- `v/1`: boolean option for verbose reporting of generated random tests (default is false).
- `pb/2`: progress bar option for executed random tests when the verbose option is false (first argument is a boolean, default is false; second argument is the tick number, a positive integer).

The `quick_check/1` predicate uses the default option values. The `quick_check/1-2` predicates print the test results and are thus better reserved for testing at the top-level interpreter. The `quick_check/3` predicate returns results in reified form:

- `passed(SequenceSeed, Discarded, Labels)`
- `failed(Goal, SequenceSeed, TestSeed)`
- `error(Error, Goal, SequenceSeed, TestSeed)`
- `broken(Why, Culprit)`

The `broken(Why, Culprit)` result only occurs when the user-defined testing setup is broken. For example, a non-callable template (e.g., a non-existing predicate), an invalid option, a problem with the pre-condition closure or with the label closure (e.g., a pre-condition that always fails or a label that fails to classify a

generated test), or errors/failures when generating tests (e.g., due to an unknown type being used in the template or a broken custom type arbitrary value generator).

The Goal argument is the random test that failed.

The SequenceSeed argument is the starting seed used to generate the sequence of random tests. The TestSeed is the seed used to generate the test that failed. Both seeds should be regarded as opaque terms. When the test seed is equal to the sequence seed, this means that the failure or error occurred while using only type edge cases. See below how to use the seeds when testing bug fixes.

The Discarded argument returns the number of generated tests that were discarded for failing to comply with a pre-condition specified using the pc/1 option. This option is specially useful when constraining or enforcing a relation between the generated arguments and is often used as an alternative to define a custom type. For example, if we define the following predicate:

```
condition(I) :-
    between(0, 127, I).
```

We can then use it to filter the generated tests:

```
| ?- lgtunit::quick_check(integer(+byte), [pc(condition)]).
% 100 random tests passed, 94 discarded
% starting seed: seed(416,18610,17023)
yes
```

The Labels argument returns a list of pairs Label-N where N is the number of generated tests that are classified as Label by a closure specified using the l/1 option. For example, assuming the following predicate definition:

```
label(I, Label) :-
    (   I mod 2 == 0 ->
        Label = even
    ;   Label = odd
    ).
```

We can try:

```
| ?- lgtunit::quick_check(integer(+byte), [l(label), n(10000)]).
% 10000 random tests passed, 0 discarded
% starting seed: seed(25513,20881,16407)
% even: 5037/10000 (50.370000%)
% odd: 4963/10000 (49.630000%)
yes
```

The label statistics are key to verifying that the generated tests provide the necessary coverage. The labeling predicates can return a single test label or a list of test labels. Labels should be ground and are typically atoms. To examine the generated tests themselves, you can use the verbose option, v/1. For example:

```
| ?- lgtunit::quick_check(integer(+integer), [v(true), n(7), pc([I]>>(I>5))]).
% Discarded: integer(0)
% Passed:    integer(786)
% Passed:    integer(590)
% Passed:    integer(165)
% Discarded: integer(-412)
% Passed:    integer(440)
% Discarded: integer(-199)
```

(continues on next page)

(continued from previous page)

```
% Passed:    integer(588)
% Discarded: integer(-852)
% Discarded: integer(-214)
% Passed:    integer(196)
% Passed:    integer(353)
% 7 random tests passed, 5 discarded
% starting seed: seed(23671,3853,29824)
yes
```

When a counter-example is found, the verbose option also prints the shrink steps. For example:

```
| ?- lgtunit::quick_check(atom(+atomic), [v(true), ec(false)]).
% Passed:    atom('dy0=Xv_MX-3b/U4KH U')
*   Failure:  atom(-198)
*   Shrunk:   atom(-99)
*   Shrunk:   atom(-49)
*   Shrunk:   atom(-24)
*   Shrunk:   atom(-12)
*   Shrunk:   atom(-6)
*   Shrunk:   atom(-3)
*   Shrunk:   atom(-1)
*   Shrunk:   atom(0)
*   quick check test failure (at test 2 after 8 shrinks):
*   atom(0)
*   starting seed: seed(3172,9814,20125)
*   test seed:    seed(7035,19506,18186)
no
```

The template can be a `(:)/2`, `(<<)/2`, or `(:)/2` qualified callable term. When the template is an unqualified callable term, it will be used to construct a goal to be called in the context of the *sender* using the `(<<)/2` debugging control construct. Another simple example is passing a template that will trigger a failed test (as the `random::random/1` predicate always returns non-negative floats):

```
| ?- lgtunit::quick_check(random::random(-negative_float)).
*   quick check test failure (at test 1 after 0 shrinks):
*   random::random(0.09230089279334841)
*   starting seed: seed(3172,9814,20125)
*   test seed:    seed(3172,9814,20125)
no
```

When QuickCheck exposes a bug in the tested code, we can use the reported counter-example to help diagnose it and fix it. As tests are randomly generated, we can use the starting seed reported with the counter-example to confirm the bug fix by calling the `quick_check/2-3` predicates with the `rs(Seed)` option. For example, assume the following broken predicate definition:

```
every_other([], []).
every_other([_, X | L], [X | R]) :-
    every_other(L, R).
```

The predicate is supposed to construct a list by taking every other element of an input list. Cursory testing may fail to notice the bug:

```
| ?- every_other([1,2,3,4,5,6], List).
List = [2, 4, 6]
yes
```

But QuickCheck will report a bug with lists with an odd number of elements with a simple property that verifies that the predicate always succeeds and returns a list of integers:

```
| ?- lgtunit::quick_check(every_other(+list(integer), -list(integer))).
* quick check test failure (at test 2 after 0 shrinks):
* every_other([0],A)
* starting seed: seed(3172,9814,20125)
* test seed: seed(3172,9814,20125)
no
```

We could fix this particular bug by rewriting the predicate:

```
every_other([], []).
every_other([H| T], L) :-
    every_other(T, H, L).

every_other([], X, [X]).
every_other([_| T], X, [X| L]) :-
    every_other(T, L).
```

By retesting with the same test seed that uncovered the bug, the same random test that found the bug will be generated and run again:

```
| ?- lgtunit::quick_check(
    every_other(+list(integer), -list(integer)),
    [rs(seed(3172,9814,20125))]
).
% 100 random tests passed, 0 discarded
% starting seed: seed(3172,9814,20125)
yes
```

Still, after verifying the bug fix, is also a good idea to re-run the tests using the sequence seed instead, as bug fixes sometimes cause regressions elsewhere.

When retesting using the `logtalk_tester` automation script, the starting seed can be set using the `-r` option. For example:

```
$ logtalk_tester -r "seed(3172,9814,20125)"
```

We could now move to other properties that the predicate should comply with (e.g., all elements in the output list being present in the input list). Often, both traditional unit tests and QuickCheck tests are used, complementing each other to ensure the required code coverage.

Another example using a Prolog module predicate:

```
| ?- lgtunit::quick_check(
    pairs:pairs_keys_values(
        +list(pair(atom, integer)),
        -list(atom),
        -list(integer)
    )
)
```

(continues on next page)

(continued from previous page)

```

    ).
% 100 random tests passed, 0 discarded
% starting seed: seed(3172,9814,20125)
yes

```

As illustrated by the examples above, properties are expressed using predicates. In the most simple cases, that can be the predicate that we are testing itself. But, in general, it will be an auxiliary predicate calling the predicate or predicates being tested and checking properties that the results must comply with.

The QuickCheck test dialects and predicates take as argument the mode template for a property, generate random values for each input argument based on the type information, and check each output argument. For common types, the implementation tries first (by default) common edge cases (e.g., empty atom, empty list, or zero) before generating arbitrary values. When the output arguments check fails, the QuickCheck implementation tries (by default) up to 64 shrink operations of the counter-example to report a simpler case to help debugging the failed test. Edge cases, generating arbitrary terms, and shrinking terms make use of the library arbitrary category via the type object (both entities can be extended by the user by defining clauses for multifile predicates).

The mode template syntax is the same as that used in the `info/2` predicate directives with an additional notation, `{}/1`, for passing argument values as-is instead of generating random values for these arguments. For example, assume that we want to verify the `type::valid/2` predicate, which takes as its first argument a type. Randomly generating random types would be cumbersome at best but the main problem is that we need to generate random values for the second argument according to the first argument. Using the `{}/1` notation, we can solve this problem for any specific type, e.g. integer, by writing:

```
| ?- lgtunit::quick_check(type::valid({integer}, +integer)).
```

We can also test all (ground, i.e. non-parametric) types with arbitrary value generators by writing:

```
| ?- forall(
    (type::type(Type), ground(Type), type::arbitrary(Type)),
    lgtunit::quick_check(type::valid({Type}, +Type))
).
```

You can find the list of the basic supported types for use in the template in the API documentation for the library entities `type` and `arbitrary`. Note that other library entities, including third-party or your own, can contribute with additional type definitions, as both `type` and `arbitrary` entities are user-extensible by defining clauses for their multifile predicates.

The user can define new types to use in the property mode templates to use with its QuickCheck tests by defining clauses for the `type` library object and the `arbitrary` library category multifile predicates. QuickCheck will use the later to generate arbitrary input arguments and the former to verify output arguments. As a toy example, assume that the property mode template has an argument of type `bit` with possible values `0` and `1`. We would then need to define:

```
:- multifile(type::type/1).
type::type(bit).

:- multifile(type::check/2).
type::check(bit, Term) :-
    once((Term == 0; Term == 1)).

:- multifile(arbitrary::arbitrary/1).
arbitrary::arbitrary(bit).
```

(continues on next page)

(continued from previous page)

```
:- multifile(arbitrary::arbitrary/2).
arbitrary::arbitrary(bit, Arbitrary) :-
    random::member(Arbitrary, [0, 1]).
```

5.13.11 Skipping tests

A test object can define the `condition/0` predicate (which defaults to `true`) to test if some necessary condition for running the tests holds. The tests are skipped if the call to this predicate fails or generates an error.

Individual tests that for some reason should be unconditionally skipped can have the test clause head prefixed with the `(-)/1` operator. For example:

```
- test(not_yet_ready) :-
    ...
```

In this case, it's a good idea to use the `test/3` dialect with a `note/1` option that briefly explains why the test is being skipped. For example:

```
- test(xyz_reset, true, [note('Feature xyz reset not yet implemented')]) :-
    ...
```

The number of skipped tests is reported together with the numbers of passed and failed tests. To skip a test depending on some condition, use the `test/3` dialect and the `condition/1` option. For example:

```
test(test_id, true, [condition(current_prolog_flag(bounded, true))]) :-
    ...
```

The test is skipped if the condition goal fails or generates an error. The conditional compilation directives can also be used in alternative, but note that in this case there will be no report on the number of skipped tests.

5.13.12 Selecting tests

While debugging an application, we often want to temporarily run just a selection of relevant tests. This is specially useful when running all the tests slows down and distracts from testing fixes for a specific issue. This can be accomplished by prefixing the clause heads of the selected tests with the `(+)/1` operator. For example:

```
:- object(tests,
    extends(lgtunit)).

cover(ack).

test(ack_1, true(Result == 11)) :-
    ack::ack(2, 4, Result).

+ test(ack_2, true(Result == 61)) :-
    ack::ack(3, 3, Result).

test(ack_3, true(Result == 125)) :-
```

(continues on next page)

(continued from previous page)

```

ack::ack(3, 4, Result).

:- end_object.

```

In this case, only the `ack_2` would run. Just be careful to remove all `(+)/1` test prefixes when done debugging the issue that prompted you to run just the selected tests. After, be sure to run all the tests to ensure there are no regressions introduced by your fixes.

5.13.13 Checking test goal results

Checking test goal results can be performed using the `test/2-3` supported outcomes such as `true(Assertion)` and `deterministic(Assertion)`. For example:

```

test(compare_3_order_less, deterministic(Order == (<))) :-
    compare(Order, 1, 2).

```

For the other test dialects, checking test goal results can be performed by calling the `assertion/1-2` utility predicates or by writing the checking goals directly in the test body. For example:

```

test(compare_3_order_less) :-
    compare(Order, 1, 2),
    ^^assertion(Order == (<)).

```

or:

```

succeeds(compare_3_order_less) :-
    compare(Order, 1, 2),
    Order == (<).

```

Using assertions is, however, preferable to directly checking test results in the test body as it facilitates debugging by printing the unexpected results when the assertions fail.

The `assertion/1-2` utility predicates are also useful for the `test/2-3` dialects when we want to check multiple assertions in the same test. For example:

```

test(dictionary_clone_4_01, true) :-
    as_dictionary([], Dictionary),
    clone(Dictionary, DictionaryPairs, Clone, ClonePairs),
    empty(Clone),
    ^^assertion(original_pairs, DictionaryPairs == []),
    ^^assertion(clone_pairs, ClonePairs == []).

```

Ground results can be compared using the standard `==/2` term equality built-in predicate. Non-ground results can be compared using the `variant/2` predicate provided by `lgtunit`. The standard `subsumes_term/2` built-in predicate can be used when testing a compound term structure while abstracting some of its arguments. Floating-point numbers can be compared using the `==~/2`, `approximately_equal/3`, `essentially_equal/3`, and `tolerance_equal/4` predicates provided by `lgtunit`. Using the `==/2` term unification built-in predicate is almost always an error, as it would mask test goals failing to bind output arguments. The `lgtunit` tool implements a linter check for the use of unification goals in test outcome assertions. In the rare cases that a unification goal is intended, wrapping the `(=)/2` goal using the `{}/1` control construct avoids the linter warning.

When the meta-argument of the `assertion/1-2` predicates is call to a local predicate (in the tests object), you need to call them using the `(:)/2` message-sending control construct instead of the `(^^)/2` *super* call

control construct. This is necessary as *super* calls preserve the *sender*, and the tests are implicitly run by the *lgtunit* object sending a message to the tests object. For example:

```
:- uses(lgtunit, [
    assertion/1
]).

test(my_test_id, true) :-
    foo(X, Y),
    assertion(consistent(X, Y)).

consistent(X, Y) :-
    ...
```

In this case, the *sender* is the tests object, and the `assertion/1` meta-predicate will call the local `consistent/2` predicate in the expected context.

5.13.14 Testing local predicates

The `<</2` debugging control construct can be used to access and test object local predicates (i.e., predicates without a scope directive). In this case, make sure that the `context_switching_calls` compiler flag is set to allow for those objects. This is seldom required, however, as local predicates are usually auxiliary predicates called by public predicates and thus tested when testing those public predicates. The code coverage support can pinpoint any local predicate clause that is not being exercised by the tests.

5.13.15 Testing non-deterministic predicates

For testing non-deterministic predicates (with a finite and manageable number of solutions), you can wrap the test goal using the standard `findall/3` predicate to collect all solutions and check against the list of expected solutions. When the expected solutions are a set, use in alternative the standard `setof/3` predicate.

If you want to check that all solutions of a non-deterministic predicate satisfy an assertion, use the `test/2` or `test/3` test dialect with the `all(Assertion)` outcome. For example:

```
test(atom_list, all(atom(Item))) :-
    member(Item, [a, b, c]).
```

See also the next section on testing *generators*.

If you want to check that a solution exists for a non-deterministic predicate that satisfies an assertion, use the `test/2` or `test/3` test dialect with the `exists(Assertion)` outcome. For example:

```
test(at_least_one_atom, exists(atom(Item))) :-
    member(Item, [1, foo(2), 3.14, abc, 42]).
```

5.13.16 Testing generators

To test all solutions of a predicate that acts as a *generator*, we can use either the `all/1` outcome or the `forall/2` predicate as the test goal with the `assertion/2` predicate called to report details on any solution that fails the test. For example:

```
test(test_solution_generator, all(test(X,Y,Z))) :-
    generator(X, Y, Z).
```

or:

```
:- uses(lgtunit, [assertion/2]).
...

test(test_solution_generator_2) :-
    forall(
        generator(X, Y, Z),
        assertion(generator(X), test(X,Y,Z))
    ).
```

While using the `all/1` outcome results in a more compact test definition, using the `forall/2` predicate allows customizing the assertion description. In the example above, we use the `generator(X)` description instead of the `test(X,Y,Z)` description implicit when we use the `all/1` outcome.

5.13.17 Testing input/output predicates

Extensive support for testing input/output predicates is provided, based on similar support found on the Prolog conformance testing framework written by Péter Szabó and Péter Szeredi.

Two sets of predicates are provided, one for testing text input/output and one for testing binary input/output. In both cases, temporary files (possibly referenced by a user-defined alias) are used. The predicates allow setting, checking, and cleaning text/binary input/output. These predicates are declared as protected and thus called using the `(^^/1)` control construct.

As an example of testing an input predicate, consider the standard `get_char/1` predicate. This predicate reads a single character (atom) from the current input stream. Some test for basic functionality could be:

```
test(get_char_1_01, true(Char == 'q')) :-
    ^^set_text_input('qwerty'),
    get_char(Char).

test(get_char_1_02, true(Assertion)) :-
    ^^set_text_input('qwerty'),
    get_char(_Char),
    ^^text_input_assertion('werty', Assertion).
```

As you can see in the above example, the testing pattern consists on setting the input for the predicate being tested, calling it, and then checking the results. It is also possible to work with streams other than the current input/output streams by using the `lgtunit` predicate variants that take a stream alias as argument. For example, when testing the standard `get_char/2` predicate, we could write:

```
test(get_char_2_01, true(Char == 'q')) :-
    ^^set_text_input(in, 'qwerty'),
    get_char(in, Char).
```

(continues on next page)

(continued from previous page)

```
test(get_char_2_02, true(Assertion)) :-
    ^^set_text_input(in, 'qwerty'),
    get_char(in, _Char),
    ^^text_input_assertion(in, 'werty', Assertion).
```

Testing output predicates follows a similar pattern by using instead the `set_text_output/1-2` and `text_output_assertion/2-3` predicates. For example:

```
test(put_char_2_02, true(Assertion)) :-
    ^^set_text_output(out, 'qwerty'),
    put_char(out, y),
    ^^text_output_assertion(out, 'qwerty', Assertion).
```

The `set_text_output/1` predicate diverts only the standard output stream (to a temporary file) using the standard `set_output/1` predicate. Most backend Prolog systems also support writing to the de facto standard error stream. But there's no standard solution to divert this stream. However, several systems provide a `set_stream/2` or similar predicate that can be used for stream redirection. For example, assume that you wanted to test a backend Prolog system warning when an `initialization/1` directive fails that is written to `user_error`. An hypothetical test could be:

```
test(singletons_warning, true(Assertion)) :-
    ^^set_text_output(''),
    current_output(Stream),
    set_stream(Stream, alias(user_error)),
    consult(broken_file),
    ^^text_output_assertion('WARNING: initialization/1 directive failed', Assertion).
```

For testing binary input/output predicates, equivalent testing predicates are provided. There is also a small set of helper predicates for dealing with stream handles and stream positions. For testing with files instead of streams, testing predicates are provided that allow creating text and binary files with given contents and check text and binary files for expected contents.

For more practical examples, check the included tests for Prolog standard conformance of built-in input/output predicates.

5.13.18 Suppressing tested predicates output

Sometimes predicates being tested output text or binary data that, at best, clutters testing logs and, at worst, can interfere with parsing of test logs. If that output itself is not under testing, you can suppress it by using the goals `^^suppress_text_output` or `^^suppress_binary_output` at the beginning of the tests. For example:

```
test(proxies_04, true(Color == yellow)) :-
    ^^suppress_text_output,
    {circle('#2', Color)}::print.
```

The `suppress_text_output/0` and `suppress_binary_output/0` predicates work by redirecting standard output to the operating-system null device. But the application may also output to e.g. `user_error` and other streams. If this output must also be suppressed, several alternatives are described next.

Output of expected warnings can be suppressed by turning off the corresponding linter flags. In this case, it is advisable to restrict the scope of the flag value changes as much as possible.

Output of expected compiler errors can be suppressed by defining suitable clauses for the `logtalk::message_hook/4` hook predicate. For example:

```
:- multifile(logtalk::message_hook/4).
:- dynamic(logtalk::message_hook/4).

% ignore expected domain error
logtalk::message_hook(compiler_error(_,_,error(domain_error(foo,bar),_)), error, core, _).
```

In this case, it is advisable to restrict the scope of the clauses as much as possible to exact exception terms. For the exact message terms, see the `core_messages` category source file. Defining this hook predicate can also be used to suppress all messages from a given component. For example:

```
:- multifile(logtalk::message_hook/4).
:- dynamic(logtalk::message_hook/4).

logtalk::message_hook(_Message, _Kind, code_metrics, _Tokens).
```

Note that there's no portable solution to suppress *all* output. However, several systems provide a `set_stream/2` or similar predicate that can be used for stream redirection. Check the documentation of the backend Prolog systems you're using for details.

5.13.19 Tests with timeout limits

There's no portable way to call a goal with a timeout limit. However, some backend Prolog compilers provide this functionality:

- B-Prolog: `time_out/3` built-in predicate
- ECLiPSe: `timeout/3` and `timeout/7` library predicates
- XVM: `call_with_timeout/2-3` built-in predicates
- SICStus Prolog: `time_out/3` library predicate
- SWI-Prolog: `call_with_time_limit/2` library predicate
- Trealla Prolog: `call_with_time_limit/2` and `time_out/3` library predicates
- XSB: `timed_call/2` built-in predicate
- YAP: `time_out/3` library predicate

Logtalk provides a timeout portability library implementing a simple abstraction for those backend Prolog compilers.

The `logtalk_tester` automation script accepts a timeout option that can be used to set a limit per-test set.

5.13.20 Setup and cleanup goals

A test object can define `setup/0` and `cleanup/0` goals. The `setup/0` predicate is called, when defined, before running the object unit tests. The `cleanup/0` predicate is called, when defined, after running all the object unit tests. The tests are skipped when the setup goal fails or throws an error. For example:

```
cleanup :-
    this(This),
    object_property(This, file(_,Directory)),
    atom_concat(Directory, serialized_objects, File),
    catch(ignore(os::delete_file(File)), _, true).
```

Per test setup and cleanup goals can be defined using the `test/3` dialect and the `setup/1` and `cleanup/1` options. The test is skipped when the setup goal fails or throws an error. Note that a broken test cleanup goal doesn't affect the test but may adversely affect any following tests. Variables in the setup and cleanup goals are shared with the test body.

5.13.21 Test annotations

It's possible to define per-unit and per-test annotations to be printed after the test results or when tests are skipped. This is particularly useful when some units or some unit tests may be run while still being developed. Annotations can be used to pass additional information to a user reviewing test results. By intercepting the unit test framework message printing calls (using the `message_hook/4` hook predicate), test automation scripts and integrating tools can also access these annotations.

Units can define a global annotation using the predicate `note/1`. To define per-test annotations, use the `test/3` dialect and the `note/1` option. For example, you can inform why a test is being skipped by writing:

```
- test(foo_1, true, [note('Waiting for Deep Thought answer')]) :-
    ...
```

Another common use is to return the execution time of one of the test sub-goals. For example:

```
test(foobar, true, [note(bar(seconds-Time))]) :-
    foo(...),
    benchmark(bar(...), Time).
```

Annotations are written, by default, between parentheses after and in the same line as the test results.

5.13.22 Test execution times and memory usage

Individual test CPU and wall execution times (in seconds) are reported by default when running the tests. Total CPU and wall execution times for passed and failed tests are reported after the tests complete. Starting and ending date and time when running a set of tests is also reported by default. The `lgtunit` object also provides several public benchmarking predicates that can be useful for e.g. reporting test sub-goals execution times using either CPU or wall clocks. When running multi-threaded code, the CPU time may or may not include all threads CPU time depending on the backend.

Be aware that the accuracy of CPU and wall time depends of the backend. Accuracy can also be different between CPU and wall time (e.g. CPU time can have nanosecond accuracy with wall time only having millisecond accuracy).

Test memory usage is not reported by default due to the lack of a portable solution to access memory data. However, several backend Prolog systems provide a `statistics/2` or similar predicate that can be used for

a custom solution. Depending on the system, individual keys may be provided for each memory area (heap, trail, atom table, ...). Aggregating keys may also be provided. As a hypothetical example, assume you're running Logtalk with a backend providing a `statistics/2` predicate with a `memory_used` key:

```
test(ack_3, true(Result == 125), [note(memory-Memory)]) :-
    statistics(memory_used, Memory0),
    ack::ack(3, 4, Result),
    statistics(memory_used, Memory1),
    Memory is Memory1 - Memory0.
```

Consult the documentation of the backend Prolog systems for actual details.

5.13.23 Working with test data files

Frequently, tests make use of test data files that are usually stored in the test set directory or in sub-directories. These data files are referenced using their relative paths. But to allow the tests to run independently of the Logtalk process current directory, the relative paths often must be expanded into an absolute path before being passed to the predicates being tested. The `file_path/2` protected predicate can be used in the test definitions to expand the relative paths. For example:

```
% check that the encoding/1 option is accepted
test(lgt_unicode_open_4_01, true) :-
    ^^file_path(sample_utf_8, Path),
    open(Path, write, Stream, [encoding('UTF-8')]),
    close(Stream).
```

The absolute path is computed relative to the path of *self*, i.e. relative to the path of the test object that received the message that runs the tests.

It's also common for tests to create temporary files and directories that should be deleted after the tests completion. The `clean_file/1` and `clean_directory/1` protected predicates can be used for this purpose. For example, assuming that the tests create a `foo.txt` text file and a `tmp` directory in the same directory as the tests object:

```
cleanup :-
    ^^clean_file('foo.txt'),
    ^^clean_directory('tmp').
```

Similar to the `file_path/2` predicate, relative paths are interpreted as relative to the path of the test object. This predicate also closes any open stream connected to the file before deleting it.

5.13.24 Flaky tests

Flaky tests are tests that pass or fail non-deterministically, usually due to external conditions (e.g., computer or network load). Thus, flaky tests often don't result from bugs in the code being tested itself but from test execution conditions that are not predictable. The `flaky/0` test option declares a test to be flaky. For example:

```
test(foo, true, [flaky]) :-
    ...
```

For backwards compatibility, the `note/1` annotation can also be used to alert that a test failure is for a flaky test when its argument is an atom containing the sub-atom `flaky`.

The testing automation support outputs the text [flaky] when reporting failed flaky tests. Moreover, the `logtalk_tester` automation script will ignore failed flaky tests when setting its exit status.

5.13.25 Mocking

Sometimes the code being tested performs complex tasks that are not feasible or desirable when running tests. For example, the code may perform a login operation requiring the user to provide a username and a password using some GUI widget. In this case, the tests may require the login operation to still be performed but using canned data (also simplifying testing automation). I.e. we want to *mock* (as in *imitate*) the login procedure. Ideally, this should be accomplished without requiring any changes to the code being tested. Logtalk provides two solutions that can be used for mocking: *term-expansion* and *hot patching*. A third solution is possible if the code we want to mock uses the *message printing mechanism*.

Using the term-expansion mechanism, we would define a *hook object* that expands the login predicate into a fact:

```
:- object(mock_login,
    implements(expanding)).

    term_expansion((login(_, _) :- _), login(jdoe, test123)).

:- end_object.
```

The tests driver file would then load the application object responsible for user management using this hook object:

```
:- initialization((
    ...,
    logtalk_load(mock_login),
    logtalk_load(user_management, [hook(mock_login)]),
    ...
)).
```

Using hot patching, we would define a *complementing category* patching the object that defines the login predicate:

```
:- category(mock_login,
    complements(user_management)).

    login(jdoe, test123).

:- end_category.
```

The tests driver file would then set the complements flag to allow and load the patch after loading application code:

```
:- initialization((
    ...,
    set_logtalk_flag(complements, allow),
    logtalk_load(application),
    logtalk_load(mock_login),
    ...
)).
```

There are pros and cons for each solution. Term-expansion works by defining hook objects that are used at compile-time, while hot patching happens at runtime. Complementing categories can also be dynamically created, stacked, and abolished. Hot patching disables static binding optimizations, but that's usually not a problem as the code being tested is often compiled in debug mode to collect code coverage data. Two advantages of the term-expansion solution are that it allows defining conditions for expanding terms and goals and can replace both predicate definitions and predicate calls. Limitations in the current Prolog standards prevent patching callers to local predicates being patched. But often both solutions can be used, with the choice depending on code clarity and user preference. See the Handbook sections on term-expansion and hot patching for more details on these mechanisms.

In those cases where the code we want to mock uses the message printing mechanism, the solution is to intercept and rewrite the messages being printed and/or the questions being asked using the `logtalk::message_hook/4` and `logtalk::question_hook/6` hook predicates.

5.13.26 Debugging messages in tests

Sometimes it is useful to write debugging or logging messages from tests when running them manually. But those messages are better suppressed when running the tests automated. A common solution is to use debug *meta-messages*. For example:

```
:- uses(logtalk, [
    print_message(debug, my_app, Message) as dbg(Message)
]).

test(some_test_id, ...) :-
    ...,
    dbg('Some intermediate value'-Value),
    ...,
    dbg([Stream]>>custom_print_goal(Stream, ...)),
    ...
```

The messages are only printed (and the user-defined printing goals are only called) when the debug flag is turned on. Note that this doesn't require compiling the tests in debug mode: you simply toggle the flag to toggle the debug messages. Also note that the `print_message/3` goals are suppressed by the compiler when compiling with the optimize flag turned on.

5.13.27 Debugging failed tests

Debugging of failed unit tests is simplified by using test assertions as the reason for the assertion failures is printed out. Thus, use preferably the `test/2-3` dialects with `true(Assertion)`, `deterministic(Assertion)`, `subsumes(Expected, Result)`, or `variant(Term1, Term2)` outcomes. If a test checks multiple assertions, you can use the predicate `assertion/2` in the test body. In the case of QuickCheck tests, the `v(true)` verbose option can be used to print the generated test case that failed if necessary.

If the assertion failures don't provide enough information, you can use the debugger tool to debug failed unit tests. Start by compiling the unit test objects and the code being tested in debug mode. Load the debugger and trace the test that you want to debug. For example, assuming your tests are defined in a `tests` object and that the identifier of the test to be debugged is `test_foo`:

```
| ?- logtalk_load(debugger(loader)).
...
| ?- debugger::trace.
```

(continues on next page)

(continued from previous page)

```
...
| ?- tests::run(test_foo).
...
```

You can also compile the code and the tests in debug mode but without using the `hook/1` compiler option for the tests compilation. Assuming that the `context_switching_calls` flag is set to allow, you can then use the `(<<)/2` debugging control construct to debug the tests. For example, assuming that the identifier of the test to be debugged is `test_foo` and that you used the `test/1` dialect:

```
| ?- logtalk_load(debugger(loader)).
...
| ?- debugger::trace.
...
| ?- tests<<test(test_foo).
...
```

In the more complicated cases, it may be worth defining `loader_debug.lgt` and `tester_debug.lgt` driver files that load code and tests in debug mode and also load the debugger.

When using the `logtalk_tester` automation script, some test sets can be reported as broken or crashed. The script creates a logs directory, by default named `logtalk_tester_logs`, in the same directory where the script was called. The log files, notably the `*.errors` files, often contain valuable information for debugging broken and crashed test sets.

5.13.28 Code coverage

If you want entity predicate clause coverage information to be collected and printed, you will need to compile the entities that you're testing using the flags `debug(on)` and `source_data(on)`. Be aware, however, that compiling in debug mode results in a performance penalty.

A single test object may include tests for one or more entities (objects, protocols, and categories). The entities being tested by a unit test object for which code coverage information should be collected must be declared using the `cover/1` predicate. For example, to collect code coverage data for the objects `foo` and `bar` include in the tests object the two clauses:

```
cover(foo).
cover(bar).
```

Code coverage is listed using the predicates clause indexes (counting from one). For example, using the `points` example in the Logtalk distribution:

```
% point: default_init_option/1 - 2/2 - (all)
% point: instance_base_name/1 - 1/1 - (all)
% point: move/2 - 1/1 - (all)
% point: position/2 - 1/1 - (all)
% point: print/0 - 1/1 - (all)
% point: process_init_option/1 - 1/2 - [1]
% point: position_/2 - 0/0 - (all)
% point: 7 out of 8 clauses covered, 87.500000% coverage
```

The numbers after the predicate indicators represent the clauses covered and the total number of clauses. E.g. for the `process_init_option/1` predicate, the tests cover 1 out of 2 clauses. After these numbers, we either get `(all)` telling us that all clauses are covered or a list of indexes for the covered clauses. E.g. only the first clause for the `process_init_option/1` predicate, `[1]`. Summary clause coverage numbers are also printed for entities and for clauses across all entities.

In the printed predicate clause coverage information, you may get a total number of clauses smaller than the covered clauses. This results from the use of dynamic predicates with clauses asserted at runtime. You may easily identify dynamic predicates in the results as their clauses often have an initial count equal to zero.

The list of indexes of the covered predicate clauses can be quite long. Some backend Prolog compilers provide a flag or a predicate to control the depth of printed terms that can be useful:

- CxProlog: `write_depth/2` predicate
- ECLiPSe: `print_depth` flag
- SICStus Prolog: `toplevel_print_options` flag
- SWI-Prolog 7.1.10 or earlier: `toplevel_print_options` flag
- SWI-Prolog 7.1.11 or later: `answer_write_options` flag
- Trealla Prolog: `answer_write_options` flag
- XSB: `set_file_write_depth/1` predicate
- XVM 3.2.0 or later: `answer_write_options` flag
- YAP: `write_depth/2-3` predicates

Code coverage is only available when testing Logtalk code. But Prolog modules can often be compiled as Logtalk objects and plain Prolog code may be wrapped in a Logtalk object. For example, assuming a `module.pl` module file, we can compile and load the module as an object by simply calling:

```
| ?- logtalk_load(module).  
...
```

The module exported predicates become object public predicates. For a plain Prolog file, say `plain.pl`, we can define a Logtalk object that wraps the code using an `include/1` directive:

```
:- object(plain).  
    :- include('plain.pl').  
:- end_object.
```

The object can also declare as public the top Prolog predicates to simplify writing the tests. In alternative, we can use the `object_wrapper_hook` provided by the `hook_objects` library:

```
| ?- logtalk_load(hook_objects(loader)).  
...  
  
| ?- logtalk_load(plain, [hook(object_wrapper_hook)]).  
...
```

These workarounds may thus allow generating code coverage data also for Prolog code by defining tests that use the `<</2` debugging control construct to call the Prolog predicates.

See also the section below on exporting code coverage results to XML files, which can be easily converted and published as e.g. HTML reports.

5.13.29 Utility predicates

The lgtunit tool provides several public utility predicates to simplify writing unit tests and for general use:

- `variant(Term1, Term2)`
To check when two terms are a variant of each other (e.g., to check expected test results against actual results when they contain variables).
- `assertion(Goal)`
To generate an exception in case the goal argument fails or throws an error.
- `assertion(Description, Goal)`
To generate an exception in case the goal argument fails or throws an error (the first argument allows assertion failures to be distinguished when using multiple assertions).
- `approximately_equal(Number1, Number2)`
For number approximate equality using the epsilon arithmetic constant value.
- `approximately_equal(Number1, Number2, Epsilon)`
For number approximate equality. Weaker equality than essential equality.
- `essentially_equal(Number1, Number2, Epsilon)`
For number essential equality. Stronger equality than approximate equality.
- `tolerance_equal(Number1, Number2, RelativeTolerance, AbsoluteTolerance)`
For number equality within tolerances.
- `Number1 == Number2`
For number (or list of numbers) close equality (usually floating-point numbers).
- `benchmark(Goal, Time)`
For timing a goal.
- `benchmark_reified(Goal, Time, Result)`
Reified version of `benchmark/2`.
- `benchmark(Goal, Repetitions, Time)`
For finding the average time to prove a goal.
- `benchmark(Goal, Repetitions, Clock, Time)`
For finding the average time to prove a goal using a cpu or a wall clock.
- `deterministic(Goal)`
For checking that a predicate succeeds without leaving a choice-point.
- `deterministic(Goal, Deterministic)`
Reified version of the `deterministic/1` predicate.

The `assertion/1-2` predicates can be used in the body of tests where using two or more assertions is convenient or in the body of tests written using the `test/1`, `succeeds/1`, and `deterministic/1` dialects to help differentiate between the test goal and checking the test goal results and to provide more informative test failure messages.

When the `assertion`, `benchmarking`, and `deterministic` meta-predicates call a local predicate of the tests object, you must call them using an implicit or explicit message instead of using a *super* call. For example, to use an implicit message to call the `assertion/1-2` meta-predicates, add the following directive to the tests object:

```
:- uses(lgtunit, [assertion/1, assertion/2]).
```

The reason this is required is that meta-predicates goals arguments are always called in the context of the *sender*, which would be the `lgtunit` object in the case of a `(^^)/2` call (as it preserves both *self* and *sender* and the tests are internally run by a message sent from the `lgtunit` object to the tests object).

As the `benchmark/2-4` predicates are meta-predicates, turning on the `optimize` compiler flag is advised to avoid runtime compilation of the meta-argument, which would add an overhead to the timing results. But this advice conflicts with collecting code coverage data, which requires compilation in debug mode. The solution is to use separate test objects for benchmarking and for code coverage. Note that the CPU and wall execution times (in seconds) for each individual test are reported by default when running the tests.

The `(==)/2` predicate is typically used by adding the following directive to the object (or category) calling it:

```
:- uses(lgtunit, [
    op(700, xfx, ==), (==)/2
]).
```

Consult the `lgtunit` object API documentation for more details on these predicates.

5.13.30 Exporting test results

xUnit XML format

To output test results in the xUnit XML format (from JUnit; see e.g. <https://github.com/windyroad/JUnit-Schema> or <https://llg.cubic.org/docs/junit/>), simply load the `xunit_output.lgt` file before running the tests. This file defines an object, `xunit_output`, that intercepts and rewrites unit test execution messages, converting them to the xUnit XML format.

To export the test results to a file using the xUnit XML format, simply load the `xunit_report.lgt` file before running the tests. A file named `xunit_report.xml` will be created in the same directory as the object defining the tests. When running a set of test suites as a single unified suite (using the `run_test_sets/1` predicate), the single xUnit report is created in the directory of the first test suite object in the set.

To use the xUnit.net v2 XML format (<https://xunit.net/docs/format-xml-v2>), load either the `xunit_net_v2_output.lgt` file or the `xunit_net_v2_report.lgt` file.

When using the `logtalk_tester` automation script, use either the `-f xunit` option or the `-f xunit_net_v2` option to generate the `xunit_report.xml` files in the test set directories.

There are several third-party xUnit report converters that can generate HTML files for easy browsing. For example:

- <https://allurereport.org/docs/> (supports multiple reports)
- <https://github.com/Zir0-93/xunit-to-html> (supports multiple test sets in a single report)
- <https://www.npmjs.com/package/xunit-viewer>
- <https://github.com/JatechUK/NUnit-HTML-Report-Generator>
- <https://plugins.jenkins.io/xunit>

TAP output format

To output test results in the TAP (Test Anything Protocol) format, simply load the `tap_output.lgt` file before running the tests. This file defines an object, `tap_output`, that intercepts and rewrites unit test execution messages, converting them to the TAP output format.

To export the test results to a file using the TAP (Test Anything Protocol) output format, load instead the `tap_report.lgt` file before running the tests. A file named `tap_report.txt` will be created in the same directory as the object defining the tests.

When using the `logtalk_tester` automation script, use the `-f tap` option to generate the `tap_report.xml` files in the test set directories.

When using the `test/3` dialect with the TAP format, a `note/1` option whose argument is an atom starting with a `TODO` or `todo` word results in a test report with a TAP `TODO` directive.

When running a set of test suites as a single unified suite, the single TAP report is created in the directory of the first test suite object in the set.

There are several third-party TAP report converters that can generate HTML files for easy browsing. For example:

- <https://github.com/Quobject/tap-to-html>
- <https://plugins.jenkins.io/tap/>

CTRF JSON format

To output test results in the CTRF (Common Test Report Format) JSON format, simply load the `ctrf_output.lgt` file before running the tests. This file defines an object, `ctrf_output`, that intercepts and rewrites unit test execution messages, converting them to the CTRF JSON format.

To export the test results to a file using the CTRF JSON format, load instead the `ctrf_report.lgt` file before running the tests. A file named `ctrf_report.json` will be created in the same directory as the object defining the tests.

When using the `logtalk_tester` automation script, use the `-f ctrf` option to generate the `ctrf_report.json` files in the test set directories.

When running a set of test suites as a single unified suite, the single CTRF report is created in the directory of the first test suite object in the set.

Subunit formats

To output test results in the Subunit v1 text streaming format, simply load the `subunit_v1_output.lgt` file before running the tests. This file defines an object, `subunit_v1_output`, that intercepts and rewrites unit test execution messages, converting them to a Subunit text stream.

To output test results in the Subunit v2 binary streaming format, load the `subunit_v2_output.lgt` file before running the tests. This file defines an object, `subunit_v2_output`, that intercepts and rewrites unit test execution messages, converting them to Subunit v2 binary packets.

These adapters write directly to the current output stream and are intended for piping test execution output to tools that parse Subunit streams.

To export test results to files using Subunit streaming formats, load instead the `subunit_v1_report.lgt` file (Subunit text stream) or the `subunit_v2_report.lgt` file (Subunit v2 binary stream) before running the tests. Files named `subunit_v1_report.txt` and `subunit_v2_report.bin` are created in the same directory as the object defining the tests.

5.13.31 Generating Allure reports

A `logtalk_allure_report.pl` Bash shell script and a `logtalk_allure_report.ps1` PowerShell script are provided for generating **Allure reports** (version 2.26.0 or later required). This requires exporting test results in xUnit XML format. A simple usage example (assuming a current directory containing tests):

```
$ logtalk_tester -p gnu -f xunit
$ logtalk_allure_report
$ allure open
```

The `allure open` command accepts `--host` and `--port` arguments in case their default values are not suitable (e.g., when running Logtalk in a remote host over a SSH connection).

The `logtalk_allure_report` script supports command-line options to pass the tests directory (i.e., the directory where the `logtalk_tester` script was run), the directory where to collect all the xUnit report files for generating the report, the directory where the report is to be saved, and the report title (see the script man page or type `logtalk_allure_report -h`). The script also supports saving the history of past test runs. In this case, a persistent location for both the results and report directories must be used.

It's also possible to use the script just to collect the xUnit report files generated by `lgtunit` and delegate the actual generation of the report to e.g. an Allure Docker container or to a Jenkins plug-in. Two examples are:

- <https://github.com/fescobar/allure-docker-service>
- <https://plugins.jenkins.io/allure-jenkins-plugin/>

In this case, we would use the `logtalk_allure_report` script option to only perform the preprocessing step:

```
$ logtalk_allure_report -p
```

The script also supports passing *environment pairs*, which are displayed in the generated Allure reports in the environment pane. This feature can be used to pass e.g. the backend name and the backend version or git commit hash. The option syntax differs, however, between the two scripts. For example, using the Bash script:

```
$ logtalk_allure_report -- Backend='GNU Prolog' Version=1.5.0
```

Or:

```
$ logtalk_allure_report -- Project='Deep Thought' Commit=$(git rev-parse --short HEAD)
```

In the case of the PowerShell script, the pairs are passed comma separated inside a string:

```
PS> logtalk_allure_report -e "Backend='GNU Prolog',Version=1.5.0"
```

Or:

```
PS> logtalk_allure_report -e "Project='Deep Thought',Commit=bf166b6"
```

To show test run trends in the report (e.g., when running the tests for each application source code commit), save the processed test results and the report data to permanent directories. For example:

```
$ logtalk_allure_report \
-i "$HOME/my_project/allure-results" \
-o "$HOME/my_project/allure-report"
$ allure open "$HOME/my_project/allure-report"
```

Note that (in Allure 2.x) the single file option is not compatible with run trends.

The generated reports can include links to the tests source code. This requires using the `logtalk_tester` shell script option that allows passing the base URL for those links. This option needs to be used together with the option to suppress the tests directory prefix so that the links can be constructed by appending the tests file relative path to the base URL. For example, assuming that you want to generate a report for the tests included in the Logtalk distribution when using the GNU Prolog backend:

```
$ cd $LOGTALKUSER
$ logtalk_tester \
  -p gnu \
  -f xunit \
  -s "$LOGTALKUSER" \
  -u "https://github.com/LogtalkDotOrg/logtalk3/tree/3e4ea295986fb09d0d4aade1f3b4968e29ef594e"
↪
```

The use of a git hash in the base URL ensures that the generated links will always show the exact versions of the tests that were run. The links include the line number for the tests in the test files (assuming that the git repo is stored in a BitBucket, GitHub, or GitLab server). But note that not all supported backends provide accurate line numbers.

It's also possible to generate single-file reports. For example:

```
$ logtalk_allure_report -s -t "My Amazing Tests Report"
```

There are some caveats when generating Allure reports that users must be aware of. First, Allure expects test names to be unique across different tests sets. If there are two tests with the same name in two different test sets, only one of them will be reported. Second, when using the `xunit` format, dates are reported as `MM/DD/YYYY`. Finally, when using the `xunit_net_v2` format, tests are reported in a random order instead of their run order, and dates are displayed as “unknown” in the overview page.

5.13.32 Exporting code coverage results

The `lgtunit` tool can export code coverage stats in three formats by loading the corresponding report object before running tests:

- `coverage_report.lgt` for a simple XML report (`coverage_report.xml`)
- `cobertura_report.lgt` for a Cobertura XML report (`cobertura.xml`)
- `lcov_report.lgt` for an LCOV report (`lcov.info`)

When using the `logtalk_tester` automation scripts, select the report format using the `-c` option with one of the values `xml`, `cobertura`, or `lcov` (or none to disable code coverage reports).

Simple XML format

To export code coverage results in simple XML format, load the `coverage_report.lgt` file before running the tests. A file named `coverage_report.xml` will be created in the same directory as the object defining the tests.

The XML file can be opened in most web browsers (with the notorious exception of Google Chrome) by copying to the same directory the `coverage_report.dtd` and `coverage_report.xsl` files found in the `tools/lgtunit` directory (when using the `logtalk_tester` script, these two files are copied automatically). In alternative, an XSLT processor can be used to generate an XHTML file instead of relying on a web browser for the transformation. For example, using the popular `xsltproc` processor:

```
$ xsltproc -o coverage_report.html coverage_report.xml
```

On Windows operating-systems, this processor can be installed using e.g. Chocolatey. On a POSIX operating-systems (e.g., Linux, macOS, ...) use the system package manager to install it if necessary.

The coverage report can include links to the source code when hosted on Bitbucket, GitHub, or GitLab. This requires passing the base URL as the value for the `url` XSLT parameter. The exact syntax depends on the XSLT processor, however. For example:

```
$ xsltproc \  
--stringparam url https://github.com/LogtalkDotOrg/logtalk3/blob/master \  
-o coverage_report.html coverage_report.xml
```

Note that the base URL should preferably be a permanent link (i.e., it should include the commit SHA1) so that the links to source code files and lines remain valid if the source code is later updated. It's also necessary to suppress the local path prefix in the generated `coverage_report.xml` file. For example:

```
$ logtalk_tester -c xml -s $HOME/logtalk/
```

Alternatively, you can pass the local path prefix to be suppressed to the XSLT processor (note that the `logtalk_tester` script suppresses the `$HOME` prefix by default):

```
$ xsltproc \  
--stringparam prefix logtalk/ \  
--stringparam url https://github.com/LogtalkDotOrg/logtalk3/blob/master \  
-o coverage_report.html coverage_report.xml
```

If you are using Bitbucket, GitHub, or GitLab hosted on your own servers, the `url` parameter may not contain a `bitbucket`, `github`, or `gitlab` string. In this case, you can use the XSLT parameter `host` to indicate which service you are running.

Cobertura XML format

To export code coverage results in Cobertura XML format, load the `cobertura_report.lgt` file before running the tests. A file named `cobertura.xml` will be created in the same directory as the object defining the tests.

When using the `logtalk_tester` automation scripts, use the option:

```
$ logtalk_tester -c cobertura
```

LCOV format

To export code coverage results in LCOV format, load the `lcov_report.lgt` file before running the tests. A file named `lcov.info` will be created in the same directory as the object defining the tests.

When using the `logtalk_tester` automation scripts, use the option:

```
$ logtalk_tester -c lcov
```

5.13.33 Automatically creating bug reports at issue trackers

To automatically create bug report issues for failed tests in GitHub or GitLab servers, see the `issue_creator` tool.

5.13.34 Minimizing test results output

To minimize the test results output, simply load the `minimal_output.lgt` file before running the tests. This file defines an object, `minimal_output`, that intercepts and summarizes the unit test execution messages.

5.13.35 Help with warnings

Load the `tutor` tool to get help with selected warnings printed by the `lgtunit` tool.

5.13.36 Known issues

Deterministic unit tests are currently not available when using Quintus Prolog as it lacks built-in support that cannot be sensibly defined in Prolog.

Parameter variables (`_VariableName_`) cannot currently be used in the definition of the `condition/1`, `setup/1`, and `cleanup/1` test options when using the `test/3` dialect. For example, the following condition will not work:

```
test(some_id, true, [condition(_ParVar_ == 42)]) :-
    ...
```

The workaround is to define an auxiliary predicate called from those options. For example:

```
test(check_xyz, true, [condition(xyz_condition)]) :-
    ...

xyz_condition :-
    _ParVar_ == 42.
```

5.14 linter

Logtalk provides a built-in linter tool that runs automatically when compiling and loading source files. The lint warnings are controlled by a `set of flags`. The default values for these flags are defined in the backend Prolog compiler adapter files and can be overridden from a settings file, from a source file (e.g., a loader file), or from an entity. These flags can be set globally using the `set_logtalk_flag/2` built-in predicate. For (source file or entity) local scope, use instead the `set_logtalk_flag/2` directive.

The linter flags can be managed as a group using the `linter` meta-flag. See the documentation for details.

Some lint checks are turned off by default, specially when computationally expensive. Still, it's a good idea to turn them on to check your code on a regular basis (e.g., in CI/CD pipelines).

Note that, in some cases, the linter may generate false warnings due to source code analysis limitations or special cases that, while valid when intended, usually result from programming issues. When a code rewrite is not a sensible solution to avoid the warning, the workaround is to turn off as locally as possible the flag that controls the warning.

5.14.1 Main linter checks

Lint checks include:

- Missing directives (including scope, meta-predicate, dynamic, discontinuous, and multifile directives)
- Duplicated directives, clauses, and grammar rules
- Missing predicates (unknown messages plus calls to non-declared and non-defined predicates)
- Calls to declared but not defined static predicates
- Non-terminals called as predicates (instead of via the phrase/2-3 built-in methods)
- Predicates called as non-terminals (instead of via the call/1 built-in method)
- Non-portable predicate calls, predicate options, arithmetic function calls, directives, flags, and flag values
- Missing arithmetic functions (with selected backends)
- Suspicious calls (syntactically valid calls that are likely semantic errors; e.g. float comparisons using the standard arithmetic comparison operators or comparing numbers using unification)
- Deprecated directives, predicates, arithmetic functions, control constructs, and flags
- References to unknown entities (objects, protocols, categories, or modules)
- Top-level shortcuts used as directives
- Unification goals that will succeed without binding any variables
- Unification goals that will succeed by creating a cyclic term
- Goals that are always true or always false
- Trivial goal failures (due to no matching predicate clause)
- Redefined built-in predicates
- Redefined standard operators
- Lambda expression unclassified variables and mixed up variables
- Lambda expression with parameter variables used elsewhere in a clause
- Singleton variables
- If-then-else and soft cut control constructs without an else part
- If-then-else and soft cut control constructs where the test is a unification between a variable and a ground term
- Missing parentheses around if-then-else and disjunction control constructs in the presence of cuts in the first argument
- Cuts in clauses for multifile predicates
- Missing cut in repeat loops
- Possible non-steadfast predicate definitions
- Non-tail recursive predicate definitions
- Redundant calls to control constructs and built-in predicates
- Calls to all-solutions predicates with existentially qualified variables not occurring in the qualified goal
- Calls to all-solutions predicates with no shared variables between template and goal

- Calls to `bagof/3` and `setof/3` where the goal argument contains singleton variables
- Calls to `findall/3` used to backtrack over all solutions of a goal without collecting them
- Calls to `catch/3` that catch all exceptions
- Calls to standard predicates that have more efficient alternatives
- Unsound calls in grammar rules
- File, entity, predicate, and variable names not following official coding guidelines
- Variable names that differ only on case
- Clauses whose body is a disjunction (and that can be rewritten as multiple clauses per coding guidelines)
- Naked meta-variables in cut-transparent control constructs
- Left-recursion in clauses and grammar rules

Additional lint checks are provided by the `lgtunit`, `lgtdoc`, `make`, and `dead_code_scanner` tools. For large projects, the data generated by the `code_metrics` tool may also be relevant in accessing code quality and suggesting code refactoring candidates.

5.14.2 Help on linter warnings

By loading the tutor tool, most lint warnings are expanded with explanations and suggestions on how to fix the reported issues. See also the [coding guidelines](#) for additional explanations.

5.14.3 Extending the linter

Experimental support for extending the linter with user-defined warnings is available using the [logtalk_linter_hook/7](#) multifile hook predicate. For example, the `format` and `list` library objects define this hook predicate to lint calls to the `format/2-3` and `append/3` predicates for common errors and misuses.

5.14.4 Linting Prolog modules

This tool can also be applied to Prolog modules that Logtalk is able to compile as objects. For example, if the Prolog module file is named `module.pl`, try:

```
| ?- logtalk_load(module, [source_data(on)]).
```

Due to the lack of standardization of module systems and the abundance of proprietary extensions, this solution is not expected to work for all cases.

5.14.5 Linting plain Prolog files

This tool can also be applied to plain Prolog code. For example, if the Prolog file is named `code.pl`, simply define an object including its code:

```
:- object(code).
   :- include('code.pl').
:- end_object.
```

Save the object to an e.g. `code.lgt` file in the same directory as the Prolog file and then load it:

```
| ?- logtalk_load(code, [source_data(on)]).
```

In alternative, use the `object_wrapper_hook` provided by the `hook_objects` library:

```
| ?- logtalk_load(hook_objects(loader)).  
...  
  
| ?- logtalk_load(code, [hook(object_wrapper_hook), source_data(on)]).
```

With either wrapping solution, pay special attention to any compilation warnings that may signal issues that could prevent the plain Prolog from being fully checked when wrapped by an object.

5.15 linter_reporter

This tool intercepts compiler linter warnings and caches them as machine-readable diagnostics. These diagnostics can be queried directly or serialized as SARIF using the standalone `sarif` tool.

5.15.1 API documentation

This tool API documentation is available at:

[../..../apis/library_index.html#linter_reporter](http://logtalk.org/..../apis/library_index.html#linter_reporter)

5.15.2 Loading

Load the tool before loading the code to be checked:

```
| ?- logtalk_load(linter_reporter(loader)).  
...
```

Enable collecting linter warnings data using default options:

```
| ?- linter_reporter::enable.  
true.
```

Or using explicit options:

```
| ?- linter_reporter::enable([explanations(true)]).  
true.
```

Load the code for which you want diagnostics collected:

```
| ?- logtalk_load(my_application(loader)).  
...
```

Disable further collecting of linter warnings:

```
| ?- linter_reporter::disable.  
true.
```

Query the cached diagnostics directly:

```
| ?- linter_reporter::diagnostics(all, Diagnostics).
...
```

Warnings originating in an included file are not always file-scoped. When the `include/1` directive appears inside an entity, the diagnostic context can be that entity; otherwise the context is typically the included file.

Or generate a SARIF report using the standalone `sarif` tool:

```
| ?- logtalk_load(sarif(loader)).
...

| ?- sarif::generate(linter_reporter, all, file('./linter_warnings.sarif'), []).
true.
```

5.15.3 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(linter_reporter(tester)).
```

The test suite reuses errors example files to exercise representative built-in linter warnings and validates both the diagnostics API and standalone SARIF generation in explanation-disabled and explanation-enabled configurations.

5.15.4 Usage

Load the tool, call `enable/0-1` before compiling the code to be checked, call `disable/0` when warning collection is finished, and then query the cached warnings using either the legacy warning predicates or the diagnostics protocol predicates. To generate SARIF from the cached diagnostics, load the standalone `sarif` tool and call `sarif::generate(linter_reporter, all, file('./linter_warnings.sarif'), [])`.

5.15.5 Options

- `explanations(Boolean)` Boolean option accepted by `enable/1`. When set to `true` (default is `false`), the tool enriches warnings with explanations from the `tutor_explanations` category provided by the tutor tool.

5.16 make

Logtalk provides a `make` tool supporting several targets using the `logtalk_make/0-1` built-in predicates. Top-level shortcuts for the targets are also provided.

5.16.1 API documentation

To consult the documentation of the `logtalk_make/0-1` built-in predicates, open in a web browser the links:

- [../refman/predicates/logtalk_make_0.html](#)
- [../refman/predicates/logtalk_make_1.html](#)

There is also a user-defined hook predicate that supports defining additional actions for the make targets (e.g., running tests automatically on make check or regenerating API documentation on make documentation):

- [../refman/predicates/logtalk_make_target_action_1.html](#)

5.16.2 Help with warnings

Load the tutor tool to get help with selected warnings printed by the make tool.

5.16.3 Known issues

The implementation of the `logtalk_make/0-1` predicates for the target `all` tries to avoid or minimize compilation warnings due to out-of-order loading of modified source files by performing a file topological sort based on the dependencies in the previous versions of the files. But it's always possible that the new versions of the files bring changes to those dependencies and thus result in compilation warnings that could possibly be avoided with a different loading order.

5.17 mutation_testing

This tool provides mutation testing support for loaded Logtalk entities. It's still under development and breaking changes should be expected.

5.17.1 API documentation

This tool API documentation is available at:

[../apis/library_index.html#mutation-testing](#)

5.17.2 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(mutation_testing(loader)).
```

5.17.3 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(mutation_testing(tester)).
```

5.17.4 Features

- Deterministic mutant discovery per entity predicate directive, predicate/non-terminal clause/rule, and mutator.
- Configurable mutator sets and campaign size limits.
- Deterministic sampled execution (`sampling(all|count(N)|rate(R))` plus seed/1).
- Mutation score computation (killed versus survived mutants).
- Mutation generation guided by code coverage stats.
- Threshold gating suitable for CI/CD checks.
- Exporting of mutation campaign reports in plain text and JSON formats.

5.17.5 Default mutators

The default mutators are:

- `fail_insertion`
Inserts failure in the selected predicate/non-terminal rule bodies.
- `body_goal_negation`
Negates the selected predicate/non-terminal clause behavior.
- `relational_operator_replacement`
Replaces a relational operator (e.g. `>`, `<`, `@=<`, `==`) with a complementary alternative.
- `arithmetic_operator_replacement`
Replaces an arithmetic operator in selected expressions.
- `truth_literal_flip`
Flips truth literals (`true <-> fail`) in selected goals.
- `head_arguments_mutation`
Mutates one compile-time bound head argument using `type::mutation/3`.
- `head_arguments_reordering`
Reorders head arguments (swapping the first two arguments).
- `clauses_reordering`
Reorders predicate/non-terminal clauses/rules by swapping a clause with its successor.
- `scope_directive_replacement`
Replaces a matching predicate/non-terminal scope directive visibility with an alternative visibility.
- `predicate_directive_suppression`
Suppresses a matching predicate/non-terminal directive.
- `uses_directive_resource_deletion`
Deletes one matching resource from a `uses/2` directive.

5.17.6 Usage

Mutation testing **requires** an existing set of tests for the code being mutated. These tests are used to verify if they are able to kill the generated mutants or if they are too weak allowing the mutants to survive.

By default, the tool looks for the tests driver file (usually, `tester.lgt`) in the directory of the entity, library, or directory being tested. If the tests driver file is not found there, the tool looks into the startup directory of the Logtalk process.

Both the name of the tests driver file and its directory can be set using the `tester_file_name/1` and `tester_directory/1` options if the defaults don't work in your particular case.

In this tool documentation, **campaign predicates** refers to the following public predicates:

- `entity/1-2`
- `predicate/2-3`
- `library/1-2`
- `directory/1-2`

These predicates run mutants and report results in one step.

Run mutation testing for one entity using defaults:

```
| ?- mutation_testing::entity(my_object).
```

List generated mutants for one entity:

```
| ?- mutation_testing::entity_mutants(my_object, Mutants).
```

Run with custom options:

```
| ?- mutation_testing::entity(my_object, [
    mutators([fail_insertion, body_goal_negation]),
    max_mutations_per_mutator(5),
    sampling(count(100)),
    seed(20260303),
    max_mutators(100),
    threshold(60.0),
    verbose(true),
    print_mutation(true)
]).
```

Run campaigns for all loaded entities from a specific library:

```
| ?- mutation_testing::library(core).
```

Run campaigns for all loaded entities from a specific directory:

```
| ?- mutation_testing::directory('/path/to/sources').
```

Pretty-print a report term using the default text format:

```
| ?- mutation_testing::report_entity(my_object, Report, []),
    mutation_testing::format_report(Report).
```

Suppress report formatting output for campaign predicates:

```
| ?- mutation_testing::entity(my_object, [format(none)]).
```

Execution uses lgtunit test sets only. When using `timeout(Timeout)`, timeout handling is best-effort.

When running mutation campaigns inside another lgtunit test run, per-mutant killed versus survived classification may be affected by nested test-run message and event handling. Campaign summary accounting remains deterministic.

5.17.7 Limitations

- Mutation points are currently mostly clause-level (one occurrence per clause per mutator), except `head_arguments_mutation` which creates one occurrence per compile-time bound head argument with supported type; internal expression-level candidate enumeration is not implemented.
- Equivalent mutant detection and duplicate mutant pruning are not implemented.
- Built-in mutation `apply/revert` currently targets loaded source entities (objects/categories), not protocols.
- Nested lgtunit runs may affect per-mutant killed vs survived status classification.

5.17.8 Options

- `include_entities(Entities)`
List of loaded entities to include (default [] meaning all).
- `exclude_entities(Entities)`
List of loaded entities to exclude (default []).
- `max_mutators(Max)`
Maximum number of discovered mutators to use when `mutators/1` is not explicitly provided (all or positive integer; default all).
- `max_mutations_per_mutator(Max)`
Maximum number of mutations generated per mutator and predicate/non-terminal (all or positive integer; default 5).
- `mutators(Mutators)`
Names of the mutators to use. When set to [] (default), mutators are auto-discovered and optionally limited by `max_mutators/1`.
- `sampling(Sampling)`
Mutant selection strategy (all, count(N), or rate(R) where $0.0 \leq R \leq 1.0$; default all). See below for full details.
- `seed(Seed)`
Pseudo-random generator seed for reproducible sampling (default 123456789; ignored when `sampling(all)`).
- `timeout(Timeout)`
Maximum time in seconds for each mutant execution (positive number; default 300). Timeout handling is best-effort.
- `threshold(Threshold)`
Minimum mutation score in range $0.0..100.0$ (default 0.0).
- `verbose(Boolean)`

Print per-mutant results (default false).

- `format(Format)` Controls report formatting output (none, text, or json; default text). When set to text or json, a report file is generated automatically for campaign predicates. The json format implements the Stryker Mutation Testing Framework report format: <https://stryker-mutator.io/>
- `report_file_name(FileName)` Report output file base name or path without extension (atom; default `mutation_test_report`). The extension is inferred from `format/1` (text -> `.txt`, json -> `.json`). When not absolute, the file is saved in the tests driver directory.
- `print_mutation(Boolean)`
When true, prints original and mutated terms with source location for mutators. This option is only effective when `verbose(true)` (default false).
- `tester_file_name(Tester)`
Name of the tests driver file for the code being tested (default `tester.lgt`).
- `tester_directory(Directory)`
Full path to the directory containing the tests driver file for the code being tested (no default).

5.17.9 Sampling semantics

The `sampling(Sampling)` option controls **which generated mutants are actually executed** in a mutation campaign. It is a post-generation selection step:

1. Generate candidate mutants for the selected entities/predicates/mutators.
2. Apply `sampling(...)` to select a subset (or all).
3. Execute only the selected mutants.

Because of this design, `sampling(...)` affects campaign predicates (entity/2, predicate/3, library/2, directory/2).

The `format(Format)` option affects only campaign predicates that execute and report in one step (entity/2, predicate/3, library/2, directory/2). It does not affect the `*_mutants` predicates or the `report_*` predicates, which only compute and return terms.

For campaign predicates, a report file is written automatically unless `format(none)` is used.

For `report_library/3` and `report_directory/3`, call `format_report/2-3` explicitly if you want to persist the computed report term. A single `format_report(..., Report)` call always generates a single output document (including JSON).

The `mutants/2-3` predicates are still useful for inspecting the generated deterministic mutant list itself; they are not execution/reporting predicates.

Sampling modes

- `sampling(all)` Execute all generated mutants.
- `sampling(count(N))` Shuffle the generated mutant list and execute up to N mutants. If N is greater than the total number of generated mutants, all mutants are executed.
- `sampling(rate(R))` Compute `round(R * TotalGeneratedMutants)`, shuffle the mutant list, and execute that many mutants. If the computed value is greater than the total number of generated mutants, all mutants are executed.

The randomization used by `count/1` and `rate/1` is controlled by the `seed/1` option, allowing reproducible selection.

Sampling usage in practice

- **Control campaign cost** Run a representative subset when the full mutant set is large.
- **Trade off speed vs. confidence** Use smaller samples for quick feedback and larger/full runs for stronger confidence before release.
- **Enable reproducible sampling** With a fixed seed/1, repeated runs with the same generated mutant set and same sampling options select the same subset. When all mutants are executed, (using `sampling(all)`) changing the seed has no effect on selection.
- **Complement other limits** `max_mutators/1` and `max_mutations_per_mutator/1` constrain mutant generation; `sampling/1` controls execution selection after generation.

5.17.10 Defining new mutators

Define a new mutator as parametric object implementing the expanding protocol, either the `clause_mutator_protocol` or the `directive_mutator_protocol` marker protocol, and importing the `mutator_common` category:

```
:- object(my_mutator(_Entity_, _Predicate_, _TargetIndex_, _Occurrence_, _PrintMutation_),
    implements((expanding, clause_mutator_protocol)),
    imports(mutator_common)).

    % mutate code when loading it using this object as a hook object
    term_expansion(Term, Mutation) :-
        ...

    % compute a term mutation; generate multiple mutations by backtracking
    mutation(Term, Mutation) :-
        ...

    % reset any required internal state
    reset :-
        ...

:- end_object.
```

New mutators are found by dynamically looking for objects that conform to either the `clause_mutator_protocol` or the `directive_mutator_protocol` marker protocols.

For implementation examples, see the default mutators:

- `mutators/fail_insertion.lgt`
- `mutators/body_goal_negation.lgt`
- `mutators/relational_operator_replacement.lgt`
- `mutators/arithmetic_operator_replacement.lgt`
- `mutators/truth_literal_flip.lgt`
- `mutators/head_arguments_mutation.lgt`
- `mutators/head_arguments_reordering.lgt`
- `mutators/clauses_reordering.lgt`
- `mutators/scope_directive_replacement.lgt`

5.18 packs

This tool provides predicates for downloading, installing, upgrading, and uninstalling third-party libraries and applications, generically known as *packs*. Collections of pack specifications are made available using *registries*. Registries can be local to a system, publicly shared, or private to a company (e.g., only available in a VPN). There is no concept of a central registry. Users can decide which registries they trust and want to use and add them using their published URLs. The tool supports both pack checksums and signatures and takes several steps to sanitize registry and pack specifications. As with other Logtalk developer tools, portability is a main goal. This tool can be used with any supported Prolog backend and run on both POSIX and Windows systems. Moreover, this tool can be used not only for handling Logtalk packs but also for Prolog only packs, thus providing a solution for sharing portable resources between multiple systems.

A list of public Logtalk and Prolog pack registries is available at:

<https://github.com/LogtalkDotOrg/pack-registries>

This tool is in the beta stage of development. Feedback is most welcome.

5.18.1 Requirements

On POSIX systems (Linux, macOS, ...), the following shell commands are required:

- sha256sum (provided by GNU coreutils)
- curl (default file downloader)
- wget (alternative file downloader)
- bsdtar (provided by libarchive or libarchive-tools)
- gpg (provided by gnupg2)
- git
- direnv (when using virtual environments)

The tool uses bsdtar instead of GNU tar so that it can uncompress .zip archives (unzip doesn't provide the desired options that allow a simple and reliable solution for ignoring the non-predictable name of the wrapper directory).

On Windows systems, the following shell commands are required:

- certutil.exe
- curl.exe (default file downloader)
- wget.exe (alternative file downloader)
- tar.exe
- gpg.exe
- git.exe
- Set-PSEnv (when using virtual environments)

In recent Windows 10 builds, only wget, gpg, git, and Set-PSEnv should require installation. You can download the GnuPG software from:

<https://www.gnupg.org/>

You can download Git from:

<https://gitforwindows.org>

You can download Wget from:

<https://eternallybored.org/misc/wget/>

You can also use Chocolatey to install the commands above:

```
> choco install gnupg git wget
```

To install `Set-PsEnv` from the PowerShell Gallery:

```
PS> Install-Module -Name Set-PsEnv
```

On macOS systems, Apple bundles both `curl` and BSD `tar` (under the name `tar`; you can simply create a `bsdtar` alias or install a more recent version). The required commands can be easily installed using MacPorts:

```
$ sudo port install coreutils wget libarchive gnupg2 git direnv
```

Or using Homebrew:

```
$ brew install coreutils wget libarchive gnupg2 git direnv
```

On Linux systems, use the distribution's own package manager to install any missing command. For example, in recent Ubuntu versions:

```
$ sudo apt update
$ sudo apt install coreutils curl wget libarchive-tools gnupg2 git direnv
```

5.18.2 API documentation

This tool API documentation is available at:

../..apis/library_index.html#packs

5.18.3 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(packs(loader)).
```

5.18.4 Testing

To run the tool tests, use the query:

```
| ?- logtalk_load(packs(tester)).
```

The tests can be run without interfering with the user packs setup.

5.18.5 Usage

The packs tool loads at startup all the currently defined registry and pack specifications (from the registries/packs storage directory; see below). When no registry/pack setup exists, a new one is automatically created.

The tool provides two main objects, registries and packs, for handling, respectively, registries and packs. Both objects accept a `help/0` message that describes the most common queries.

5.18.6 Programmatic query API

Besides the user-facing predicates that print information (e.g., `describe/1-2`), the packs object also provides a query API intended for programmatic use by tools and tests. This API exposes resolved installed-pack metadata and current-session load participation as data instead of formatted text.

The main predicates are:

- `pack_metadata/4`
- `pack_property/4`
- `pack_object/3`
- `loaded_pack/3`
- `loaded_pack_file/4`
- `pack_dependency/6`
- `loaded_pack_dependency/6`

The `pack_metadata/4` predicate enumerates installed packs and returns a resolved metadata term with the shape:

```
metadata(Name, Description, License, Home, SourceURL, Checksum, Dependencies, Portability,
↳Directory, Pinned, Installed, Loaded)
```

The term combines metadata declared by the pack specification object with the metadata selected from the installed version. Missing optional metadata is normalized to `none`. The `Installed` field is always `true` on success. The `Loaded` field reports if the installed pack contributed at least one currently loaded file to the session.

For example:

```
| ?- packs::pack_metadata(Registry, Pack, Version, Metadata).
...
```

The `pack_property/4` predicate provides selective access to resolved metadata using property terms. Supported properties are:

- `name(Name)`
- `description(Description)`
- `license(License)`
- `home(Home)`
- `source_url(URL)`
- `checksum(Checksum)`

- dependencies(Dependencies)
- portability(Portability)
- directory(Directory)
- pinned(Boolean)
- installed(Boolean)
- loaded(Boolean)

For example:

```
| ?- packs::pack_property(local_1_d, foo, Version, license(License)).
License = 'Apache-2.0'.
```

The pack_object/3 predicate is a low-level bridge that returns the pack specification object implementing the pack_protocol for a given registry and pack. Most clients should prefer pack_metadata/4 or pack_property/4 unless direct access to the pack object is required.

For example:

```
| ?- packs::pack_object(local_1_d, foo, PackObject).
PackObject = foo_pack.
```

The loaded_pack/3 and loaded_pack_file/4 predicates expose current-session pack participation. A pack is considered loaded iff at least one currently loaded file belongs to its installation directory. The first predicate is the high-level query, while the second predicate can be used to inspect the files that justify that status.

For example:

```
| ?- packs::loaded_pack(Registry, Pack, Version).
...

| ?- packs::loaded_pack_file(Registry, Pack, Version, File).
...
```

The pack_dependency/6 predicate enumerates resolved direct pack-to-pack dependencies for installed pack versions using the same version comparison and dependency selection semantics used by the installation logic. Version ranges, alternatives, and conjunctions are resolved canonically. Non-pack dependencies on Logtalk, Prolog backends, and operating systems are ignored by this predicate.

For example:

```
| ?- packs::pack_dependency(local_1_d, foo, 2:0:0, DependencyRegistry, DependencyPack, ↵
↵DependencyVersion).
DependencyRegistry = local_2_d,
DependencyPack = baz,
DependencyVersion = 1:0:0.
```

The loaded_pack_dependency/6 predicate is the loaded-session counterpart of pack_dependency/6. It only succeeds when both the source pack and the dependency pack contributed loaded files to the current session.

For example:

```
| ?- packs::loaded_pack_dependency(Registry, Pack, Version, DependencyRegistry, ↵
↵DependencyPack, DependencyVersion).
...
```

These predicates are especially useful for tools such as `sbom`, which need to query resolved pack metadata, determine which installed packs contributed code to the current session, and enumerate relationships between those packs without reconstructing that information manually.

5.18.7 Registries and packs storage

The tool uses a directory specified using the `logtalk_packs` library alias when defined (in a settings file or in a backend Prolog initialization file). When this library alias is not defined, the tool uses the value of the `LOGTALKPACKS` environment variable when defined. Otherwise, it defaults to the `~/logtalk_packs` directory. The actual directory can be retrieved by the query:

```
| ?- packs::logtalk_packs(Directory).  
...
```

This directory holds sub-directories for registries, packs, and archives. These sub-directories are automatically created when loading the packs tool if they don't exist. Users should not manually modify the contents of these directories. Multiple and independent registry/pack setups are possible using *virtual environments* as explained next.

Your registries and packs setup can be saved and restored (e.g., in a different system) by using the `packs::save/1-2` and `packs::restore/1-2` predicates, as explained in the next section about virtual environments. If necessary, before restoring, the `packs::reset/0` predicate can be called to delete any defined registries and installed packs.

5.18.8 Virtual environments

An application may require specific pack versions. These requirements may differ between applications. Different applications may also have conflicting requirements. Therefore, a *virtual environment* where an application requirements are fulfilled may be required to develop and/or run it. A virtual environment is essentially a registries/packs storage directory.

Defining the `logtalk_packs` library alias in a settings file or defining the `LOGTALKPACKS` environment variable before starting Logtalk allows easy creation and switching between virtual environments. By using a per-application settings file (or a per-application environment variable definition), each application can thus use its own virtual environment. The `settings.lgt` file can define the `logtalk_packs` library alias using code such as:

```
:- initialization((  
    logtalk_load_context(directory, Directory),  
    assertz(logtalk_library_path(logtalk_packs, Directory))  
)).
```

The definition of the `logtalk_packs` library alias **must** always be an atom and thus never use library notation (i.e., it must never depend on other library aliases).

When a virtual environment also requires a specific Logtalk version (e.g., the version used to test and certify it), this can be installed as a pack from the official [talkshow](#) registry and used by (re)defining the `LOGTALKHOME` and `LOGTALKUSER` environment variables to point to its pack directory (which can be queried by using the `packs::directory/2` message).

Experimental `lgtenv.sh` and `lgtenv.ps1` scripts are included to simplify creating virtual environments. For example:

```
$ lgtenv -d /my_venv -c -p logtalk_packs
$ cd /my_venv
direnv: loading /my_venv/.envrc
direnv: export +LOGTALKPACKS
```

Type `lgtenv -h` for details on the script options.

These scripts require, respectively, `direnv` and `Set-PsEnv` to be installed. These utilities load and unload environment variables when changing the current directory. On Windows systems, when using the `lgtenv.ps1` script, you also need to redefine the PowerShell prompt in a profile file (e.g., `$HOME\Documents\PowerShell\Profile.ps1`) to mimic the functionality of `direnv` of automatically loading an existing `.env` file when changing to its directory. For example:

```
function prompt {
    Set-PsEnv
    'PS ' + $(Get-Location) + '> '
}
```

A virtual environment setup (i.e., the currently defined registries and installed packs) can be saved into a file (e.g., `requirements.lgt`) using the `packs::save/1-2` predicates. For example:

```
| ?- packs::save('requirements.lgt').
...
```

This query saves a listing of all the installed packs and their registries. Using the saved file, the virtual environment setup can then be restored using the `packs::restore/1-2` predicates. The file uses a simple format with `registry/2`, `pack/3`, `pinned_registry/1`, and `pinned_pack/1` facts (in this order) and can be manually created or edited if necessary. For example:

```
registry(talkshow, 'https://github.com/LogtalkDotOrg/talkshow.git').
pack(talkshow, logtalk, 3:45:0).
pack(talkshow, lflat, 2:1:0).
```

These files can be distributed with applications so that users can easily fulfill application requirements by running the query once:

```
| ?- packs::restore('requirements.lgt').
```

Subsequently, the application `loader.lgt` file can then load the required packs using their loader files:

```
:- initialization((
    % load required packs
    logtalk_load(foo(loader)),
    logtalk_load(bar(loader)),
    ...
    % load application files
    ...
)).
```

Note that restoring encrypted registries or encrypted packs requires entering the required passphrases. Although the `restore/2` predicate accepts a list of options that include the `gpg/1` option, this only allows specifying a single and common passphrase when interactive entering of passphrases is not convenient or possible.

5.18.9 Lock files

For deterministic installs and CI/CD reproducibility when working with virtual environments, the `save/2` and `restore/2` predicates support a `lock(true)` option. Saving with this option writes lock extension facts in addition to the requirements facts:

```
lockfile_version(1).
lock_registry_commit(talkshow, '0123456789abcdef0123456789abcdef01234567').
lock_integrity(talkshow, lflat, 2:1:0, sha256, '8774b3863efc03bb6c85dcf34f69f1156d2496a3').
```

The `lock_registry_commit/2` facts are generated only for registries defined from git URLs. Restoring with `lock(true)` enforces lock extension facts, including exact pack versions, integrity hashes, and git registry commits.

5.18.10 Registry specification

A registry is a git remote repo that can be cloned, a downloadable or local archive, or a local directory containing a Logtalk loader file that loads source files defining the registry itself and the packs it provides. The registry name is ideally a valid unquoted atom. The registry directory must contain at least two Logtalk source files:

- A file defining an object named after the registry with a `_registry` suffix, implementing the `registry_protocol`. This naming convention helps prevent name conflicts.
- A loader file (named `loader.lgt` or `loader.logtalk`) that loads the registry object file and all pack object files.

An example of a registry specification object would be:

```
:- object(jdoe_awesome_packs_registry,
    implements(registry_protocol)).

    :- info([
        version is 1:0:0,
        author is 'John Doe',
        date is 2021-10-18,
        comment is 'John Doe awesome packs registry spec.'
    ]).

    name(jdoe_awesome_packs).

    description('John Doe awesome packs').

    home('https://example.com/jdoe_awesome_packs').

    clone('https://github.com/jdoe/jdoe_awesome_packs.git').

    archive('https://github.com/jdoe/jdoe_awesome_packs/archive/main.zip').

:- end_object.
```

Optionally, the registry object can also define a `note(Action, Note)` predicate. The `Action` argument is an atom: `add`, `update`, or `delete`. The `Note` argument is also an atom. The tool will print any available notes when executing one of the registry actions. See the `registry_protocol` documentation for more details.

The registry directory should also contain `LICENSE` and `README.md` files (individual packs can use a different license, however). The path to the `README.md` file is printed when the registry is added. It can also be queried using the `registries::directory/2` predicate. The `NOTES.md` file name can also be used in alternative to the recommended `README.md` file name.

Summarizing the required directory structure using the above example (note that the registry and pack specification files are named after the objects):

```
jdoe_awesome_packs
  LICENSE
  README.md
  jdoe_awesome_packs_registry.lgt
  loader.lgt
  foo_pack.lgt
  bar_pack.lgt
  ...
```

With the contents of the `loader.lgt` file being:

```
:- initialization((
    logtalk_load(jdoe_awesome_packs_registry),
    logtalk_load(foo_pack),
    logtalk_load(bar_pack),
    ...
)).
```

It would be, of course, possible to have all objects in a single source file. But having a file per-object and a loader file helps maintenance, and it's also a tool requirement for applying safety procedures to the source file contents and thus successfully loading the registry and pack specs.

As registries are git repos in the most common case, and thus adding them performs a git repo cloning, they should only contain the strictly required files.

5.18.11 Registry handling

Registries can be added using the `registries::add/1-3` predicates, which take a registry URL. Using the example above:

```
| ?- registries::add('https://github.com/jdoe/jdoe_awesome_packs.git').
```

HTTPS URLs must end with either a `.git` extension or an archive extension (same valid extensions as for pack archives, including gpg encrypted). Git cloning URLs are preferred as they simplify updating registries. But a registry can also be made available via a local directory (using a `file://` URL) or a downloadable archive (using a `https://` URL).

For registries made available using an archive, the `registries::add/2-3` predicates **must** be used as the registry name cannot in general be inferred from the URL basename or from the archived directory name. The registry argument must also be the declared registry name in the registry specification object. For example:

```
| ?- registries::add(
    jdoe_awesome_packs,
    'https://github.com/jdoe/jdoe_awesome_packs/archive/main.zip'
).
```

When a registry may be already defined, you can use the `update(true)` option to ensure that the registry will be updated to its latest definition:

```
| ?- registries::add(
    jdoe_awesome_packs,
    'https://github.com/jdoe/jdoe_awesome_packs/archive/main.zip',
    [update(true)]
).
```

The added registries can be listed using the `registries::list/0` predicate:

```
| ?- registries::list.

% Defined registries:
%   jdoe_awesome_packs (git)
%   ...
```

The `registries::describe/1` predicate can be used to print the details of a registry:

```
| ?- registries::describe(jdoe_awesome_packs).

% Registry:    jdoe_awesome_packs
% Description: John Doe awesome packs
% Home:        https://example.com/jdoe_awesome_packs
% Cloning URL: https://github.com/jdoe/jdoe_awesome_packs.git
% Archive URL: https://github.com/jdoe/jdoe_awesome_packs/archive/main.zip
```

To update all registries, use the `registries::update/0` predicate. To update a single registry, use the `registries::update/1-2` predicates. After updating, you can use the `packs::outdated/0-1` predicates to list any outdated packs.

Registries can also be deleted using the `registries::delete/1-2` predicate. By default, any registries with installed packs cannot be deleted. If you force deletion (by using the `force(true)` option), you can use the `packs::orphaned/0` predicate to list any orphaned packs that are installed.

See the tool API documentation on the [registries](#) object for other useful predicates.

5.18.12 Registry development

To simplify registry development and testing, use a local directory and a `file://` URL when calling the `registries::add/1` predicate. For example:

```
| ?- registries::add('file:///home/jdoe/work/my_pack_collection').
```

If the directory is a git repo, the tool will clone it when adding it. Otherwise, the files in the directory are copied to the registry definition directory. This allows the registry to be added and deleted without consequences for the original registry source files.

The `registries::add/3` and `registries::update/2` predicates also accept a `commit(Hash)` option. This option checks out a git registry at an exact commit hash. Using this option with non-git registries fails with an error message.

To check your registry specifications, use the `registries::lint/0-1` predicates after adding the registry.

5.18.13 Pack specification

A pack is specified using a Logtalk source file defining an object that implements the `pack_protocol`. The source file should be named after the pack with a `_pack` suffix. This naming convention helps prevent name conflicts, notably with the pack's own objects. The file must be available from a declared pack registry (by having the registry loader file loading it). The pack name is preferably a valid unquoted atom. An example of a pack specification object would be:

```
:- object(lflat_pack,
    implements(pack_protocol)).

    :- info([
        version is 1:0:0,
        author is 'Paulo Moura',
        date is 2021-10-18,
        comment is 'L-FLAT - Logtalk Formal Language and Automata Toolkit pack spec.'
    ]).

    name(lflat).

    description('L-FLAT - Logtalk Formal Language and Automata Toolkit').

    license('MIT').

    home('https://github.com/l-flat/lflat').

    version(
        2:1:0,
        stable,
        'https://github.com/l-flat/lflat/archive/refs/tags/v2.1.0.tar.gz',
        sha256 - '9c298c2a08c4e2a1972c14720ef1498e7f116c7cd8bf7702c8d22d8ff549b6a1',
        [logtalk @>= 3:42:0],
        all
    ).

    version(
        2:0:2,
        stable,
        'https://github.com/l-flat/lflat/archive/refs/tags/v2.0.2.tar.gz',
        sha256 - '8774b3863efc03bb6c284935885dcf34f69f115656d2496a33a446b6199f3e19',
        [logtalk @>= 3:36:0],
        all
    ).

:- end_object.
```

The `license/1` argument must be an atom and should, whenever possible, be a license identifier as specified in the [SPDX standard](#).

Optionally, the pack object can also define a `note(Action, Version, Note)` predicate. The `Action` argument is an atom: `install`, `update`, or `uninstall`. The `Note` argument is also an atom. The tool will print any available notes when executing one of the registry actions. See the `pack_protocol` documentation for more details.

The pack sources must be available either as a local directory (when using a `file://` URL) or for download-

ing as a supported archive. The checksum for the archive must use the SHA-256 hash algorithm (sha256). The pack may optionally be signed. Supported archive formats and extensions are:

- .zip
- .tgz, .tar.gz
- .tbz2, .tar.bz2

Also, for encrypted packs, all the extensions above with a .gpg suffix (e.g., .zip.gpg).

The pack sources should contain LICENSE, README.md (or NOTES.md), and loader.lgt (or loader.logtalk) files. Ideally, it should also contain a tester.lgt (tester.logtalk) file. The path to the README.md file is printed when the pack is installed or updated. It can also be queried using the `packs::directory/2` predicate.

5.18.14 Encrypted packs

Packs can be gpg encrypted, with a choice of passphrase-based encryption, key-based encryption, or both. Encrypted pack archives must always have a .gpg extension. For example, to encrypt a pack archive with a symmetric cipher using a passphrase:

```
$ tar -cvzf - my_pack | gpg -c --cipher-algo AES256 > v1.2.1.tar.gz.gpg
```

In this case, the passphrase would need to be securely communicated to any users installing or updating the pack.

See the gpg documentation for full details on encrypting and decrypting archives. If you get a “gpg: problem with the agent: Inappropriate ioctl for device” error message with the command above, try:

```
$ export GPG_TTY=$(tty)
```

5.18.15 Signed packs

Packs can be gpg signed. Detached signature files are assumed and expected to share the name of the archive and use .asc or .sig extensions. For example, if the pack archive name is `v1.0.0.tar.gz`, the signature file must be named `v1.0.0.tar.gz.asc` or `v1.0.0.tar.gz.sig`. When the `checksig(true)` option is used, the signature file is automatically downloaded using a URL constructed from the pack archive URL. When both .asc and .sig files exist, the .asc file is used. An example of signing a pack and creating the .asc file (assuming the default signing key) is:

```
$ gpg --armor --detach-sign v1.0.0.tar.gz
```

To create instead a .sig file:

```
$ gpg --detach-sign v1.0.0.tar.gz
```

See the gpg documentation for full details on signing archives and sharing the public keys required to verify the signatures.

5.18.16 Pack URLs and Single Sign-On

Typically, pack archive download URLs are HTTPS URLs and handled using `curl`. It's also possible to use `git archive` to download pack archives, provided that the server supports it (as of this writing, Bitbucket and GitLab public hosting services support it but not GitHub). Using `git archive` is specially useful when the packs registry is hosted on a server using Single Sign-On (SSO) for authentication. In this case, HTTPS URLs can only be handled by `curl` by passing a token (see below for an example). When the user has setup SSH keys to authenticate to the packs registry server, `git archive` simplifies pack installation, providing a better user experience. For example:

```
version(
  1:0:1,
  stable,
  'git@gitlab.com:me/foo.git/v1.0.1.zip',
  sha256 - '0894c7cdb8968b6bbcf00e3673c1c16cfa98232573af30ceddda207b20a7a207',
  [logtalk @>= 3:36:0],
  all
).
```

The pseudo-URL must be the concatenation of the SSH repo cloning URL with the archive name. The archive name must be the concatenation of a valid tag with a supported archive extension. SSH repo cloning URLs use the format:

```
git@<hostname>:path/to/project.git
```

They can usually be easily copied from the hosting service repo webpage. To compute the checksum, you must first download the archive. For example:

```
$ git archive --output=foo-v1.0.1.zip --remote=git@gitlab.com:me/foo.git v1.0.1
$ openssl sha256 foo-v1.0.1.zip
```

Be sure to use a format that is supported by both the packs tool and the `git archive` command (the format is inferred from the `--output` option). Do not download the archive from the web interface of the git hosting service in order to compute the checksum. Different implementations of the archiving and compressing algorithms may be used, resulting in mismatched checksums.

Users installing packs available using `git archive` URLs are advised to run a SSH agent to avoid being prompted for passwords when installing or updating packs. They must also upload their SSH public keys to the pack provider hosts.

5.18.17 Multiple pack versions

A pack may specify multiple versions. Each version is described using a `version/6` predicate clause as illustrated in the example above. The versions must be listed in order from newest to oldest. For details, see the `pack_protocol` API documentation.

Listing multiple versions allows the pack specification to be updated (by updating its registry) without forcing existing users into installing (or updating to) the latest version of the pack. It allows different applications depending on different pack versions to continue to be built and deployed.

The pack version is complemented by the pack status. Valid values are `stable`, `rc`, `beta`, `alpha`, `experimental`, and `deprecated`. Packs with a `experimental` or `deprecated` status are **never** installed by default when using the `install` and `update` predicates unless their version is explicitly specified. When updating packs, we can restrict the valid status of the updates using the `status/1` option. For example, we can ensure that we only update to new stable pack versions by using the option `status([stable])`.

5.18.18 Pack dependencies

Pack dependencies on other packs can be specified using a list of `Registry::Pack Operator Version` terms where `Operator` is a standard term comparison operator:

- `Registry::Pack @>= Version` - the pack requires a dependency with a version equal or above the specified one. For example, `logtalk @>= 3:36:0` means that the pack requires Logtalk 3.36.0 or a later version.
- `Registry::Pack @<= Version` - the pack requires a dependency with a version up to the specified one. For example, `common::bits @<= 2:1` means that the pack requires a `common::bits` pack up to 2.1. This includes all previous versions and also all patches for version 2.1 (e.g., 2.1.7, 2.1.8, ...) but not version 2.2 or newer.
- `Registry::Pack @< Version` - the pack requires a dependency with version older than the specified one. For example, `common::bits @< 3` means that the pack requires a `common::bits` 2.x or older version.
- `Registry::Pack @> Version` - the pack requires a dependency with version newer than the specified one. For example, `common::bits @> 2:4` means that the pack requires a `common::bits` 2.5 or newer version.
- `Registry::Pack == Version` - the pack requires a dependency with a specific version. For example, `common::bits == 2:1` means that the pack requires a `common::bits` pack version 2.1.x (thus, from version 2.1.0 to the latest patch for version 2.1).
- `Registry::Pack \== Version` - the pack requires a dependency with any version other than the one specified. For example, `common::bits \== 2.1` means that the pack requires a `common::bits` pack version other than any 2.1.x version.

To specify *range* dependencies by using two consecutive elements with the lower bound followed by the upper bound. For example, `common::bits @>= 2, common::bits @< 3` means all `common::bits` 2.x versions but not older or newer major versions.

It's also possible to specify *alternative* dependencies using the `(;)/2` operator. For example, `(common::bits == 1:9; common::bits @>= 2:3)` means either `common::bits` 1.9.x versions or 2.3.x and later versions. Alternatives should be listed in decreasing order of preference.

When a pack also depends on a Logtalk or backend version, the name `logtalk` or the backend identifier atom can be used in place of `Registry::Pack` (see below for the table of backend specifiers). For example, `logtalk @>= 3.36.0`.

When a pack also depends on an operating-system version (e.g., a pack containing shared libraries with executable code), the `os(Name,Machine)` compound term can also be used in place of `Registry::Pack`. For example, `os('Darwin',x86_64) @>= '23.0.0'`. Note that, in this case, the release is an atom. The operating-system data (name, machine, and release) is queried using the corresponding `os` library predicates (see the library documentation for details).

5.18.19 Pack portability

Ideally, packs are fully portable and can be used with all Logtalk-supported Prolog backends. This can be declared by using the atom `all` in the last argument of the `version/6` predicate (see example above).

When a pack can only be used with a subset of the Prolog backends, the last argument of the `version/6` predicate is a list of backend identifiers (atoms):

- B-Prolog: `b`
- Ciao Prolog: `ciao`
- CxProlog: `cx`

- ECLiPSe: eclipse
- GNU Prolog: gnu
- JIProlog: ji
- Quintus Prolog: quintus
- SICStus Prolog: sicstus
- SWI-Prolog: swi
- Tau Prolog: tau
- Trealla Prolog: trealla
- XSB: xsb
- XVM: xvm
- YAP: yap

5.18.20 Pack development

To simplify pack development and testing, define a local registry and add to it a pack specification with the development version available from a local directory. For example:

```
version(
  0:11:0,
  beta,
  'file:///home/jdoe/work/my_awesome_library',
  none,
  [],
  all
).
```

If the directory is a git repo, the tool will clone it when installing the pack. Otherwise, the files in the directory are copied to the pack installation directory. This allows the pack to be installed, updated, and uninstalled without consequences for the pack source files.

You can also use a local archive instead of a directory. For example:

```
version(
  1:0:0,
  stable,
  'file:///home/jdoe/work/my_awesome_library/v1.0.0.tar.gz',
  sha256 - '1944773afba1908cc6194297ff6b5ac649a844ef69a69b2bcd267cfa8bfce1e',
  [],
  all
).
```

Packs that are expected to be fully portable should always be checked by loading them with the portability flag set to warning.

To check your pack manifest files, use the `packs::lint/0-2` predicates after adding the registry that provides the packs.

5.18.21 Pack handling

Packs must be available from a defined registry. To list all packs that are available for installation, use the `packs::available/0` predicate:

```
| ?- packs::available.
```

To list all installed packs, call the `packs::installed/0` predicate:

```
| ?- packs::installed.
```

To list only the installed packs from a specific registry, call instead the `packs::installed/1` predicate. For example:

```
| ?- packs::installed(talkshow).
```

To know more about a specific pack, use the `packs::describe/1-2` predicates. For example:

```
| ?- packs::describe(bar).
```

The `packs::describe/2` predicate can be used when two or more registries provide packs with the same name. For example:

```
| ?- packs::describe(reg, bar).
```

To install the latest version of a pack, we can use the `packs::install/1-4` predicates. In the most simple case, when a pack name is unique among registries, we can use the `packs::install/1` predicate. For example:

```
| ?- packs::install(bar).
```

Any pack dependencies are also checked and installed or updated if necessary. Other install predicates are available to disambiguate between registries and to install a specific pack version.

Packs become available for loading immediately after successful installation (no restarting of the Logtalk session is required). For example, after the pack `bar` is installed, you can load it at the top-level by typing:

```
| ?- {bar(loader)}.
```

or load it from a loader file using the goal `logtalk_load(bar(loader))`.

After updating the defined registries, outdated packs can be listed using the `packs::outdated/0` predicate. You can update all outdated packs by calling the `packs::update/0` predicate or update a single pack using the `packs::update/1-2` predicates. For example:

```
| ?- packs::update(bar).
```

By default, updating a pack fails if it would break any dependent pack (the `force(true)` option, described below, can be used to force updating in this case).

The tool provides versions of the pack install, update, and uninstall predicates that accept a list of options:

- `verbose(Boolean)` (default is false)
- `clean(Boolean)` (default is false)
- `update(Boolean)` (default is false)
- `force(Boolean)` (default is false)
- `compatible(Boolean)` (default is true)

- checksum(Boolean) (default is true)
- checksig(Boolean) (default is false)
- git(Atom) (extra command-line options; default is '')
- downloader(Atom) (downloader utility; default is curl)
- curl(Atom) (extra command-line options; default is '')
- wget(Atom) (extra command-line options; default is '')
- gpg(Atom) (extra command-line options; default is '')
- tar(Atom) (extra command-line options; default is '')

Note that, by default, only compatible packs can be installed. To install a pack that is incompatible with the current Logtalk version, backend version, or operating-system version, use the `install/4` or `update/3` predicates with the option `compatible(false)`.

When installing large packs over unreliable network conditions, you may try switching the default downloader utility from `curl` to `wget`.

When a pack may already be installed, you can use the `update(true)` option to ensure that the installation will be updated to the specified version:

```
| ?- packs::install(reg, bar, 1:1:2, [update(true)]).
```

When using a `checksig(true)` option to check a pack signature, it is strongly advised that you also use the `verbose(true)` option. For example:

```
| ?- packs::install(reg, bar, 1:1:2, [verbose(true), checksig(true)]).
```

Note that the public key used to sign the pack archive must already be present in your local system.

Downloading pack archives may require passing extra command-line options to `curl` for authentication. A common solution is to use a personal access token. The details depend on the server software. An example when using GitHub:

```
| ?- packs::install(reg, bar, 1:1:2, [curl('--header "Authorization: token foo42"')]).
```

Another example when using GitLab:

```
| ?- packs::install(reg, bar, 1:1:2, [curl('--header "PRIVATE-TOKEN: foo42"')]).
```

Pack archives may be `gpg` encrypted. Encryption can be passphrase-based, key-based, or both. When using only passphrase-based encryption, the archive passphrase must be entered (if not cached) when installing or updating a pack. In this case, the passphrase can be entered interactively or using the `gpg/1` option. For example:

```
| ?- packs::install(reg, bar, 1:1:2, [gpg('--batch --passphrase test123')]).
```

See the `gpg` documentation for details. When using the `gpg/1` option, you should be careful to not leak passphrases in, e.g., the query history.

To uninstall a pack that you no longer need, use the `packs::uninstall/1-2` predicates. By default, only packs with no dependent packs can be uninstalled. You can print or get a list of the packs that depend on a given pack by using the `packs::dependents/1-3` predicates. For example:

```
| ?- packs::dependents(reg, bar, Dependents).
```

See the tool API documentation on the [packs](#) object for other useful predicates.

5.18.22 Pack documentation

The path to the pack README.md file is printed when the pack is installed or updated. It can also be retrieved at any time by using the `readme/2` predicate. For example:

```
| ?- packs::readme(lflat, Path).
```

Additional documentation may also be available from the pack home page, which can be printed by using the `describe/1-2` predicates. For example:

```
| ?- packs::describe(lflat).

% Registry:    ...
% Pack:        lflat
% Description: L-FLAT - Logtalk Formal Language and Automata Toolkit
% License:     MIT
% Home:        https://github.com/l-flat/lflat
% Versions:
...
```

The pack API documentation can be generated using the `lgtdoc` tool library and directory predicates (depending on the pack source files organization). For example:

```
| ?- {lflat(loader)},
    {lgtdoc(loader)},
    logtalk::expand_library_path(lflat, Path),
    lgtdoc::rdirectory(Path).
...
```

This query creates a `xml_docs` directory in the current directory. The XML documentation files can then be converted into a final format, e.g., HTML, using one of the scripts provided by the `lgtdoc` tool. For example:

```
$ cd xml_docs
$ lgt2html
```

For more details and alternatives, see the `lgtdoc` tool documentation.

It is also possible to add API documentation and diagrams for all the installed packs to the Logtalk distribution API documentation and diagrams by calling the `build` and `update_svg_diagrams` scripts in the `docs/apis/sources` directory with the `-i` option. See the scripts documentation for more details.

5.18.23 Pinning registries and packs

Registries and packs can be *pinned* after installation to prevent accidental updating or deleting, e.g., when using the batch `update/0` predicate. This is useful when your application requires a specific version or for security considerations (see below). For example, if we want the `bar` pack to stay at its current installed version:

```
| ?- packs::pin(bar).
yes
```

After, any attempt to update or uninstall the pack will fail with an error message:

```
| ?- packs::update(bar).
!    Cannot update pinned pack: bar
no

| ?- packs::uninstall(bar).
!    Cannot uninstall pinned pack: bar
no
```

To enable the pack to be updated or uninstalled, the pack must first be unpinned. Alternatively, the `force(true)` option can be used. Note that if you force update a pinned pack, the new version will be unpinned.

It's also possible to pin (or unpin) all defined registries or installed packs at once by using the `pin/0` (or `unpin/0`) predicates. But note that registries added after or packs installed after will not be automatically pinned.

5.18.24 Testing packs

Logtalk packs (as most Logtalk libraries, tools, and examples) are expected to have a `tester.lgt` or `tester.logtalk` tests driver file at the root of their directory, which can be used for both automated and manual testing. For example, after installing the `foo` pack:

```
| ?- {foo(tester)}.
```

To test all installed packs, you can use the `logtalk_tester` automation script from the installed packs directory, which you can query using the goal:

```
| ?- packs::prefix(Directory).
```

Note that running the packs tests, like simply loading the pack, can result in calling arbitrary code, which can potentially harm your system. Always take into account the security considerations discussed below.

5.18.25 Security considerations

New pack registries should be examined before being added, specially if public and from a previously unknown source. The same precautions should be taken when adding or updating a pack. Note that a registry can always index third-party packs.

Pack checksums are checked by default. But pack signatures are only checked if requested, as packs are often unsigned. Care should be taken when adding public keys for pack signers to your local system.

Registry and pack spec files, plus the registry loader file, are compiled by term-expanding them so that only expected terms are actually loaded and only expected `logtalk_load/2` goals with expected relative file paths are allowed. Predicates defining URLs are discarded if the URLs are neither `https://` nor `file://` URLs or if they contain non-allowed characters (currently, only alpha-numeric ASCII characters plus the ASCII `/`, `.`, `-`, and `_` characters are accepted). But note that this tool makes no attempt to audit pack source files themselves.

Registries and packs can always be pinned so that they are not accidentally updated to a version that you may not have had the chance to audit.

5.18.26 Best practices

- Make available a new pack registry as a git repo. This simplifies updating the registry and rolling back to a previous version.
- Use registry and pack names that are valid unquoted atoms, thus simplifying usage. Use descriptive names with underscores if necessary to link words.
- Name the registry and pack specification objects after their names with a `_registry` or `_pack` suffix. Save the objects in files named after the objects.
- Create new pack versions from git tags.
- If the sources of a pack are available from a git repo, consider using signed commits and signed tags for increased security.
- When a new pack version breaks backwards compatibility, list both the old and the new versions on the pack specification file.
- Pin registries and packs when specific versions are critical for your work so that you can still easily batch update the remaining packs and registries.
- Include the `$LOGTALKPACKS` directory (or the default `~/logtalk_packs` directory) on your regular back-ups.

5.18.27 Installing Prolog packs

This tool can also be used to install Prolog packs that don't use Logtalk. After installing a `pl_pack` Prolog pack from a `pl_reg` registry, it can be found in the `$LOGTALKPACKS/packs/pl_reg/pl_pack` directory. When the `LOGTALKPACKS` environment variable is not defined, the pack directory is by default `~/logtalk_packs/packs/pl_reg/pl_pack`.

Different Prolog systems provide different solutions for locating Prolog code. For example, several Prolog systems adopted the Quintus Prolog `file_search_path/2` hook predicate. For these systems, a solution could be to add a fact to this predicate for each installed Prolog pack. For example, assuming a `pl_pack` Prolog pack:

```
:- multifile(file_search_path/2).
:- dynamic(file_search_path/2).

file_search_path(library, '$LOGTALKPACKS/packs/pl_pack').
```

If the Prolog system also supports reading an initialization file at startup, the above definition could be added there.

5.18.28 Help with warnings

Load the tutor tool to get help with selected warnings printed by the packs tool.

5.18.29 Known issues

Using the `verbose(true)` option on Windows systems may not provide the shell commands output depending on the backend.

On Windows systems, the `reset`, `delete`, and `uninstall` predicates may fail to delete all affected folders and files due to a operating-system bug. Depending on the backend, this bug may cause some of the tests to fail. For details on this bug, see:

<https://github.com/microsoft/terminal/issues/309>

The workaround is to use the Windows File Explorer to delete the leftover folders and files.

When using Ciao Prolog 1.20.0, a workaround is used for this system non-standard support for multifile predicates.

When using GNU Prolog 1.5.0 as the backend on Windows, you may get an error on `directory_files/2` calls. For details and a workaround, see:

<https://github.com/didoudiaz/gprolog/issues/4>

This issue is fixed in the latest GNU Prolog git version.

Using SICStus Prolog as the backend on Windows doesn't currently work in version 4.7.0 and earlier versions. The underlying issues are fixed in the SICStus Prolog 4.7.1 version.

XSB has an odd bug (likely in its parser) when reading files that may cause a pack installed version to be reported as the `end_of_file` atom.

Some tests fail on Windows when using ECLiPSe or XSB due to file path representation issues.

5.19 ports_profiler

This tool counts and reports the number of times each port in the *procedure box model* is traversed during the execution of queries. It can also report the number of times each clause (or grammar rule) is used. It is inspired by the ECLiPSe `port_profiler` tool.

The procedure box model is the same as the one used in the debugger tool. This is an extended version of the original Byrd's four-port model. Besides the standard `call`, `exit`, `fail`, and `redo` ports, Logtalk also defines two post-unification ports, `fact` and `rule`, and an exception port. This tool can also distinguish between deterministic exits (reported in the `exit` column in the profiling result tables) and exits that leave choice-points (reported in the `*exit` column).

5.19.1 API documentation

This tool API documentation is available at:

[../apis/library_index.html#ports-profiler](http://logtalk.org/..../apis/library_index.html#ports-profiler)

For sample queries, please see the `SCRIPT.txt` file in the tool directory.

5.19.2 Loading

```
| ?- logtalk_load(ports_profiler(loader)).
```

5.19.3 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(ports_profiler(tester)).
```

5.19.4 Compiling source files for port profiling

To compile source files for port profiling, simply compile them in debug mode and with the `source_data` flag turned on. For example:

```
| ?- logtalk_load(my_source_file, [debug(on), source_data(on)]).
```

Alternatively, you can also simply turn on the `debug` and `source_data` flags globally before compiling your source files:

```
| ?- set_logtalk_flag(debug, on), set_logtalk_flag(source_data, on).
```

Be aware, however, that loader files (e.g., library loader files) may override the default flag values, and thus the loaded files may not be compiled in debug mode. In this case, you will need to modify the loader files themselves.

5.19.5 Generating profiling data

After loading this tool and compiling the source files that you want to profile in debug mode, simply call the `ports_profiler::start` goal followed by the goals to be profiled. Use the `ports_profiler::stop` goal to stop profiling.

Note that the `ports_profiler::start/0` predicate implicitly selects the `ports_profiler` tool as the active debug handler. If you have additional debug handlers loaded (e.g., the debugger tool), those would no longer be active (there can be only one active debug handler at any given time).

5.19.6 Printing profiling data reports

After calling the goals that you want to profile, you can print a table with all profile data by typing:

```
| ?- ports_profiler::data.
```

To print a table with data for a single entity, use the query:

```
| ?- ports_profiler::data(Entity).
```

To print a table with data for a single entity predicate, use the query:

```
| ?- ports_profiler::data(Entity, Predicate).
```

In this case, the second argument must be either a predicate indicator, Name/Arity, or a non-terminal indicator, Name//Arity.

The profiling data can be reset using the query:

```
| ?- ports_profiler::reset.
```

To reset only the data for a specific entity, use the query:

```
| ?- ports_profiler::reset(Entity).
```

To illustrate the tool output, consider the family example in the Logtalk distribution:

```
| ?- {ports_profiler(loader)}.
...
yes

| ?- set_logtalk_flag(debug, on).
yes

| ?- logtalk_load(family(loader)).
...
yes

| ?- ports_profiler::start.
yes

| ?- addams::sister(Sister, Sibling).
Sister = wednesday,
Sibling = pubert ;
Sister = wednesday,
Sibling = pugsley ;
Sister = wednesday,
Sibling = pubert ;
Sister = wednesday,
Sibling = pugsley ;
no

| ?- ports_profiler::data.
-----
Entity      Predicate  Fact Rule Call Exit *Exit Fail Redo Error
-----
addams      female/1   2      0      1      1      1      0      1      0
addams      parent/2   8      0      4      3      5      1      5      0
relations   sister/2   0      1      1      0      4      1      4      0
-----
yes

| ?- ports_profiler::data(addams).
-----
Predicate   Fact Rule Call Exit *Exit Fail Redo Error
-----
female/1     2      0      1      1      1      0      1      0
parent/2     8      0      4      3      5      1      5      0
```

(continues on next page)

(continued from previous page)

```

-----
yes

| ?- ports_profiler::data(addams, parent/2).
-----
Clause  Count
-----
      1      1
      2      1
      3      2
      4      1
      5      1
      6      2
-----
yes

```

5.19.7 Interpreting profiling data

Some useful information that can be inferred from the profiling data include:

- which predicates are called more often (from the `call` port)
- unexpected failures (from the `fail` port)
- unwanted non-determinism (from the `*exit` port)
- performance issues due to backtracking (from the `*exit` and `redo` ports)
- predicates acting like a generator of possible solutions (from the `*exit` and `redo` ports)
- inefficient indexing of predicate clauses (from the `fact`, `rule`, and `call` ports)
- clauses that are never used or seldom used

The profiling data should be analyzed by taking into account the expected behavior for the profiled predicates.

5.19.8 Profiling Prolog modules

This tool can also be applied to Prolog modules that Logtalk is able to compile as objects. For example, if the Prolog module file is named `module.pl`, try:

```
| ?- logtalk_load(module, [debug(on), source_data(on)]).
```

Due to the lack of standardization of module systems and the abundance of proprietary extensions, this solution is not expected to work for all cases.

5.19.9 Profiling plain Prolog code

This tool can also be applied to plain Prolog code. For example, if the Prolog file is named `code.pl`, simply define an object including its code and declaring as public any predicates that you want to use as messages to the object. For example:

```
:- object(code).

   :- public(foo/2).
   :- include('code.pl').

:- end_object.
```

Save the object to an e.g. `code.lgt` file in the same directory as the Prolog file and then load it in debug mode:

```
| ?- logtalk_load(code, [debug(on), source_data(on)]).
```

In alternative, use the `object_wrapper_hook` provided by the `hook_objects` library:

```
| ?- logtalk_load(hook_objects(loader)).
...

| ?- logtalk_load(
    code,
    [hook(object_wrapper_hook), debug(on),
     source_data(on), context_switching_calls(allow)]
).
```

In this second alternative, you can then use the `<<>/2` context switch control construct to call the wrapped predicates. E.g.

```
| ?- code<<foo(X, Y).
```

With either wrapping solution, pay special attention to any compilation warnings that may signal issues that could prevent the plain Prolog code from working as-is when wrapped by an object. Often any required changes are straightforward (e.g., adding `use_module/2` directives for called module library predicates).

5.19.10 Known issues

Determinism information is currently not available when using Quintus Prolog as the backend compiler.

5.20 profiler

This tool contains simple wrappers for selected Prolog profiler tools.

5.20.1 Loading

This tool can be loaded using the query:

```
?- logtalk_load(profiler(loader)).
```

For sample queries, please see the `SCRIPT.txt` file in the tool directory.

5.20.2 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(profiler(tester)).
```

5.20.3 Supported backend Prolog compilers

Currently, this tool supports the profilers provided with SICStus Prolog 4, SWI-Prolog, and YAP. The tool includes two files:

- `yap_profiler.lgt`
simple wrapper for the YAP count profiler
- `sicstus_profiler.lgt`
simple wrapper for the SICStus Prolog 4 profiler

Logtalk also supports the YAP tick profiler (using the latest YAP development version) and the SWI-Prolog XPCE profiler. When using the XPCE profiler, you can avoid profiling the Logtalk compiler (which is invoked, e.g., when you use the `(:)/2` message-sending operator at the top-level interpreter) by compiling your code with the `optimize` flag turned on:

```
?- set_logtalk_flag(optimize, on).
true.

?- use_module(library(statistics)).
true.

?- profile(... :: ...).
...
```

Given that `prolog_statistics:profile/1` is a meta-predicate, Logtalk will compile its argument before calling it thanks to the `goal_expansion/2` hook predicate definitions in the adapter file. Without this hook definition, you would need to use instead (to avoid profiling the compiler itself):

```
?- logtalk << (prolog_statistics:profile(... :: ...)).
...
```

In either case, don't forget, however, to load the `prolog_statistics` module *before* using or compiling calls to the `profile/1` to allow the Logtalk compiler to access its meta-predicate template.

The profiler support attempts to conceal internal Logtalk compiler/runtime predicates and the generated entity predicates that implement predicate inheritance. Calls to internal compiler and runtime predicates have functors starting with `$lgt_`. Calls to predicates with functors such as `_def`, `_dcl`, or `_super`, used to implement inheritance, may still be listed in a few cases. Note that the time and the number of calls/redos of concealed predicates are added to the caller predicates.

5.20.4 Compiling source code for profiling

To get the user-level object and predicate names instead of the compiler-generated internal names when using the SWI-Prolog and YAP profilers, you must set `code_prefix` flag to a character other than the default `$` before compiling your source code. For example:

```
?- set_logtalk_flag(code_prefix, '.').
```

See also the `samples/settings-sample.lgt` file for automating the necessary setup at Logtalk startup.

5.21 sarif

The `sarif` tool serializes diagnostics produced by tools implementing the `tool_diagnostics_protocol` protocol into SARIF 2.1.0 reports.

5.21.1 API documentation

This tool API documentation is available at:

../apis/library_index.html#sarif

5.21.2 Loading

Load the tool using:

```
| ?- logtalk_load(sarif(loader)).  
...
```

5.21.3 Testing

To test this tool, load the `tester.lgt` file:

```
| ?- logtalk_load(sarif(tester)).
```

The test suite validates SARIF generation for single diagnostics producers and explicit aggregate reports, including JSON Schema validation against the SARIF 2.1.0 schema.

5.21.4 Usage

Use the `term/4` and `generate/4` predicates to generate a report for a single diagnostics producer, target, and options combination. These predicates generate a SARIF document with a single run.

For example:

```
| ?- sarif::generate(dead_code_scanner, entity(my_object), file('./report.sarif'), []).  
true.
```

Use the `term/2` and `generate/2` predicates to generate an explicit aggregate report from a list of specifications. Each specification must be a `tool_spec(Tool, Target, Options)` term and produces a single SARIF run. Runs are emitted in the same order as the specifications list.

For example:

```
| ?- sarif::generate([
    tool_spec(linter_reporter, all, []),
    tool_spec(dead_code_scanner, entity(my_object), [])
], file('./aggregate.sarif')).
true.
```

Each specification uses the same target and options accepted by the corresponding diagnostics producer. The aggregate API is explicit: the `sarif` tool does not infer producers or merge options across specifications.

5.22 sbom

This tool generates a Software Bill of Materials (SBOM) for an application, exported as a JSON file in either the ISO/IEC 5962:2021 standard SPDX 2.3 JSON format or the CycloneDX 1.6 JSON format:

<https://www.iso.org/standard/81870.html>

<https://cyclonedx.org/specification/overview/>

5.22.1 API documentation

This tool API documentation is available at:

[../..../apis/library_index.html#sbom](http://localhost:8080/..../apis/library_index.html#sbom)

5.22.2 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(sbm(loader)).
```

5.22.3 Testing

To run the tool tests, use the query:

```
| ?- logtalk_load(sbm(tester)).
```

5.22.4 Usage

The tool inspects the current Logtalk session and generates a JSON SBOM document describing:

- the loaded application
- the Logtalk version
- the backend Prolog system and version
- the installed packs that contributed loaded files to the current session

When exactly one loaded object conforms to the `application_protocol` protocol from the application library, the tool also uses that object as the default source for application metadata. If no conforming object exists, or if multiple conforming objects are loaded, the tool falls back to the explicit `sbom` options and built-in defaults.

For single-value application metadata such as `name/1`, `version/1`, `license/1`, `built_date/1`, `release_date/1`, `valid_until_date/1`, `supplier/1`, and `originator/1`, explicit `sbom` options override the values declared by the application object. For `creators/1`, explicit `creators/1` options override the application object creators list. Application external references declared by the object are combined with any `application_external_reference/2` options.

For CycloneDX exports, the generated BOM also includes a `serialNumber` and records the Logtalk `sbom` generator itself under `metadata.tools.components`. The BOM document license is exported as `metadata.licenses` with the SPDX identifier `CC0-1.0`. The CycloneDX BOM itself also exports top-level `externalReferences` for the Logtalk website and the Logtalk git repository.

- `bom_external_reference(Type, URL)` Adds a top-level CycloneDX `externalReferences` entry to the generated BOM. This option is ignored for SPDX exports. It can be repeated to export multiple references. The built-in Logtalk website and git repository references are still exported by default. Runtime components are exported with explicit CycloneDX scope `required`, while the `sbom` generator listed under `metadata.tools.components` is exported with scope `excluded`. When available from existing metadata, CycloneDX components also export `externalReferences`: Logtalk and the `sbom` tool export a website reference, the backend component exports a website reference from the bundled backend table, and loaded packs export their `home/1` and `version` source URL values as website and distribution references. For SPDX exports, the default `creationInfo.creators` entry is a versioned tool identifier derived from the `sbom` tool version. When URL metadata is available, SPDX packages also export `externalRefs`: the Logtalk and backend packages export website references, loaded packs export website and distribution references, and `application_external_reference/2` options are mapped to SPDX package external references for the application package.

The public predicates are:

- `document/1`
- `document/2`
- `export/1`
- `export/2`

The `document/1-2` predicates return a JSON term. The `export/1-2` predicates write the JSON document to any sink accepted by the `json::generate/2` predicate, including `atom(Atom)` and `file(Path)`. Follows the list of supported options.

Global/application options:

- `name(Name)` Sets the application name. Exported as the SPDX application package name and as the CycloneDX `metadata.component.name`. When exactly one object conforming to `application_protocol` is loaded and declares `name/1`, that value is used by default. Otherwise, the default is `loaded-application`.

- `format(Format)` Selects the export format. Possible values are `spdx` and `cyclonedx`. Default is `spdx`.
- `version(Version)` Sets the application version. Exported as the SPDX application package `versionInfo` and as the CycloneDX `metadata.component.version`. When exactly one object conforming to `application_protocol` is loaded and declares `version/1`, that value is used by default. Otherwise, the default is `0.0.0`.
- `application_license(License)` Sets the application license. Exported as the SPDX application package `licenseConcluded` and `licenseDeclared` fields and, unless the value is `NOASSERTION`, as the CycloneDX `metadata.component.licenses` entry, using `license.id` for SPDX license identifiers, expression for SPDX license expressions, and `license.name` otherwise. When exactly one object conforming to `application_protocol` is loaded and declares `license/1`, that value is used by default. Otherwise, the default is `NOASSERTION`.
- `application_built_date(Date)` Sets the application build date. Exported as the SPDX application package `builtDate` field and as the CycloneDX custom property `logtalk:sbom:built_date`. When exactly one object conforming to `application_protocol` is loaded and declares `built_date/1`, that value is used by default. Otherwise, the default is not exporting this information.
- `application_release_date(Date)` Sets the application release date. Exported as the SPDX application package `releaseDate` field and as the CycloneDX custom property `logtalk:sbom:release_date`. When exactly one object conforming to `application_protocol` is loaded and declares `release_date/1`, that value is used by default. Otherwise, the default is not exporting this information.
- `application_valid_until_date(Date)` Sets the application validity limit date. Exported as the SPDX application package `validUntilDate` field and as the CycloneDX custom property `logtalk:sbom:valid_until_date`. When exactly one object conforming to `application_protocol` is loaded and declares `valid_until_date/1`, that value is used by default. Otherwise, the default is not exporting this information.
- `application_supplier(Supplier)` Sets the application supplier. Exported as the SPDX application package `supplier` field. For CycloneDX, exported as `metadata.component.supplier` when using the `Organization: Name` convention and also as the custom property `logtalk:sbom:supplier`. When exactly one object conforming to `application_protocol` is loaded and declares `supplier/1`, that value is used by default. Otherwise, the default is not exporting this information.
- `application_originator(Originator)` Sets the application originator. Exported as the SPDX application package `originator` field. For CycloneDX, exported as `metadata.component.manufacturer` when using the `Organization: Name` convention, or as `metadata.component.authors` when using the `Person: Name` convention, and also as the custom property `logtalk:sbom:originator`. When exactly one object conforming to `application_protocol` is loaded and declares `originator/1`, that value is used by default. Otherwise, the default is not exporting this information.
- `application_external_reference(Type, Locator)` Adds application reference metadata. For SPDX exports, this becomes an application package `externalRefs` entry. For CycloneDX exports, URL-based references are exported under `metadata.component.externalReferences`, while package and provenance identifiers use the dedicated component identity fields. This option can be repeated to export multiple references. When exactly one object conforming to `application_protocol` is loaded, its declared `external_reference/2` metadata is also exported. The `homepage/1`, `distribution/1`, `package/1`, `repository/1`, `git_object_identifier/1`, and `software_heritage_identifier/1` predicates are mapped to `website`, `distribution`, `purl`, `vcs`, `gitoid`, and `swh` SBOM reference types, respectively. SPDX categories are inferred from the reference type, with `purl` exported under `PACKAGE-MANAGER` and `gitoid` and `swh` exported under `PERSISTENT-ID`. For CycloneDX exports, `purl` is written to `metadata.component.purl`, `gitoid` values to `metadata.component.omniborId`, and `swh` values to `metadata.component.swhid`.
- `namespace(Namespace)` Sets the base document namespace URI. A process and timestamp suffix is added automatically to guarantee uniqueness. This option only applies to SPDX exports

(documentNamespace) and is ignored for CycloneDX exports. Default is <https://logtalk.org/spdxdocs/logtalk-sbom>.

- `creators(Creators)` Adds all atoms in the list `Creators` to the SPDX `creationInfo.creators` list and the CycloneDX `metadata.authors` list. When no creator option is provided, and exactly one object conforming to `application_protocol` is loaded with a declared `creators/1` predicate, that list is used. Otherwise, the default for SPDX exports is the versioned tool identifier `Tool: Logtalk SBOM generator-<version>` and the default for CycloneDX exports is `Logtalk SBOM generator`.
- `validate_export(Boolean)` When true, validates the generated document against the bundled schema for the selected format before exporting it. Default is false.

Logtalk options:

- `logtalk_license(License)` Sets the Logtalk component license. Exported as the SPDX Logtalk package `licenseConcluded` and `licenseDeclared` fields and, unless the value is `NOASSERTION`, as the CycloneDX component licenses entry, using `license.id` for SPDX license identifiers, expression for SPDX license expressions, and `license.name` otherwise. Default is `Apache-2.0`.
- `logtalk_built_date(Date)` Sets the Logtalk build date. Exported as the SPDX Logtalk package `builtDate` field and as the CycloneDX custom property `logtalk:sbom:built_date`. Default is not exporting this information.
- `logtalk_release_date(Date)` Sets the Logtalk release date. Exported as the SPDX Logtalk package `releaseDate` field and as the CycloneDX custom property `logtalk:sbom:release_date`. Default is not exporting this information.
- `logtalk_valid_until_date(Date)` Sets the Logtalk validity limit date. Exported as the SPDX Logtalk package `validUntilDate` field and as the CycloneDX custom property `logtalk:sbom:valid_until_date`. Default is not exporting this information.
- `logtalk_supplier(Supplier)` Sets the Logtalk supplier. Exported as the SPDX Logtalk package `supplier` field. For CycloneDX, exported as the component supplier when using the `Organization: Name` convention and also as the custom property `logtalk:sbom:supplier`. Default is not exporting this information.
- `logtalk_originator(Originator)` Sets the Logtalk originator. Exported as the SPDX Logtalk package `originator` field. For CycloneDX, exported as the component manufacturer or authors entry, depending on the value convention, and also as the custom property `logtalk:sbom:originator`. Default is not exporting this information.

Backend options:

- `backend_license(License)` Sets the backend component license. Exported as the SPDX backend package `licenseConcluded` and `licenseDeclared` fields and, unless the value is `NOASSERTION`, as the CycloneDX component licenses entry, using `license.id` for SPDX license identifiers, expression for SPDX license expressions, and `license.name` otherwise. Default is the license specified in the backend/3 table.
- `backend_built_date(Date)` Sets the backend build date. Exported as the SPDX backend package `builtDate` field and as the CycloneDX custom property `logtalk:sbom:built_date`. Default is not exporting this information.
- `backend_release_date(Date)` Sets the backend release date. Exported as the SPDX backend package `releaseDate` field and as the CycloneDX custom property `logtalk:sbom:release_date`. Default is not exporting this information.
- `backend_valid_until_date(Date)` Sets the backend validity limit date. Exported as the SPDX backend package `validUntilDate` field and as the CycloneDX custom property `logtalk:sbom:valid_until_date`. Default is not exporting this information.

- `backend_supplier(Supplier)` Sets the backend supplier. Exported as the SPDX backend package supplier field. For CycloneDX, exported as the component supplier when using the Organization: Name convention and also as the custom property `logtalk:sbom:supplier`. Default is not exporting this information.
- `backend_originator(Originator)` Sets the backend originator. Exported as the SPDX backend package originator field. For CycloneDX, exported as the component manufacturer or authors entry, depending on the value convention, and also as the custom property `logtalk:sbom:originator`. Default is not exporting this information.

Pack options:

- `pack_license(Pack, License)` Sets the license for a loaded pack named `Pack`. Exported as the SPDX pack package `licenseConcluded` and `licenseDeclared` fields and, unless the value is `NOASSERTION`, as the CycloneDX component licenses entry, using `license.id` for SPDX license identifiers, expression for SPDX license expressions, and `license.name` otherwise. Default for packs without an explicit option is the result of sending the pack specification object the message `license(License)`, falling back to `NOASSERTION` when no license is available. Loaded packs also export a SPDX package `downloadLocation` from the pack specification `version/6` third argument and, when available, a SPDX homepage from the pack specification `home/1` predicate. Pack `home/1` and `version/6` URLs are also exported as SPDX package `externalRefs` with `website` and `distribution referenceType` values. Pack checksums are exported as SPDX package checksums and CycloneDX component hashes when the pack specification defines them in the `version/6` predicate fourth argument.
- `pack_built_date(Pack, Date)` Sets the build date for the loaded pack named `Pack`. Exported as the SPDX pack package `builtDate` field and as the CycloneDX custom property `logtalk:sbom:built_date`. Default is not exporting this information.
- `pack_release_date(Pack, Date)` Sets the release date for the loaded pack named `Pack`. Exported as the SPDX pack package `releaseDate` field and as the CycloneDX custom property `logtalk:sbom:release_date`. Default is not exporting this information.
- `pack_valid_until_date(Pack, Date)` Sets the validity limit date for the loaded pack named `Pack`. Exported as the SPDX pack package `validUntilDate` field and as the CycloneDX custom property `logtalk:sbom:valid_until_date`. Default is not exporting this information.
- `pack_supplier(Pack, Supplier)` Sets the supplier for the loaded pack named `Pack`. Exported as the SPDX pack package supplier field. For CycloneDX, exported as the component supplier when using the Organization: Name convention and also as the custom property `logtalk:sbom:supplier`. Default is not exporting this information.
- `pack_originator(Pack, Originator)` Sets the originator for the loaded pack named `Pack`. Exported as the SPDX pack package originator field. For CycloneDX, exported as the component manufacturer or authors entry, depending on the value convention, and also as the custom property `logtalk:sbom:originator`. Default is not exporting this information.

Examples:

```
| ?- sbom::document(Document).  
  
| ?- sbom::document(Document, [name(my_app), version('1.2.3')]).  
  
| ?- sbom::export(file('sbom.spdx.json')).  
  
| ?- sbom::export(file('sbom.cdx.json'), [format(cdx)]).  
  
| ?- sbom::export(atom(Atom), [  
    format(cdx),
```

(continues on next page)

(continued from previous page)

```

name(my_app),
version('1.2.3'),
application_license('MIT'),
logtalk_license('Apache-2.0'),
backend_license('BSD-2-Clause'),
application_built_date('2026-03-23T00:00:00Z'),
application_release_date('2026-03-23T00:00:00Z'),
application_valid_until_date('2027-03-23T00:00:00Z'),
application_supplier('Organization: Example Application'),
application_originator('Person: Application Maintainer'),
bom_external_reference(documentation, 'https://example.com/my_app/sbom'),
application_external_reference(website, 'https://example.com/my_app'),
application_external_reference(vcs, 'https://example.com/my_app.git'),
logtalk_supplier('Organization: Logtalk.org'),
backend_supplier('Organization: Backend Vendor'),
pack_license(my_pack, 'MIT'),
pack_supplier(my_pack, 'Organization: Pack Maintainer'),
creators(['Tool: My build pipeline', 'Person: Release Manager', 'Organization: ↵
↵Example, Inc.']),
validate_export(true)
]).

```

Use the `.spdx.json` extension for SPDX exports and the `.cdx.json` extension for CycloneDX exports.

See the `sbom-example.spdx.json` file for a representative SPDX export. See the `sbom-example.cdx.json` file for a representative CycloneDX export.

5.22.5 Known issues

The ECLiPSe and GNU Prolog backends fail several sbom tests and cannot be used for CycloneDX exports. The root cause is that both backends are limited to the US-ASCII charset, which prevents processing the SPDX/CycloneDX schema data required for CycloneDX license validation and export validation.

5.23 tool_diagnostics

This tool provides the shared `tool_diagnostics_protocol` protocol and the `tool_diagnostics_common` category used by developer tools such as `dead_code_scanner`, `lgt doc`, `lgt unit`, and `linter_reporter` to expose machine-readable diagnostics.

The protocol is designed so that producer tools can keep their own internal semantics while still providing a common interface for querying diagnostics directly or for generating SARIF reports using the `sarif` tool.

5.23.1 API documentation

This tool API documentation is available at:

../apis/library_index.html#tool_diagnostics_protocol

5.23.2 Loading

Load the protocol and common category definitions using:

```
| ?- logtalk_load(tool_diagnostics(loader)).  
...
```

5.23.3 Protocol overview

The protocol defines predicates for reporting:

- tool metadata using `diagnostics_tool/5`
- supported target patterns using `diagnostic_target/1`
- rule descriptors using `diagnostic_rule/5` and `diagnostic_rules/1`
- diagnostics using `diagnostic/2-3` and `diagnostics/2-3`
- summaries using `diagnostics_summary/2-3`
- preflight issues using `diagnostics_preflight/2-3`

Common target terms include `all`, `file(File)`, `directory(Directory)`, `rdirectory(Directory)`, `library(Library)`, `rlibrary(Library)`, `entity(Entity)`, and `tests(Object)`.

Predicates taking an `Options` argument must accept the common option explanations(`Boolean`). Tools may use it to include `explanation(Text)` properties when available, but implementations that do not provide additional explanation text must still accept the option as a no-op.

5.23.4 Term conventions

Tool metadata uses terms of the form:

```
diagnostics_tool(Id, Name, Version, InformationURI, Properties)
```

The `Id` and `Name` arguments identify the tool. The `Version` argument should be derived reflectively from the implementing object `info/1` metadata using `object_property/2` instead of being hardcoded. The `InformationURI` argument is a stable tool home or documentation URI and is exported by the `sarif` tool as the `SARIF tool.driver.informationUri` value.

The `Properties` list is intended for exporter and post-processing hints. The current `sarif` tool interprets the following entries:

- `guid(Guid)` sets the `SARIF tool.driver.guid` value when present.
- `fingerprint_algorithm(Algorithm)` sets the exported `run fingerprintAlgorithm` property and selects the canonical result fingerprint key. The current implementation distinguishes `canonical_finding_v1` from the default warning-oriented path and computes result identities from normalized relative locations instead of raw absolute file paths.

- `automation_id(target)` requests that the SARIF `run.automationDetails.id` value be derived from the queried target term, allowing different targets of the same tool to produce distinct run identities.
- `automation_id(tool)` denotes a tool-scoped automation identifier. In the current sarif implementation this is redundant because omitting `automation_id(target)` already falls back to the tool identifier.
- `include_invocations(true)` requests a SARIF `invocations` section, including `toolExecutionNotifications` when preflight issues exist.
- `include_git_metadata(true)` requests SARIF `gitBranch` and `gitCommitHash` run properties when a stable Git context can be inferred. When this succeeds, the sarif tool also normalizes artifact locations and fingerprint inputs relative to the repository root. The `gitBranch` value is descriptive metadata; the `gitCommitHash` value provides the stable revision identity used to disambiguate repository-relative paths.
- `include_version_control_provenance(true)` requests SARIF `versionControlProvenance` data when repository details can be inferred from diagnostic file paths.
- `count_key(Key)` selects the property name used for the exported count, e.g. `diagnosticsCount` or `totalWarnings`.

When `include_git_metadata(true)` is not present or a stable Git context cannot be inferred, the sarif tool falls back to application-root-relative artifact locations and fingerprint inputs instead.

Rule descriptors use terms of the form:

```
diagnostic_rule(RuleId, ShortDescription, FullDescription, DefaultSeverity, Properties)
```

The `diagnostic_rules/1` predicate should return these descriptors in a stable order. Rule descriptor properties may contain terms such as `guid(Guid)`, `help(Text)`, `help_uri(URI)`, `tags(Tags)`, and `precision(Precision)`.

Diagnostics use terms of the form:

```
diagnostic(RuleId, Severity, Confidence, Message, Context, File, Lines, Properties)
```

The `Context` argument is represented using the term `context(Kind, Identifier)` where `Kind` is a tool-defined atom such as `file`, `object`, `category`, `protocol`, or `test_set`. For file-scoped diagnostics, `Identifier` is the source file path when available and the empty atom otherwise. Diagnostics originating in an included file may still use an entity context when the `include/1` directive appears inside that entity; otherwise they are file-scoped.

The `File` argument is a source file path or the empty atom when unavailable. The `Lines` argument uses the Start-End notation. The value `0-0` means that no precise line information is available.

Supported severity values are `error`, `warning`, and `note`. Supported confidence values are `high`, `medium`, `low`, and `not_applicable`.

The `Properties` argument should include stable structured metadata such as `raw_term(Term)`, `explanation(Text)`, `flag(Flag)`, `directive(Directive)`, `indicator(Indicator)`, `key(Key)`, `test(Test)`, `option(Option)`, and any other tool-specific facts useful for post-processing.

Preflight issues use terms of the form:

```
preflight_issue(Id, Severity, Message, Context, File, Lines, Properties)
```

Preflight issues describe prerequisites or analysis quality issues instead of findings.

Summaries use terms of the form:

```
diagnostics_summary(Target, TotalContexts, TotalDiagnostics, Breakdown, ContextSummaries)
```

The Breakdown argument uses a `diagnostic_breakdown(RuleCounts, SeverityCounts, ConfidenceCounts)` term and `ContextSummaries` is a list of `context_summary(Context, DiagnosticsCount, Breakdown)` terms.

Rule counts use `rule_count(RuleId, Count)` terms, severity counts use `severity_count(Severity, Count)` terms, and confidence counts use `confidence_count(Confidence, Count)` terms.

5.23.5 Usage

Tools implementing this protocol can be queried directly for diagnostics after compiling code that results in the production of diagnostics. For example:

```
| ?- logtalk_load(dead_code_scanner(loader)).  
...  
  
| ?- logtalk_load(my_library(loader)).  
...  
  
| ?- dead_code_scanner::diagnostics(library(my_library), Diagnostics).  
...
```

To serialize diagnostics as SARIF, load the standalone sarif tool and generate a report from a diagnostics producer:

```
| ?- logtalk_load(sarif(loader)).  
...  
  
| ?- sarif::generate(dead_code_scanner, entity(my_object), file('./report.sarif'), []).  
true.
```

To generate an explicit aggregate SARIF report with multiple diagnostics producers, pass a list of `tool_spec(Tool, Target, Options)` terms:

```
| ?- sarif::generate([  
|     tool_spec(linter_reporter, all, []),  
|     tool_spec(dead_code_scanner, entity(my_object), [])  
| ], file('./aggregate.sarif')).  
true.
```

5.24 tutor

This tool adds explanations and suggestions for selected warning and error messages from the compiler/runtime and the developer tools. It's specially useful for new users not yet familiar with the warning and error messages and looking for advice on how to solve the reported issues.

5.24.1 API documentation

This tool API documentation is available at:

../apis/library_index.html#tutor

5.24.2 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(tutor(loader)).
```

5.24.3 Usage

Simply load the tool at startup (e.g., from a settings file). As an example, with this tool loaded, instead of terse compiler warnings such as:

```
*      No matching clause for goal: baz(a)
*      while compiling object main_include_compiler_warning
*      in file logtalk/examples/errors/include_compiler_warning.lgt between lines 37-38
*
*      Duplicated clause: b(one)
*      first found at or above line 45
*      while compiling object main_include_compiler_warning
*      in file logtalk/examples/errors/include_compiler_warning.lgt at or above line 48
```

the user will get:

```
*      No matching clause for goal: baz(a)
*      while compiling object main_include_compiler_warning
*      in file logtalk/examples/errors/include_compiler_warning.lgt between lines 37-38
*      Calls to locally defined predicates without a clause with a matching head
*      fail. Typo in a predicate argument? Predicate definition incomplete?
*
*      Duplicated clause: b(one)
*      first found at or above line 45
*      while compiling object main_include_compiler_warning
*      in file logtalk/examples/errors/include_compiler_warning.lgt at or above line 48
*      Duplicated clauses are usually a source code editing error and can
*      result in spurious choice-points, degrading performance. Delete or
*      correct the duplicated clause to fix this warning.
```

5.25 wrapper

This is a prototype tool to help port a plain Prolog application to Logtalk. It can also be used to enable applying other Logtalk developer tools, such as the documenting and diagramming tools, to plain Prolog code.

The tool takes a directory of Prolog files or a list of Prolog files, loads and wraps the code in each file using an object wrapper, and advises on missing directives to be added to those objects by using the compiler lint checker and the reflection API. The user can then either save the generated wrapper objects or copy and

paste the printed advice into the Prolog files (updating them to Logtalk files by adding the object opening and closing directives to the Prolog files). The wrapper objects use `include/1` directives to include the Prolog files and can be loaded for testing and for use with other tools. The wrapped Prolog files are not modified and thus require only read permission.

5.25.1 API documentation

This tool API documentation is available at:

../apis/library_index.html#wrapper

5.25.2 Loading

This tool can be loaded using the query:

```
| ?- logtalk_load(wrapper(loader)).
```

5.25.3 Workflows

The typical porting workflow is:

```
| ?- wrapper::rdirectory(root_directory_of_prolog_code).  
...  
| ?- wrapper::save.  
...
```

See the next section on how to customize the API calls for more flexible processing.

5.25.4 Customization

The tool can be customized by extending the wrapper object. A common scenario is when wrapping plain Prolog code just to take advantage, for example, of the documenting tool or for generating cross-referencing diagrams. In this case, we can workaround any compiler errors by specializing the inherited definitions for the `term_expansion/2` and `goal_expansion/2` predicates and then load the wrapper objects for further processing by using the `include_wrapped_files(false)` option described below.

The API predicates also accept a set of options for customization:

- `prolog_extensions(Extensions)`
list of file name extensions used to recognize Prolog source files (default is `['.pl', '.pro', '.prolog']`)
- `logtalk_extension(Extension)`
Logtalk file name extension to be used for the generated wrapper files (default is `'.lgt'`)
- `exclude_files(Files)`
list of Prolog source files to exclude (default is `[]`)
- `exclude_directories(Directories)`
list of sub-directories to exclude (default is `[]`)
- `include_wrapped_files(Boolean)`
generate `include/1` directives for the wrapped Prolog source files (default is `true`)

The use, by default, of `include/1` directives to wrap the code in Prolog source files facilitates running in parallel the Logtalk port and the original Prolog code. This is specially useful when the Prolog code being ported lacks a comprehensive set of tests that could be adapted to verify the Logtalk port. The generated Logtalk files can also take advantage of `uses/2` directives, including predicate aliases and predicate short-hands to minimize the changes to the original Prolog code and help verify if replacement calls to Logtalk library predicates provide the same semantics as the original calls.

5.25.5 Current limitations

- The tool cannot deal with syntax errors in the Prolog files. These errors usually occur when using a backend Prolog system different from the one used to compile the original plain Prolog code. A common cause of syntax errors is operator definitions. These can often be solved by defining those operators for the Prolog backend used to run Logtalk and this tool. An alternative is to preload the Prolog files where those operators are declared. Preloading the plain Prolog application can also help in wrapping it by ensuring that its dependencies are also loaded.
- The tool assumes that all files to be wrapped have different names (even if found in different directories). If that is not the case, the name conflicts must be manually solved before using the tool.
- There's only preliminary support for dealing with meta-predicates and advise on missing meta-predicate directives.

LIBRARIES

The documentation of each library can also be found in the library directory in the `NOTES.md` file.

6.1 Overview

This folder contains libraries of useful objects, categories, and protocols. The currently available libraries can be grouped by scope as follows (the grouping is thematic only, as some libraries naturally span multiple areas):

- Application metadata, configuration, and logging: `application`, `expand_library_alias_paths`, `command_line_options`, `logging`, and `options`.
- Types, collections and generic data processing: `basic_types`, `types`, `assignvars`, `deque`s, `dictionaries`, `nested_dictionaries`, `graphs`, `heaps`, `hierarchies`, `intervals`, `queues`, `sets`, `subsequences`, `union_find`, and `zipper`s.
- Combinatorics: `arrangements`, `cartesian_products`, `combinations`, `derangements`, `multisets`, `partitions`, and `permutations`.
- Meta-programming: `meta` and `meta_compiler`
- Monoid implementations: `expecteds`, `optionals`, and `validations`.
- Control, state, and developer support: `dependents`, `edcg`, `events`, `listing`, `reader`, and `term_io`.
- Term- and goal-expansion: `hook_flows` and `hook_objects`.
- Portability: `coroutining`, `dif`, `format`, `process`, `recorded_database`, and `timeout`.
- Dates and time: `dates`, `dates_tz`, `time_scales`, and `tzif`.
- Geospatial data: `crs_projections`, `geojson`, `geohash`, `geospatial`, `nmea`, `tle_orbits`, and `wkt_wkb`
- Space communications and telemetry: `ccsds_frames`, `ccsds_link_profiles`, `ccsds_packet_services`, `ccsds_packetization`, `ccsds_packets`, `ccsds_tc_services`, `ccsds_time_codes`, `ccsds_time_fields`,
- Text and NLP: `character_sets`, `grammars`, `stemming`, `string_distance`, and `strings`.
- Web: `html`, `mime_types`, and `url`.
- Identifiers: `cuid2`, `genint`, `gensym`, `ids`, `ksuid`, `nanoid`, `snowflakeid`, `ulid`, and `uuid`.
- Interchange formats and wire protocols: `amqp`, `avro`, `base32`, `base58`, `base64`, `base85`, `cbor`, `csv`, `json`, `json_ld`, `json_lines`, `json_pointer`, `json_rpc`, `json_schema`, `message_pack`, `mcp_server`, `protobuf`, `stomp`, `toml`, `toon`, `tsv`, and `yaml`.
- Coordination and data stores: `linda`, `memcached`, and `redis`.
- System and external integration: `git`, `java`, `os`, and `sockets`.

- Logic and symbolic computing: `datalog`.
- Security and integrity: `hashes` and `hmac`.
- Randomness: `arbitrary`, `mutations`, and `random`.
- Mathematics, statistics, and optimization: `ieee_754`, `linear_algebra`, `simulated_annealing`, and `statistics`.
- Machine learning:
 - Classification: `classification_protocols`, `adaptive_boosting_classifier`, `c45_classifier`, `gradient_boosting_classifier`, `kernel_svm_classifier`, `knn_classifier`, `lda_classifier`, `linear_svm_classifier`, `logistic_regression_classifier`, `naive_bayes_classifier`, `nearest_centroid_classifier`, `qda_classifier`, `random_forest_classifier`, and `sgd_classifier`.
 - Anomaly detection: `anomaly_detection_protocols`, `cusum_anomaly_detector`, `ewma_anomaly_detector`, `iqr_anomaly_detector`, `isolation_forest_anomaly_detector`, `knn_distance_anomaly_detector`, `lof_anomaly_detector`, `modified_z_score_anomaly_detector`, and `z_score_anomaly_detector`.
 - Regression: `regression_protocols`, `bayesian_ridge_regression`, `elastic_net_regression`, `gradient_boosting_regression`, `gaussian_process_regression`, `knn_regression`, `lasso_regression`, `linear_regression`, `random_forest_regression`, `regression_tree`, and `ridge_regression`.
 - Ranking: `ranking_protocols`, `borda_ranker`, `bradley_terry_ranker`, `colley_ranker`, `copeland_ranker`, `elo_ranker`, `glicko2_ranker`, `glicko2_periodic_ranker`, `hodge_rank`, `kemeny_young_ranker`, `massey_ranker`, `plackett_luce_ranker`, `plackett_luce_last_ranker`, `rank_centrality`, `ranked_pairs`, `regularized_bradley_terry_ranker`, `schulze_ranker`, and `thurstone_mosteller_ranker`.
 - Clustering: `clustering_protocols`, `agglomerative_clusterer`, `dbscan_clusterer`, `gaussian_mixture_clusterer`, `hdbscan_clusterer`, `hierarchical_clustering`, `kcenters_clusterer`, `kmeans_clusterer`, `kmedians_clusterer`, `kmedoids_clusterer`, `kmodes_clusterer`, `kprototypes_clusterer`, and `optics_clusterer`.
 - Dimension reduction: `dimension_reduction_protocols`, `ica_projection`, `kernel_pca_projection`, `lda_projection`, `nmf_projection`, `pca_projection`, `pls_projection`, `probabilistic_pca_projection`, and `random_projection`.
 - Pattern mining: `pattern_mining_protocols`, `frequent_pattern_mining_protocols`, `sequential_pattern_mining_protocols`, `apriori_pattern_miner`, `clo_span_pattern_miner`, `eclat_pattern_miner`, `fp_growth_pattern_miner`, `gsp_pattern_miner`, `prefix_span_pattern_miner`, and `spade_pattern_miner`.

In addition to the loader-based libraries, this directory also contains a small number of standalone reusable entities, namely `attributes`, `cloning`, `counters`, and `streamvars`.

Specific notes about individual libraries can be found in the corresponding library directory `NOTES.md` files.

A plain Prolog version of the Unicode 6.2 standard is also included in the `unicode_data` folder. See its `README.md` file for details.

A `parallel_logtalk_processes_setup.pl` Prolog file is also provided with sample code for selected backend Prolog compilers for initializing Logtalk processes such that each process uses a unique scratch directory, therefore allowing parallel process execution (e.g., for usage at continuous integration servers). Starting with Logtalk 3.48.0, this setup is only required in general when running with the `clean` flag turned off. See the comments in the file itself for usage instructions.

6.1.1 Library documentation

Specific notes about each library can be found in the corresponding NOTES.md files. HTML documentation for each library API can be found on the docs directory (open the ../docs/handbook/index.html file with your web browser).

6.1.2 Loading libraries

All the individual libraries can be loaded using the <library name>(loader) notation as argument for the compiling and loading predicates. For example:

```
| ?- logtalk_load(random(loader)).
```

There is also a file named all_loader.lgt that will load all libraries. Simply type the goal:

```
| ?- logtalk_load(library(all_loader)).
```

As a general rule, always use the corresponding loader file to load a library. Most library entities are part of small hierarchies or depend on other libraries and thus cannot be loaded and compiled separately (e.g., the list object implements the listp protocol and is part of a basic types hierarchy). Using the loader files takes care of all dependencies and also ensures compilation in optimized mode.

6.1.3 Testing libraries

Most of the libraries include unit tests in their directory, together with a tester.lgt file for running them. For example, to run the tests for the random library, we can use the goal:

```
| ?- logtalk_load(random(tester)).
```

To run all library tests, we can use the logtalk_tester automation script from the library directory at the root of the Logtalk distribution. For example, assuming the Logtalk user directory is ~/logtalk and that we want to run the tests using ECLiPSe as the backend Prolog compiler:

```
$ cd ~/logtalk/library
$ logtalk_tester -p eclipse
```

6.1.4 Credits

Some code in this library is based on public domain Prolog code, in particular, code adopted from the Edinburgh Prolog library. The definition of the predicate reverse/2 in the list object is from Richard O’Keefe and can be found in its book “The Craft of Prolog”.

Some elements of this library are inspired by Richard O’Keefe library proposal available at:

<http://www.cs.otago.ac.nz/staffpriv/ok/pllib.htm>

Some libraries, or parts of libraries, are either ports of Prolog system libraries or inspired by Prolog system libraries. See the individual library notes for details. See also the NOTICE.txt file at the root of the Logtalk distribution for copyright information on third-party source code.

6.1.5 Other notes

Some files contained in this directory represent work in progress and are not loaded by default by any loader utility file.

6.2 adaptive_boosting_classifier

Adaptive Boosting (aka AdaBoost) classifier using C4.5 decision trees as base learners. Implements the SAMME (Stagewise Additive Modeling using a Multi-class Exponential loss function) variant, which supports multi-class classification by adjusting the weight update formula to account for the number of classes. Builds an ensemble of weighted decision trees where each subsequent tree focuses on the examples misclassified by previous trees: after each iteration, the weights of misclassified examples are increased so that subsequent learners focus more on difficult cases.

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning an ensemble classifier from a dataset, using it to make predictions (with class probabilities), and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. See `test_files` directory for examples.

6.2.1 API documentation

Open the ../docs/library_index.html#adaptive_boosting_classifier link in a web browser.

6.2.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(adaptive_boosting_classifier(loader)).
```

6.2.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(adaptive_boosting_classifier(tester)).
```

6.2.4 Features

- **Adaptive Boosting:** Iteratively trains weighted decision trees, focusing on misclassified examples
- **SAMME Algorithm:** Supports multi-class classification via the SAMME variant of AdaBoost
- **C4.5 Base Learners:** Uses C4.5 decision trees as weak learners for each boosting round
- **Weighted Voting:** Final predictions determined by weighted voting where more accurate learners have higher weight
- **Probability Estimation:** Provides confidence scores based on weighted vote proportions
- **Early Stopping:** Training stops early if a perfect classifier is found or if a base learner is worse than random

- **Configurable Options:** Number of estimators (boosting rounds) configurable via predicate options
- **Classifier Export:** Learned classifiers can be exported as predicate clauses

6.2.5 Options

The following options can be passed to the `learn/3` predicate:

- `number_of_estimators(N)`: Number of boosting rounds / weak learners (default: 10)

6.2.6 Classifier representation

The learned classifier is represented as a compound term:

```
ab_classifier(WeightedTrees, ClassValues, Options)
```

Where:

- **WeightedTrees**: List of `weighted_tree(Alpha, C45Tree, AttributeNames)` elements
- **ClassValues**: List of possible class values
- **Options**: List of options used during learning

When exported using `export_to_clauses/4` or `export_to_file/4`, this classifier term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.2.7 References

1. Freund, Y. and Schapire, R.E. (1997). "A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting". *Journal of Computer and System Sciences*, 55(1), 119-139.
2. Zhu, J., Zou, H., Rosset, S., and Hastie, T. (2009). "Multi-class AdaBoost". *Statistics and Its Interface*, 2(3), 349-360.
3. Quinlan, J.R. (1993). "C4.5: Programs for Machine Learning". Morgan Kaufmann.

6.2.8 Usage

Learning a Classifier

```
% Learn an AdaBoost classifier with default options (10 estimators)
| ?- adaptive_boosting_classifier::learn(play_tennis, Classifier).
...

% Learn with custom options
| ?- adaptive_boosting_classifier::learn(play_tennis, Classifier, [number_of_
↪estimators(20)]).
...
```

Making Predictions

```
% Predict class for a new instance
| ?- adaptive_boosting_classifier::learn(play_tennis, Classifier),
      adaptive_boosting_classifier::predict(Classifier, [outlook-sunny, temperature-hot,
↳humidity-high, wind-weak], Class).
Class = no
...

% Get probability distribution from weighted voting
| ?- adaptive_boosting_classifier::learn(play_tennis, Classifier),
      adaptive_boosting_classifier::predict_probabilities(Classifier, [outlook-overcast,
↳temperature-mild, humidity-normal, wind-weak], Probabilities).
Probabilities = [yes-0.9, no-0.1]
...
```

Exporting the Classifier

```
% Export as predicate clauses
| ?- adaptive_boosting_classifier::learn(play_tennis, Classifier),
      adaptive_boosting_classifier::export_to_clauses(play_tennis, Classifier, my_boost,
↳Clauses).
...

% Export to a file
| ?- adaptive_boosting_classifier::learn(play_tennis, Classifier),
      adaptive_boosting_classifier::export_to_file(play_tennis, Classifier, my_boost, 'boost.
↳pl').
...
```

Using a Saved Classifier

```
% Load and use a previously saved classifier
| ?- logtalk_load('boost.pl'),
      my_boost(Classifier),
      adaptive_boosting_classifier::predict(Classifier, [outlook-sunny, temperature-cool,
↳humidity-normal, wind-weak], Class).
Class = yes
...
```

Printing the Classifier

```
% Print a summary of the AdaBoost classifier
| ?- adaptive_boosting_classifier::learn(play_tennis, Classifier),
      adaptive_boosting_classifier::print_classifier(Classifier).

AdaBoost Classifier
=====

Learning options: [number_of_estimators(10)]

Class values: [yes,no]
Number of estimators: 10

Weighted trees:
  Estimator 1 (alpha=1.2345, features: [outlook,temperature,humidity,wind]):
    -> tree rooted at outlook
    ...
  ...
```

6.3 agglomerative_clusterer

Agglomerative clusterer.

The library implements the `clusterer_protocol` defined in the `clustering_protocols` library. It provides predicates for learning a clusterer from a dataset, assigning new instances to clusters, and exporting the learned clusterer as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol from the `clustering_protocols` library.

6.3.1 API documentation

Open the ../apis/library_index.html#agglomerative_clusterer link in a web browser.

6.3.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(agglomerative_clusterer(loader)).
```

6.3.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(agglomerative_clusterer(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(agglomerative_clusterer(tester_performance)).
```

6.3.4 Features

- **Bottom-Up Clustering:** Uses deterministic bottom-up agglomerative_clusterer clustering and stops when the requested number of clusters is reached.
- **Continuous Datasets:** Accepts datasets containing only continuous attributes.
- **Linkage Strategies:** Supports single, complete, and average linkage.
- **Distance Metrics:** Supports euclidean and manhattan distances.
- **Optional Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Linkage-Aware Prediction:** New instances are assigned to the nearest learned cluster using the selected linkage strategy and distance metric applied to the learned cluster members.
- **Deterministic Ordering:** Equal-distance merges are broken using node-id order and final clusters are ordered by minimum training example id so equivalent dataset permutations keep the same cluster ids.
- **Cached Distances:** Inter-cluster distances are cached and incrementally updated after each merge instead of being fully recomputed from cluster members at every iteration.
- **Priority-Queue Merge Selection:** Candidate merges are tracked in a min-heap keyed by distance and node-id order, allowing stale entries to be discarded lazily while keeping merge selection deterministic.
- **Rich Diagnostics:** Diagnostics report the training example count, performed merge count, initial pair count, maximum heap size, stale-pair discard count, deterministic pair-selection strategy, and linkage-aware prediction strategy.
- **Fail-Fast Consistency Checks:** Internal heap, active-node, and cached-distance inconsistencies raise explicit `agglomerative_error/2` exceptions instead of failing silently.
- **Portable Export:** Learned clusterers can be exported as clauses or files and reused later.

6.3.5 Options

The following options can be passed to the `learn/3` predicate:

- `k(K)`: Number of clusters to retain after merging. Default is 2.
- `linkage(Linkage)`: Linkage strategy to use. Options: single, complete, or average (default).
- `distance_metric(Metric)`: Distance metric to use. Options: euclidean (default) or manhattan.
- `feature_scaling(FeatureScaling)`: Whether to standardize continuous attributes before clustering. Options: on (default) or off.

6.3.6 Clusterer representation

The learned clusterer is represented as a compound term with the functor chosen by the user when exporting the clusterer and arity 4. For example:

```
agglomerative_clusterer(Encoders, Clusters, Prototypes, Options, Diagnostics)
```

Where:

- Encoders: List of continuous attribute encoders storing attribute name, mean, and scale.
- Clusters: List of cluster(Id, Points) terms in cluster-id order.
- Prototypes: List of average vectors used for display, diagnostics, and export metadata.
- Options: Effective training options used to learn the clusterer.
- Diagnostics: Training metadata including heap and prediction details.

6.3.7 Diagnostics

The diagnostics/2 predicate returns metadata including:

- training_example_count/1
- merge_count/1
- initial_pair_count/1
- maximum_heap_size/1
- stale_pair_discard_count/1
- pair_selection(priority_queue)
- prediction_strategy(cluster_member_linkage_distance)
- tie_breaking(node_id_order)
- options/1

6.4 amqp

Portable AMQP 0-9-1 (Advanced Message Queuing Protocol) client implementation. This library uses the sockets library and supports all backend Prolog systems supported by that library: ECLiPSe, GNU Prolog, SICStus Prolog, SWI-Prolog, Trealla Prolog, and XVM.

<https://www.rabbitmq.com/amqp-0-9-1-reference.html>

AMQP 0-9-1 is a binary wire-level protocol for message-oriented middleware. It is the de facto standard for message broker interoperability and is supported by RabbitMQ, Apache Qpid, Apache ActiveMQ, and other brokers.

6.4.1 API documentation

Open the ../apis/library_index.html#amqp link in a web browser.

6.4.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(amqp(loader)).
```

6.4.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(amqp(tester)).
```

Note: Integration tests require a running AMQP 0-9-1 server (e.g., RabbitMQ). RabbitMQ AMQP listens on port 5672 by default with guest/guest credentials.

6.4.4 Features

- Full AMQP 0-9-1 protocol support with binary frame encoding/decoding
- Connection management with heartbeat negotiation
- Automatic reconnection with configurable retry attempts and delays
- Connection pooling with automatic connection management
- Multiple concurrent channels over a single connection
- Exchange operations: declare, delete, bind, unbind
- Queue operations: declare, delete, bind, unbind, purge
- Basic messaging: publish, consume, get, ack, nack, reject
- Quality of service (QoS) with prefetch settings
- Transaction support: tx.select, tx.commit, tx.rollback
- Publisher confirms (RabbitMQ extension)
- SASL PLAIN authentication
- Custom message properties and headers
- Content properties: content-type, delivery-mode, priority, etc.

6.4.5 Usage

Connecting to an AMQP server

Basic connection to RabbitMQ with default settings:

```
?- amqp::connect(localhost, 5672, Connection, []).
```

Connection with custom credentials and virtual host:

```
?- amqp::connect(localhost, 5672, Connection, [
    username('myuser'),
    password('mypassword'),
    virtual_host('/myvhost'),
    heartbeat(30)
]).
```

Connection with automatic reconnection enabled:

```
?- amqp::connect(localhost, 5672, Connection, [
    reconnect(true),
    reconnect_attempts(5),
    reconnect_delay(2)
]).
```

This will attempt to connect up to 5 times with a 2 second delay between attempts. If all attempts fail, an `amqp_error(reconnect_failed)` error is thrown. The reconnection options are:

- `reconnect(Boolean)` - Enable automatic reconnection (default: `false`)
- `reconnect_attempts(N)` - Maximum number of connection attempts (default: 3)
- `reconnect_delay(Seconds)` - Delay between attempts in seconds (default: 1)

Opening a channel

AMQP operations require a channel. Open one on the connection:

```
?- amqp::channel_open(Connection, 1, Channel).
```

You can open multiple channels (with different numbers) on a single connection.

Declaring exchanges and queues

Declare a direct exchange:

```
?- amqp::exchange_declare(Channel, 'my.exchange', [
    type(direct),
    durable(true)
]).
```

Declare a queue:

```
?- amqp::queue_declare(Channel, 'my.queue', [
    durable(true),
    auto_delete(false)
]).
```

Declare a queue with server-generated name:

```
?- amqp::queue_declare(Channel, Queue, [exclusive(true)]).
% Queue will be unified with the generated name
```

Bind a queue to an exchange:

```
?- amqp::queue_bind(Channel, 'my.queue', 'my.exchange', [
    routing_key('my.routing.key')
]).
```

Publishing messages

Publish a simple message:

```
?- amqp::basic_publish(Channel, 'my.exchange', 'Hello, World!', [
    routing_key('my.routing.key')
]).
```

Publish with properties:

```
?- amqp::basic_publish(Channel, 'my.exchange', '{"data": "json"}', [
    routing_key('my.routing.key'),
    content_type('application/json'),
    delivery_mode(2), % persistent
    correlation_id('abc123'),
    reply_to('reply.queue')
]).
```

Publish with custom headers:

```
?- amqp::basic_publish(Channel, 'my.exchange', 'Message', [
    routing_key('my.key'),
    headers([
        'x-custom-header'-longstr('value'),
        'x-priority'-int(5)
    ])
]).
```

Consuming messages

Start a consumer:

```
?- amqp::basic_consume(Channel, 'my.queue', [
    consumer_tag('my-consumer'),
    no_ack(false)
]).
```

Receive messages:

```
?- amqp::receive(Channel, Message, [timeout(5000)]).
```

Extract message data:

```
?- amqp::message_body(Message, Body),
    amqp::message_delivery_tag(Message, DeliveryTag),
    amqp::message_property(Message, content_type, ContentType).
```

Acknowledge a message:

```
?- amqp::basic_ack(Channel, DeliveryTag, []).
```

Reject a message:

```
?- amqp::basic_reject(Channel, DeliveryTag, [requeue(true)]).
```

Synchronous get

Get a single message synchronously:

```
?- amqp::basic_get(Channel, 'my.queue', [no_ack(false)]).
```

Quality of Service

Set prefetch count to limit unacknowledged messages:

```
?- amqp::basic_qos(Channel, [prefetch_count(10)]).
```

Transactions

Enable transactions on a channel:

```
?- amqp::tx_select(Channel).
```

Publish messages within a transaction:

```
?- amqp::basic_publish(Channel, 'exchange', 'Msg1', [routing_key('key')]),
    amqp::basic_publish(Channel, 'exchange', 'Msg2', [routing_key('key')]),
    amqp::tx_commit(Channel).
```

Rollback a transaction:

```
?- amqp::tx_rollback(Channel).
```

Publisher Confirms (RabbitMQ extension)

Enable publisher confirms:

```
?- amqp::confirm_select(Channel).
```

Closing connections

Close a channel:

```
?- amqp::channel_close(Channel).
```

Close the connection:

```
?- amqp::close(Connection).
```

6.4.6 Connection Pooling

The library provides connection pooling through the `amqp_pool` category. To create a connection pool, define an object that imports this category:

Defining a pool

```
:- object(my_pool,  
    imports(amqp_pool)).  
:- end_object.
```

Initializing the pool

Initialize the pool with configuration options:

```
?- my_pool::initialize([  
    host(localhost),  
    port(5672),  
    min_size(2),  
    max_size(10),  
    connection_options([  
        username('guest'),  
        password('guest'),  
        virtual_host('/')  
    ])  
]).
```

Pool configuration options:

- `host(Host)` - AMQP server hostname (default: `localhost`)

- `port(Port)` - AMQP server port (default: 5672)
- `min_size(N)` - Minimum connections to maintain (default: 1)
- `max_size(N)` - Maximum connections allowed (default: 10)
- `connection_options(Options)` - Options passed to `amqp::connect/4` (default: [])

Acquiring and releasing connections

Manually acquire and release connections:

```
?- my_pool::acquire(Connection),
   amqp::channel_open(Connection, 1, Channel),
   % ... use the channel ...
   amqp::channel_close(Channel),
   my_pool::release(Connection).
```

Using with_connection/1

Use `with_connection/1` for automatic connection management:

```
?- my_pool::with_connection(do_work).

do_work(Connection) :-
    amqp::channel_open(Connection, 1, Channel),
    amqp::basic_publish(Channel, '', 'Hello!', [routing_key('my.queue')]),
    amqp::channel_close(Channel).
```

The connection is automatically released even if the goal fails or throws an exception.

Pool statistics

Get pool statistics:

```
?- my_pool::stats(stats(Available, InUse, Total, MinSize, MaxSize)).
```

Resizing the pool

Resize the pool at runtime:

```
?- my_pool::resize(5, 20).
```

Destroying the pool

Close all connections and clear pool state:

```
?- my_pool::destroy.
```

Creating pools dynamically

Pools can also be created at runtime using `create_object/4`:

```
?- create_object(dynamic_pool, [imports(amqp_pool)], [], []),  
   dynamic_pool::initialize([host(localhost), port(5672)]).
```

6.4.7 Binary Frame Encoding

AMQP 0-9-1 is a binary protocol. The library provides low-level encoding and decoding predicates for working with raw frames:

Frame structure

An AMQP frame consists of:

- Type (1 byte): 1=method, 2=header, 3=body, 8=heartbeat
- Channel (2 bytes): Channel number (0 for connection frames)
- Size (4 bytes): Payload size
- Payload (Size bytes): Frame-specific data
- Frame end (1 byte): 0xCE marker

Encoding/decoding frames

Encode a frame to bytes:

```
?- amqp::encode_frame(Frame, Bytes).
```

Decode bytes to a frame:

```
?- amqp::decode_frame(Bytes, Frame).
```

6.4.8 Data Types

The library handles AMQP data types automatically:

- **octet**: 8-bit unsigned integer
- **short**: 16-bit unsigned integer (big-endian)
- **long**: 32-bit unsigned integer (big-endian)
- **longlong**: 64-bit unsigned integer (big-endian)

- **shortstr**: Short string (length ≤ 255)
- **longstr**: Long string
- **table**: Field table (key-value pairs)
- **array**: Field array

Field values in tables use type tags:

- `bool(true/false)` - Boolean
- `byte(V)` - Signed 8-bit
- `short(V)` - Signed 16-bit
- `int(V)` - Signed 32-bit
- `long(V)` - Signed 64-bit
- `float(V)` - 32-bit float
- `double(V)` - 64-bit float
- `longstr(V)` - Long string
- `table(Pairs)` - Nested table
- `array(Values)` - Array
- `timestamp(V)` - Timestamp
- `void` - No value

6.4.9 Float/Double Encoding/Decoding

The library implements proper IEEE 754 encoding and decoding for floating-point values:

Float (32-bit IEEE 754 single precision):

- Format: 1 sign bit + 8 exponent bits (bias 127) + 23 mantissa bits
- Encoded as 4 bytes in big-endian order
- Used for `float(Value)` field values

Double (64-bit IEEE 754 double precision):

- Format: 1 sign bit + 11 exponent bits (bias 1023) + 52 mantissa bits
- Encoded as 8 bytes in big-endian order
- Used for `double(Value)` field values

Special values representation:

The library represents IEEE 754 special values using Prolog compound terms:

- `@infinity` - Positive infinity
- `@negative_infinity` - Negative infinity
- `@not_a_number` - NaN (Not a Number)

These special values are automatically encoded to and decoded from their standard IEEE 754 binary representations.

Known limitation:

Most Prolog backends cannot distinguish between `-0.0` and `0.0` when comparing floating-point values. While the library correctly encodes and decodes both positive and negative zero according to IEEE 754, these values will typically compare as equal in Prolog arithmetic operations. This limitation is unlikely to affect typical AMQP messaging applications.

6.4.10 Error Handling

The library throws structured errors:

- `amqp_error(connection_failed)` - TCP connection failed
- `amqp_error(auth_failed)` - Authentication failed
- `amqp_error(protocol_error(Msg))` - Protocol violation
- `amqp_error(channel_error(Msg))` - Channel-level error
- `amqp_error(exchange_error(Msg))` - Exchange operation failed
- `amqp_error(queue_error(Msg))` - Queue operation failed
- `amqp_error(basic_error(Msg))` - Basic operation failed
- `amqp_error(tx_error(Msg))` - Transaction error

6.4.11 Comparison with STOMP

AMQP 0-9-1 and STOMP are both messaging protocols, but differ significantly:

Feature	AMQP 0-9-1	STOMP
Protocol type	Binary	Text
Complexity	High	Low
Exchange types	Multiple	Broker-dependent
Routing	Flexible	Simple
Transactions	Native	Native
QoS	Native prefetch	Limited
Performance	Higher	Lower

Use AMQP when you need:

- Fine-grained routing control
- High performance
- Advanced message patterns
- Exchange-based routing

Use STOMP when you need:

- Simple text-based protocol
- Easy debugging
- Protocol simplicity

6.4.12 Future Work

- Logtalk protocols for messaging patterns (request/reply, pub/sub, etc.)
- Categories for common message transformations
- Async message handlers

6.4.13 Known Limitations and Issues

- SSL/TLS connections not yet supported (use stunnel or similar)
- Heartbeat sending must be done manually via `send_heartbeat/1`
- Most Prolog backends cannot distinguish between `-0.0` and `0.0` (see the “Float/Double Encoding/Decoding” section above for details)
- GNU Prolog support is partial as it doesn’t currently support null characters in atoms

6.4.14 AMQP 1.0 vs AMQP 0-9-1

This library implements AMQP 0-9-1 only. Despite the similar name, **AMQP 1.0 is a fundamentally different protocol** and is not supported by this library. The two versions are not wire-compatible and have different conceptual models:

Aspect	AMQP 0-9-1	AMQP 1.0
Design	Broker-centric	Peer-to-peer
Routing	Exchanges + queues	Link-based addressing
Frame structure	Custom binary framing	Layered performatives
Data encoding	Custom type system	CBOR-like encoding

Key differences:

- AMQP 0-9-1 uses exchanges, bindings, and queues as core protocol concepts
- AMQP 1.0 abstracts these as broker-specific “nodes” and uses links
- Frame encoding, handshake, and message format are completely different
- Minimal code could be shared between the two implementations

AMQP 1.0 support would require a separate library with its own distinct API reflecting the link-based model, rather than the exchange/queue model of AMQP 0-9-1.

Practical note: RabbitMQ and most message brokers continue to primarily use AMQP 0-9-1, making this library suitable for the vast majority of use cases.

6.5 anomaly_detection_protocols

This library provides protocols used in the implementation of machine learning anomaly-detection algorithms. Datasets are represented as objects implementing the `anomaly_dataset_protocol` protocol. Anomaly detectors are represented as objects importing the `anomaly_detector_common` category which imports the `anomaly_detector_protocol` protocol. The category provides shared `learn/2`, `predict/3-4`, `diagnostics/2`, `diagnostic/2`, `anomaly_detector_options/2`, file export, baseline training-selection helpers, and dataset helper predicates. It keeps threshold-based prediction and export behavior separate from the algorithm-specific learning, scoring, clause export, pretty-printing, and diagnostics metadata code.

The shared category also provides reusable protected predicates for baseline-only training workflows. Libraries can use the `baseline_class_values/1` and `baseline_selection_policy/1` options via a single helper instead of reimplementing class-label validation and baseline-example filtering or rejection logic locally.

Learned detector terms can be validated explicitly using the shared `check_anomaly_detector/1` and `valid_anomaly_detector/1` predicates. This validation API is never called implicitly by scoring, prediction, printing, or export predicates.

This library also provides a reusable shared category, anomaly benchmark datasets, and a small family smoke-test suite.

6.5.1 Export header format

The shared exporter in the `anomaly_detector_common` category writes a header before the exported clauses in the following format:

```
% exported anomaly detector predicate: Functor/Arity
% training dataset: Dataset
% options: Options
% Functor(Detector)
Functor(Detector)
```

The exported clauses serialize the learned detector term as a single predicate argument so that loading the file gives a detector term that can be passed directly to the `predict/3-4` and `score_all/3` predicates.

When exporting a serialized detector term, using a noun such as `detector/1` or `model/1` is recommended.

6.5.2 API documentation

Open the ../apis/library_index.html#anomaly-detection-protocols link in a web browser.

6.5.3 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(anomaly_detection_protocols(loader)).
```

6.5.4 Testing

To run the library smoke tests, shared category tests, and dataset checks, load the `tester.lgt` file:

```
| ?- logtalk_load(anomaly_detection_protocols(tester)).
```

6.5.5 Test datasets

Several sample datasets are included in the `test_datasets` directory:

- `gaussian_anomalies.lgt` — A synthetic 2D anomaly detection dataset with 48 examples and 2 continuous attributes (x, y). Normal points are sampled from a standard normal distribution centered at the origin. Anomalous points are placed far from the cluster center. Inspired by the canonical test case used in the Extended Isolation Forest paper by Hariri et al. (2019).
- `malformed_anomalies.lgt` — A negative fixture with invalid class labels for testing family-level dataset validation.
- `mixed_anomalies.lgt` — A small mixed-feature anomaly dataset with 16 examples, 2 continuous attributes (age, income), and 2 categorical attributes (student, credit_rating). Includes missing values and uncommon feature combinations to exercise anomaly-detection code on heterogeneous data.
- `mixed_distance_behaviors.lgt` — A compact mixed-feature anomaly fixture with 8 examples, 2 continuous attributes (size, weight), and 2 categorical attributes (color, shape). Intended for smoke-testing continuous plus categorical distance behavior and basic mixed-data handling.
- `sensor_anomalies.lgt` — A synthetic industrial sensor anomaly dataset with 40 examples and 3 continuous attributes (temperature, pressure, vibration). Contains missing values (14 examples with missing values, represented using anonymous variables). Normal readings cluster around typical operating ranges. Anomalous readings show extreme values indicating equipment malfunction.
- `shuttle_anomalies.lgt` — A subset of the Statlog Shuttle dataset with 50 examples and 9 continuous attributes representing sensor readings from the NASA Space Shuttle. Class 1 (Rad Flow) is the majority class (normal), while all other classes are treated as anomalies. Originally from Catlett, J. (1991). Available from the UCI Machine Learning Repository: <https://archive.ics.uci.edu/dataset/148/statlog+shuttle>
- `water_potability.lgt` — A water potability dataset with 48 examples and 9 continuous attributes (pH, hardness, solids, chloramines, sulfate, conductivity, organic carbon, trihalomethanes, turbidity). Normal instances represent potable water samples within acceptable ranges. Anomalous instances represent water samples with hazardous contamination levels. Based on the publicly available Water Quality dataset (Kadiwal, A., 2020, Kaggle).

6.6 application

This library provides the `application_common` category and the `application_protocol` protocol for declaring application metadata, including optional git-related facts such as repository URL, branch, commit, author, and commit message, plus optional package and archive identifiers. Application metadata is typically consumed by tools such as `sbom`.

The library distinguishes two kinds of information:

- release metadata, such as name, version, description, license, distribution location, package identifier, release date, and validity date
- optional source provenance metadata, such as repository URL, branch, commit, commit date, author, commit message, and archive identifiers

Source provenance predicates are explicit facts declared by the application. They are not a reflection of the current status of a local git checkout and should not be assumed to identify a released artifact unless the application author chooses to state them for that purpose.

6.6.1 API documentation

Open the `../apis/library_index.html#application` link in a web browser.

6.6.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(application(loader)).
```

6.6.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(application(tester)).
```

6.6.4 Usage

Define an application metadata object importing the `application_common` category and declaring the meta-data that is known explicitly. For example:

```
:- object(my_application,
    imports(application_common)).

    name(my_application).
    version('1.2.3').
    description('Example application metadata object').
    license('Apache-2.0').
    homepage('https://example.com/my_application').
    distribution('https://example.com/my_application/releases/download/v1.2.3/my_application.
↳tgz').
    package('pkg:generic/my_application@1.2.3').
    creators(['Tool: Build pipeline', 'Person: Alice']).
    supplier('Organization: Example Application').
    originator('Organization: Upstream Project').
    repository('https://example.com/my_application.git').
    repository_branch(main).
    repository_commit('0123456789abcdef0123456789abcdef01234567').
    git_object_identifier('gitoid:commit:sha1:0123456789abcdef0123456789abcdef01234567').
    software_heritage_identifier('swh:1:rev:0123456789abcdef0123456789abcdef01234567').
    repository_commit_author('Alice').

:- end_object.
```

The library distinguishes between creators and originator using two predicates:

- `creators/1` identifies the people, organizations, or tools credited with creating the application or preparing its release metadata
- `originator/1` identifies the original source of the software when that is relevant and distinct from the creators

For example, an internal application might use the same party for both `creators/1` and `originator/1`. A packaged or redistributed application may use `creators/1` for the team or toolchain preparing the release metadata and `originator/1` for the original upstream source.

The imported category provides a default definition for the `loader_file/1` predicate and derived external references that use the same vocabulary as the corresponding first-class predicates:

- `external_reference(homepage, URL)` from `homepage/1`
- `external_reference(distribution, URL)` from `distribution/1`
- `external_reference(package, Identifier)` from `package/1`
- `external_reference(repository, URL)` from `repository/1`
- `external_reference(git_object_identifier, Identifier)` from `git_object_identifier/1`
- `external_reference(software_heritage_identifier, Identifier)` from `software_heritage_identifier/1`

The following predicates are intended to help tools such as sbom export stronger standardized references:

- `package/1` Stores an application package identifier as a PURL. This is distinct from `distribution/1`, which stores a download location URL. Tools such as sbom can use `package/1` to export a package identity reference instead of only a release download location. Example: `package('pkg:generic/my_application@1.2.3')`.
- `git_object_identifier/1` Stores a standardized Git object identifier as a gitoid. This is distinct from `repository_commit/1`, which stores the raw commit hash as provenance metadata. Tools such as sbom can use `git_object_identifier/1` to export a stronger provenance reference. Example: `git_object_identifier('gitoid:commit:sha1:0123456789abcdef0123456789abcdef01234567')`.
- `software_heritage_identifier/1` Stores a Software Heritage identifier (SWHID) for an archived release or revision. Tools such as sbom can use it to export a stable archived source provenance reference when one is known. Example: `software_heritage_identifier('swh:1:rev:0123456789abcdef0123456789abcdef01234567')`.

Git-related predicates are optional source provenance facts, not reflection over the current state of a local repository checkout.

Release-oriented metadata typically includes:

- `name/1`
- `version/1`
- `description/1`
- `license/1`
- `homepage/1`
- `distribution/1`
- `package/1`
- `creators/1`
- `supplier/1`

- `originator/1`
- `built_date/1`
- `release_date/1`
- `valid_until_date/1`
- `external_reference/2`

Optional source provenance metadata typically includes:

- `repository/1`
- `repository_branch/1`
- `repository_commit/1`
- `repository_commit_abbreviated/1`
- `repository_commit_date/1`
- `repository_commit_author/1`
- `repository_commit_message/1`
- `git_object_identifier/1`
- `software_heritage_identifier/1`

6.7 `apriori_pattern_miner`

Apriori frequent itemset miner for transaction datasets. The library depends on the `frequent_pattern_mining_protocols` support library, implements the `generic_pattern_miner_protocol` defined in the `pattern_mining_protocols` core library, and mines frequent itemsets using deterministic level-wise candidate generation and anti-monotone pruning with one transaction rescan per candidate level using a candidate hash tree backed by keyed bucket dictionaries. Requires a dataset implementing `transaction_dataset_protocol` with transactions represented as canonical sorted lists of unique declared items.

6.7.1 API documentation

Open the ../apis/library_index.html#apriori_pattern_miner link in a web browser.

6.7.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(apriori_pattern_miner(loader)).
```

6.7.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(apriori_pattern_miner(tester)).
```

6.7.4 Features

- **Deterministic Level-Wise Mining:** Builds frequent itemsets level by level by generating deterministic candidate combinations, pruning candidates whose subsets are infrequent, and rescanning transactions once per level to compute support counts for all candidates using a candidate hash tree.
- **Candidate Hash Tree Counting:** Counts supports for an entire candidate level by traversing a hash tree with keyed bucket and item dictionaries instead of linearly scanning bucket lists for every transaction.
- **Library Hashing:** Uses the hashes library `fnv1a_32` object to hash candidate items instead of relying on an ad hoc local hash function.
- **Apriori Join Step:** Generates level candidates by pairwise joins of the previous frequent itemsets with shared prefixes.
- **Apriori Pruning:** Rejects candidate itemsets whose immediate subsets are not all frequent using ordered subset checks over the previous level.
- **Canonical Transactions:** Validates that transactions are sorted, duplicate-free, and restricted to declared items.
- **Flexible Support Thresholds:** Supports relative minimum support and absolute minimum support count.
- **Model Export:** Mined pattern collections can be exported as predicate clauses or written to a file.

6.7.5 Options

The `mine/3` predicate accepts the following options:

- `minimum_support/1`: Relative minimum support threshold in the interval `]0.0, 1.0]`. The default is `0.5`.
- `minimum_support_count/1`: Absolute minimum support count. When both support options are provided, this option takes precedence.
- `maximum_pattern_length/1`: Maximum itemset length to mine. The default is `1000`, which is effectively capped by the longest transaction in the dataset.
- `minimum_pattern_length/1`: Minimum itemset length retained in the mined result. The default is `1`.

6.7.6 Pattern miner representation

The mined pattern miner result is represented by a compound term with the functor chosen by the implementation and arity 3. For example:

```
apriori_pattern_miner(ItemDomain, Patterns, Options)
```

Where:

- ItemDomain: Canonical sorted list of declared dataset items.
- Patterns: List of `itemset(Items, SupportCount)` terms ordered first by pattern length and then lexicographically.
- Options: Effective mining options used to mine the frequent itemsets.

6.7.7 References

1. Agrawal, R. and Srikant, R. (1994) - “Fast algorithms for mining association rules in large databases”.

6.8 arbitrary

The arbitrary library defines an arbitrary category providing predicates for generating random values for selected types to the type object, complementing its type checking predicates. Both the object and the category predicates can be extended by the user with definitions for new types by defining clauses for multifile predicates. This library is notably used in the QuickCheck implementation by the `lgtunit` tool. See also the documentation of the mutations library for related functionality.

6.8.1 API documentation

Open the ../apis/library_index.html#arbitrary link in a web browser.

6.8.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(arbitrary(loader)).
```

6.8.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(arbitrary(tester)).
```

Several of the provided tests are generic and verify the correct behavior of all pre-defined and loaded user-defined generators and shrinkers for all ground types.

6.8.4 Pre-defined types

This library defines random generators for the most common Logtalk and Prolog types. See the [API documentation](#) for a listing of all the pre-defined types.

6.8.5 Usage

The arbitrary category complements the type object and thus its predicates are accessed via this object. For example:

```
| ?- type::arbitrary(integer, Arbitrary).
Arbitrary = -816
yes
```

6.8.6 Defining new generators and shrinkers

To define a generator of arbitrary values for a type, define a clause for the `arbitrary::arbitrary/1` multifile predicate specifying the type and a clause for the `arbitrary::arbitrary/2` multifile predicate generating an arbitrary term of the specified type. As a simple example, assume that we want to define an “odd integer type”. We start by defining both the type checker and the arbitrary generator:

```
:- multifile(type::type/1).
type::type(odd).

:- multifile(type::check/2).
type::check(odd, Term) :-
    ( var(Term) ->
        throw(instantiation_error)
    ; integer(Term),
      Term mod 2 == 1 ->
        true
    ; throw(type_error(odd, Term))
    ).

:- multifile(arbitrary::arbitrary/1).
arbitrary::arbitrary(odd).

:- multifile(arbitrary::arbitrary/2).
arbitrary::arbitrary(odd, Arbitrary) :-
    type::arbitrary(integer, Arbitrary0),
    ( Arbitrary0 mod 2 == 1 ->
        Arbitrary = Arbitrary0
    ; Arbitrary is Arbitrary0 + 1
    ).
```

The arbitrary library also provides *meta-types* that can simplify the definition of new generators. For example, the odd type above can also be defined using the `constrain/2` meta-type to only generate values that satisfy a closure:

```
arbitrary::arbitrary(odd, Arbitrary) :-
    arbitrary(constrain(integer, [Arbitrary]>>(Arbitrary mod 2 == 1)), Arbitrary).
```

Another example is using the `transform/2` meta-type to transform generated values for a base type using a closure. Assuming that we want to generate a sorted list of random integers, we can write:

```
arbitrary::arbitrary(sorted_integer_list, Arbitrary) :-
  arbitrary(transform(list(integer), sort), Arbitrary).
```

We can also define a clause for the `arbitrary::shrinker/1` multifile predicate to declare a new shrinker and a `arbitrary::shrink/3` multifile predicate for shrinking arbitrary values for QuickCheck usage:

```
:- multifile(arbitrary::shrinker/1).
arbitrary::shrinker(odd).

:- multifile(arbitrary::shrink/3).
arbitrary::shrink(odd, Large, Small) :-
  integer(Large),
  (   Large < -1 ->
      Small is Large + 2
  ;   Large > 1,
      Small is Large - 2
  ).
```

Definitions for the `shrink/3` predicate should either succeed or fail but never throw an exception. The `shrink_sequence/3` predicate can be used to help test that shrinking a value results in a finite sequence of values.

It is also possible to define edge cases for a given type for use with QuickCheck implementations. For example:

```
:- multifile(arbitrary::edge_case/2).
arbitrary::edge_case(odd, 1).
arbitrary::edge_case(odd, -1).
```

Edge cases are tried before resorting to generating arbitrary values for a type.

A more complex example is generating arbitrary values for a recursive type. A simple example of a recursive type is a binary tree. Assuming that we are working with a binary tree holding integers where each node is represented by a `node(Left, Right)` compound term, we can define a `node(Depth)` type where `Depth` is the maximum depth of the tree. This argument allows us to prevent excessively deep trees:

```
:- category(binary_tree).

:- multifile(type::type/1).
type::type(node(_)).

:- multifile(type::check/2).
type::check(node(_), Term) :-
  (   check(Term) ->
      true
  ;   var(Term) ->
      throw(instantiation_error)
  ;   throw(type_error(node(_), Term))
  ).

check(Term) :-
  (   integer(Term) ->
```

(continues on next page)

(continued from previous page)

```

    true
;   compound(Term),
    Term = node(Left, Right),
    check(Left),
    check(Right)
).

:- multifile(arbitrary::arbitrary/1).
arbitrary::arbitrary(node(_)).

:- multifile(arbitrary::arbitrary/2).
arbitrary::arbitrary(node(Depth), Arbitrary) :-
(   Depth > 1 ->
    NewDepth is Depth - 1,
    type::arbitrary(
        types_frequency([
            integer - 1,
            compound(
                node,
                [
                    types([node(NewDepth), integer]),
                    types([node(NewDepth), integer])
                ]
            ) - 3
        ]),
        Arbitrary
    )
;   type::arbitrary(
    integer, Arbitrary
).

:- end_category.

```

In this second example, we use some of the pre-defined types provided by the library. The `types_frequency(Pairs)` type supports generating random terms for a type in the Type-Frequency pairs list where the type is randomly chosen after the types relative frequency. The `compound(Name, Types)` type supports generating compound terms with a given name and random arguments after the given types:

```

| ?- type::arbitrary(node(4), Arbitrary).
Arbitrary = 907
yes

| ?- type::arbitrary(node(4), Arbitrary).
Arbitrary = node(node(node(522, 509), node(83, 453)), node(454, -197))
yes

| ?- type::arbitrary(node(4), Arbitrary).
Arbitrary = node(node(-875, -866), -254)
yes

| ?- type::arbitrary(node(4), Arbitrary).
Arbitrary = node(-133, -831)

```

(continues on next page)

(continued from previous page)

yes

The source code of these examples can be found in the `test_files` directory. Other examples of arbitrary term generators can be found in the implementation of the `optionals` and `expecteds` libraries.

6.8.7 Scoped generators and shrinkers

Declaring a new generator and possibly a shrinker for a custom type raises the possibility of a conflict with third-party-defined generators and shrinkers. An alternative is to use the `(::)/2` meta-type to define scoped generators and shrinkers. For example:

```
:- object(scoped).

% the same predicate is used for both generating and validating
:- public(custom/1).
custom(Term) :-
    (   var(Term) ->
        % assume predicate used as a generator
        random::random(Term)
      ;   % assume predicate used as a validator
        float(Term)
    ).

% a predicate with the same name is used for shrinking
:- public(custom/2).
custom(Larger, Small) :-
    Small is Larger / 2.

:- end_object.
```

Some sample calls:

```
| ?- type::arbitrary(scoped::custom, Arbitrary).
Arbitrary = 0.5788130906607927
yes

| ?- type::valid(scoped::custom, foo).
no

| ?- type::check(scoped::custom, _).
ERROR: type_error(instantiation_error)

| ?- type::check(scoped::custom, foo).
ERROR: type_error(scoped::custom, foo)

| ?- type::shrink(scoped::custom, 0.42, Smaller).
Smaller = 0.21
yes
```

The source code of this example can be found in the `test_files` directory.

6.8.8 Reproducing sequences of arbitrary terms

The arbitrary category provides access to the pseudo-random generator it uses via the `get_seed/1` and `set_seed/1`. This allows sequences of arbitrary values to be reproduced. For example:

```
| ?- type::get_seed(Seed).
Seed = seed(3172, 9814, 20125)
yes

| ?- type::arbitrary(integer, Arbitrary).
Arbitrary = -816
yes

| ?- type::arbitrary(integer, Arbitrary).
Arbitrary = -113
yes

| ?- type::arbitrary(integer, Arbitrary).
Arbitrary = 446

| ?- type::set_seed(seed(3172, 9814, 20125)).
yes

| ?- type::arbitrary(integer, Arbitrary).
Arbitrary = -816
yes

| ?- type::arbitrary(integer, Arbitrary).
Arbitrary = -113
yes

| ?- type::arbitrary(integer, Arbitrary).
Arbitrary = 446
yes
```

The seed should be regarded as an opaque term and handled using the `get_seed/1` and `set_seed/1` predicates. These predicates are notably used in the QuickCheck implementation provided by the `lgtunit` tool.

6.8.9 Default size of generated terms

The library uses the value 42 for the default size of generated terms for types where size is meaningful and implicit. To override this default value, define a clause for the `arbitrary::max_size/1` multifile predicate. The new default size must be a positive integer. For example:

```
:- multifile(arbitrary::max_size/1).
arbitrary::max_size(7).
```

When multiple definitions exist, the first valid one found is used. When no definition is valid, the default value of 42 is used.

6.8.10 Known issues

Some Prolog systems either don't support the null character or provide buggy results when calling `char_code/2` with a code of zero. When that's the case, the null character is excluded when generating arbitrary characters or character codes.

Generating arbitrary Unicode characters (instead of Unicode codepoints) is inherently problematic as the process first generates codepoints and then tries to use the standard `char_code/2` to convert them to characters. But, depending on the backend Prolog system and its internal (if any) Unicode normalization, it may not be possible to convert a codepoint to a single character.

6.9 arrangements

This library provides predicates for generating and querying arrangements over lists. Arrangements are ordered selections of a given length with element repetition allowed. The following categories of predicates are provided:

- **Generation operations** - Predicates for generating arrangements.
- **Ordering variants** - Predicates that support an additional order argument (default or lexicographic) for controlling output order.
- **Distinct-value generation** - Predicates for generating arrangements while deduplicating equal-valued results.
- **Indexed access** - Predicates for direct access to arrangements at specific positions, including distinct arrangements.
- **Lexicographic stepping** - Predicates for navigating arrangements in lexicographic order.
- **Counting operations** - Predicates for counting arrangements and distinct arrangements.
- **Random selection** - Predicates for randomly selecting and sampling arrangements and distinct arrangements.

Dedicated `cartesian_products`, `permutations`, `combinations`, `multisets`, `derangements`, `partitions`, and `subsequences` libraries are also available for focused APIs on related operations.

6.9.1 API documentation

Open the ../apis/library_index.html#arrangements link in a web browser.

6.9.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(arrangements(loader)).
```

6.9.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(arrangements(tester)).
```

6.10 assignvars

The `assignvarsp` protocol declares the predicates used for logical assignment of Prolog terms developed by Nobukuni Kino.

The `assignvars` object provides a declarative implementation of the `assignvarsp` protocol. It can be used with any backend Prolog compiler.

The `nd_assignvars` object provides a non-declarative but faster implementation of the `assignvarsp` protocol. It can be used with the following backend Prolog compilers: B-Prolog, CxProlog, ECLiPSe, GNU Prolog, SICStus Prolog, SWI-Prolog, and YAP.

For more information on `assignvars`, please consult the URL:

https://web.archive.org/web/20160818050049/http://www.kprolog.com/en/logical_assignment/

The representation of assignable terms should be regarded as an opaque term and only accessed using the library predicates.

6.10.1 API documentation

Open the `../..apis/library_index.html#assignvars` link in a web browser.

6.10.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(assignvars(loader)).
```

6.10.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(assignvars(tester)).
```

6.11 avro

The `avro` library implements predicates for reading (parsing) and writing (generating) data in the Apache Avro binary format:

- <https://avro.apache.org/docs/current/specification/>

This library requires a backend supporting unbounded integer arithmetic.

6.11.1 API documentation

Open the `../..apis/library_index.html#avro` link in a web browser.

6.11.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(avro(loader)).
```

6.11.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(avro(tester)).
```

6.11.4 Schema representation

Schema files use the `.avsc` extension and their contents are always valid JSON values. They can be parsed using the `json(curl,dash,atom)` library object.

6.11.5 Data representation

Data to be serialized to Avro binary format files can be represented in the JSON Lines format (using the `.jsonl` extension) and parsed using the `json_lines(curl, dash, atom)` library object.

6.11.6 Schema Examples

Primitive type schemas:

- `null`
- `boolean`
- `int`
- `long`
- `float`
- `double`
- `bytes`
- `string`

Array schema example represented as a JSON term:

```
{type-array, items-int}
```

6.11.7 Encoding

Encoding is accomplished using the `generate/3` or `generate/4` predicates. For example, assuming the schema is just `int`:

```
| ?- avro::generate(bytes(Bytes), int, 42).
Bytes = [84]
yes
```

Or an array of `int`:

```
| ?- avro::generate(bytes(Bytes), {type-array,items-int}, [42,37,13,17]).
Bytes = [8,84,74,26,34,0]
yes
```

To include the schema in the output (as an Avro Object Container File), use the `generate/4` predicate with the second argument set to `true`. For example:

```
| ?- avro::generate(file('output.avro'), true, {type-array,items-int}, [42,37,13,17]).
yes
```

6.11.8 Decoding

Decoding is accomplished using the `parse/2` or `parse/3` predicates.

When parsing a file that includes a schema (Avro Object Container File), use `parse/2` which returns a Schema-Data pair. For example:

```
| ?- avro::parse(file('input.avro'), Schema-Data).
```

When the schema is not present in the file, Schema is unified with `false`.

When parsing with a known schema, use instead the `parse/3` predicate. For example:

```
| ?- avro::parse(bytes([84]), int, Data).
Data = 42
yes
```

6.12 base32

The `base32` library provides predicates for encoding and decoding data in the Base32 format as per the specification found at:

<https://www.rfc-editor.org/rfc/rfc4648>

6.12.1 API documentation

Open the `../apis/library_index.html#base32` link in a web browser.

6.12.2 Loading

To load all entities in this library, load the `loader.lgt` utility file:

```
| ?- logtalk_load(base32(loader)).
```

6.12.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(base32(tester)).
```

6.12.4 Encoding

Encoding a list of bytes in Base32 format is accomplished by the `base32::generate/2` predicate. For example:

```
| ?- atom_codes('Hello!', Bytes),
    base32::generate(atom(Base32), Bytes).
Base32 = 'JBSWY3DP EE====='
Bytes = [72,101,108,108,111,33]
yes

| ?- atom_codes('Hello!', Bytes),
    base32::generate(codes(Base32), Bytes).
Base32 = [74,66,83,87,89,51,68,80,69,69,61,61,61,61,61]
Bytes = [72,101,108,108,111,33]
yes
```

The Base32 result can also be represented using a list of chars, written to a file or to a stream. See the API documentation for details.

6.12.5 Decoding

Decoding of Base32 data is accomplished using the `base32::parse/2` predicate. For example:

```
| ?- base32::parse(atom('JBSWY3DP EE====='), Bytes),
    atom_codes(Atom, Bytes).
Atom = 'Hello!'
Bytes = [72,101,108,108,111,33]
yes

| ?- base32::parse(chars(['J','B','S','W','Y','3','D','P','E','E','=','','=','','=','','=','=']),
    ↪Bytes),
    atom_codes(Atom, Bytes).
Atom = 'Hello!'
```

(continues on next page)

(continued from previous page)

```
Bytes = [72,101,108,108,111,33]
yes
```

The `base32::parse/2` predicate accepts other input sources such as a file or a stream. See the API documentation for details. Both uppercase and lowercase letters are accepted when decoding.

6.13 base58

The `base58` library provides predicates for encoding and decoding data in the Base58 format using the Bitcoin alphabet variant. Base58 is commonly used in Bitcoin addresses and other cryptocurrency applications. The Bitcoin alphabet excludes visually ambiguous characters:

- 0 (zero), O (uppercase o)
- I (uppercase i), l (lowercase L)

Alphabet: 123456789ABCDEFGHJKLMNPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz

For more details, see for example:

<https://bitcoinwiki.org/wiki/base58>

This library requires a backend supporting unbounded integer arithmetic.

6.13.1 API documentation

Open the `../apis/library_index.html#base58` link in a web browser.

6.13.2 Loading

To load all entities in this library, load the `loader.lgt` utility file:

```
| ?- logtalk_load(base58(loader)).
```

6.13.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(base58(tester)).
```

6.13.4 Encoding

Encoding a list of bytes in Base58 format is accomplished by the `base58::generate/2` predicate. For example:

```
| ?- atom_codes('Hello World', Bytes),
    base58::generate(atom(Base58), Bytes).
Base58 = 'JxF12TrwUP45BMd'
Bytes = [72,101,108,108,111,32,87,111,114,108,100]
yes
```

Leading zero bytes are preserved and encoded as '1' characters:

```
| ?- base58::generate(atom(Base58), [0, 0, 0, 1, 2, 3]).
Base58 = '111Ldp'
yes
```

6.13.5 Decoding

Decoding of Base58 data is accomplished using the `base58::parse/2` predicate. For example:

```
| ?- base58::parse(atom('JxF12TrwUP45Bmd'), Bytes),
    atom_codes(Atom, Bytes).
Atom = 'Hello World'
Bytes = [72,101,108,108,111,32,87,111,114,108,100]
yes
```

6.14 base64

The `base64` library provides predicates for parsing and generating data in the Base64 and Base64URL formats as per the specification found at:

<https://www.rfc-editor.org/rfc/rfc4648>

6.14.1 API documentation

Open the `../..apis/library_index.html#base64` link in a web browser.

6.14.2 Loading

To load all entities in this library, load the `loader.lgt` utility file:

```
| ?- logtalk_load(base64(loader)).
```

6.14.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(base64(tester)).
```

6.14.4 Encoding

Encoding a list of bytes in Base64 format is accomplished by the `base64::generate/2` predicate. For example:

```
| ?- atom_codes('Hello world!', Bytes),
    base64::generate(atom(Base64), Bytes).
Base64 = 'SGVsbG8gd29ybGQh'
Bytes = [72,101,108,108,111,32,119,111,114,108,100,33]
yes

| ?- atom_codes('Hello world!', Bytes),
    base64::generate(codes(Base64), Bytes).
Base64 = [83,71,86,115,98,71,56,103,100,50,57,121,98,71,81,104]
Bytes = [72,101,108,108,111,32,119,111,114,108,100,33]
yes
```

The Base64 result can also be represented using a list of chars, written to a file or to a stream. See the API documentation for details.

For safe encoding of URLs, use instead the Base64URL format. For example:

```
| ?- base64url::generate(atom(Base64URL), 'https://logtalk.org').
Base64URL == 'aHR0cHM6Ly9sb2d0YWxrLm9yZW'
yes
```

The Base64URL can also be represented using a list of chars or a list of codes. The input URL should be in the same format.

6.14.5 Decoding

Decoding of Base64 data is accomplished using the `base64::parse/2` predicate. For example:

```
| ?- base64::parse(atom('SGVsbG8gd29ybGQh'), Bytes),
    atom_codes(Atom, Bytes).
Atom = 'Hello world!'
Bytes = [72,101,108,108,111,32,119,111,114,108,100,33]
yes

| ?- base64::parse(chars(['S','G','V',s,b,'G','8',g,d,'2','9',y,b,'G','Q',h]), Bytes),
    atom_codes(Atom, Bytes).
Atom = 'Hello world!'
Bytes = [72,101,108,108,111,32,119,111,114,108,100,33]
yes
```

The `base64::parse/2` predicate accepts other input source such as a file or a stream. See the API documentation for details.

For decoding of URLs in the Base64URL format, use the `base64url::parse/2` predicate. For example:

```
| ?- base64url::parse(atom('aHR0cHM6Ly9sb2d0YWxrLm9yZW'), URL).
URL == 'https://logtalk.org'
yes
```

The `base64url::parse/2` predicate also accepts a list of chars or a list of codes as input. See the API documentation for details.

6.15 base85

The base85 library provides predicates for encoding and decoding data in the Base85 (Ascii85) format. Ascii85 is used in PostScript and PDF files. The encoding uses printable ASCII characters from ! (33) to u (117). A special shortcut z represents four zero bytes. For more details, see for example:

<https://en.wikipedia.org/wiki/Ascii85>

6.15.1 API documentation

Open the `../..apis/library_index.html#base85` link in a web browser.

6.15.2 Loading

To load all entities in this library, load the `loader.lgt` utility file:

```
| ?- logtalk_load(base85(loader)).
```

6.15.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(base85(tester)).
```

6.15.4 Encoding

Encoding a list of bytes in Base85 format is accomplished by the `base85::generate/2` predicate. For example:

```
| ?- atom_codes('Man ', Bytes),  
    base85::generate(atom(Base85), Bytes).  
Base85 = '9jqo^'  
Bytes = [77,97,110,32]  
yes
```

Four zero bytes are encoded as the special character z:

```
| ?- base85::generate(atom(Base85), [0, 0, 0, 0]).  
Base85 = 'z'  
yes
```

6.15.5 Decoding

Decoding of Base85 data is accomplished using the `base85::parse/2` predicate. For example:

```
| ?- base85::parse(atom('9jqo^'), Bytes),
    atom_codes Atom, Bytes).
Atom = 'Man '
Bytes = [77,97,110,32]
yes
```

The parser also accepts input with optional `<~` and `~>` delimiters:

```
| ?- base85::parse(atom('<~9jqo^~>'), Bytes),
    atom_codes Atom, Bytes).
Atom = 'Man '
Bytes = [77,97,110,32]
yes
```

6.16 basic_types

The `basic_types` library is a virtual library that loads only basic types from the `types` library:

- `comparingp`
- `term`, `term`
- `atomic`, `atom`, `number`, `float`, `integer`
- `compound`, `listp`, `list`
- `type`

6.16.1 API documentation

Open the ../apis/library_index.html#types link in a web browser.

6.16.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(basic_types(loader)).
```

6.16.3 Testing

To test this library predicates, load the `tester.lgt` file for the `types` library:

```
| ?- logtalk_load(types(tester)).
```

6.17 bayesian_ridge_regression

Bayesian ridge regression regressor supporting continuous and mixed-feature datasets. The library implements the `regressor_protocol` defined in the `regression_protocols` library and learns a Bayesian linear model using evidence maximization for the global weight and noise precisions together with Gamma hyperpriors over both precision terms.

6.17.1 API documentation

Open the ../apis/library_index.html#bayesian_ridge_regression link in a web browser.

6.17.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(bayesian_ridge_regression(loader)).
```

6.17.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(bayesian_ridge_regression(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(bayesian_ridge_regression(tester_performance)).
```

6.17.4 Features

- **Continuous and Mixed Features:** Supports continuous attributes and categorical attributes encoded using reference-level dummy coding from the declared dataset attribute values.
- **Automatic Hyperparameter Tuning:** Learns the global coefficient precision and observation-noise precision using MacKay-style evidence maximization with configurable Gamma hyperpriors instead of a user-supplied ridge penalty.
- **Posterior Uncertainty:** Exposes predictive Gaussian distributions using coefficient posterior uncertainty plus observation noise, matching the usual scikit-learn `BayesianRidge` treatment where the intercept is not probabilistic. Posterior coefficient variances are also exposed.
- **Feature Scaling:** Continuous attributes can be standardized using z-score scaling before fitting and prediction.
- **Stable Posterior Solves:** Evidence-maximization updates clamp the learned weight and noise precisions to configurable `precision_bounds`(Min, Max) to avoid degenerate zero or infinite precision estimates. Posterior solves use Cholesky factorization of positive-definite precision matrices, diagnostics report any diagonal jitter applied when factorization retries are needed, and the evidence-maximization loop computes the effective degrees of freedom from a one-time eigenspectrum of the centered Gram surrogate while still switching to a sample-space solve when the active encoded feature count exceeds the number of training rows.

- **Missing Values:** Missing numeric and categorical values represented using anonymous variables are encoded using explicit missing-value indicator features.
- **Unknown Values:** Prediction requests containing categorical values that are not declared by the dataset raise a domain error.
- **Zero-Variance Features:** Encoded columns with zero variance are excluded from posterior updates and assigned zero mean and zero posterior variance.
- **Diagnostics Metadata:** Learned regressors record model name, target, training example count, Cholesky stabilization attempts and applied jitter, Gamma hyperpriors for both precision terms, effective precision bounds, learned precisions, learned noise variance, final log evidence, the full log-evidence score trace, active feature count, posterior variances, intercept treatment, convergence metric and status, encoded feature count, and effective options.
- **Model Export:** Learned regressors can be exported as predicate clauses or written to a file.

6.17.5 Regressor representation

The learned regressor is represented by default as:

- `bayesian_ridge_regressor(Encoders, Bias, Weights, ActiveFlags, PosteriorCovariance, NoiseVariance, Diagnostics)`

The exported predicate clauses therefore use the shape:

- `Functor(Encoders, Bias, Weights, ActiveFlags, PosteriorCovariance, NoiseVariance, Diagnostics)`

6.17.6 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(bayesian_ridge_regression),
  target(Target),
  training_example_count(TrainingExampleCount),
  options(Options),
  solver(cholesky_factorization),
  stabilization_attempts(StabilizationAttempts),
  stabilization_jitter(StabilizationJitter),
  precision_bounds(MinimumPrecision, MaximumPrecision),
  weight_precision_hyperprior(gamma(LambdaShape, LambdaRate)),
  noise_precision_hyperprior(gamma(AlphaShape, AlphaRate)),
  weight_precision(Alpha),
  noise_precision(Beta),
  noise_variance(NoiseVariance),
  log_evidence(LogEvidence),
  scores(Scores),
  active_feature_count(ActiveFeatureCount),
  weight_prior(isotropic_zero_mean_gaussian),
  intercept_treatment(non_probabilistic),
  bias_variance(BiasVariance),
  weight_variances(WeightVariances),
  convergence_metric(coefficient_l1),
```

(continues on next page)

(continued from previous page)

```
convergence(Convergence),  
iterations(Iterations),  
final_delta(FinalDelta),  
encoded_feature_count(FeatureCount)  
]
```

The `scores/1` diagnostic is analogous to `scikit-learn scores_`: it stores the log marginal likelihood at the initial hyperparameters followed by the value after each evidence-maximization update. The final element is identical to `log_evidence/1`.

The `bias_variance/1` diagnostic is always `0.0` because the intercept is treated as a deterministic centering adjustment rather than as a probabilistic parameter.

Use the `regression_protocols` `diagnostic/2` and `regressor_options/2` helper predicates when you only need a single metadata term or the effective options.

6.17.7 Options

The `learn/3` predicate accepts the following options:

- `maximum_iterations/1`: Maximum number of evidence-maximization updates. The default is `300`.
- `tolerance/1`: Convergence tolerance on the L1 change between consecutive active coefficient vectors across evidence-maximization updates. The default is `1.0e-6`.
- `initial_weight_precision/1`: Positive initial value for the shared coefficient precision. The default is `1.0`.
- `initial_noise_precision/1`: Positive initial value for the observation-noise precision or `auto` to derive it from the target variance. The default is `auto`.
- `alpha_1/1`: Non-negative shape hyperparameter of the Gamma prior over the learned observation-noise precision. The default is `1.0e-6`.
- `alpha_2/1`: Non-negative rate hyperparameter of the Gamma prior over the learned observation-noise precision. The default is `1.0e-6`.
- `lambda_1/1`: Non-negative shape hyperparameter of the Gamma prior over the learned coefficient precision. The default is `1.0e-6`.
- `lambda_2/1`: Non-negative rate hyperparameter of the Gamma prior over the learned coefficient precision. The default is `1.0e-6`.
- `feature_scaling/1`: Controls z-score standardization of continuous attributes before training and prediction. Accepted values are `true` and `false`. The default is `true`.
- `precision_bounds/2`: Lower and upper positive bounds used to clamp the learned weight and noise precisions during evidence maximization for numerical stability. The default is `precision_bounds(1.0e-12, 1.0e12)`.

6.18 borda_ranker

Borda grouped-ranking ranker. Ranks each item by summing, across groups, the number of same-group items with strictly lower relevance.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from grouped relevance judgments, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The training dataset must declare each group once, use only declared groups and items in relevance judgments, and assign non-negative integer relevance values.

6.18.1 API documentation

Open the ../apis/library_index.html#borda_ranker link in a web browser.

6.18.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(borda_ranker(loader)).
```

6.18.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(borda_ranker(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(borda_ranker(tester_performance)).
```

6.18.4 Features

- **Grouped Relevance Learning:** Learns one deterministic item score from grouped ranking or relevance-judgment datasets.
- **Portable Borda Scoring:** Computes scores using only non-negative integer grouped relevance judgments and standard Logtalk library predicates. Within each group, an item receives one point for every same-group item with strictly lower relevance when using `tie_scoring(standard)` and the average of the minimum and maximum tied positions when using `tie_scoring(fractional)`.
- **Deterministic Ranking:** Orders candidate items by learned score with deterministic tie-breaking. Ranking ties are broken deterministically using the standard term order of the item identifiers after sorting by descending score.
- **Missing relevance semantics:** Missing relevance facts are treated as zero by default using the `missing_relevance(zero)` option and can be rejected using `missing_relevance(error)`.
- **Strict Dataset Validation:** Rejects duplicate groups, duplicate items within a group, undeclared groups or items in relevance judgments, and non-integer or negative relevance values.

- **Explicit Semantics Options:** The `learn/3` predicate exposes the current tie and missing-relevance policies using the `tie_scoring/1` and `missing_relevance/1` options.
- **Benchmark Coverage:** Includes a dedicated performance suite for a large grouped dataset benchmark.
- **Training Diagnostics:** Learned rankers include dataset summary metadata that can be accessed using the `diagnostics/2` predicate.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.
- **Shared Ranking Infrastructure:** Uses the common `ranking_protocols` helper predicates for option processing, dataset validation, diagnostics, export, and candidate ranking.

6.18.5 Scoring semantics

This implementation uses a grouped Borda count variant over the declared items of each group. With the default `tie_scoring(standard)` option, an item receives one point for every same-group item with strictly lower relevance. Tied items therefore receive the same per-group contribution, because equal relevance values do not add or subtract points.

With the `tie_scoring(fractional)` option, each tied relevance class receives the average of the minimum and maximum per-group Borda points available to that tie block. For example, when two items tie above a single lower-ranked item, both tied items receive 1.5 points instead of the 1 point assigned by the default policy.

Missing relevance facts are treated as relevance 0 only for items that are declared in the group when using the default `missing_relevance(zero)` option. This allows grouped datasets to omit explicit zero judgments while keeping the score computation deterministic. Use `missing_relevance(error)` to reject grouped datasets that omit a declared item relevance.

6.18.6 Usage

Learning a ranker

```
% Learn from a grouped ranking dataset object
| ?- borda_ranker::learn(my_dataset, Ranker).
...

% Learn with an explicit empty options list
| ?- borda_ranker::learn(my_dataset, Ranker, []).
...

% Learn while requiring every declared group item to have a relevance fact
| ?- borda_ranker::learn(my_dataset, Ranker, [missing_relevance(error)]).
...

% Learn using fractional tie scoring for tied relevance levels
| ?- borda_ranker::learn(my_dataset, Ranker, [tie_scoring(fractional)]).
...
```

The current implementation accepts the `missing_relevance/1` and `tie_scoring/1` options described below.

Inspecting diagnostics

```
% Inspect model and dataset summary metadata
| ?- borda_ranker::learn(my_dataset, Ranker),
    borda_ranker::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
...
```

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- borda_ranker::learn(my_dataset, Ranker),
    borda_ranker::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
...
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- borda_ranker::learn(my_dataset, Ranker),
    borda_ranker::export_to_clauses(my_dataset, Ranker, my_ranker, Clauses).
Clauses = [my_ranker(borda_ranker(...))]
...

% Export to a file
| ?- borda_ranker::learn(my_dataset, Ranker),
    borda_ranker::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.pl').
...
```

6.18.7 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(borda_ranker),
  options(Options),
  dataset_summary(DatasetSummary)
]
```

Where:

- `model(borda_ranker)` identifies the learning algorithm that produced the ranker.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.

- `dataset_summary(DatasetSummary)` stores a summary list describing the validated training dataset.

The current `dataset_summary/1` payload has the form:

```
[
  groups(NumberOfGroups),
  items(NumberOfItems),
  relevance_judgments(NumberOfJudgments)
]
```

Use the `ranking_protocols` `diagnostic/2` and `ranker_options/2` helper predicates when you only need a single metadata term or the effective options.

6.18.8 Options

The following options can be passed to the `learn/3` predicate:

- `missing_relevance(Policy)`: Controls how declared group items without an explicit relevance fact are handled. The supported values are `zero` (default) and `error`.
- `tie_scoring(Policy)`: Controls the grouped Borda tie semantics. The current implementation supports `standard` (minimum tied-block score) and `fractional` (average tied-block score).

6.18.9 Ranker representation

The learned ranker is represented by a compound term of the form:

```
borda_ranker(Items, Scores, Diagnostics)
```

Where:

- `Items`: List of ranked items.
- `Scores`: List of Item-Score pairs.
- `Diagnostics`: List of metadata terms, including the effective options and dataset summary.

When exported using `export_to_clauses/4` or `export_to_file/4`, this ranker term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.18.10 References

1. de Borda, J.-C. (1781). Mémoire sur les élections au scrutin. *Histoire de l'Académie Royale des Sciences*.

6.19 bradley_terry_ranker

Bradley-Terry pairwise preference ranker. Uses a deterministic minorization-maximization update to estimate one relative strength parameter per item from weighted pairwise wins and losses.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from pairwise preferences, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The training dataset must declare each ranked item once, enumerate positive-weight pairwise preferences between distinct declared items, induce a connected undirected comparison graph, and induce a strongly connected directed win graph so that a finite Bradley-Terry maximum-likelihood estimate exists.

6.19.1 API documentation

Open the ../apis/library_index.html#bradley_terry_ranker link in a web browser.

6.19.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(bradley_terry_ranker(loader)).
```

6.19.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(bradley_terry_ranker(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(bradley_terry_ranker(tester_performance)).
```

6.19.4 Features

- **Pairwise Preference Learning:** Learns relative item strengths from weighted head-to-head outcomes.
- **Original MM Fidelity:** Rejects datasets without a finite Bradley-Terry maximum-likelihood estimate instead of masking them with implicit strength flooring or other hidden regularization.
- **Deterministic Ranking:** Orders candidate items by learned strength with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items, self-preferences, non-positive weights, disconnected comparison graphs, and pairwise datasets that do not admit a finite Bradley-Terry maximum-likelihood estimate.
- **Training Diagnostics:** Learned rankers include convergence and dataset summary metadata that can be accessed using the `diagnostics/2` predicate.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.

- **Sparse Comparison Processing:** Training aggregates weighted comparisons into sparse adjacency lists so iteration cost scales with observed pairwise comparisons instead of the full dense item cross-product.

6.19.5 Dataset requirements

This implementation requires more than undirected connectedness. In order to admit a finite Bradley-Terry maximum-likelihood estimate, the directed win graph induced by the preferences must be strongly connected. Datasets that leave one or more items isolated, split the undirected comparison graph into multiple components, or create dominance partitions with no directed path back to stronger items are rejected instead of producing degenerate strengths.

For a related MAP formulation that keeps the same pairwise-preference setting but uses an explicit Gamma prior to admit connected datasets whose directed win graph is not strongly connected, see the `regularized_bradley_terry_ranker` library.

6.19.6 Usage

Learning a ranker

```
% Learn from a pairwise ranking dataset object
| ?- bradley_terry_ranker::learn(my_dataset, Ranker).
...

% Learn with custom iteration and convergence options
| ?- bradley_terry_ranker::learn(my_dataset, Ranker, [maximum_iterations(500), tolerance(1.
↪0e-7)]).
...
```

Inspecting diagnostics

```
% Inspect convergence and dataset summary metadata
| ?- bradley_terry_ranker::learn(my_dataset, Ranker),
    bradley_terry_ranker::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
...
```

6.19.7 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(bradley_terry_ranker),
  options(Options),
  convergence(Status),
  iterations(Iterations),
  final_delta(FinalDelta),
  dataset_summary(DatasetSummary)
]
```

Where:

- `model(bradley_terra_ranker)` identifies the learning algorithm that produced the ranker.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `convergence(Status)` records the training stop condition. The current values are converged and `maximum_iterations_exhausted`.
- `iterations(Iterations)` stores the number of update iterations that were executed.
- `final_delta(FinalDelta)` stores the maximum absolute strength update in the last iteration.
- `dataset_summary(DatasetSummary)` stores a summary list describing the validated training dataset.

The current `dataset_summary/1` payload has the form:

```
[
  items(NumberOfItems),
  preferences(NumberOfPreferences),
  connected_components(NumberOfComponents),
  isolated_items(IsolatedItems)
]
```

Where `IsolatedItems` is the list of declared items that have no comparisons. For valid Bradley-Terry training datasets this list is expected to be empty, because disconnected datasets are rejected.

For example, learning from the `head_to_head` test dataset currently returns diagnostics with the structure:

```
[
  model(bradley_terra_ranker),
  options([maximum_iterations(5000), tolerance(1.0e-6)]),
  convergence(converged),
  iterations(...),
  final_delta(...),
  dataset_summary([
    items(4),
    preferences(6),
    connected_components(1),
    isolated_items([])
  ])
]
```

Use the `ranking_protocols diagnostic/2` and `ranker_options/2` helper predicates when you only need a single metadata term or the effective options.

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- bradley_terra_ranker::learn(my_dataset, Ranker),
    bradley_terra_ranker::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
...
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- bradley_terry_ranker::learn(my_dataset, Ranker),
    bradley_terry_ranker::export_to_clauses(my_dataset, Ranker, my_ranker, Clauses).
Clauses = [my_ranker(bt_ranker(...))]
...

% Export to a file
| ?- bradley_terry_ranker::learn(my_dataset, Ranker),
    bradley_terry_ranker::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.pl').
...
```

6.19.8 Options

The following options can be passed to the learn/3 predicate:

- maximum_iterations(MaximumIterations): Positive integer iteration bound.
- tolerance(Tolerance): Positive convergence tolerance.

6.19.9 Ranker representation

The learned ranker is represented by a compound term of the form:

```
bt_ranker(Items, Strengths, Diagnostics)
```

Where:

- Items: List of ranked items.
- Scores: List of Item-Strength pairs.
- Diagnostics: List of metadata terms, including the effective options, convergence status, iteration count, final update delta, and dataset summary.

When exported using export_to_clauses/4 or export_to_file/4, this ranker term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.19.10 References

1. Bradley, R. A. and Terry, M. E. (1952). Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, 39(3/4), 324-345.
2. Hunter, D. R. (2004). MM algorithms for generalized Bradley-Terry models. *The Annals of Statistics*, 32(1), 384-406.

6.20 c45_classifier

This library implements the C4.5 decision tree learning algorithm. C4.5 is an extension of the ID3 algorithm that uses information gain ratio instead of information gain for attribute selection, which avoids bias towards attributes with many values (see below for implementation details).

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning a decision tree from a dataset, optionally prune it, using it to make predictions, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. See `test_files` directory for examples.

6.20.1 API documentation

Open the [../apis/library_index.html#c45_classifier](http://logtalk.org/apic/library_index.html#c45_classifier) link in a web browser.

6.20.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(c45_classifier(loader)).
```

6.20.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(c45_classifier(tester)).
```

6.20.4 Implemented features

- Information gain ratio for attribute selection (avoids bias towards attributes with many values)
- Handling of discrete (categorical) attributes with multi-way splits
- Handling of continuous (numeric) attributes with binary threshold splits (selects the threshold with the highest gain ratio from midpoints between consecutive sorted values)
- Handling of missing attribute values (represented using anonymous variables): examples with missing values for an attribute are distributed to all branches during tree construction; gain ratio computation uses only examples with known values for the attribute being evaluated; prediction with missing values uses majority voting by exploring all possible branches and selecting the most common class
- Tree pruning using pessimistic error pruning (PEP): estimates error rates using the upper confidence bound of the binomial distribution (Wilson score interval) with a configurable confidence factor (default 0.25) and minimum instances per leaf (default 2); replaces subtrees with leaf nodes when doing so would not increase the estimated error; helps reduce overfitting and improve generalization
- Export of learned decision trees as predicate clauses
- Pretty-printing of learned decision trees

6.20.5 Learned tree representation

For discrete attributes, the learned decision tree is represented as `leaf(Class)` for leaf nodes and `tree(Attribute, Subtrees)` for internal nodes with discrete attributes, where `Subtrees` is a list of `Value-Subtree` pairs.

For continuous (numeric) attributes, the tree uses binary threshold splits represented as `tree(Attribute, threshold(Threshold), LeftSubtree, RightSubtree)` where `LeftSubtree` corresponds to values \leq `Threshold` and `RightSubtree` to values $>$ `Threshold`.

6.20.6 Limitations

- No incremental learning (the tree must be rebuilt from scratch when new examples are added)

6.20.7 References

- Quinlan, J.R. (1986). Induction of Decision Trees. *Machine Learning*, 1(1), 81-106. <https://doi.org/10.1007/BF00116251>
- Quinlan, J.R. (1993). *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers. <https://doi.org/10.1016/C2009-0-27846-9>
- Mitchell, T.M. (1997). *Machine Learning*. McGraw-Hill. Chapter 3: Decision Tree Learning.

6.20.8 Usage

To learn a decision tree from a dataset:

```
| ?- c45_classifier::learn(play_tennis, Tree).
```

To prune a learned tree with the default parameters (confidence factor 0.25, minimum instances per leaf 2):

```
| ?- c45_classifier::learn(breast_cancer, Tree),  
    c45_classifier::prune(breast_cancer, Tree, PrunedTree).
```

To prune a learned tree with both custom confidence factor and minimum instances per leaf:

```
| ?- c45_classifier::learn(breast_cancer, Tree),  
    c45_classifier::prune(breast_cancer, Tree, 0.1, 3, PrunedTree).
```

To export the tree as a list of predicate clauses:

```
| ?- c45_classifier::learn(play_tennis, Tree),  
    c45_classifier::tree_to_clauses(play_tennis, Tree, classify, Clauses).
```

To export the tree to a file:

```
| ?- c45_classifier::learn(play_tennis, Tree),  
    c45_classifier::tree_to_file(play_tennis, Tree, classify, 'tree.pl').
```

To print the tree to the current output:

```
| ?- c45_classifier::learn(play_tennis, Tree),  
    c45_classifier::print_tree(Tree).
```

To predict the class for a new instance (as a list of attribute-value pairs):

```
| ?- c45_classifier::learn(play_tennis, Tree),
    c45_classifier::predict(Tree, [outlook-sunny, temperature-hot, humidity-high, wind-
    ↪weak], Class).
```

6.21 cartesian_products

This library provides predicates for generating and querying Cartesian products over lists. The following categories of predicates are provided:

- **Generation operations** - Predicates for generating Cartesian products and tuples.
- **Ordering variants** - Predicates that support an additional order argument (default or lexicographic) for controlling output order.
- **Distinct-value generation** - Predicates for generating tuples while deduplicating repeated values in factor lists.
- **Indexed access** - Predicates for direct access to tuples at specific positions.
- **Lexicographic stepping** - Predicates for navigating tuples in lexicographic order.
- **Counting operations** - Predicates for counting tuples without generating them.
- **Random selection** - Predicates for randomly selecting and sampling tuples.

Dedicated arrangements, permutations, combinations, multisets, derangements, partitions, and subsequences libraries are also available for focused APIs on related operations.

6.21.1 API documentation

Open the ../apis/library_index.html#cartesian_products link in a web browser.

6.21.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(cartesian_products(loader)).
```

6.21.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(cartesian_products(tester)).
```

6.22 cbor

The cbor library implements predicates for importing and exporting data in the Concise Binary Object Representation (CBOR) format:

- <https://www.rfc-editor.org/rfc/rfc8949.html>
- <http://cbor.io/>

This library is a work-in-progress. Currently it requires a backend supporting unbounded integer arithmetic.

6.22.1 Representation

- Maps are represented using curly-bracketed terms, {Pairs}, where each pair uses the representation Key-Value.
- Arrays are represented using lists.
- Byte strings uses bytes(List) compound terms.
- Text strings can be represented as atoms, chars(List), or codes(List). The default when decoding is to use atoms when using the cbor object. To decode text strings into lists of chars or code, use the cbor/1 with the parameter bound to chars or codes. For example:

```
| ?- cbor::parse([0x65,0x68,0x65,0x6c,0x6c,0x6f], Term).  
Term = hello  
yes  
  
| ?- cbor(atom)::parse([0x65,0x68,0x65,0x6c,0x6c,0x6f], Term).  
Term = hello  
yes  
  
| ?- cbor(chars)::parse([0x65,0x68,0x65,0x6c,0x6c,0x6f], Term).  
Term = chars([h,e,l,l,o])  
yes  
  
| ?- cbor(codes)::parse([0x65,0x68,0x65,0x6c,0x6c,0x6f], Term).  
Term = codes([104,101,108,108,111])  
yes
```

- Tagged data uses tag(Tag, Data) compound terms.
- Simple values can be represented using simple(Simple) compound terms.
- The CBOR elements false, true, null, and undefined are represented by, respectively, the @false, @true, @null, and @undefined compound terms.
- The compound terms @infinity, @negative_infinity, and @not_a_number are used to represent the corresponding CBOR elements.
- Only some backends distinguish between positive zero and negative zero. The compound terms @zero and @negative_zero can be used as an alternative for encoding. The decoder, however, produces the 0.0 and -0.0 floats.

6.22.2 Encoding

Encoding is accomplished using the `generate/2` predicate. For example:

```
| ?- cbor::generate([a,{b-c}], Encoding).
Encoding = [0x9f,0x61,0x61,0xbf,0x61,0x62,0x61,0x63,0xff,0xff]
yes
```

The encoding of arrays and maps uses indefinite-length encoding. All floats are currently encoded using decimal fractions. Encoding indicators and big floats are not currently supported.

6.22.3 Decoding

Decoding is accomplished using the `parse/2` predicate. For example:

```
| ?- cbor::parse([0x9f,0x61,0x61,0xbf,0x61,0x62,0x61,0x63,0xff,0xff], Term).
Term = [a,{b-c}]
yes
```

6.22.4 API documentation

Open the ../apis/library_index.html#cbor link in a web browser.

6.22.5 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(cbor(loader)).
```

6.22.6 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(cbor(tester)).
```

6.23 character_sets

This library provides a `character_set_protocol` protocol plus concrete objects for converting between lists of character codes and lists of bytes. It also provides metadata predicates `preferred_mime_name/1`, `name/1`, `alias/1`, and `mibenum/1` based on the IANA character set registry:

<https://www.iana.org/assignments/character-sets/character-sets.xhtml>

The currently provided objects are:

- `us_ascii`
- `iso_8859_1`
- `iso_8859_2`

- iso_8859_3
- iso_8859_4
- iso_8859_9
- iso_8859_10
- iso_8859_13
- iso_8859_14
- iso_8859_15
- iso_8859_16
- windows_1250
- windows_1251
- windows_1252
- windows_1253
- windows_1254
- windows_1257
- utf_8
- utf_16le
- utf_16be
- utf_32le
- utf_32be

Object names are derived from the preferred IANA MIME names by lowercasing them and replacing hyphens with underscores. When a registry entry has no distinct preferred MIME alias, the registered IANA name is used instead. A compatibility alias object named `utf16be` is also provided for `utf_16be`.

The Unicode character set objects work with Unicode scalar values and do not emit or consume a byte order mark (BOM).

This library intentionally does not currently provide `Shift_JIS` or `GB18030` objects because portable mapping tables for those multibyte encodings are not yet included.

No input validation is performed when converting between character codes and bytes. When necessary, use the types library validation and checking predicates before calling the `codes_to_bytes/2` and `bytes_to_codes/2` predicates.

6.23.1 API documentation

Open the ../apis/library_index.html#character_sets link in a web browser.

6.23.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(character_sets(loader)).
```

6.23.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(character_sets(tester)).
```

6.23.4 Usage

The UTF, ISO 8859, and Windows character set objects are grouped in three main files:

- `utf_character_sets.lgt`
- `iso_8859_character_sets.lgt`
- `windows_character_sets.lgt`

This allows some customization of the character set objects loaded by your application. Note that the `character_set_protocol.lgt` and `character_sets.lgt` base files must always be loaded (they include the `us_ascii` character set, which is thus always loaded).

6.24 ccsds_frames

The `ccsds_frames` library provides support for CCSDS transfer frames. The current implementation includes CCSDS telemetry transfer frames and CCSDS telecommand transfer frames, plus CCSDS advanced orbiting systems transfer frames, all using fixed frame lengths and optional mission-profile fields configured at the object level.

6.24.1 Available entities

- `ccsds_frame_protocol` Common protocol for transfer frame format objects.
- `ccsds_frames` Facade object for introspecting frame terms and bridging raw frame data fields with CCSDS space packets.
- `ccsds_tm_frames(FrameLength, SecondaryHeaderLength, HasFECF)` Parametric object implementing parsing, generation, and accessors for CCSDS telemetry transfer frames.
- `ccsds_tc_frames(FrameLength, SegmentHeaderLength, HasFECF)` Parametric object implementing parsing, generation, and accessors for CCSDS telecommand transfer frames.
- `ccsds_aos_frames(FrameLength, InsertZoneLength, HasOCF, HasFECF)` Parametric object implementing parsing, generation, and accessors for CCSDS advanced orbiting systems transfer frames.
- `ccsds_frames_types` Type definitions for TM, TC, and AOS transfer frame byte encodings and terms.

6.24.2 Representation

Telemetry transfer frames are represented using the compound term:

```
tm_transfer_frame(  
    Version,  
    SpacecraftId,  
    VirtualChannelId,  
    OCFFlag,  
    MasterChannelFrameCount,  
    VirtualChannelFrameCount,  
    SecondaryHeaderFlag,  
    SynchronizationFlag,  
    PacketOrderFlag,  
    SegmentLengthIdentifier,  
    FirstHeaderPointer,  
    SecondaryHeader,  
    DataField,  
    OCF,  
    FECF  
)
```

Where:

- SecondaryHeader is either none or secondary_header(Bytes)
- OCF is either none or ocf(Bytes)
- FECF is either none or fecf(Bytes)

Telecommand transfer frames are represented using the compound term:

```
tc_transfer_frame( Version, BypassFlag, ControlCommandFlag, SpacecraftId, VirtualChannelId, SequenceNumber, SegmentHeader, DataField, FECF )
```

Where:

- SegmentHeader is either none or segment_header(Bytes)
- FECF is either none or fecf(Bytes)

Advanced orbiting systems transfer frames are represented using the compound term:

```
aos_transfer_frame( Version, SpacecraftId, VirtualChannelId, VirtualChannelFrameCount, SignalingField, InsertZone, DataField, OCF, FECF )
```

Where:

- SignalingField is signaling_field(ReplayFlag, FrameCountUsageFlag, SpacecraftIdExtension, FrameCountCycle)
- InsertZone is either none or insert_zone(Bytes)
- OCF is either none or ocf(Bytes)
- FECF is either none or fecf(Bytes)

6.24.3 Parsing and generating

To parse a fixed 16-octet telemetry transfer frame without a telemetry transfer frame secondary header and without a frame error control field:

```
| ?- ccstds_tm_frames(16, 0, false)::parse(bytes([0x02, 0xA7, 0x10, 0x20, 0x18, 0x00, 1, 2, 3,
↪ 4, 5, 6, 0xDE, 0xAD, 0xBE, 0xEF]), Frames).
```

To generate the original byte sequence from the parsed term:

```
| ?- ccstds_tm_frames(16, 0, false)::generate(bytes(Bytes), Frames).
```

To parse a fixed 10-octet telecommand transfer frame without a segment header:

```
| ?- ccstds_tc_frames(10, 0, true)::parse(bytes([0x20, 0x2A, 0x0C, 0x09, 0x07, 1, 2, 3, 0x44, 0x6D]),
Frames).
```

To generate the original byte sequence from the parsed term:

```
| ?- ccstds_tc_frames(10, 0, true)::generate(bytes(Bytes), Frames).
```

When FECF support is enabled for a frame object, parsing verifies the incoming FECF, `valid/1` checks that it matches the remaining frame fields, and `generate/2-3` recomputes it from the remaining frame content. The parametric frame objects also provide `update_fecf/2` and `verify_fecf/1` for explicit integrity handling:

```
| ?- ccstds_tc_frames(10, 0, true)::update_fecf(tc_transfer_frame(0, 1, 0, 42, 3, 7, none, [1,2,3], none),
Frame).
```

```
| ?- ccstds_tc_frames(10, 0, true)::verify_fecf(Frame).
```

To parse a fixed 12-octet AOS transfer frame without an insert zone, OCF, or FECF:

```
| ?- ccstds_aos_frames(12, 0, false, false)::parse(bytes([0x4A, 0x83, 0x12, 0x34, 0x56, 0x80, 1, 2, 3, 4, 5,
6]), Frames).
```

To generate the original byte sequence from the parsed term:

```
| ?- ccstds_aos_frames(12, 0, false, false)::generate(bytes(Bytes), Frames).
```

6.24.4 Types

The `ccstds_frames_types` category defines the following types:

- `ccstds_tm_frame(FrameLength, SecondaryHeaderLength, HasFECF)` for byte encodings
- `ccstds_tm_frame_term(FrameLength, SecondaryHeaderLength, HasFECF)` for frame terms
- `ccstds_tc_frame(FrameLength, SegmentHeaderLength, HasFECF)` for byte encodings
- `ccstds_tc_frame_term(FrameLength, SegmentHeaderLength, HasFECF)` for frame terms
- `ccstds_aos_frame(FrameLength, InsertZoneLength, HasOCF, HasFECF)` for byte encodings
- `ccstds_aos_frame_term(FrameLength, InsertZoneLength, HasOCF, HasFECF)` for frame terms

For example:

```
| ?- type::check(ccstds_tm_frame(16, 0, false), [0x02, 0xA7, 0x10, 0x20, 0x18, 0x00, 1, 2, 3, ↪
↪ 4, 5, 6, 0xDE, 0xAD, 0xBE, 0xEF]).

| ?- type::check(ccstds_tm_frame_term(16, 0, false), tm_transfer_frame(0, 42, 3, 1, 16, 32, 0,
```

(continues on next page)

(continued from previous page)

```

→ 0, 0, 3, 0, none, [1,2,3,4,5,6], ocf([0xDE,0xAD,0xBE,0xEF]), none)).

| ?- type::check(ccsds_tc_frame(10, 0, true), [0x20, 0x2A, 0x0C, 0x09, 0x07, 1, 2, 3, 0x44,
→ 0x6D]).

| ?- type::check(ccsds_tc_frame_term(10, 0, true), tc_transfer_frame(0, 1, 0, 42, 3, 7, none,
→ [1,2,3], fecf([0x44,0x6D]))).

| ?- type::check(ccsds_aos_frame(12, 0, false, false), [0x4A, 0x83, 0x12, 0x34, 0x56, 0x80,
→ 1, 2, 3, 4, 5, 6]).

| ?- type::check(ccsds_aos_frame_term(12, 0, false, false), aos_transfer_frame(1, 42, 3,
→ 0x123456, signaling_field(1, 0, 0, 0), none, [1,2,3,4,5,6], none, none)).

```

6.24.5 Facade helpers

The `ccsds_frames` facade object provides generic accessors that work across TM, TC, and AOS frame terms:

- `frame_type/2`
- `version/2`
- `spacecraft_id/2`
- `virtual_channel_id/2`
- `data_field/2`
- `ocf/2`
- `fecf/2`
- `update_fecf/2`
- `verify_fecf/1`

It also provides raw payload bridge helpers for CCSDS space packets:

- `extract_packets(Frame, SecondaryHeaderLength, Packets)` parses the frame data field bytes using `ccsds_packets(SecondaryHeaderLength)`.
- `insert_packets(Packets, SecondaryHeaderLength, Frame, UpdatedFrame)` generates packet bytes using `ccsds_packets(SecondaryHeaderLength)` and replaces the frame data field while leaving the remaining frame fields unchanged.

For example:

```

| ?- ccsds_frames::extract_packets(tm_transfer_frame(0, 42, 3, 0, 16, 32, 0, 0, 0, 3, 0,
→ none, [0x08,0x01,0xC0,0x00,0x00,0x03,0xDE,0xAD,0xBE,0xEF], none, none), 0, Packets).

| ?- ccsds_frames::insert_packets([ccsds_packet(0, 0, 1, 1, 3, 0, none, [0xDE,0xAD,0xBE,
→ 0xEF])], 0, tm_transfer_frame(0, 42, 3, 0, 16, 32, 0, 0, 0, 3, 0, none, [], none, none),
→ UpdatedFrame).

| ?- ccsds_frames::update_fecf(tc_transfer_frame(0, 1, 0, 42, 3, 7, none, [1,2,3], fecf([0x00,0x00])), UpdatedFrame).

| ?- ccsds_frames::verify_fecf(UpdatedFrame).

```

6.24.6 API documentation

Open the `../apis/library_index.html#ccsds_frames` link in a web browser.

6.24.7 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ccsds_frames(loader)).
```

6.24.8 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ccsds_frames(tester)).
```

6.25 ccsds_link_profiles

The `ccsds_link_profiles` library provides an ergonomic wrapper layer over the existing parametric CCSDS frame objects.

Instead of selecting one of the raw `ccsds_tm_frames(...)`, `ccsds_tc_frames(...)`, or `ccsds_aos_frames(...)` objects directly at each call site, callers can work with explicit profile terms and generic wrapper predicates.

6.25.1 Representation

Link profiles are represented using the compound terms:

```
tm_profile(FrameLength, SecondaryHeaderLength, HasFECF)
tc_profile(FrameLength, SegmentHeaderLength, HasFECF)
aos_profile(FrameLength, InsertZoneLength, HasOCF, HasFECF)
```

Where:

- `FrameLength` is the fixed frame length in octets
- `SecondaryHeaderLength`, `SegmentHeaderLength`, and `InsertZoneLength` are the mission-profile field lengths in octets
- `HasOCF` and `HasFECF` are the atoms `true` or `false`

6.25.2 Public API

The current implementation provides the predicates:

- `valid_profile/1`
- `parse_frame/3`
- `parse_frames/3`
- `generate_frame/3`
- `generate_frames/3`
- `valid_reassembly_state/1`
- `initial_reassembly_state/1`
- `pending_fragments/2`
- `valid_discontinuity_policy/1`
- `extract_packets/4`
- `insert_packets/5`
- `reassemble_packets/6-8`
- `reassemble_frames/6-8`

6.25.3 Parsing and generating

To parse exactly one telemetry transfer frame using a profile term:

```
| ?- ccstds_link_profiles::parse_frame(bytes([0x02, 0xA7, 0x10, 0x20, 0x18, 0x00, 1, 2, 3, 4, ↵
↵5, 6, 0xDE, 0xAD, 0xBE, 0xEF]), tm_profile(16, 0, false), Frame).
```

To generate a telecommand transfer frame using a profile term:

```
| ?- ccstds_link_profiles::generate_frame(bytes(Bytes), tc_profile(10, 0, true), tc_transfer_
↵frame(0, 1, 0, 42, 3, 7, none, [1,2,3], fecf([0x44,0x6D]))).
```

`parse_frame/3` expects the source to contain exactly one frame. If the source contains zero frames or more than one frame, the predicate throws a `domain_error(ccstds_single_frame_source, Source)` exception.

Parsing and generation inherit FECF verification and regeneration from the underlying TM, TC, and AOS frame objects selected by the profile.

To parse or generate multi-frame sources and sinks, use `parse_frames/3` and `generate_frames/3`.

6.25.4 Packet helpers

Packet extraction and insertion are provided for TM and AOS profiles using the high-level packet-zone semantics from `ccstds_packet_services`:

```
| ?- ccstds_link_profiles::extract_packets(tm_profile(18, 0, false), 0, tm_transfer_frame(0, ↵
↵42, 3, 0, 16, 32, 0, 0, 0, 3, 2, none, [0xAA,0xBB,0x08,0x01,0xC0,0x00,0x00,0x03,0xDE,0xAD,
↵0xBE,0xEF], none, none), PacketZone).
```

(continues on next page)

(continued from previous page)

```
| ?- ccstds_link_profiles::insert_packets(aos_profile(20, 0, false, false), 0, packet_
↳zone([0xAA,0xBB], [ccstds_packet(0, 0, 1, 1, 3, 0, none, [0xDE,0xAD,0xBE,0xEF])], []), aos_
↳transfer_frame(1, 42, 3, 0x123456, signaling_field(1, 0, 0, 0), none, [0,0,0,0,0,0,0,0,0,
↳0,0,0,0], none, none), UpdatedFrame).
```

For cross-frame TM and AOS packet reconstruction, initialize a reassembly state and call the generic profile wrapper predicates:

```
| ?- ccstds_link_profiles::initial_reassembly_state(State), ccstds_link_profiles::reassemble_
↳frames(tm_profile(16, 0, false), 0, Frames, resynchronize, State, Packets, UpdatedState).
```

Telecommand profiles are intentionally rejected by `extract_packets/4` and `insert_packets/5`, `reassemble_packets/6-8`, and `reassemble_frames/6-8`, which throw a `domain_error(ccstds_packet_link_profile, tc_profile(...))` exception.

6.25.5 API documentation

Open the ../apis/library_index.html#ccstds_link_profiles link in a web browser.

6.25.6 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ccstds_link_profiles(loader)).
```

6.25.7 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ccstds_link_profiles(tester)).
```

6.26 ccstds_packet_services

The `ccstds_packet_services` library provides service-aware helpers that sit above the raw `ccstds_packets` and `ccstds_frames` codec libraries.

The current implementation supports telemetry and advanced orbiting systems packet-zone handling above the raw `ccstds_frames` payload bridge:

- splitting a packet-zone byte sequence using a TM transfer frame first header pointer
- parsing the complete packets that start in that zone
- preserving any leading continuation bytes and trailing incomplete bytes
- rebuilding a packet zone and reinserting it into a TM transfer frame
- splitting an AOS M_PDU packet-service data field into its first-header-pointer header and packet zone
- representing AOS idle-only packet zones distinctly from AOS no-start packet zones using the same `packet_zone/3` term

- rebuilding an AOS M_PDU packet-service data field and reinserting it into an AOS transfer frame
- carrying trailing packet fragments across multiple TM or AOS frames using an explicit packet reassembly state term
- reassembling complete packets across frame sequences while preserving any still-incomplete trailing fragment for the next batch
- tracking TM and AOS reassembly state per frame type, spacecraft identifier, and virtual channel identifier
- detecting dropped or out-of-order TM and AOS frames by checking the expected virtual-channel frame counts during reassembly
- applying configurable discontinuity recovery policies at the frame level, choosing between throwing an error, dropping the discontinuous frame, or resynchronizing on the current frame
- returning explicit recovery event streams from the richer frame reassembly predicates so callers can observe dropped fragments, skipped frames, and resynchronization points

6.26.1 Representation

Packet zones are represented using the compound term:

```
packet_zone(PrefixData, Packets, SuffixData)
```

Where:

- PrefixData is a list of bytes preceding the first complete packet
- Packets is a list of complete `ccsds_packet(...)` terms
- SuffixData is a list of bytes trailing the last complete packet

When a telemetry transfer frame first header pointer indicates that no packet starts in the frame data field, the full data field is represented as:

```
packet_zone(DataField, [], [])
```

For AOS packet service data fields, the same term is used after decoding the two-octet M_PDU first header pointer:

- `packet_zone(PrefixData, [], [])` corresponds to first header pointer 2047 and means that no packet starts in this frame packet zone
- `packet_zone([], [], SuffixData)` corresponds to first header pointer 2046 and means that the packet zone contains idle data only

Cross-frame packet reassembly state is represented using the compound term:

```
packet_reassembly_state(PendingData)
```

Where PendingData is a list of bytes belonging to an incomplete packet whose header started in an earlier frame and whose remaining bytes are expected in a later frame.

Frame-channel reassembly state is represented using the compound term:

```
channel_reassembly_state(Channels)
```

Where Channels is a list of terms of the form:

```
reassembly_channel(FrameType, SpacecraftId, VirtualChannelId, CounterModulus,
↳ExpectedFrameCount, PendingData)
```

The `FrameType` argument is either `tm` or `aos`. The keyed state is used by the frame-level reassembly predicates so that interleaved virtual channels and TM versus AOS streams do not collide.

Discontinuity recovery policies are represented by the atoms:

```
throw
drop
resynchronize
```

The policies have the following meaning when a frame-count discontinuity is detected for a keyed channel:

- `throw`: raise a `domain_error/2` exception and leave recovery to the caller
- `drop`: discard any buffered fragment for that keyed channel and skip the discontinuous frame payload entirely
- `resynchronize`: discard any buffered fragment for that keyed channel but still process the current frame, ignoring any pre-pointer continuation bytes

Recovery event streams are returned by the seven-argument TM and AOS frame reassembly predicates as lists of terms of the form:

```
dropped_fragment(FrameType, SpacecraftId, VirtualChannelId, Reason, PendingData)
skipped_frame(FrameType, SpacecraftId, VirtualChannelId, Discontinuity)
resynchronized(FrameType, SpacecraftId, VirtualChannelId, Discontinuity)
```

Where `Reason` is either:

- `discontinuity(StoredCounterModulus, CounterModulus, ExpectedFrameCount, VirtualChannelFrameCount)`
- `idle_only(VirtualChannelFrameCount)` for AOS idle-only M_PDUs that clear a buffered fragment

6.26.2 Packet-zone helpers

To split a packet zone using a first header pointer and parse the complete packets that start there:

```
| ?- ccstds_packet_services::split_packet_zone([0xAA,0xBB,0x00,0x00,0xC0,0x00,0x00,0x00,0x42,
↳0xCC], 2, 0, PacketZone).
```

To rebuild a packet zone and recover the corresponding first header pointer:

```
| ?- ccstds_packet_services::join_packet_zone(packet_zone([0xAA,0xBB], [ccstds_packet(0,0,0,0,
↳3,0,none,[0x42])], [0xCC]), 0, Bytes, FirstHeaderPointer).
```

6.26.3 Telemetry transfer frame helpers

To extract a packet zone from a telemetry transfer frame term:

```
| ?- ccstds_packet_services::extract_tm_packets(tm_transfer_frame(0, 42, 3, 0, 16, 32, 0, 0,
↪0, 3, 2, none, [0xAA,0xBB,0x00,0x00,0xC0,0x00,0x00,0x00,0x42,0xCC], none, none), 0,
↪PacketZone).
```

To update a telemetry transfer frame from a packet zone term:

```
| ?- ccstds_packet_services::insert_tm_packets(packet_zone([0xAA,0xBB], [ccstds_packet(0,0,0,
↪3,0,none,[0x42]]), [0xCC]), 0, tm_transfer_frame(0, 42, 3, 0, 16, 32, 0, 0, 0, 3, 0, none,
↪[], none, none), UpdatedFrame).
```

If the input TM frame already carries a fecf/1 term, insert_tm_packets/4 refreshes it from the updated frame bytes. Frames with none keep none.

6.26.4 Advanced orbiting systems helpers

To split an AOS M_PDU packet-service data field:

```
| ?- ccstds_packet_services::split_aos_packet_zone([0x00,0x02,0xAA,0xBB,0x00,0x00,0xC0,0x00,
↪0x00,0x00,0x42,0xCC], 0, PacketZone).
```

To rebuild an AOS M_PDU packet-service data field from a packet zone term:

```
| ?- ccstds_packet_services::join_aos_packet_zone(packet_zone([0xAA,0xBB], [ccstds_packet(0,0,
↪0,0,3,0,none,[0x42]]), [0xCC]), 0, Bytes).
```

To extract a packet zone from an AOS transfer frame term:

```
| ?- ccstds_packet_services::extract_aos_packets(aos_transfer_frame(1, 42, 3, 16, signaling_
↪field(0,0,0,0), none, [0x00,0x02,0xAA,0xBB,0x00,0x00,0xC0,0x00,0x00,0x00,0x42,0xCC], none,
↪none), 0, PacketZone).
```

To update an AOS transfer frame from a packet zone term:

```
| ?- ccstds_packet_services::insert_aos_packets(packet_zone([0xAA,0xBB], [ccstds_packet(0,0,0,
↪3,0,none,[0x42]]), [0xCC]), 0, aos_transfer_frame(1, 42, 3, 16, signaling_field(0,0,0,0),
↪none, [], none, none), UpdatedFrame).
```

If the input AOS frame already carries a fecf/1 term, insert_aos_packets/4 refreshes it from the updated frame bytes. Frames with none keep none.

6.26.5 Cross-frame packet reassembly

To get the initial low-level packet reassembly state:

```
| ?- ccstds_packet_services::initial_reassembly_state(State).
```

To get the initial keyed frame-channel reassembly state:

```
| ?- ccstds_packet_services::initial_channel_reassembly_state(State).
```

To reassemble a packet zone against a prior trailing fragment:

```
| ?- ccstds_packet_services::reassemble_packet_zone(packet_zone([0x00,0x00,0x42], [ccstds_
→packet(0,0,0,0,3,1,none,[0x43]]), []), 0, packet_reassembly_state([0x00,0x00,0xC0,0x00]),
→Packets, UpdatedState).
```

To reassemble complete packets across a sequence of telemetry transfer frames:

```
| ?- ccstds_packet_services::reassemble_tm_frames([tm_transfer_frame(0, 42, 3, 0, 16, 32, 0,
→0, 0, 3, 0, none, [0x00,0x00,0xC0,0x00]), none, none), tm_transfer_frame(0, 42, 3, 0, 16,
→33, 0, 0, 0, 3, 3, none, [0x00,0x00,0x42,0x00,0x00,0xC0,0x01,0x00,0x00,0x43]), none, none)],
→0, channel_reassembly_state([], Packets, UpdatedState).
```

To reassemble telemetry transfer frames while resynchronizing automatically on a detected discontinuity:

```
| ?- ccstds_packet_services::reassemble_tm_packets(tm_transfer_frame(0, 42, 3, 0, 16, 34, 0,
→0, 0, 3, 3, none, [0x00,0x00,0x42,0x00,0x00,0xC0,0x02,0x00,0x00,0x44]), none, none), 0,
→resynchronize, channel_reassembly_state([reassemble_channel(tm,42,3,256,33,[0x00,0x00,0xC0,
→0x00])]), Packets, UpdatedState).
```

To reassemble a telemetry transfer frame and also receive explicit recovery events:

```
| ?- ccstds_packet_services::reassemble_tm_packets(tm_transfer_frame(0, 42, 3, 0, 16, 34, 0,
→0, 0, 3, 3, none, [0x00,0x00,0x42,0x00,0x00,0xC0,0x02,0x00,0x00,0x44]), none, none), 0,
→resynchronize, channel_reassembly_state([reassemble_channel(tm,42,3,256,33,[0x00,0x00,0xC0,
→0x00])]), Packets, UpdatedState, Events).
```

To reassemble complete packets across a sequence of AOS transfer frames:

```
| ?- ccstds_packet_services::reassemble_aos_frames([aos_transfer_frame(1, 42, 3, 16,
→signaling_field(0,0,0,0), none, [0x00,0x00,0x00,0x00,0xC0,0x00]), none, none), aos_transfer_
→frame(1, 42, 3, 17, signaling_field(0,0,0,0), none, [0x00,0x03,0x00,0x00,0x42,0x00,0x00,
→0xC0,0x01,0x00,0x00,0x43]), none, none)], 0, channel_reassembly_state([], Packets,
→UpdatedState).
```

To reassemble AOS transfer frames while dropping a discontinuous frame and resetting the keyed buffered fragment:

```
| ?- ccstds_packet_services::reassemble_aos_packets(aos_transfer_frame(1, 42, 3, 18,
→signaling_field(0,0,0,0), none, [0x00,0x03,0x00,0x00,0x42,0x00,0x00,0xC0,0x02,0x00,0x00,
→0x44]), none, none), 0, drop, channel_reassembly_state([reassemble_channel(aos,42,3,
→16777216,17,[0x00,0x00,0xC0,0x00])]), Packets, UpdatedState).
```

To reassemble AOS transfer frames and receive explicit recovery events across a frame sequence:

```
| ?- ccstds_packet_services::reassemble_aos_frames([aos_transfer_frame(1, 42, 3, 16,
→signaling_field(0,0,0,0), none, [0x00,0x00,0x00,0x00,0xC0,0x00]), none, none), aos_transfer_
→frame(1, 42, 3, 18, signaling_field(0,0,0,0), none, [0x00,0x03,0x00,0x00,0x42,0x00,0x00,
→0xC0,0x02,0x00,0x00,0x44]), none, none)], 0, drop, channel_reassembly_state([], Packets,
→UpdatedState, Events).
```

To inspect the currently buffered keyed pending fragments:

```
| ?- ccstds_packet_services::pending_fragments(channel_reassembly_state([reassemble_
→channel(tm,42,3,256,33,[0x00,0x00,0xC0,0x00])]), PendingFragments).
```

For AOS packet service, idle-only M_PDUs clear any pending fragment state. This matches the packet-zone interpretation where first header pointer 2046 denotes idle data rather than continuation bytes.

If a TM or AOS frame arrives with a virtual-channel frame count different from the keyed state expected count, the frame-level reassembly predicates apply the selected discontinuity recovery policy. The five-argument variants default to the throw policy. The seven-argument variants additionally report the recovery decisions as an explicit event stream.

6.26.6 Scope

This milestone covers TM transfer frame packet zones, AOS packet-service M_PDUs, low-level packet-fragment stitching, and keyed TM/AOS cross-frame packet reassembly with continuity checks and configurable discontinuity recovery. Future milestones can extend the same library with segmentation policies, idle packet generation, richer per-channel recovery controls, and higher level packetization policies.

6.26.7 API documentation

Open the ../apis/library_index.html#ccsds_packet_services link in a web browser.

6.26.8 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ccsds_packet_services(loader)).
```

6.26.9 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ccsds_packet_services(tester)).
```

6.27 ccsds_packetization

The `ccsds_packetization` library provides the inverse service-layer helpers to the existing `ccsds_packet_services` library.

This first implementation slice covers:

- fitting packet streams into TM and AOS packet-service regions
- preserving partially emitted packet bytes across frame boundaries
- tracking queued packets per frame type, spacecraft identifier, and virtual channel identifier
- generating idle packets deliberately when there is remaining frame capacity
- reporting explicit packetization events for buffered fragments and generated idle packets

6.27.1 Representation

Packetizer state is represented using the compound term:

```
packetizer_state(Channels)
```

Where Channels is a list of terms of the form:

```
packetizer_channel(FrameType, SpacecraftId, VirtualChannelId, PendingBytes, QueuedPackets)
```

Where:

- FrameType is either tm or aos
- SpacecraftId is the keyed spacecraft identifier
- VirtualChannelId is the keyed virtual channel identifier
- PendingBytes is a list of packet bytes already started in an earlier frame but not yet fully emitted
- QueuedPackets is a list of complete ccsds_packet(...) terms waiting to be packetized for that keyed channel

Idle packets are generated as telemetry packets using APID 2047, standalone sequence flags, the requested sequence count, a zero-filled secondary header of the requested length, and zero-filled user data.

6.27.2 Packetization

To get the initial packetizer state:

```
| ?- ccsds_packetization::initial_state(State).
```

To inspect the currently buffered pending bytes and queued packets:

```
| ?- ccsds_packetization::pending_packets(packetizer_state([packetizer_channel(tm, 42, 3, _
↳ [0x22,0x33], [ccsds_packet(0,0,0,1,3,7,none,[0x44])])), PendingPackets).
```

To packetize packets into a TM transfer frame:

```
| ?- ccsds_packetization::packetize_tm_packets(tm_transfer_frame(0, 42, 3, 0, 16, 32, 0, 0, _
↳ 0, 3, 2047, none, [0,0,0,0,0,0,0], none, none), 0, packetizer_state([]), [ccsds_packet(0, _
↳ 0, 0, 0, 3, 0, none, [0x42])), UpdatedFrame, UpdatedState).
```

When the input TM frame already carries a fecf/1 term, packetize_tm_packets/6-7 refresh it from the updated frame bytes. Frames with none keep none.

To packetize packets into an AOS transfer frame:

```
| ?- ccsds_packetization::packetize_aos_packets(aos_transfer_frame(1, 42, 3, 16, signaling_
↳ field(0,0,0,0), none, [0,0,0,0,0,0,0,0], none, none), 0, packetizer_state([]), [ccsds_
↳ packet(0, 0, 0, 0, 3, 0, none, [0x42])), UpdatedFrame, UpdatedState).
```

When the input AOS frame already carries a fecf/1 term, packetize_aos_packets/6-7 refresh it from the updated frame bytes. Frames with none keep none.

To packetize packets across a sequence of TM transfer frames:

```
| ?- ccstds_packetization::packetize_tm_frames([Frame1, Frame2], 0, packetizer_state([]),
↪Packets, UpdatedFrames, RemainingPackets, UpdatedState).
```

6.27.3 Events

The richer single-frame and multi-frame packetization predicates return event lists using the terms:

```
buffered_packet_fragment(FrameType, SpacecraftId, VirtualChannelId, PendingBytes)
generated_idle_packet(FrameType, SpacecraftId, VirtualChannelId, Packet)
```

Where:

- buffered_packet_fragment/4 reports trailing packet bytes buffered for the next frame
- generated_idle_packet/4 reports an idle packet synthesized to consume remaining frame capacity

6.27.4 Idle packets

To generate a telemetry idle packet with APID 2047:

```
| ?- ccstds_packetization::generate_idle_packet(0, 7, 2, Packet).
```

6.27.5 API documentation

Open the `../apis/library_index.html#ccstds_packetization` link in a web browser.

6.27.6 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ccstds_packetization(loader)).
```

6.27.7 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ccstds_packetization(tester)).
```

6.28 ccstds_packets

The `ccstds_packets` library implements predicates for parsing and generating CCSDS (Consultative Committee for Space Data Systems) Space Packets following the CCSDS 133.0-B-2 standard (Space Packet Protocol).

Reference documentation:

- <https://public.ccsds.org/Pubs/133x0b2e1.pdf>
- <https://jastoolbox.sandia.gov/topic/communication-specification/jas-packets/ccstds-telecommand-and-telemetry-format-packet-standard/>

6.28.1 Packet Structure

A CCSDS Space Packet consists of:

1. **Primary Header (6 bytes):**

- Version Number (3 bits) - Always 000 for Space Packets
- Packet Type (1 bit) - 0=Telemetry, 1=Telecommand
- Secondary Header Flag (1 bit) - 0=absent, 1=present
- Application Process ID (APID) (11 bits) - 0-2047
- Sequence Flags (2 bits) - 00=continuation, 01=first, 10=last, 11=standalone
- Packet Sequence Count (14 bits) - 0-16383
- Packet Data Length (16 bits) - Number of octets in data field minus 1

2. **User Data Field** - Variable length payload

6.28.2 Representation

Packets are represented using the compound term:

```
ccsds_packet(Version, Type, SecHeaderFlag, APID, SeqFlags, SeqCount, SecHeader, UserData)
```

Where:

- Version is an integer (0-7, typically 0)
- Type is an integer (0=telemetry, 1=telecommand)
- SecHeaderFlag is an integer (0 or 1)
- APID is an integer (0-2047)
- SeqFlags is an integer (0-3)
- SeqCount is an integer (0-16383)
- SecHeader is either none or secondary_header(Bytes) where Bytes is a list of bytes
- UserData is a list of bytes

Note that the DataLength field from the wire format is not stored in the term representation as it can be computed from SecHeader and UserData. The data_length/2 accessor predicate computes and returns this value when needed.

6.28.3 Parsing

The parse/2 predicate accepts a source term as its first argument. The source can be file(File), stream(Stream), or bytes(Bytes). All source types return a list of packets for uniformity.

To parse packets from a list of bytes:

```
| ?- ccsds_packets::parse(bytes([0x08, 0x01, 0xC0, 0x00, 0x00, 0x03, 0xDE, 0xAD, 0xBE,
↪ 0xEF]), Packets).
Packets = [ccsds_packet(0, 0, 1, 1, 3, 0, none, [222, 173, 190, 239])]
yes
```

To parse packets from a binary file:

```
| ?- ccstds_packets::parse(file('telemetry.bin'), Packets).
```

To parse packets from a binary stream:

```
| ?- ccstds_packets::parse(stream(Stream), Packets).
```

When the packets include a secondary header, the secondary header length must be known. In this case, use the `ccstds_packets(SecondaryHeaderLength)` object instead of the `ccstds_packets` object. For example, to parse packets with a secondary header of 6 bytes:

```
| ?- ccstds_packets(6)::parse(bytes([0x08, 0x01, 0xC0, 0x00, 0x00, 0x07, 0x01, 0x02, 0x03,
↳ 0x04, 0x05, 0x06, 0xAA, 0xBB]), Packets).
```

6.28.4 Generating

The `generate/2` predicate accepts a sink term as its first argument and a list of packet terms as the second argument. The sink can be `file(File)`, `stream(Stream)`, or `bytes(Bytes)`.

To generate bytes from a list of packet terms:

```
| ?- ccstds_packets::generate(bytes(Bytes), [ccstds_packet(0, 0, 1, 1, 3, 0, none, [0xDE, 0xAD,
↳ 0xBE, 0xEF])]).
Bytes = [8, 1, 192, 0, 0, 3, 222, 173, 190, 239]
yes
```

To write the bytes generated from a list of packet terms to a binary file:

```
| ?- ccstds_packets::generate(file('output.bin'), Packets).
```

To write the bytes generated from a list of packet terms to a binary stream:

```
| ?- ccstds_packets::generate(stream(Stream), Packets).
```

6.28.5 Accessor Predicates

The library provides convenient accessor predicates for extracting packet fields:

```
| ?- ccstds_packets::apid(Packet, APID).

% Returns telemetry or telecommand
| ?- ccstds_packets::type(Packet, Type).

% Returns continuation, first, last, or standalone
| ?- ccstds_packets::sequence_flags(Packet, Flags).

% Returns the secondary header as a list of bytes or none
| ?- ccstds_packets::secondary_header(Packet, SecHeader).

% Parses a self-describing secondary-header time field when Descriptor is a variable.
| ?- ccstds_packets(7)::secondary_header_time(Packet, Descriptor, Time).
```

(continues on next page)

(continued from previous page)

```
% Parses raw secondary-header T-field bytes using an explicit descriptor term.
| ?- ccstds_packets(10)::secondary_header_time(Packet, cuc_descriptor(5, 5, ccstds_epoch),_
↳Time).

| ?- ccstds_packets::user_data(Packet, Data).

% Returns the packet data length (computed from secondary header and user data)
| ?- ccstds_packets::data_length(Packet, Length).
```

6.28.6 Types and arbitrary generators

The library includes a `ccstds_packets_types` category that provides `ccstds_packet` and `ccstds_packet(SecondaryHeaderLength)` types and arbitrary generators for CCSDS packets. For example:

```
| ?- type::check(ccstds_packet, Bytes).

| ?- type::arbitrary(ccstds_packet(42), Bytes).
```

It also provides a `ccstds_packets(N)` and `ccstds_packets(SecondaryHeaderLength, N)` types for generating a list with N packets. For example:

```
| ?- type::arbitrary(ccstds_packets(10), Bytes).

| ?- type::check(ccstds_packets(42, 10), Bytes).
```

For a term representation of a packet, use the `ccstds_packet_term` and `ccstds_packet_term(SecondaryHeaderLength)` types and arbitrary generators. For example:

```
| ?- type::check(ccstds_packet_term, Packet).

| ?- type::arbitrary(ccstds_packet_term(42), Packet).
```

6.28.7 API documentation

Open the `../apis/library_index.html#ccstds_packets` link in a web browser.

6.28.8 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ccstds_packets(loader)).
```

6.28.9 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ccsds_packets(tester)).
```

To test the performance of the library parsing predicates, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(ccsds_packets(tester_performance)).
```

6.29 ccsds_tc_services

The `ccsds_tc_services` library provides service-aware helpers that sit above raw `ccsds_frames` telecommand transfer frame terms.

This first implementation slice covers:

- extracting telecommand segments from CCSDS telecommand transfer frames
- encoding telecommand segments back into telecommand transfer frames
- tracking segmented telecommand reassembly state per spacecraft identifier and virtual channel identifier
- keeping pending segmented data separately per MAP identifier inside each virtual channel
- detecting telecommand frame-sequence discontinuities during reassembly
- applying configurable discontinuity recovery policies during reassembly
- returning explicit recovery event streams from the richer frame reassembly predicates
- returning provenance-aware reassembled service-unit terms that preserve the ordered per-frame mission-specific segment-header suffixes contributing to a complete telecommand service unit
- grouping complete telecommand service units into MAP-aware dispatch buckets so callers can route them by MAP identifier without rebuilding that bookkeeping

6.29.1 Representation

Telecommand segments extracted by this library are represented using the compound term:

```
tc_segment(SequenceFlags, MapId, HeaderSuffix, Data)
```

Where:

- `SequenceFlags` is one of `continuation`, `first`, `last`, or `unsegmented`
- `MapId` is the multiplexer access point identifier in the range 0-63
- `HeaderSuffix` is the list of mission-specific segment-header octets that follow the standard first octet
- `Data` is the segment data as a list of bytes

The compact term `tc_segment(SequenceFlags, MapId, Data)` is also accepted by predicates such as `valid_segment/1` and `insert_tc_segment/3` as shorthand for an empty header suffix.

Reassembled service units that are reconstructed from multiple transfer frames use the empty list as `HeaderSuffix`, as there is no single mission-specific suffix value that unambiguously represents the merged result.

When callers need that provenance instead of the compact compatibility view, the companion reassembly predicates return terms of the form:

```
tc_reassembled_segment(MapId, HeaderSuffixes, Data)
```

Where HeaderSuffixes is the ordered list of the per-frame mission-specific segment-header suffixes that contributed to the complete service unit.

MAP-aware dispatch helpers group complete telecommand service units into terms of the form:

```
map_dispatch(MapId, ServiceUnits)
```

Where ServiceUnits is a list of complete telecommand service-unit terms for a single MapId, preserving the original service-unit order inside each bucket and the first-seen order of the MAP buckets.

Telecommand reassembly state is represented using the compound term:

```
tc_reassembly_state(Channels)
```

Where Channels is a list of terms of the form:

```
tc_reassembly_channel(SpacecraftId, VirtualChannelId, ExpectedSequenceNumber, ↵
↵PendingSegments)
```

The PendingSegments list stores the currently buffered segmented data by MAP identifier using terms of the form:

```
tc_pending_segment(MapId, PendingData)
```

Pending buffered fragments can be queried using pending_fragments/2, which returns flattened terms of the form:

```
pending_fragment(SpacecraftId, VirtualChannelId, MapId, PendingData)
```

Discontinuity recovery policies are represented by the atoms:

```
throw
drop
resynchronize
```

The policies have the following meaning when a frame sequence discontinuity is found for a telecommand virtual channel:

- throw: raise a domain_error/2 exception
- drop: discard all pending MAP fragments for that virtual channel and skip the discontinuous frame payload
- resynchronize: discard all pending MAP fragments for that virtual channel and still process the current frame payload

Recovery event streams are returned by reassemble_tc_frame/6 and reassemble_tc_frames/6 as lists of terms of the form:

```
dropped_fragment(SpacecraftId, VirtualChannelId, MapId, Reason, PendingData)
skipped_frame(SpacecraftId, VirtualChannelId, Discontinuity)
resynchronized(SpacecraftId, VirtualChannelId, Discontinuity)
```

Where Discontinuity is represented by:

```
discontinuity(ExpectedSequenceNumber, SequenceNumber)
```

Provenance-aware reassembly is provided by the companion predicates `reassemble_tc_frame_with_provenance/5-6` and `reassemble_tc_frames_with_provenance/5-6`.

MAP-aware dispatch is provided by `dispatch_service_units_by_map/2-3`, which accept both compact telecommand segment terms and provenance-aware reassembled telecommand service-unit terms.

6.29.2 Scope

This first milestone interprets the standard telecommand MAP and sequence semantics from the first segment-header octet, where the high two bits encode the sequence flags and the low six bits encode the MAP identifier. Transfer frames with longer mission-specific segment headers are also supported: additional segment-header octets are exposed in extracted segment terms as the `HeaderSuffix` component and are preserved when updating frames.

6.29.3 Examples

To extract a telecommand segment from a telecommand transfer frame term:

```
| ?- ccstds_tc_services::extract_tc_segment(tc_transfer_frame(0, 1, 0, 42, 3, 7, segment_
↳header([0xC5]), [1,2,3], none), Segment).
```

To update a telecommand transfer frame from a telecommand segment term:

```
| ?- ccstds_tc_services::insert_tc_segment(tc_segment(first, 42, [1,2,3]), tc_transfer_
↳frame(0, 1, 0, 42, 3, 7, none, [], none), UpdatedFrame).
```

To inspect the mission-specific segment-header suffix from an extracted telecommand segment term:

```
| ?- ccstds_tc_services::segment_header_suffix(tc_segment(first, 5, [0xAA,0xBB], [1,2,3]),
↳HeaderSuffix).
```

Mission-specific segment-header suffix octets are preserved when present in the input frame term:

```
| ?- ccstds_tc_services::insert_tc_segment(tc_segment(first, 42, [1,2,3]), tc_transfer_
↳frame(0, 1, 0, 42, 3, 7, segment_header([0x00,0xAA,0xBB]), [], none), UpdatedFrame).
```

To reassemble a segmented telecommand service unit across a sequence of frames:

```
| ?- ccstds_tc_services::reassemble_tc_frames([tc_transfer_frame(0, 1, 0, 42, 3, 7, segment_
↳header([0x45]), [1,2], none), tc_transfer_frame(0, 1, 0, 42, 3, 8, segment_header([0x85]),
↳[3,4], none)], tc_reassembly_state([]), Segments, UpdatedState).
```

To reassemble a segmented telecommand service unit across a sequence of frames and preserve the ordered per-frame mission-specific header-suffix provenance:

```
| ?- ccstds_tc_services::reassemble_tc_frames_with_provenance([tc_transfer_frame(0, 1, 0, 42,
↳3, 7, segment_header([0x45,0xAA]), [1,2], none), tc_transfer_frame(0, 1, 0, 42, 3, 8,
↳segment_header([0x85,0xBB]), [3,4], none)], tc_reassembly_state([]), ReassembledSegments,
↳UpdatedState, Events).
```

To group complete telecommand service units into MAP-aware dispatch buckets:

```
| ?- ccstds_tc_services::dispatch_service_units_by_map([tc_segment(unsegmented, 7, [],  
→[0x01]), tc_reassembled_segment(6, [[0xBB]], [0x03]), tc_reassembled_segment(7, [[0xAA]],  
→[0x02])), Dispatches).
```

To resynchronize on a discontinuous telecommand frame while receiving explicit recovery events:

```
| ?- ccstds_tc_services::reassemble_tc_frame(tc_transfer_frame(0, 1, 0, 42, 3, 9, segment_  
→header([0xC6]), [9], none), resynchronize, tc_reassembly_state([tc_reassembly_channel(42,  
→3, 8, [tc_pending_segment(5, [1])])), Segments, UpdatedState, Events).
```

6.29.4 API documentation

Open the ../apis/library_index.html#ccstds_tc_services link in a web browser.

6.29.5 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ccstds_tc_services(loader)).
```

6.29.6 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ccstds_tc_services(tester)).
```

6.30 ccstds_time_codes

The `ccstds_time_codes` library provides support for CCSDS time codes. The current implementation includes CCSDS Unsegmented Time Code (CUC) and CCSDS Day Segmented Time Code (CDS), and CCSDS Calendar Segmented Time Code (CCS), keeping the package small and reusable by the existing `ccstds_packets` packet library and by future frame-level libraries.

6.30.1 Available entities

- `ccstds_time_code_protocol` Common protocol for time code format objects.
- `ccstds_time_codes` Facade object for introspecting time code terms.
- `ccstds_cuc`(CoarseOctets, FineOctets, Epoch) Parametric object implementing parsing, generation, and conversion for CUC time codes.
- `ccstds_cds`(DaySegmentOctets, SubmillisecondOctets, Epoch) Parametric object implementing parsing, generation, and conversion for CDS time codes.
- `ccstds_ccs`(CalendarVariant, FractionOctets) Parametric object implementing parsing, generation, and conversion for CCS time codes.
- `ccstds_time_codes_types` Type definitions and arbitrary generators for CUC terms and byte encodings.

6.30.2 Representation

CUC values are represented using the compound term:

```
cuc_time(Coarse, Fine)
```

Where:

- Coarse is the integer value stored in the coarse time octets
- Fine is the integer value stored in the fine time octets

The number of octets used for each field is selected by the parametric object.

CDS values are represented using the compound terms:

```
cds_time(Days, Milliseconds)
cds_time(Days, Milliseconds, Submilliseconds)
```

Where:

- Days is the integer value stored in the day segment
- Milliseconds is the millisecond-of-day value
- Submilliseconds is the optional submillisecond value when the object uses a submillisecond segment; 2-octet objects store microseconds and 4-octet objects store picoseconds within the millisecond

CCS values are represented using the compound terms:

```
ccs_calendar_time(Year, Month, Day, Hour, Minute, Second, Fraction)
ccs_ordinal_time(Year, DayOfYear, Hour, Minute, Second, Fraction)
```

Where:

- Year is the four-digit BCD year
- Month and Day are used by the calendar variant
- DayOfYear is used by the day_of_year variant
- Fraction stores the optional BCD fractional digits encoded by the object

6.30.3 Parsing and generating

To parse a 4-octet coarse and 2-octet fine CUC value using the CCSDS epoch:

```
| ?- ccstds_cuc(4, 2, ccstds_epoch)::parse(bytes([0x00, 0x00, 0x01, 0x00, 0x00, 0x80]),_
↳TimeCode).
TimeCode = cuc_time(256, 128)
yes
```

To generate the corresponding byte sequence:

```
| ?- ccstds_cuc(4, 2, ccstds_epoch)::generate(bytes(Bytes), cuc_time(256, 128)).
Bytes = [0, 0, 1, 0, 0, 128]
yes
```

To parse a 16-bit day-segment CDS value without the optional submillisecond segment:

```
| ?- ccstds_cds(2, 0, ccstds_epoch)::parse(bytes([0x00, 0x01, 0x00, 0x00, 0x07, 0xD0]),  
↪TimeCode).  
TimeCode = cds_time(1, 2000)  
yes
```

To parse a CDS value with a 16-bit submillisecond segment:

```
| ?- ccstds_cds(2, 2, unix_epoch)::parse(bytes([0x00, 0x01, 0x00, 0x00, 0x07, 0xD0, 0x01,  
↪0xF4]), TimeCode).  
TimeCode = cds_time(1, 2000, 500)  
yes
```

To parse a CDS value with a 32-bit submillisecond segment:

```
| ?- ccstds_cds(3, 4, unix_epoch)::parse(bytes([0x00, 0x00, 0x01, 0x00, 0x00, 0x07, 0xD0,  
↪0x1D, 0xCD, 0x65, 0x00]), TimeCode).  
TimeCode = cds_time(1, 2000, 500000000)  
yes
```

To parse a CCS calendar-segmented value without fractional digits:

```
| ?- ccstds_ccs(calendar, 0)::parse(bytes([0x20, 0x26, 0x05, 0x08, 0x14, 0x30, 0x45]),  
↪TimeCode).  
TimeCode = ccs_calendar_time(2026, 5, 8, 14, 30, 45, 0)  
yes
```

To parse an ordinal CCS value with one fractional BCD octet:

```
| ?- ccstds_ccs(day_of_year, 1)::parse(bytes([0x20, 0x26, 0x01, 0x28, 0x14, 0x30, 0x45,  
↪0x67]), TimeCode).  
TimeCode = ccs_ordinal_time(2026, 128, 14, 30, 45, 67)  
yes
```

6.30.4 Unix time conversion

To convert a CUC value that uses the Unix epoch into Unix seconds:

```
| ?- ccstds_cuc(4, 2, unix_epoch)::unix_seconds(cuc_time(1, 32768), Seconds).  
Seconds = 1.5  
yes
```

To convert Unix seconds back into a CUC value:

```
| ?- ccstds_cuc(4, 2, unix_epoch)::from_unix_seconds(1.5, TimeCode).  
TimeCode = cuc_time(1, 32768)  
yes
```

To convert a CDS value that uses the Unix epoch into Unix seconds:

```
| ?- ccstds_cds(2, 2, unix_epoch)::unix_seconds(cds_time(1, 2000, 500), Seconds).  
Seconds = 86402.0005  
yes
```

To convert Unix seconds back into a CDS value:

```
| ?- ccstds_cds(2, 2, unix_epoch)::from_unix_seconds(86402.0005, TimeCode).
TimeCode = cds_time(1, 2000, 500)
yes
```

To convert a CCS calendar value into Unix seconds:

```
| ?- ccstds_ccs(calendar, 0)::unix_seconds(ccs_calendar_time(1970, 1, 1, 0, 0, 0, 0),
↪Seconds).
Seconds = 0
yes
```

To convert Unix seconds back into an ordinal CCS value:

```
| ?- ccstds_ccs(day_of_year, 1)::from_unix_seconds(1.5, TimeCode).
TimeCode = ccs_ordinal_time(1970, 1, 0, 0, 1, 50)
yes
```

6.30.5 Types and arbitrary generators

The `ccstds_time_codes_types` category defines the following types:

- `ccstds_cuc(CoarseOctets, FineOctets, Epoch)` for byte encodings
- `ccstds_cuc_time(CoarseOctets, FineOctets)` for `cuc_time/2` terms
- `ccstds_cds(DaySegmentOctets, SubmillisecondOctets, Epoch)` for byte encodings
- `ccstds_cds_time(DaySegmentOctets, SubmillisecondOctets)` for CDS terms
- `ccstds_ccs(CalendarVariant, FractionOctets)` for byte encodings
- `ccstds_ccs_time(CalendarVariant, FractionOctets)` for CCS terms

For example:

```
| ?- type::check(ccstds_cuc(4, 2, ccstds_epoch), [0, 0, 1, 0, 0, 128]).
| ?- type::check(ccstds_cuc_time(4, 2), cuc_time(256, 128)).
| ?- type::check(ccstds_cds(2, 0, ccstds_epoch), [0, 1, 0, 0, 7, 208]).
| ?- type::check(ccstds_cds_time(2, 2), cds_time(1, 2000, 500)).
| ?- type::check(ccstds_ccs(calendar, 0), [0x20, 0x26, 0x05, 0x08, 0x14, 0x30, 0x45]).
| ?- type::check(ccstds_ccs_time(day_of_year, 1), ccs_ordinal_time(2026, 128, 14, 30, 45,
↪67)).
```

6.30.6 API documentation

Open the ../apis/library_index.html#ccsds_time_codes link in a web browser.

6.30.7 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ccsds_time_codes(loader)).
```

6.30.8 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ccsds_time_codes(tester)).
```

6.31 ccsds_time_fields

The `ccsds_time_fields` library adds the self-describing time-field layer that is missing between raw embedded time bytes and the format-specific objects in `ccsds_time_codes`.

This library covers CCSDS binary P-fields for:

- CUC time fields using 1 to 7 coarse octets and 0 to 10 fine octets, including the two-octet extended P-field encoding
- CDS time fields using 16-bit or 24-bit day segments and either no submillisecond segment, a 16-bit microsecond segment, or a 32-bit picosecond segment
- CCS time fields using either calendar or day-of-year encoding and up to 6 BCD fraction octets

6.31.1 Representation

Descriptors are represented using the compound terms:

```
cuc_descriptor(CoarseOctets, FineOctets, Epoch)
cds_descriptor(DaySegmentOctets, SubmillisecondOctets, Epoch)
ccs_descriptor(CalendarVariant, FractionOctets)
```

Where:

- Epoch is either `ccsds_epoch` or `unix_epoch`
- CalendarVariant is either `calendar` or `day_of_year`

This library uses the existing `ccsds_time_codes` objects to parse and generate the T-field bytes once the descriptor has been decoded.

For the current scope, the agency-defined epoch bit in CUC and CDS descriptors is mapped to the existing `unix_epoch` atom so callers can keep using the same time-code objects and terms already provided by `ccsds_time_codes`.

6.31.2 Parsing and generating

To parse a self-describing CUC time field:

```
| ?- ccstds_time_fields::parse(bytes([0x1E, 0x00, 0x00, 0x01, 0x00, 0x00, 0x80]), Descriptor, _  
↪ TimeCode).
```

To parse a self-describing extended CUC time field:

```
| ?- ccstds_time_fields::parse(bytes([0x9F, 0x28, 0x00, 0x00, 0x00, 0x01, 0x00, 0x00, 0x00, _  
↪ 0x00, 0x00, 0x80]), Descriptor, TimeCode).
```

To generate the same time field from a descriptor and a time-code term:

```
| ?- ccstds_time_fields::generate(bytes(Bytes), cuc_descriptor(4, 2, ccstds_epoch), cuc_  
↪ time(256, 128)).
```

To parse a self-describing CDS time field:

```
| ?- ccstds_time_fields::parse(bytes([0x40, 0x00, 0x01, 0x00, 0x00, 0x07, 0xD0]), Descriptor, _  
↪ TimeCode).
```

To parse a self-describing CDS time field with a 32-bit submillisecond segment:

```
| ?- ccstds_time_fields::parse(bytes([0x4E, 0x00, 0x00, 0x01, 0x00, 0x00, 0x07, 0xD0, 0x1D, _  
↪ 0xCD, 0x65, 0x00]), Descriptor, TimeCode).
```

To generate a self-describing CCS day-of-year time field:

```
| ?- ccstds_time_fields::generate(bytes(Bytes), ccs_descriptor(day_of_year, 1), ccs_ordinal_  
↪ time(2026, 128, 14, 30, 45, 67)).
```

6.31.3 Descriptor helpers

The library also provides helper predicates for descriptor introspection:

- valid_descriptor/1
- format/2
- epoch/2

For example:

```
| ?- ccstds_time_fields::format(cds_descriptor(3, 4, unix_epoch), Format).  
  
| ?- ccstds_time_fields::epoch(ccs_descriptor(calendar, 1), Epoch).
```

6.31.4 API documentation

Open the `../apis/library_index.html#ccsds_time_fields` link in a web browser.

6.31.5 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ccsds_time_fields(loader)).
```

6.31.6 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ccsds_time_fields(tester)).
```

6.32 classification_protocols

This library provides protocols used in the implementation of machine learning classifier algorithms. Datasets are represented as objects implementing the `dataset_protocol` protocol. Classifiers are represented as objects implementing the `classifier_protocol` protocol.

This library also provides reusable shared categories, smoke tests, and test datasets. See below for details.

Logtalk currently provides several classifiers including `c45_classifier`, `knn_classifier`, `linear_svm_classifier`, `logistic_regression_classifier`, `naive_bayes_classifier`, `nearest_centroid_classifier`, and `random_forest_classifier`. See these libraries documentation for details.

6.32.1 Shared category

The library includes one reusable category intended to be imported by classifier algorithm implementations:

- `classifier_common` — shared diagnostics accessors and classifier export helpers.

This category keeps diagnostics access and file export behavior separate from the algorithm-specific learning, prediction, and pretty-printing code.

6.32.2 Diagnostics

The `classifier_common` category provides shared accessor predicates such as `diagnostics/2`, `diagnostic/2`, and `classifier_options/2`. These predicates make it possible to inspect learned-classifier metadata without depending on the exact term representation used by a particular classifier implementation.

The detailed contents of the diagnostics data are classifier-dependent. For example, some classifiers report effective training options, while others report structural metadata such as attribute names, feature types, or the number of training examples or models.

6.32.3 Export header format

The shared classifier exporter in the `classifier_common` category writes a header before the exported clauses in the following format:

```
% exported classifier predicate: Functor/Arity
% training dataset: Dataset
% dataset prediction schema: Functor(Attribute1, ..., AttributeN, Class)
% diagnostics: Diagnostics
% Functor(Classifier)
Functor(Classifier)
```

The dataset prediction schema line always uses an ASCII-only title case conversion for the attribute names and class. This line documents the dataset-level prediction interface for readability, even when the exported clauses serialize a model term instead of an executable predictor relation.

When exporting a serialized classifier term, using a noun such as `classifier/1` or `model/1` is recommended.

6.32.4 API documentation

Open the `../apis/library_index.html#classification_protocols` link in a web browser.

6.32.5 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(classification_protocols(loader)).
```

6.32.6 Testing

To test this library predicates, shared categories, and datasets, load the `tester.lgt` file:

```
| ?- logtalk_load(classification_protocols(tester)).
```

6.32.7 Test datasets

Several sample datasets are included in the `test_datasets` directory:

- `play_tennis.lgt` — The classic weather/tennis dataset with 14 examples and 4 discrete attributes (outlook, temperature, humidity, wind). Originally from Quinlan (1986) and widely used in machine learning textbooks including Mitchell (1997). Also available from the UCI Machine Learning Repository: <https://archive.ics.uci.edu/dataset/349/tennis+major+tournament+match+statistics>
- `contact_lenses.lgt` — A dataset with 24 examples and 4 discrete attributes (age, spectacle prescription, astigmatism, tear production rate) for deciding the type of contact lenses to prescribe. Originally from Cendrowska, J. (1987). PRISM: An algorithm for inducing modular rules. *International Journal of Man-Machine Studies*, 27(4), 349-370. Available from the UCI Machine Learning Repository: <https://archive.ics.uci.edu/dataset/58/lenses>
- `iris.lgt` — The classic Iris flower dataset with 150 examples and 4 continuous attributes (sepal length, sepal width, petal length, petal width) for classifying iris species (setosa, versicolor, virginica). Originally from Fisher, R.A. (1936). The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7(2), 179-188. Available from the UCI Machine Learning Repository: <https://archive.ics.uci.edu/dataset/53/iris>

- `breast_cancer.lgt` — A dataset with 286 examples and 9 discrete attributes (age, menopause, tumor size, inv-nodes, node-caps, degree of malignancy, breast, breast quadrant, irradiation) for predicting breast cancer recurrence events. Contains missing values (9 examples with missing values in the node-caps and breast-quad attributes, represented using anonymous variables). Originally from the Institute of Oncology, University Medical Centre, Ljubljana, Yugoslavia. Donors: Ming Tan and Jeff Schlimmer. Available from the UCI Machine Learning Repository: <https://archive.ics.uci.edu/dataset/14/breast+cancer>

6.33 clo_span_pattern_miner

CloSpan closed sequential pattern miner for sequence datasets. The library depends on the `sequential_pattern_mining_protocols` support library, implements the generic `pattern_miner_protocol` defined in the `pattern_mining_protocols` core library, and mines closed frequent sequential patterns directly using closure-aware projected-database search with explicit projected-database-equivalence backward pruning that skips equivalent branches before recursion, together with a same-support closed frontier that merges branch results during the search without a final post-filter.

Requires a dataset implementing `sequence_dataset_protocol` with sequences represented as ordered lists of canonical sorted itemsets over a declared item domain.

6.33.1 API documentation

Open the `../apis/library_index.html#clo_span_pattern_miner` link in a web browser.

6.33.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(clo_span_pattern_miner(loader)).
```

6.33.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(clo_span_pattern_miner(tester)).
```

6.33.4 Features

- **Closed Pattern Mining:** Retains only frequent sequential patterns that have no superpattern with the same support.
- **Backward and Closure Pruning:** Uses projected-database equivalence to skip equivalent backward-growth branches before recursion, then maintains a same-support closed frontier to suppress dominated patterns during the search instead of post-filtering another miner output.
- **Canonical Sequences:** Uses the shared sequential validation logic from the support library.
- **Flexible Support Thresholds:** Supports relative minimum support and absolute minimum support count. When both are given, the absolute-count threshold takes precedence.
- **Model Export:** Closed pattern collections can be exported as predicate clauses or written to a file.

6.33.5 Options

The `mine/3` predicate accepts the following options:

- `minimum_support/1`: Relative minimum support threshold in the interval `[0.0, 1.0]`. The default is `0.5`.
- `minimum_support_count/1`: Absolute minimum support count. When both support options are provided, this option takes precedence.
- `maximum_pattern_length/1`: Maximum total number of items in a mined sequential pattern. The default is `1000`, effectively capped by the longest sequence in the dataset.
- `minimum_pattern_length/1`: Minimum total number of items retained in the mined result. The default is `1`.

6.33.6 Pattern miner representation

The mined pattern miner result is represented by a compound term with the functor chosen by the implementation and arity 3. For example:

```
clo_span_pattern_miner(ItemDomain, Patterns, Options)
```

Where:

- `ItemDomain`: Canonical sorted list of declared dataset items.
- `Patterns`: List of `sequence_pattern(Pattern, SupportCount)` terms ordered first by total item count and then lexicographically.
- `Options`: Effective mining options used to mine the closed sequential patterns.

6.33.7 References

1. Yan, X., Han, J., and Afshar, R. (2003) - “CloSpan: Mining closed sequential patterns in large datasets”.

6.34 clustering_protocols

This library provides protocols used in the implementation of machine learning clustering algorithms. Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol. Clusterers are represented as objects implementing the `clusterer_protocol` protocol.

This library also provides test datasets and a small smoke-test suite. The shared test suite also includes cross-library comparison tests for clusterers that implement the same protocols, allowing common datasets and validation failures to be checked in one place.

Concrete clustering algorithms are intentionally out of scope for this package. The goal is to provide a portable foundation for future libraries such as `kmeans_clusterer`.

6.34.1 API documentation

Open the ../apis/library_index.html#clustering_protocols link in a web browser.

6.34.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(clustering_protocols(loader)).
```

6.34.3 Testing

To run the library smoke tests and shared comparison tests, load the `tester.lgt` file:

```
| ?- logtalk_load(clustering_protocols(tester)).
```

6.34.4 Test datasets

Several sample datasets are included in the `test_datasets` directory:

- `all_noise.lgt` — A synthetic 2D continuous dataset with 4 examples and 2 continuous attributes (x , y). The examples are well separated, making the dataset useful for all-noise tests where a density-based clusterer should reject every point under a small epsilon radius.
- `bridge_noise.lgt` — A synthetic 2D continuous dataset with 10 examples and 2 continuous attributes (x , y). The examples form two dense blobs connected by two sparse bridge points, making the dataset suitable for testing density-based algorithms that should keep the blobs separate while treating the bridge as noise.
- `dead_component_blobs.lgt` — A synthetic 2D continuous dataset with 6 examples and 2 continuous attributes (x , y). The examples form two tiny ordered blobs sized so an over-specified Gaussian mixture with `first_k` initialization can drive one component fully dead, making the dataset useful for dead-component policy regression tests.
- `duplicate_points.lgt` — A synthetic 2D continuous dataset with 6 examples and 2 continuous attributes (x , y). The examples include repeated coordinates forming a dense local cluster plus one isolated outlier, making the dataset useful for duplicate-point and density-threshold tests.
- `imbalanced_three_modes.lgt` — A synthetic 2D continuous dataset with 9 examples and 2 continuous attributes (x , y). The examples form one dense blob plus two much sparser distant modes, making the dataset useful for imbalanced-cluster and Gaussian mixture stress tests.
- `iris_unlabeled.lgt` — A compact Iris-derived dataset with 9 examples and 4 continuous attributes (`sepal_length`, `sepal_width`, `petal_length`, `petal_width`). It is derived from the classic Iris dataset but intentionally omits species labels so it can be used with unsupervised algorithms.
- `large_two_blobs.lgt` — A synthetic 2D continuous dataset with 100 examples and 2 continuous attributes (x , y). The examples form two dense 5x10 grids that are well separated, making the dataset useful for performance benchmarks that need a larger deterministic density-based clustering workload than the small smoke-test datasets.
- `mixed_profiles.lgt` — A mixed-feature dataset with 6 examples, 2 continuous attributes (`age`, `income`), and 2 discrete attributes (`channel`, `region`). This dataset is intended for clustering algorithms that support both numeric and categorical features.

- `scaling_bands.lgt` — A synthetic 2D continuous dataset with 6 examples and 2 continuous attributes (x, y). The examples form two horizontal bands with large variation along x , making the dataset useful for tests that compare clustering behavior with feature scaling turned on versus off.
- `shopping_profiles.lgt` — A categorical dataset with 6 examples and 4 discrete attributes (channel, region, loyalty, device). The examples form two clear shopping-profile segments suitable for categorical clustering smoke tests.
- `single_blob.lgt` — A synthetic 2D continuous dataset with 6 examples and 2 continuous attributes (x, y). The examples form a single compact blob suitable for one-cluster smoke tests.
- `two_blobs.lgt` — A synthetic 2D continuous dataset with 8 examples and 2 continuous attributes (x, y). The examples form two compact, well-separated blobs suitable for deterministic clustering smoke tests.

6.35 colley_ranker

Colley pairwise preference ranker.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from pairwise preferences, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. Requires well-formed connected pairwise datasets so that the learned rankings remain globally comparable across all ranked items.

6.35.1 API documentation

Open the ../apis/library_index.html#colley_ranker link in a web browser.

6.35.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(colley_ranker(loader)).
```

6.35.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(colley_ranker(tester)).
```

6.35.4 Features

- **Pairwise Preference Learning:** Learns one deterministic rating per item from aggregated pairwise outcomes.
- **Colley Matrix Fidelity:** Solves the standard Colley linear system with diagonal regularization $2 + \text{games}_i$ and off-diagonal entries $-\text{games}_{ij}$.
- **Numerically Hardened Solver:** Uses Gaussian elimination with partial pivoting plus residual checks before accepting the learned ratings.
- **Conservative Ratings:** Produces ratings in the interval $[0, 1]$ with the neutral prior centered at 0.5 .
- **Deterministic Ranking:** Orders candidate items by learned rating with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items, self-preferences, non-positive weights, and disconnected comparison graphs.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.
- **Shared Ranking Infrastructure:** Reuses the `ranking_protocols` helpers for dataset validation, diagnostics, export, and candidate ranking.

6.35.5 Scoring semantics

This implementation aggregates pairwise preferences into matchup totals and then solves the Colley system

```
C r = b
```

where $C_{ii} = 2 + \text{games}_i$, $C_{ij} = -\text{games}_{ij}$ for $i \neq j$, and $b_i = 1 + (\text{wins}_i - \text{losses}_i) / 2$.

The linear system is solved using deterministic Gaussian elimination with partial pivoting. The implementation also verifies that the recovered solution has a small residual and only clamps negligible floating-point boundary noise at 0.0 or 1.0 ; materially invalid out-of-range values are rejected instead of being silently accepted.

The resulting ratings remain close to 0.5 when evidence is weak and move toward 1.0 or 0.0 only when repeated pairwise results justify doing so.

6.35.6 Usage

Learning a ranker

```
% Learn from a pairwise ranking dataset object
| ?- colley_ranker::learn(my_dataset, Ranker).
...

% Learn with an explicit empty options list
| ?- colley_ranker::learn(my_dataset, Ranker, []).
...
```

The current implementation accepts only the empty options list `[]`. Any non-empty options list is rejected.

Inspecting diagnostics

```
% Inspect model and dataset summary metadata
| ?- colley_ranker::learn(my_dataset, Ranker),
      colley_ranker::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
...
```

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- colley_ranker::learn(my_dataset, Ranker),
      colley_ranker::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
...
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- colley_ranker::learn(my_dataset, Ranker),
      colley_ranker::export_to_clauses(my_dataset, Ranker, my_ranker, Clauses).
Clauses = [my_ranker(colley_ranker(...))]
...

% Export to a file
| ?- colley_ranker::learn(my_dataset, Ranker),
      colley_ranker::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.pl').
...
```

6.35.7 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(colley_ranker),
  options(Options),
  dataset_summary(DatasetSummary)
]
```

6.35.8 Ranker representation

The learned ranker is represented by a compound term of the form:

```
colley_ranker(Items, Ratings, Diagnostics)
```

Where:

- **Items:** List of ranked items.
- **Scores:** List of Item-Rating pairs.
- **Diagnostics:** List of metadata terms, including the effective options and dataset summary.

6.35.9 References

1. Colley, W. N. (2002). *Colley's bias free college football ranking method: The Colley Matrix explained*.

6.36 combinations

This library provides predicates for generating and querying combinations over lists. The following categories of predicates are provided:

- **Generation operations** - Predicates for generating combinations.
- **Ordering variants** - Predicates that support an additional order argument (default, lexicographic, or shortlex) for controlling output order.
- **Distinct-value generation** - Predicates for generating combinations while deduplicating equal-valued results.
- **Indexed access** - Predicates for direct access to combinations at specific positions, including distinct combinations.
- **Counting operations** - Predicates for counting combinations and distinct combinations.
- **Random selection** - Predicates for randomly selecting and sampling combinations and distinct combinations.
- **Stepping operations** - Predicates for moving to the next or previous combination in lexicographic order.

Dedicated arrangements, cartesian_products, multisets, permutations, derangements, partitions, and subsequences libraries are also available for focused APIs on related operations. The multisets library is the repetition-allowed counterpart to this library.

6.36.1 API documentation

Open the ../apis/library_index.html#combinations link in a web browser.

6.36.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(combinations(loader)).
```

6.36.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(combinations(tester)).
```

6.37 command_line_options

This library provides command-line parsing predicates where command-line options are defined as objects that import the `command_line_option` category. Predicates are also provided for generating help text.

Parsed options are returned by default as a list of `Name(Value)` terms. A parser option, `output_functor(Functor)`, allows returning options as `Functor(Key, Value)` terms instead.

Adapted with significant changes from the SWI-Prolog `optparse` library by Marcus Uneson. Most documentation, parsing options, and help options retained from the original library. Command line options specification changed to an object-based representation. Parsing simplified, no longer using DCGs. Comprehensive set of tests added, based on documentation examples. Full portability to all supported backends.

6.37.1 API documentation

Open the `../apis/library_index.html#command_line_options` file in a web browser.

6.37.2 Loading

To load all entities in this library, load the `loader.lgt` utility file:

```
| ?- logtalk_load(command_line_options(loader)).
```

6.37.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(command_line_options(tester)).
```

6.37.4 Terminology

The terminology is partly borrowed from Python (see <http://docs.python.org/library/optparse.html#terminology>). *Arguments* are what you provide on the command line and for many Prolog systems show up as a list of atoms `Args` in `current_prolog_flag(argv, Args)`. They can be divided into:

- *Runtime arguments*, which control the Prolog runtime; conventionally terminated by `--`.
- *Options*, which are key-value pairs (with a boolean value possibly implicit) intended to control your program.
- *Positional arguments*, which are what remains after all runtime arguments and options have been removed (with implicit arguments `true/false` for booleans filled in).

Positional arguments are used for mandatory arguments without sensible defaults (e.g., input file names). Options offer flexibility by letting you change a default setting. This library has no notion of mandatory or required options.

6.37.5 Defining option objects

Define command-line options by creating objects that import the `command_line_option` category. Each object represents one command-line option and overrides the predicates to specify its properties.

Create an object that imports `command_line_option` and override the following predicates:

- `key/1`: The key used to identify this option in the result. This predicate must be overridden.
- `short_flags/1`: List of single-character short flags (e.g., `[v]` for `-v`). Default: `[]`.
- `long_flags/1`: List of long flags (e.g., `[verbose]` for `--verbose`). Default: `[]`.
- `type/1`: Option value type. One of `boolean`, `atom`, `integer`, `float`, or `term`. Default: `term`.
- `default/1`: Default value. Optional.
- `meta/1`: Metasyntactic variable name for help text. Default: `''`.
- `help/1`: Help text (atom or list of atoms for pre-broken lines). Default: `''`.

6.37.6 Validating option definitions

Option objects provide two predicates for validating their definitions:

- `check/0`: Validates the option definition, throwing an error if invalid. The validation checks:
 - The `key/1` predicate is defined (not just using the default).
 - All short flags are single characters.
 - The type is a known type (supported by the `type` library).
 - If a default value is defined and the type is not `term`, the default value matches the declared type.
- `valid/0`: Succeeds deterministically if the option definition is valid, fails otherwise. This is useful when you want to check validity without handling exceptions.

The `parse/4-5` and `help/2-3` predicates automatically call `check/0` on all option objects before processing, ensuring invalid definitions are caught early with clear error messages.

Example:

```
| ?- verbose_option::check.
yes

| ?- verbose_option::valid.
yes
```

6.37.7 Example option objects

```
:- object(verbose_option,
    imports(command_line_option)).

    key(verbose).
    short_flags([v]).
    long_flags([verbosity]).
    type(integer).
    default(2).
    meta('V').
    help('verbosity level, 1 <= V <= 3').

:- end_object.

:- object(mode_option,
    imports(command_line_option)).

    key(mode).
    short_flags([m]).
    long_flags([mode]).
    type(atom).
    default('SCAN').
    help(['data gathering mode, one of',
        '  SCAN: do this',
        '  READ: do that',
        '  MAKE: fabricate numbers',
        '  WAIT: do nothing']).

:- end_object.

:- object(cache_option,
    imports(command_line_option)).

    key(cache).
    short_flags([r]).
    long_flags(['rebuild-cache']).
    type(boolean).
    default(true).
    help('rebuild cache in each iteration').

:- end_object.
```

(continues on next page)

(continued from previous page)

```

:- object(outfile_option,
    imports(command_line_option)).

    key(outfile).
    short_flags([o]).
    long_flags(['output-file']).
    type(atom).
    meta('FILE').
    help('write output to FILE').

:- end_object.

% Configuration parameter without command-line flags
:- object(path_option,
    imports(command_line_option)).

    key(path).
    default('/some/file/path/').

:- end_object.

```

6.37.8 Parsing command-line arguments

Use the `command_line_options` object `parse/4-5` predicates:

```

| ?- command_line_options::parse(
    [verbose_option, mode_option, cache_option, outfile_option],
    ['-v', '5', '-m', 'READ', '-ooutput.txt', 'input.txt'],
    Options,
    PositionalArguments
).

Options = [verbose(5), mode('READ'), cache(true), outfile('output.txt')],
PositionalArguments = ['input.txt'].

```

6.37.9 Generating help text

Use the `command_line_options::help/2` predicate for default formatting:

```

| ?- command_line_options::help([verbose_option, mode_option, cache_option], Help).

```

Or use `command_line_options::help/3` with custom help options:

```

| ?- command_line_options::help([verbose_option, mode_option], Help, [line_width(100)]).

```

6.37.10 Help options

The `help/3` predicate supports the following help options:

- `line_width(Width)`: Maximum line width for help text. Default: 80.
- `min_help_width(Width)`: Minimum width for help text column. Default: 40.
- `break_long_flags(Boolean)`: If true, break long flags across multiple lines. Default: false.
- `suppress_empty_meta(Boolean)`: If true (default), suppress empty metasyntactic variables in help output.

6.37.11 Parse options

The `parse/5` predicate supports the following parse options:

- `output_functor(Functor)`: When defined, options are returned as `Functor(Key, Value)` terms instead of `Key(Value)` terms. No default.
- `duplicated_flags(Keep)`: How to handle duplicate options. One of `keepfirst`, `keeplast`, `keepall`. Default: `keeplast`.
- `allow_empty_flag_spec(Boolean)`: If true (default), options without flags are allowed (useful for configuration parameters). Set to false to raise errors on empty flags.

6.37.12 Notes and tips

- The default type is `term`, which subsumes `integer`, `float`, and `atom`. However, always specifying types is recommended for reliable parsing and clear error messages.
- `-sbar` is taken as `-s bar`, not as `-s -b -a -r` (no clustering).
- `-s=foo` is disallowed.
- Duplicate flags default to `keeplast` (controllable via the `duplicated_flags` parse option).
- Unknown flags (not in the specification) will raise errors.

6.38 core

This library consists of built-in entities.

6.38.1 API documentation

See this Handbook “Objects”, “Protocols”, and “Categories” sections.

6.38.2 Loading

All entities in this library are automatically loaded at Logtalk startup.

6.38.3 Testing

Tests for this library can be found in the `tests/logtalk/entities` directory.

6.39 copeland_ranker

Copeland pairwise preference ranker. Ranks each item by its number of matchup wins minus losses after aggregating weighted pairwise preferences per observed opponent pair.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from pairwise preference judgments, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The training dataset must declare each ranked item once, use only declared items in preferences, assign positive weights to preferences between distinct items, and induce a connected undirected comparison graph.

6.39.1 API documentation

Open the ../apis/library_index.html#copeland_ranker link in a web browser.

6.39.2 Loading

To load this library, load the loader `.lgt` file:

```
| ?- logtalk_load(copeland_ranker(loader)).
```

6.39.3 Testing

To test this library predicates, load the tester `.lgt` file:

```
| ?- logtalk_load(copeland_ranker(tester)).
```

6.39.4 Features

- **Pairwise Preference Learning:** Learns one deterministic score per item from pairwise preference datasets.
- **Portable Copeland Scoring:** Computes Copeland scores from aggregated head-to-head outcomes using only standard Logtalk library predicates.
- **Integer Score Semantics:** Learned scores are restricted to integers, because each observed aggregated matchup contributes exactly +1, -1, or 0 to an item's total score.

- **Deterministic Ranking:** Orders candidate items by learned score with deterministic tie-breaking using the standard term order of the item identifiers after sorting by descending score.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items in preferences, self-preferences, non-positive weights, and disconnected comparison graphs.
- **Training Diagnostics:** Learned rankers include dataset summary metadata that can be accessed using the `diagnostics/2` predicate.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.
- **Shared Ranking Infrastructure:** Uses the common `ranking_protocols` helper predicates for option processing, dataset validation, diagnostics, export, and candidate ranking.

6.39.5 Scoring semantics

This implementation uses a Copeland score defined over aggregated observed matchups. For each unordered pair of observed opponents, the preference data is aggregated into total wins for the left and right item. The item with the higher aggregated total receives +1 for that matchup, the item with the lower total receives -1, and both receive 0 when the aggregated totals tie.

Only observed opponent pairs contribute to the learned scores. Unobserved pairs are ignored rather than treated as implicit ties. This makes the implementation a sparse-data Copeland variant suitable for incomplete pairwise datasets.

Because each observed matchup contributes only +1, -1, or 0, every learned Copeland score is an integer. The ranker validation logic enforces this invariant when consuming serialized or exported ranker terms so that malformed non-integer score payloads are rejected instead of silently accepted.

6.39.6 Usage

Learning a ranker

```
% Learn from a pairwise ranking dataset object
| ?- copeland_ranker::learn(my_dataset, Ranker).
...

% Learn with an explicit empty options list
| ?- copeland_ranker::learn(my_dataset, Ranker, []).
...
```

The current implementation accepts only the empty options list `[]`. Any non-empty options list is rejected.

Inspecting diagnostics

```
% Inspect model and dataset summary metadata
| ?- copeland_ranker::learn(my_dataset, Ranker),
    copeland_ranker::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
...
```

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- copeland_ranker::learn(my_dataset, Ranker),
    copeland_ranker::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
...
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- copeland_ranker::learn(my_dataset, Ranker),
    copeland_ranker::export_to_clauses(my_dataset, Ranker, my_ranker, Clauses).
Clauses = [my_ranker(copeland_ranker(...))]
...

% Export to a file
| ?- copeland_ranker::learn(my_dataset, Ranker),
    copeland_ranker::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.pl').
...
```

6.39.7 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(copeland_ranker),
  options(Options),
  dataset_summary(DatasetSummary)
]
```

Where:

- `model(copeland_ranker)` identifies the learning algorithm that produced the ranker.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `dataset_summary(DatasetSummary)` stores a summary list describing the validated training dataset.

The current `dataset_summary/1` payload has the form:

```
[
  items(NumberOfItems),
  preferences(NumberOfPreferences),
  connected_components(NumberOfComponents),
  isolated_items(IsolatedItems)
]
```

Use the `ranking_protocols` `diagnostic/2` and `ranker_options/2` helper predicates when you only need a single metadata term or the effective options.

6.39.8 Options

The current `learn/3` implementation does not define any user options beyond the default empty list. Non-empty options lists are rejected.

6.39.9 Ranker representation

The learned ranker is represented by a compound term of the form:

```
copeland_ranker(Items, Scores, Diagnostics)
```

Where:

- *Items*: List of ranked items.
- *Scores*: List of Item-Score pairs.
- *Diagnostics*: List of metadata terms, including the effective options and dataset summary.

The *Scores* payload is expected to contain integer Copeland scores only. This restriction is not just an implementation preference but a direct consequence of the scoring semantics: each observed aggregated matchup changes an item's score by one of the three integer values +1, -1, or 0.

When exported using `export_to_clauses/4` or `export_to_file/4`, this ranker term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.39.10 References

1. Copeland, A. H. (1951). A reasonable social welfare function. *Seminar on Mathematics in the Social Sciences, University of Michigan*.

6.40 coroutining

The coroutining object provides a portable abstraction over how common coroutining predicates are made available by the supported backend Prolog systems (ECLiPSe, SICStus Prolog, SWI-Prolog, Trealla Prolog, XVM, and YAP) that provide them. Partial support for XSB is provided (the predicate `frozen/2` is not available and calls to it fail).

Calls to the library predicates are inlined when compiled with the `optimize` flag turned on. In this case, there is no overhead compared with calling the abstracted predicates directly.

See also the `dif` library.

6.40.1 API documentation

Open the ../apis/library_index.html#coroutining link in a web browser.

6.40.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(coroutining(loader)).
```

6.40.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(coroutining(tester)).
```

6.40.4 Usage

Load this library from your application loader file. To call the coroutining predicates using implicit message-sending, add the following directive to any object or category calling the predicates (adjust the list to the predicates actually called):

```
:- uses(coroutining, [
    dif/2, dif/1, freeze/2, frozen/2, when/2
]).
```

6.41 crs_projections

This library provides a `crs_projections_protocol` protocol and a `crs_projections` object for working with a small set of common coordinate reference systems and transforming coordinates between them.

Both 2D and explicit 3D WGS84 geodetic coordinates are supported.

6.41.1 API documentation

Open the ../apis/library_index.html#crs_projections link in a web browser.

6.41.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(crs_projections(loader)).
```

6.41.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(crs_projections(loader)).
```

6.41.4 Supported coordinate reference systems

The library currently supports:

- wgs84 geographic coordinates represented as `geographic(Latitude,Longitude)` in degrees
- wgs84_3d geographic coordinates represented as `geographic(Latitude,Longitude,EllipsoidalHeight)` in degrees and meters
- web_mercator projected coordinates represented as `projected(X,Y)` in meters
- world_mercator projected coordinates represented as `projected(X,Y)` in meters
- ecef geocentric coordinates represented as `ecef(X,Y,Z)` in meters
- enu(OriginCoordinate) local tangent-plane coordinates represented as `enu(East,North,Up)` in meters
- lambert_azimuthal_equal_area(OriginCoordinate) equal-area projected coordinates represented as `projected(X,Y)` in meters
- utm(Zone,Hemisphere) projected coordinates represented as `grid(Easting,Northing)` in meters

EPSG helpers are available for these reference systems:

- 4326 for wgs84
- 4979 for wgs84_3d
- 3857 for web_mercator
- 3395 for world_mercator
- 4978 for ecef
- 32601-32660 for northern-hemisphere UTM zones
- 32701-32760 for southern-hemisphere UTM zones

Additional metadata helpers are provided for querying CRS names, kinds, coordinate units, and dimensions.

Additional projection helpers are provided for working directly with local ENU tangent-plane coordinates and Lambert azimuthal equal-area coordinates.

6.41.5 Usage

Load the library:

```
| ?- logtalk_load(crs_projections(loader)).
```

Validate CRS terms and look up EPSG codes:

```
| ?- crs_projections::valid_crs(utm(29, north)).  
| ?- crs_projections::crs_name(wgs84_3d, Name).  
| ?- crs_projections::crs_name(ecef, Name).
```

(continues on next page)

(continued from previous page)

```
| ?- crs_projections::crs_kind(world_mercator, Kind).
| ?- crs_projections::crs_units(wgs84, Units).
| ?- crs_projections::crs_units(wgs84_3d, Units).
| ?- crs_projections::crs_dimensions(ecef, Dimensions).
| ?- crs_projections::crs_dimensions(wgs84_3d, Dimensions).
| ?- crs_projections::crs_epsg(web_mercator, EPSG).
| ?- crs_projections::crs_epsg(wgs84_3d, EPSG).
| ?- crs_projections::epsg_crs(32629, CRS).
| ?- crs_projections::epsg_crs(4979, CRS).
| ?- crs_projections::valid_crs(enu(geographic(38.7223, -9.1393))).
| ?- crs_projections::valid_crs(lambert_azimuthal_equal_area(geographic(38.7223, -9.1393))).
```

Infer the native UTM CRS for a WGS84 coordinate:

```
| ?- crs_projections::utm_zone(geographic(38.7223, -9.1393), Zone).
| ?- crs_projections::utm_crs(geographic(38.7223, -9.1393), CRS).
```

Transform WGS84 coordinates to projected systems and back:

```
| ?- crs_projections::transform(wgs84, web_mercator, geographic(38.7223, -9.1393),
↪Projected).
| ?- crs_projections::transform(wgs84, wgs84_3d, geographic(38.7223, -9.1393), Geodetic3D).
| ?- crs_projections::transform(wgs84_3d, ecef, geographic(38.7223, -9.1393, 123.45),
↪Geocentric3D).
| ?- crs_projections::transform(ecef, wgs84_3d, Geocentric3D, Geodetic3D).
| ?- crs_projections::transform(web_mercator, wgs84, Projected, Coordinate).
| ?- crs_projections::transform(wgs84, world_mercator, geographic(38.7223, -9.1393),
↪Mercator).
| ?- crs_projections::transform(wgs84, ecef, geographic(38.7223, -9.1393), Geocentric).
| ?- crs_projections::transform(wgs84_3d, enu(geographic(38.7223, -9.1393, 100.0)),
↪geographic(38.7223, -9.1393, 125.0), Local3D).
| ?- crs_projections::transform(enu(geographic(38.7223, -9.1393, 100.0)), wgs84_3d, Local3D,
↪Geodetic3D).
| ?- crs_projections::transform(wgs84, enu(geographic(38.7223, -9.1393)), geographic(38.7233,
↪-9.1383), Local).
| ?- crs_projections::transform(wgs84, lambert_azimuthal_equal_area(geographic(38.7223, -9.
↪1393)), geographic(38.7233, -9.1383), EqualArea).
| ?- crs_projections::transform(wgs84, utm(29, north), geographic(38.7223, -9.1393), UTM).
| ?- crs_projections::transform(utm(29, north), wgs84, UTM, Coordinate).
```

Use direct helper predicates for local tangent-plane and equal-area workflows:

```
| ?- crs_projections::local_tangent_plane(geographic(38.7223, -9.1393), geographic(38.7233, -
↪9.1383), Local).
| ?- crs_projections::local_tangent_plane_inverse(geographic(38.7223, -9.1393), Local,
↪Coordinate).
| ?- crs_projections::lambert_azimuthal_equal_area(geographic(38.7223, -9.1393),
↪geographic(38.7233, -9.1383), EqualArea).
| ?- crs_projections::lambert_azimuthal_equal_area_inverse(geographic(38.7223, -9.1393),
↪EqualArea, Coordinate).
```

6.41.6 Notes

The `utm_zone/2`, `utm_crs/2`, and UTM transformation predicates only succeed for latitudes in the `[-80.0, 84.0[` range, matching the standard UTM coverage.

Transformations between projected coordinate reference systems are implemented by converting through `wgs84`.

When converting from `wgs84` to `wgs84_3d`, the ellipsoidal height defaults to `0.0`. When converting from 2D projected or geographic coordinate reference systems to `wgs84_3d`, the resulting ellipsoidal height is also `0.0` because the source coordinate reference system does not carry height information.

Transforms between `ecef` and projected coordinate reference systems are also implemented through `wgs84`.

Transforms involving `enu(OriginCoordinate)` preserve the local Up component when converting directly to or from `ecef` and `wgs84_3d`.

6.42 csv

The `csv` library provides predicates for reading and writing CSV files and streams:

<https://www.rfc-editor.org/rfc/rfc4180.txt>

The main object, `csv/4`, is a parametric object allowing passing options for the handling of the header of the file, the fields separator, the handling of double-quoted fields, and comments handling. When comments handling is `true`, lines starting with the `#` character are skipped when reading files and streams.

The `csv/3` parametric object is kept for backward compatibility and extends `csv/4` by setting the comments option to `false`. The `csv` object extends the `csv/4` parametric object using default option values.

The library also includes predicates to guess the separator and guess the number of columns in a given CSV file.

Files and streams can be read into a list of rows (with each row being represented by a list of fields) or asserted using a user-defined dynamic predicate. Reading can be done by first loading the whole file (using the `read_file/2-3` predicates) into memory or line by line (using the `read_file_by_line/2-3` predicates). Reading line by line is usually the best option for parsing large CSV files.

Data can be saved to a CSV file or stream by providing the object and predicate for accessing the data plus the name of the destination file or the stream handle or alias.

6.42.1 API documentation

Open the [../apis/library_index.html#csv](http://localhost:8080/..../apis/library_index.html#csv) link in a web browser.

6.42.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(csv(loader)).
```

6.42.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(csv(tester)).
```

6.42.4 Usage

The `csv(Header, Separator, IgnoreQuotes, Comments)` parametric object allows passing the following options:

1. Header: possible values are `missing`, `skip`, and `keep`.
2. Separator: possible values are `comma`, `tab`, `semicolon`, and `colon`.
3. IgnoreQuotes: possible values are `true` to ignore double quotes surrounding field data and `false` to preserve the double quotes.
4. Comments: possible values are `false` (default) and `true` to skip lines starting with `#`.

The `csv` object uses the default values `keep`, `comma`, `false`, and `false`.

When writing CSV files or streams, set the `quoted fields` option to `false` to write all non-numeric fields double-quoted (i.e., escaped).

The library objects can also be used to guess the separator used in a CSV file if necessary. For example:

```
| ?- csv::guess_separator('test_files/crlf_ending.csv', Separator).
Is this the proper reading of a line of this file (y/n)? [aaa,bb,ccc]
|> y.

Separator = comma ?
```

This information can then be used to read the CSV file returning a list of rows:

```
| ?- csv(keep, comma, true)::read_file('test_files/crlf_ending.csv', Rows).

Rows = [[aaa,bbb,ccc],[zzz,yyy,xxx]] ?

| ?- csv(keep, comma, true, true)::read_file('test_files/comments.csv', Rows).

Rows = [[aaa,bbb,ccc],[zzz,yyy,xxx]] ?
```

Alternatively, the CSV data can be saved using a public and dynamic object predicate (that must be previously declared). For example:

```
| ?- assertz(p(_,_,_)), retractall(p(_,_,_)).
yes

| ?- csv(keep, comma, true)::read_file('test_files/crlf_ending.csv', user, p/3).
yes

| ?- p(A,B,C).

A = aaa
B = bbb
```

(continues on next page)

(continued from previous page)

```
C = ccc ? ;

A = zzz
B = yyy
C = xxx
```

Given a predicate representing a table, the predicate data can be written to a file or stream. For example:

```
| ?- csv(keep, comma, true)::write_file('output.csv', user, p/3).
yes
```

leaving the content just as the original file thanks to the use of true for the IgnoreQuotes option:

```
aaa,bbb,ccc
zzz,yyy,xxx
```

Otherwise:

```
| ?- csv(keep, comma, false)::write_file('output.csv', user, p/3).
yes
```

results in the following file content:

```
"aaa","bbb","ccc"
"zzz","yyy","xxx"
```

The guess_arity/2 method, to identify the arity, i. e. the number of fields or columns per record in a given CSV file, for example:

```
| ?- csv(keep, comma, false)::guess_arity('test_files/crlf_ending.csv', Arity).
Is this the proper reading of a line of this file (y/n)? [aaa,bbb,ccc]
|> y.

Arity = 3
```

6.43 cuid2

This library generates random Cuid2 identifiers:

<https://github.com/paralleldrive/cuid2>

By default, identifiers are represented as atoms with 24 symbols and use a lowercase alphanumeric alphabet:

```
abcdefghijklmnopqrstuvwxyz0123456789
```

Custom size, alphabet, and representation (atoms, lists of characters, or lists of character codes) are supported using a parametric object.

The generation of random identifiers uses the /dev/urandom random number generator when available. This includes macOS, Linux, *BSD, and other POSIX operating systems. On Windows, a pseudo-random generator is used, randomized using the current wall time.

See also the ids, nanoid, ksuid, snowflakeid, uuid, and ulid libraries.

6.43.1 API documentation

Open the `../..apis/library_index.html#cuid2` link in a web browser.

6.43.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(cuid2(loader)).
```

6.43.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(cuid2(tester)).
```

6.43.4 Usage

To generate an identifier using the default configuration:

```
| ?- cuid2::generate(Cuid2).
Cuid2 = 'k4f9mdd51t2r9y53i8h4j1bz'
yes
```

To generate a 10-symbol identifier represented as a list of characters:

```
| ?- cuid2(chars, 10, 'abcdef012345')::generate(Cuid2).
Cuid2 = ['a', '2', 'f', 'e', '5', 'c', '1', 'd', '0', 'b']
yes
```

6.44 cusum_anomaly_detector

CUSUM (Cumulative Sum Control Chart) anomaly detector for continuous sequence-like datasets. This is a statistical anomaly-detection method based on a two-sided CUSUM control chart. Declared continuous attributes are interpreted as ordered monitoring steps.

The library implements the `anomaly_detector_protocol` defined in the `anomaly_detection_protocols` library. It learns a detector from a continuous dataset, computes anomaly scores for new instances, predicts normal or anomaly, and exports learned detectors as clauses or files.

Datasets are represented as objects implementing the `anomaly_dataset_protocol` protocol from the `anomaly_detection_protocols` library. Declared continuous attributes are interpreted as ordered monitoring steps in a sequence. See the `cusum_anomaly_detector/tests.lgt` file for example datasets.

6.44.1 API documentation

Open the ../apis/library_index.html#cusum_anomaly_detector link in a web browser.

6.44.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(cusum_anomaly_detector(loader)).
```

6.44.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(cusum_anomaly_detector(tester)).
```

6.44.4 Features

- **Statistical method:** implements anomaly detection based on a two-sided CUSUM control chart, using learned per-step population means and standard deviations for continuous attributes.
- **Ordered sequence interpretation:** declared continuous attributes are treated as ordered monitoring steps. For each known step value x_t , the library computes $z_t = (x_t - \mu_t) / \sigma_t$ and updates the positive and negative CUSUM recurrences along that attribute order. The learned detector stores a precomputed attribute schema so that this ordering does not need to be rebuilt for every scoring call.
- **CUSUM recurrences:** the positive and negative cumulative sums are updated as $C+_t = \max(0, C+_{t-1} + z_t - k)$ and $C-_t = \max(0, C-_{t-1} - z_t - k)$, where k is the learn-time allowance. The raw anomaly score is the maximum excursion over all positive and negative cumulative sums.
- **Continuous features only:** accepts datasets whose declared attributes are all continuous.
- **Baseline training selection:** supports learn-time `baseline_class_values(ClassValues)` and `baseline_selection_policy(Policy)` options. The default baseline class values are `[normal]`. The default reject policy throws an error if any non-baseline training example is found. The filter policy removes non-baseline examples before fitting the baseline statistics.
- **Missing-value tolerant:** ignores missing values when fitting per-step statistics and skips them during scoring. Queries must still provide at least one known value.
- **Bounded scoring:** maps the raw CUSUM excursion to $[0.0, 1.0)$ using $\text{Score} = \text{Raw} / (1 + \text{Raw})$.
- **CUSUM control parameters:** supports learn-time `allowance/1` and `decision_interval/1` options. The default `allowance(0.5)` and `decision_interval(5.0)` correspond to a common standardized CUSUM setup.
- **Default threshold:** the default `anomaly_threshold(0.8333333333333334)` corresponds to the default raw decision interval `5.0`. If a custom `decision_interval/1` is passed to `learn/3` without an explicit `anomaly_threshold/1`, the stored anomaly threshold is derived automatically as $H / (1 + H)$.
- **Learn-time control parameters:** `allowance/1` and `decision_interval/1` are recorded in the learned detector and reused for subsequent scoring and prediction. Passing them to `predict/4` does not override the learned values. Only `anomaly_threshold/1` can be overridden at predict time.

- **All-missing queries rejected:** scoring and prediction throw a `domain_error(non_empty_known_values, AttributeNames)` exception when every declared step is missing in the query.
- **Featureless datasets rejected:** datasets must declare at least one continuous feature; otherwise `learn/2-3` throws a `domain_error(non_empty_features, Dataset)` exception.
- **Detector export:** learned detectors can be exported as predicate clauses.
- **Explicit validation and diagnostics:** supports the shared `check_anomaly_detector/1`, `valid_anomaly_detector/1`, `diagnostics/2`, `diagnostic/2`, and `anomaly_detector_options/2` predicates.

6.44.5 Options

The following options are supported by the public API:

- `anomaly_threshold(Threshold)`: Threshold for `predict/3-4` (default: `0.8333333333333334`)
- `allowance(Allowance)`: Learn-time CUSUM allowance k (default: `0.5`)
- `baseline_class_values(ClassValues)`: Learn-time class labels that are admissible for baseline fitting (default: `[normal]`)
- `baseline_selection_policy(Policy)`: Learn-time handling of examples whose class is not listed in `baseline_class_values/1`. Supported values are `reject` and `filter` (default: `reject`)
- `decision_interval(DecisionInterval)`: Learn-time raw decision interval H (default: `5.0`). If no explicit `anomaly_threshold/1` is passed to `learn/3`, the stored threshold is derived from this value as $H / (1 + H)$.

6.44.6 Detector representation

The learned detector is represented by default as:

```
cusum_detector(TrainingDataset, AttributeSchema, Encoders, Diagnostics)
```

Where:

- `TrainingDataset`: training dataset object identifier
- `AttributeSchema`: precomputed attribute ordering metadata used to validate and reorder query step values efficiently during scoring
- `Encoders`: list of `cusum_encoder(Attribute, Mean, Scale)` records
- `Diagnostics`: learned metadata terms including `model/1`, `training_dataset/1`, `attribute_names/1`, `feature_count/1`, `example_count/1`, and `options/1`. The `example_count/1` value is the effective number of training examples after applying the selected baseline selection policy.

When exported using `export_to_clauses/4` or `export_to_file/4`, this detector term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.44.7 Notes

Scoring has three stages. First, the detector computes one standardized deviation $z_t = (x_t - \mu_t) / \sigma_t$ for each known monitoring step. Second, those deviations are processed sequentially using the positive and negative CUSUM recurrences with the learned allowance/1 value. Third, the maximum raw excursion is mapped to the interval $[0.0, 1.0)$ using $\text{Score} = \text{Raw} / (1 + \text{Raw})$.

The allowance/1 option changes the CUSUM update rule itself by controlling how much drift must accumulate before the chart grows. Larger values make the detector less sensitive to small shifts. The decision_interval/1 option does not change scoring; it only affects the default threshold stored when learning a detector.

The baseline_class_values/1 option declares which dataset class labels are admissible for baseline fitting. The baseline_selection_policy/1 option then controls what happens when other labels are present in the training data. The default reject policy raises a `domain_error(baseline_only_training_data, Dataset)` exception when any non-baseline example is found. The filter policy removes non-baseline examples before fitting and raises a `domain_error(non_empty_baseline_training_data, Dataset)` exception if no training examples remain after filtering.

Attributes with zero observed dispersion are assigned a fallback scale of 1.0. This keeps the detector well-defined for singleton datasets or constant steps while still yielding zero score for matching values and positive scores for deviating values.

6.45 datalog

This library provides a portable Datalog and incremental rule engine for educational purposes and non-demanding applications. This library is work-in-progress and changes, possibly breaking backwards compatibility, are to be expected.

6.45.1 API documentation

Open the ../apis/library_index.html#datalog link in a web browser.

6.45.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(datalog(loader)).
```

6.45.3 Testing

To run the library unit tests, load the `tester.lgt` file:

```
| ?- logtalk_load(datalog(tester)).
```

6.45.4 Scope

- positive and stratified negation rules (neg/1 body literals)
- conservative aggregates (agg(Op, Template, Goals, Result)) with lower-strata dependencies (Op in count, sum, min, max)
- safe rules (Head and negated literal variables must appear in positive Body literals)
- fixpoint materialization
- incremental updates using support-count based propagation
- simple explanations for derived facts
- rule management (add_rule/3, remove_rule/1)
- transaction support (begin/0, commit/0, rollback/0)
- stratum introspection (predicate_stratum/3, strata/1)
- rule body normalization (positive ground literals first, then positive non-ground, then aggregates, then negation)

6.45.5 Public API overview

- lifecycle: clear/0, load_program/1, materialize/0
- rule management: add_rule/3, remove_rule/1, rules/1
- fact management: assert_fact/1, retract_fact/1, facts/1
- querying and explanations: query/1, query/2, explain/2
- updates: update/3
- transactions: begin/0, commit/0, rollback/0
- stratification introspection: predicate_stratum/3, strata/1

Rules and facts are represented as terms:

- rule(Id, Head, Body) where Body is a list of Literal terms:
 - positive literal: Term
 - negative literal: neg(Term)
 - aggregate literal: agg(Op, Template, Goals, Result) where Op is one of count, sum, min, max
- fact(Fact) for extensional database (EDB) facts

Aggregate notes:

- Aggregate dependencies are required to be in lower strata.
- min and max require at least one matching value; otherwise no fact is derived.
- Incremental updates involving negation or aggregates currently fallback to full rematerialization for correctness.

6.45.6 Basic usage

```
| ?- logtalk_load(datalog(loader)).  
...  
  
| ?- Program = [  
    rule(path_base, path(X,Y), [edge(X,Y)]),  
    rule(path_rec, path(X,Z), [edge(X,Y), path(Y,Z)]),  
    fact(edge(a,b)),  
    fact(edge(b,c))  
],  
    datalog::load_program(Program),  
    datalog::query(path(a,c)).
```

6.45.7 Limitations

This first version intentionally keeps the implementation simple and portable. Negation support is currently limited to stratified programs. Aggregates and cost-based optimization are planned future enhancements. Aggregate support is currently limited to count, sum, min, and max.

6.46 dates

The date object implements some useful calendar date predicates.

It also provides portable predicates for date-time handling, including:

- Unix epoch conversions (`date_time_to_unix/2`, `unix_to_date_time/2`)
- date-time arithmetic with durations (`add_duration/3`, `subtract_duration/3`, `duration_between/3`)
- UTC/local conversion using explicit offsets (`utc_to_local/3`, `local_to_utc/3`)
- date-time formatting with explicit offsets (`format_date_time/4`)
- calendar utilities (`day_of_year/2`, `week_of_year_iso/2`, `weekday/2`)
- reverse and positional calendar utilities (`day_of_year_date/3`, `month_weekday_date/5`)
- date-time normalization and validation (`normalize_date_time/2`, `valid_date_time/1`)
- date-time comparison (`before/2`, `after/2`, `same_instant/2`, `compare_date_time/3`)

The duration predicates accept two duration representations:

- `duration(Days, Hours, Minutes, Seconds)` — a fixed-length duration converted to seconds
- `duration(Years, Months, Days, Hours, Minutes, Seconds)` — a calendar-aware period where the year and month delta is applied first using calendar arithmetic, clamping the day to the last valid day of the resulting month when necessary (e.g. adding one month to January 31 gives February 28 or 29 depending on the year), with the remaining day and time delta applied via fixed-length arithmetic. The `duration_between/3` predicate returns this form when called with a 6-arity skeleton such as `duration(Yr, Mo, Da, Hr, Mi, Se)`.

Date-time values are represented using the `date_time/6` compound term:

- `date_time(Year, Month, Day, Hours, Minutes, Seconds)`

The `format_date_time/4` predicate currently supports the named formats `rfc3339`, `iso8601`, `atom`, `rfc2822`, `rfc5322`, `rss`, `http_date`, `rfc1123`, `unix_date`, `common_log`, `date_short`, `date_medium`, `date_long`, `date_full`, `time_short`, `time_medium`, `time_long`, `time_full`, `date_time_short`, `date_time_medium`, `date_time_long`, and `date_time_full`.

The explicit offset argument is given in seconds. Formats that include numeric offsets require offsets representable in whole minutes. The `http_date` and `rfc1123` formats always normalize the output to GMT.

The style presets are currently English-only presentation formats:

- `date_short` -> 2026-04-08
- `date_medium` -> 8 Apr 2026
- `date_long` -> April 8, 2026
- `date_full` -> Tuesday, April 8, 2026
- `time_short` -> 13:45
- `time_medium` -> 13:45:30
- `time_long` -> 13:45:30 +01:00
- `time_full` -> 1:45:30 PM +01:00
- `date_time_short` -> 2026-04-08 13:45
- `date_time_medium` -> 8 Apr 2026 13:45:30
- `date_time_long` -> April 8, 2026 1:45:30 PM +01:00
- `date_time_full` -> Tuesday, April 8, 2026 1:45:30 PM +01:00

The time object implements some useful time predicates.

Please note that the functionality of these objects depends on the chosen Prolog support for accessing the operating system time and date.

6.46.1 API documentation

Open the [../apis/library_index.html#dates](http://logtalk.org/..../apis/library_index.html#dates) link in a web browser.

6.46.2 Loading

To load all entities in this library, load the `loader.lgt` utility file:

```
| ?- logtalk_load(dates(loader)).
```

6.46.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(dates(tester)).
```

6.47 dates_tz

Bridge library linking the dates and tzif libraries for zone-aware date-time conversions. Provides predicates that convert UTC date-times to civil local date-times and vice-versa, handling DST transitions including ambiguous (fold) and non-existent (gap) times. Also provides cross-zone conversion that routes through UTC as an intermediate representation.

This library requires unbounded integer arithmetic support from the backend Prolog compiler (as does the tzif library it depends on). It is not available on backends where the bounded Prolog flag is true.

6.47.1 API documentation

Open the ../apis/library_index.html#dates_tz link in a web browser.

6.47.2 Loading

To load all entities in this library, load the loader.lgt utility file:

```
| ?- logtalk_load(dates_tz(loader)).
```

6.47.3 Testing

To test this library predicates, load the tester.lgt file:

```
| ?- logtalk_load(dates_tz(tester)).
```

6.47.4 Usage

Zone data must be loaded separately into the tzif cache before calling any dates_tz predicates. For example:

```
:- initialization((
    logtalk_load(dates_tz(loader)),
    tzif::load(file('/usr/share/zoneinfo/America/New_York', 'America/New_York')),
    tzif::load(directory('/usr/share/zoneinfo'))
)).
```

6.47.5 Examples

Convert a UTC date_time/6 compound to the civil local date-time in the named IANA zone. The UTC instant always has exactly one local representation.

```
| ?- dates_tz::utc_to_local_tz(date_time(2024,1,15,12,0,0), 'America/New_York', Local).
Local = date_time(2024,1,15,7,0,0).
```

Convert a civil local date_time/6 compound in the named zone to UTC using strict interpretation. Fails silently if the local time falls in a DST gap (non-existent) or a DST fold (ambiguous). Throws an error if the zone is not cached.

```
| ?- dates_tz::local_to_utc_tz(date_time(2024,1,15,7,0,0), 'America/New_York', UTC).
UTC = date_time(2024,1,15,12,0,0).
```

Convert a civil local date-time to UTC with an explicit resolution mode. The mode can be:

- `strict` — fail unless exactly one interpretation exists
- `first` — prefer the earliest valid interpretation (earliest UTC instant)
- `second` — prefer the latest valid interpretation (latest UTC instant)
- `all` — enumerate all valid interpretations in chronological order (non-deterministic)

Fall-back fold in New York: 2024-11-03 01:30 occurs twice:

```
| ?- dates_tz::local_to_utc_tz_with_resolution(
    date_time(2024,11,3,1,30,0), 'America/New_York', first, UTC).
UTC = date_time(2024,11,3,5,30,0). % EDT interpretation

| ?- dates_tz::local_to_utc_tz_with_resolution(
    date_time(2024,11,3,1,30,0), 'America/New_York', second, UTC).
UTC = date_time(2024,11,3,6,30,0). % EST interpretation
```

Convert a civil local date-time from one IANA zone to another, routing through UTC. Uses strict interpretation for the source zone: fails if the source local time is ambiguous or non-existent. Requires both zones to be cached.

```
| ?- dates_tz::convert_zones(
    date_time(2024,1,15,7,0,0), 'America/New_York', 'Asia/Kathmandu', Result).
Result = date_time(2024,1,15,17,45,0).
```

Convert a civil local date-time between zones with an explicit resolution mode applied to the source zone. Useful when the source local time may be ambiguous due to a DST transition.

```
| ?- dates_tz::convert_zones_with_resolution(
    date_time(2024,11,3,1,30,0), 'America/New_York', first, 'Asia/Kathmandu', Result).
Result = date_time(2024,11,3,11,15,0).
```

6.47.6 Notes

- Date-time values are represented as `date_time(Year, Month, Day, Hour, Minute, Second)` compounds, consistent with the `dates` library.
- UTC offset arithmetic is performed via integer seconds using `date::add_duration/3` and `date::subtract_duration/3`.
- This library does not provide a system clock or current-time predicate. Use the `dates` library's `date::now/6` or equivalent for the current time.
- For performance when converting many instants in the same zone, load all required zones once at startup and keep them in the `tzif` cache.

6.48 dbscan_clusterer

DBSCAN clusterer. Uses deterministic density-based clustering based on epsilon neighborhoods and minimum point counts. Supports continuous attributes only.

The library implements the `clusterer_protocol` defined in the `clustering_protocols` library. It provides predicates for learning a clusterer from a dataset, assigning new instances to clusters, and exporting the learned clusterer as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol from the `clustering_protocols` library.

6.48.1 API documentation

Open the ../apis/library_index.html#dbscan_clusterer link in a web browser.

6.48.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(dbscan_clusterer(loader)).
```

6.48.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(dbscan_clusterer(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(dbscan_clusterer(tester_performance)).
```

6.48.4 Features

- **Density-Based Clustering:** Learns density-connected clusters using epsilon neighborhoods and minimum point counts.
- **Adaptive Neighborhood Indexing:** Builds a low-dimensional epsilon-grid index when it is likely to be cheaper and otherwise falls back to a deterministic metric tree. The metric-tree path uses adaptive leaf sizing, balanced tie-aware subtree splits, selectable exact or heuristic pivot scoring, and lower-allocation range-query traversal.
- **Continuous Datasets:** Accepts datasets containing only continuous attributes.
- **Distance Metrics:** Supports euclidean and manhattan distances.
- **Optional Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Reachable-Core Prediction:** New instances are assigned to the cluster of the nearest reachable core point within the learned epsilon radius; otherwise the atom noise is returned.
- **Portable Export:** Learned clusterers can be exported as clauses or files and reused later.

6.48.5 Options

The following options can be passed to the `learn/3` predicate:

- `epsilon(Epsilon)`: Neighborhood radius used to determine density connectivity. Default is `1.0`.
- `minimum_points(MinimumPoints)`: Minimum number of points in an epsilon neighborhood for a core point. Default is `2`.
- `distance_metric(Metric)`: Distance metric to use. Options: `euclidean` (default) or `manhattan`.
- `feature_scaling(FeatureScaling)`: Whether to standardize continuous attributes before clustering. Options: `on` (default) or `off`.
- `pivot_scoring(PivotScoring)`: Metric-tree pivot scoring strategy. Options: `heuristic` (default, single-pass dispersion scoring with one final sort) or `exact` (more expensive gap-and-range scoring that sorts each candidate profile).

6.48.6 Clusterer representation

The learned clusterer is represented as a compound term with the functor chosen by the user when exporting the clusterer and arity 4. For example:

```
dbscan_clusterer(Encoders, Clusters, Noise, Options)
```

Where:

- `Encoders`: List of continuous attribute encoders storing attribute name, mean, and scale.
- `Clusters`: List of `cluster(Id, CorePoints, BorderPoints)` terms in cluster-id order.
- `Noise`: List of encoded training points classified as noise.
- `Options`: Effective training options used to learn the clusterer.

6.48.7 References

1. Ester, Kriegel, Sander, and Xu (1996) - "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise". KDD, 226-231.

6.49 dependents

The observer and subject categories implement the Smalltalk dependents handling mechanism. This mechanism can be used as an alternative to Logtalk system-wide support for event-driven programming.

6.49.1 API documentation

Open the ../apis/library_index.html#dependents link in a web browser.

6.49.2 Loading

To load all entities in this library, load the `loader.lgt` utility file:

```
| ?- logtalk_load(dependents(loader)).
```

6.50 dequeues

This library implements double-ended queues, dequeues. The queue representation should be regarded as an opaque term and only accessed using this library predicates.

6.50.1 API documentation

Open the ../apis/library_index.html#dequeues link in a web browser.

6.50.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(deques(loader)).
```

6.50.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(deques(tester)).
```

6.50.4 Usage

To create a new deque, use the `new/1` predicate:

```
| ?- deque::new(Deque).  
Deque = ...  
yes
```

Elements can be added and removed both at the front and at the back of the deque using the `push_front/3`, `push_back/3`, `pop_front/3`, and `pop_back/3` predicates. For example:

```
| ?- deque::(as_deque([1,2,3], Deque0), push_front(0, Deque0, Deque)).  
Deque0 = ...,  
Deque = ...  
yes
```

We can also peek at the front and back of the deque using the `peek_front/2` and `peek_back/2` predicates. For example:

```
| ?- deque::(as_deque([1,2,3], Deque), peek_back(Deque, Element)).
Deque = ...,
Element = 3
yes
```

Mapping a closure over all elements of a deque can be done using the `map/2` and `map/3` predicates. For example:

```
| ?- deque::(as_deque([1,2,3], Deque), map(write, Deque)).
123
Deque = ...
yes
```

For details on these and other provided predicates, consult the library API documentation.

6.51 derangements

This library provides predicates for generating and querying derangements over lists. Derangements are full-length permutations where no result element is identical to the input element at the same position. The following categories of predicates are provided:

- **Generation operations** - Predicates for generating derangements.
- **Ordering variants** - Predicates that support an additional order argument (default, lexicographic, or shortlex) for controlling output order.
- **Distinct-value generation** - Predicates for generating derangements while deduplicating equal-valued results.
- **Indexed access** - Predicates for direct access to derangements at specific positions, including distinct derangements.
- **Lexicographic stepping** - Predicates for navigating derangements in lexicographic order.
- **Counting operations** - Predicates for counting derangements and distinct derangements.
- **Random selection** - Predicates for randomly selecting and sampling derangements and distinct derangements.

Dedicated permutations, arrangements, cartesian_products, combinations, multisets, partitions, and subsequences libraries are also available for focused APIs on related operations.

6.51.1 API documentation

Open the ../apis/library_index.html#derangements link in a web browser.

6.51.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(derangements(loader)).
```

6.51.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(derangements(tester)).
```

6.52 dictionaries

This library provides a dictionary (also known as associative array, map, or symbol table) protocol and binary tree, AVL tree, Red-Black tree, splay tree, and 2-3 tree implementations. The different representations of a dictionary should be regarded as opaque terms and only accessed using the library predicates.

6.52.1 API documentation

Open the ../apis/library_index.html#dictionaries link in a web browser.

6.52.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(dictionaries(loader)).
```

6.52.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(dictionaries(tester)).
```

6.52.4 Usage

First, select the dictionary implementation that you want to use. For cases where the number of elements is relatively small and performance is not critical, `bintree` can be a good choice. For other cases, `avltree`, `rbtree`, or `two3tree` are likely better choices. The `splaytree` implementation can be a good choice when recently accessed elements are likely to be accessed again, as it moves accessed elements closer to the root. If you want to compare the performance of the implementations, either define an object alias or use a `uses/2` directive so that you can switch between implementations by simply changing the alias definition or the first argument of the directive. Note that you can switch between implementations at runtime without code changes by using a parameter variable in the first argument of a `uses/2` directive.

Dictionary keys should preferably be ground terms. If the keys contain variables, the user must ensure that any instantiation of those variables when calling this library predicates will not affect the key ordering.

To create a new dictionary, you can use the `new/1` predicate. For example:

```
| ?- avltree::new(Dictionary).
Dictionary = ...
yes
```

You can also create a new dictionary from a list of key-value pairs by using the `as_dictionary/2` predicate. For example:

```
| ?- avltree::as_dictionary([a-1,c-3,b-2], Dictionary).
Dictionary = ...
yes
```

Several predicates are provided for inserting key-value pairs, lookup key-value pairs updating the value associated with a key, and deleting key-value pairs. For example:

```
| ?- avltree::(
    new(Dictionary0),
    insert(Dictionary0, a, 1, Dictionary1),
    update(Dictionary1, a, 2, Dictionary2),
    lookup(a, Value, Dictionary2)
).
Dictionary0 = ...,
Dictionary1 = ...,
Dictionary2 = ...,
Value = 2
yes
```

For details on these and other provided predicates, consult the library API documentation.

6.52.5 Credits

The AVL tree implementation is an adaptation to Logtalk of the `assoc` SWI-Prolog library authored by R.A.O’Keefe, L.Damas, V.S.Costa, Glenn Burgess, Jiri Spitz, and Jan Wielemaker. Additional predicates authored by Paulo Moura.

The Red-Black tree implementation is an adaptation to Logtalk of the `rbtrees` Prolog library authored by Vitor Santos Costa.

The 2-3 tree implementation is a contribution by Michael T. Richter.

6.53 dif

The `dif` object provides a portable abstraction over how the `dif/2` predicate is made available by the supported backend Prolog systems that implement it (B-Prolog, ECLiPSe, SICStus Prolog, SWI-Prolog, Trealla Prolog, XSB, XVM, and YAP).

Calls to the library predicates are inlined when compiled with the `optimize` flag turned on. In this case, there is no overhead compared with calling the abstracted predicate directly.

See also the coroutining library.

6.53.1 API documentation

Open the `../apis/library_index.html#dif` link in a web browser.

6.53.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(dif(loader)).
```

6.53.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(dif(tester)).
```

6.53.4 Usage

Load this library from your application loader file. To call the `dif/1-2` predicates using implicit message-sending, add the following directive to any object or category calling the predicates:

```
:- uses(dif, [  
    dif/2, dif/1  
]).
```

6.54 dimension_reduction_protocols

This library provides protocols and reusable support predicates used in the implementation of machine learning dimension reduction algorithms. Unsupervised datasets are represented as objects implementing the `dimension_reduction_dataset_protocol` protocol. Labeled datasets for supervised reducers such as Linear Discriminant Analysis are represented as objects implementing the `supervised_dimension_reduction_dataset_protocol` protocol. Target-valued datasets for reducers such as partial least squares projections are represented as objects implementing the `target_supervised_dimension_reduction_dataset_protocol` protocol. Dimension reducers are represented as objects importing the `dimension_reducer_common` category.

This library also provides reusable test datasets and a small smoke-test suite.

Concrete dimension reduction algorithms are intentionally out of scope for this package. The goal is to provide a portable foundation for libraries such as `pca_projection`, `truncated_svd_projection`, `random_projection`, `nmf_projection`, `lda_projection`, and `pls_projection`, including shared reducer helpers for common tasks such as feature encoding, projection, reducer validation, diagnostics, export, and pretty-printing.

Exported reducers are serialized protocol-wide as single-argument predicates such as `reducer(Reducer)`, allowing the saved reducer term to be loaded and passed directly to `transform/3`.

All reducers importing the shared category are also expected to record a diagnostics list containing at least `model/1` and `options/1` terms. The shared protocol surface exposes this metadata using the `diagnostics/2`, `diagnostic/2`, and `dimension_reducer_options/2` predicates and validates serialized reducers using `check_dimension_reducer/1`.

6.54.1 API documentation

Open the ../apis/library_index.html#dimension_reduction_protocols link in a web browser.

6.54.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(dimension_reduction_protocols(loader)).
```

6.54.3 Common options

The `dimension_reducer_common` category supports the `feature_scaling/1` option used by importing reducers to control continuous feature normalization before projection:

- `feature_scaling(true)` standardizes each continuous attribute using its training-set mean and standard deviation.
- `feature_scaling(false)` only centers each continuous attribute using its training-set mean.

The current `pca_projection`, `random_projection`, `lda_projection`, and `pls_projection` libraries all define `feature_scaling(true)` as their default.

Reducers can also specialize preprocessing behavior. For example, the `nmf_projection` library only supports `center(false)` because the learned feature representation must remain non-negative.

6.54.4 Testing

To run the library smoke tests, load the `tester.lgt` file:

```
| ?- logtalk_load(dimension_reduction_protocols(tester)).
```

6.54.5 Test datasets

Several sample datasets are included in the `test_datasets` directory:

- `correlated_plane.lgt` — A compact continuous dataset with 8 examples and 3 continuous attributes (`x`, `y`, `z`) where the features are strongly correlated. It is intended for PCA and projection smoke tests.
- `high_dimensional_measurements.lgt` — A small continuous dataset with 10 examples and 6 continuous attributes (`f1` through `f6`). It is intended for testing projected dimensionality and reducer output shape.
- `low_rank_rectangular.lgt` — A compact continuous dataset with 4 examples and 3 continuous attributes whose third feature is the sum of the first two. It is intended for testing matrix-rank truncation and singular-value-based reducers such as `truncated_svd_projection`.
- `labeled_measurements.lgt` — A compact labeled continuous dataset with 9 examples, 4 continuous attributes (`length`, `width`, `height`, `weight`), and 3 class labels (`alpha`, `beta`, `gamma`). It is intended for testing supervised reducers such as `lda_projection`.
- `parts_based_measurements.lgt` — A compact non-negative continuous dataset with 5 examples and 4 continuous attributes (`f1` through `f4`) built from two latent additive parts. It is intended for testing parts-based reducers such as `nmf_projection`.

- `target_latent_measurements.lgt` — A compact target-valued continuous dataset with 6 examples, 4 continuous attributes (f1 through f4), and a numeric target attribute (score) influenced by two latent directions. It is intended for testing target-supervised reducers such as `pls_projection`.
- Target-supervised smoke coverage is currently provided by the `target_measurements` object in the library smoke tests, which exercises the `target_supervised_dimension_reduction_dataset_protocol` contract.

6.55 eclat_pattern_miner

Eclat frequent itemset miner for transaction datasets. Normalizes dataset transaction identifiers to internal ascending ordinals, builds vertical tidsets for frequent singleton items, and recursively extends them by lexicographic suffix joins and tidset intersections.

This library depends on the `frequent_pattern_mining_protocols` support library, implements the generic `pattern_miner_protocol` defined in the `pattern_mining_protocols` core library.

Requires a dataset implementing `transaction_dataset_protocol` with transactions represented as canonical sorted lists of unique declared items. External transaction identifiers are treated as opaque metadata and normalized to internal ordinals before tidset construction.

6.55.1 API documentation

Open the ../apis/library_index.html#eclat_pattern_miner link in a web browser.

6.55.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(eclat_pattern_miner(loader)).
```

6.55.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(eclat_pattern_miner(tester)).
```

6.55.4 Features

- **Vertical Tidsets:** Represents each frequent extension by the sorted list of transaction identifiers containing it.
- **Transaction Ordinal Normalization:** Treats external transaction identifiers as opaque metadata and normalizes them to internal ascending ordinals before building vertical tidsets.
- **Depth-First Mining:** Extends frequent prefixes recursively using tidset intersections.
- **Canonical Transactions:** Validates that transactions are sorted, duplicate-free, and restricted to declared items.
- **Flexible Support Thresholds:** Supports relative minimum support and absolute minimum support count. When both are given, the absolute-count threshold takes precedence.

- **Model Export:** Mined pattern collections can be exported as predicate clauses or written to a file.

6.55.5 Options

The `mine/3` predicate accepts the following options:

- `minimum_support/1`: Relative minimum support threshold in the interval $]0.0, 1.0]$. The default is `0.5`.
- `minimum_support_count/1`: Absolute minimum support count. When both support options are provided, this option takes precedence.
- `maximum_pattern_length/1`: Maximum itemset length to mine. The default is `1000`, which is effectively capped by the longest transaction in the dataset.
- `minimum_pattern_length/1`: Minimum itemset length retained in the mined result. The default is `1`.

6.55.6 Pattern miner representation

The mined pattern miner result is represented by a compound term with the functor chosen by the implementation and arity 3. For example:

```
eclat_pattern_miner(ItemDomain, Patterns, Options)
```

Where:

- `ItemDomain`: Canonical sorted list of declared dataset items.
- `Patterns`: List of `itemset(Items, SupportCount)` terms ordered first by pattern length and then lexicographically.
- `Options`: Effective mining options used to mine the frequent itemsets.

6.55.7 References

1. Zaki, M. J. (2000) - "Scalable algorithms for association mining".

6.56 edcg

This library provides a Logtalk port of Peter Van Roy's extended DCG implementation. For full documentation on EDCGs, see:

<https://webperso.info.ucl.ac.be/~pvr/edcg.html>

This Logtalk version defines a hook object, `edcg`. Source files defining EDCGs must be compiled using the compiler option `hook(edcg)`:

```
| ?- logtalk_load(source, [hook(edcg)]).
```

Alternatively, the following directive can be added at the beginning of the source file containing the EDCGs:

```
:- set_logtalk_flag(hook, edcg).
```

The hook object automatically adds the EDCGs -->> infix operator scoped to the source file.

This port has simplified by copying and then modifying Michael Hendricks's edcg repo at:

<https://github.com/mndrix/edcg>

A notable difference is that Michael's version declares Peter's original predicates for declaring accumulators and predicates using the hidden arguments as multifile predicates. But this is risky as two independent EDCGs may use e.g. the same accumulator names and introduce conflicts. The Logtalk version uses instead the edcg hook object internal state to temporarily save those predicates in order to parse the corresponding EDCGs.

6.56.1 API documentation

Open the `../apis/library_index.html#edcg` link in a web browser.

6.56.2 Loading

To load all entities in this library, load the `loader.lgt` utility file:

```
| ?- logtalk_load(edcg(loader)).
```

6.56.3 Testing

To test this library predicates, load the `tester.lgt` file of the edcgs example:

```
| ?- logtalk_load(edcgs(tester)).
```

6.56.4 Usage

Follows the usage documentation written by Michael Hendricks (with a contribution from Peter Ludemann), used here with permission, with the necessary changes for the Logtalk port.

```
% declare accumulators
acc_info(adder, X, In, Out, integer::plus(X,In,Out)).

% declare predicates using these hidden arguments
pred_info(len,0,[adder,dcg]).
pred_info(increment,0,[adder]).

increment -->>
    % add one to the accumulator
    [1]:adder.

len(Xs,N) :-
    len(0,N,Xs,[]).

len -->>
    % 'dcg' accumulator has an element
    [_,
```

(continues on next page)

(continued from previous page)

```
!,
% increment the 'adder' accumulator
increment,
len.
len -->>
[].
```

6.56.5 Introduction

DCG notation gives us a single, hidden accumulator. Extended DCG notation (implemented by this library) lets predicates have arbitrarily many hidden accumulators. As demonstrated by the synopsis above, those accumulators can be implemented with arbitrary goals (like `integer::plus/3`).

Benefits of this library:

- avoid tedium and errors from manually threading accumulators through your predicates
- add or remove accumulators with a single declaration
- change the accumulator implementation with a single declaration (ex, switching from ordsets to rb-trees)

6.56.6 Syntax

Extended DCG syntax is very similar to DCG notation. An EDCG is created with clauses whose neck is the `-->>` operator. The following syntax is supported inside an EDCG clause:

- `{Goal}` - don't expand any hidden arguments of `Goal`
- `Goal` - expand all hidden arguments of `Goal` that are also in the head. Those hidden arguments not in the head are given default values.
- `Goal:L` - If `Goal` has no hidden arguments then force the expansion of all arguments in `L` in the order given. If `Goal` has hidden arguments then expand all of them, using the contents of `L` to override the expansion. `L` is either a term of the form `Acc`, `Acc(Left,Right)`, `Pass`, `Pass(Value)`, or a list of such terms. When present, the arguments `Left`, `Right`, and `Value` override the default values of arguments not in the head.
- `List:Acc` - Accumulate a list of terms in the accumulator `Acc`
- `List` - Accumulate a list of terms in the accumulator `dcg`
- `X/Acc` - Unify `X` with the left term for the accumulator `Acc`
- `Acc/X` - Unify `X` with the right term for the accumulator `Acc`
- `X/Acc/Y` - Unify `X` with the left and `Y` with the right term for the accumulator `Acc`
- `insert(X,Y):Acc` - Insert the arguments `X` and `Y` into the chain implementing the accumulator `Acc`. This is useful when the value of the accumulator changes radically because `X` and `Y` may be the arguments of an arbitrary relation
- `insert(X,Y)` - Insert the arguments `X` and `Y` into the chain implementing the accumulator `dcg`. This inserts the difference list `X-Y` into the accumulated list

6.56.7 Declaration of Predicates

Predicates are declared with facts of the following form:

```
pred_info(Name, Arity, List).
```

The predicate `Name/Arity` has the hidden parameters given in `List`. The parameters are added in the order given by `List` and their names must be atoms.

6.56.8 Declaration of Accumulators

Accumulators are declared with facts in one of two forms. The short form is:

```
acc_info(Acc, Term, Left, Right, Joiner).
```

The long form is:

```
acc_info(Acc, Term, Left, Right, Joiner, LStart, RStart).
```

In most cases the short form gives sufficient information. It declares the accumulator `Acc`, which must be an atom, along with the accumulating function, `Joiner`, and its arguments `Term`, the term to be accumulated, and `Left` & `Right`, the variables used in chaining.

The long form of `acc_info` is useful in more complex programs. It contains two additional arguments, `LStart` and `RStart`, that are used to give default starting values for an accumulator occurring in a body goal that does not occur in the head. The starting values are given to the unused accumulator to ensure that it will execute correctly even though its value is not used. Care is needed to give correct values for `LStart` and `RStart`. For DCG-like list accumulation both may remain unbound.

Two conventions are used for the two variables used in chaining depending on which direction the accumulation is done. For forward accumulation, `Left` is the input and `Right` is the output. For reverse accumulation, `Right` is the input and `Left` is the output.

6.56.9 Declaration of Passed Arguments

Passed arguments are conceptually the same as accumulators with `=/2` as the joiner function. Passed arguments are declared as facts in one of two forms. The short form is:

```
pass_info(Pass).
```

The long form is:

```
pass_info(Pass, PStart).
```

In most cases the short form is sufficient. It declares a passed argument `Pass`, that must be an atom. The long form also contains the starting value `PStart` that is used to give a default value for a passed argument in a body goal that does not occur in the head. Most of the time this situation does not occur.

6.56.10 Additional documentation

Peter Van Roy's page: [Declarative Programming with State](#)

Technical Report UCB/CSD-90-583 [Extended DCG Notation: A Tool for Applicative Programming in Prolog](#) by Peter Van Roy

- The Tech Report's PDF is [here](#)

A short [Wikipedia article](#) on DCGs and extensions.

6.57 elastic_net_regression

Elastic net regression regressor supporting continuous and mixed-feature datasets. The library implements the regressor_protocol defined in the regression_protocols library and learns a linear model using cyclic coordinate descent with coefficient-wise soft-thresholding updates and an additional L2 term in order to minimize mean squared error plus a standard elastic net penalty.

6.57.1 API documentation

Open the [../apis/library_index.html#elastic_net_regression](#) link in a web browser.

6.57.2 Loading

To load this library, load the loader.lgt file:

```
| ?- logtalk_load(elastic_net_regression(loader)).
```

6.57.3 Testing

To test this library predicates, load the tester.lgt file:

```
| ?- logtalk_load(elastic_net_regression(tester)).
```

The unit test suite covers the default mixed-penalty behavior together with the boundary cases `l1_ratio(0)` and `l1_ratio(1.0)`.

To run the performance benchmark suite, load the tester_performance.lgt file:

```
| ?- logtalk_load(elastic_net_regression(tester_performance)).
```

6.57.4 Features

- **Continuous and Mixed Features:** Supports continuous attributes and categorical attributes.
- **Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Missing Values:** Missing numeric and categorical values represented using anonymous variables are encoded using explicit missing-value indicator features.
- **Unknown Values:** Prediction requests containing categorical values that are not declared by the dataset raise a domain error.

- **Mixed Penalty:** Combines coefficient-wise L1 shrinkage with an L2 penalty controlled by the `regularization/1` and `l1_ratio/1` options, including the ridge-like `l1_ratio(0.0)` and lasso-like `l1_ratio(1.0)` endpoints.
- **Standard Encoding Semantics:** Categorical attributes are encoded using reference-level dummy coding derived from the declared dataset attribute values, with a trailing missing-value indicator feature, and each encoded coefficient is regularized independently.
- **Diagnostics Metadata:** Learned regressors record model name, target, training example count, optimization stop reason, completed iterations, final parameter delta, encoded feature count, and effective options, accessible using the shared regression diagnostics predicates.
- **Model Export:** Learned regressors can be exported as predicate clauses or written to a file.
- **Reference Benchmarks:** Includes a dedicated performance suite reporting training time, RMSE, and MAE for both zero-penalty baselines and explicit elastic-net runs on sparse, mixed, categorical, and wide mixed datasets.

6.57.5 Regressor representation

The learned regressor is represented by default as:

- `elastic_net_regressor(Encoders, Bias, Weights, Diagnostics)`

The exported predicate clauses therefore use the shape:

- `Functor(Encoders, Bias, Weights, Diagnostics)`

6.57.6 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(elastic_net_regression),
  target(Target),
  training_example_count(TrainingExampleCount),
  options(Options),
  convergence(Status),
  iterations(Iterations),
  final_delta(FinalDelta),
  encoded_feature_count(FeatureCount)
]
```

Where:

- `model(elastic_net_regression)` identifies the learning algorithm that produced the regressor.
- `target(Target)` stores the target attribute name declared by the training dataset.
- `training_example_count(TrainingExampleCount)` stores the number of examples used during training.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `convergence(Status)` records the optimization stop condition. The current values are `tolerance` when the maximum Karush-Kuhn-Tucker optimality violation across the intercept and all encoded features is within the configured tolerance and `maximum_iterations_exhausted` when training stops because the iteration cap is reached.

- `iterations(Iterations)` stores the number of coordinate-descent sweeps completed during training.
- `final_delta(FinalDelta)` stores the maximum Karush-Kuhn-Tucker optimality violation measured during the final optimization check.
- `encoded_feature_count(FeatureCount)` stores the number of numeric features induced by the encoder list, including missing-value indicator features.

Use the `regression_protocols diagnostic/2` and `regressor_options/2` helper predicates when you only need a single metadata term or the effective options.

6.57.7 Options

The `learn/3` predicate accepts the following options:

- `maximum_iterations/1`: Maximum number of coordinate-descent sweeps to run before stopping even if the tolerance criterion has not been met. The default is 2000.
- `tolerance/1`: Convergence threshold for the maximum Karush-Kuhn-Tucker optimality violation in a full coordinate-descent sweep. Training stops early when both the intercept condition and all encoded-feature subgradient conditions are satisfied within this value. The default is $1.0e-7$.
- `regularization/1`: Overall penalty coefficient applied during optimization. Higher values increase shrinkage and can reduce overfitting. The default is 0.01.
- `l1_ratio/1`: Fraction of the overall penalty assigned to the L1 part of the elastic net penalty. The remaining fraction is assigned to the L2 part. Accepted values are floats in the interval $[0.0, 1.0]$, where 0.0 gives the ridge endpoint and 1.0 gives the lasso endpoint. The default is 0.5.
- `feature_scaling/1`: Controls z-score standardization of continuous attributes before training and prediction. Accepted values are `true` and `false`. The default is `true`.

6.58 elo_ranker

Elo pairwise preference ranker. Processes the pairwise preference stream in dataset enumeration order using the standard Elo expected-score formula and symmetric rating updates after each observed result.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from pairwise preferences, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The current implementation requires a well-formed connected pairwise dataset so that learned ratings remain globally comparable across all ranked items.

6.58.1 API documentation

Open the `../apis/library_index.html#elo_ranker` link in a web browser.

6.58.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(elo_ranker(loader)).
```

6.58.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(elo_ranker(tester)).
```

6.58.4 Features

- **Pairwise Preference Learning:** Learns one deterministic rating per item from pairwise outcomes.
- **Deterministic Batch Elo Semantics:** Replays the dataset preference stream in enumeration order using the standard Elo expected-score formula. Because the current pairwise dataset protocol does not record historical timestamps, the implementation is a deterministic batch interpretation of that enumeration order rather than a reconstruction of a literal chronological competition log.
- **Configurable Rating Parameters:** Exposes the initial rating, K-factor, and rating scale as user options.
- **Integer Weight Fidelity:** Preference weights must be positive integers and are replayed as repeated unit outcomes in dataset enumeration order.
- **Deterministic Ranking:** Orders candidate items by learned rating with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items, self-preferences, non-positive weights, and disconnected comparison graphs.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.
- **Shared Ranking Infrastructure:** Reuses the `ranking_protocols` helpers for option processing, dataset validation, diagnostics, export, and candidate ranking.

6.58.5 Rating semantics

This implementation uses a deterministic batch Elo interpretation over the pairwise preference stream. Each preference is processed in dataset enumeration order. For a winner with rating R_w and a loser with rating R_l , the expected winner score is computed as

```
1 / (1 + 10^((R_l - R_w)/Scale))
```

and the winner receives a rating update of

```
K * (1 - ExpectedWinnerScore)
```

with the loser receiving the symmetric negative update.

Because the current pairwise dataset protocol does not encode timestamps, this is a deterministic batch Elo variant rather than a literal historical competition-log replay.

Positive integer preference weights are replayed as repeated unit outcomes. Datasets using non-integer preference weights are rejected because they do not map cleanly to standard Elo update semantics.

6.58.6 Usage

Learning a ranker

```
% Learn from a pairwise ranking dataset object
| ?- elo_ranker::learn(my_dataset, Ranker).
...

% Learn with custom Elo parameters
| ?- elo_ranker::learn(my_dataset, Ranker, [initial_rating(1400.0), k_factor(24.0), rating_
↪scale(200.0)]).
...
```

Inspecting diagnostics

```
% Inspect model, options, and dataset summary metadata
| ?- elo_ranker::learn(my_dataset, Ranker),
    elo_ranker::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
...
```

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- elo_ranker::learn(my_dataset, Ranker),
    elo_ranker::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
...
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- elo_ranker::learn(my_dataset, Ranker),
    elo_ranker::export_to_clauses(my_dataset, Ranker, my_ranker, Clauses).
Clauses = [my_ranker(elo_ranker(...))]
...

% Export to a file
| ?- elo_ranker::learn(my_dataset, Ranker),
    elo_ranker::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.pl').
...
```

6.58.7 Options

The following options can be passed to the learn/3 predicate:

- `initial_rating(Rating)`: Initial rating assigned to every item.
- `k_factor(KFactor)`: Positive Elo K-factor.
- `rating_scale(Scale)`: Positive rating-scale denominator used in the expected-score formula.

Datasets supplied to the ranker must use positive integer preference weights. Non-integer weights are rejected.

6.58.8 Ranker representation

The learned ranker is represented by a compound term of the form:

```
elo_ranker(Items, Ratings, Diagnostics)
```

Where:

- `Items`: List of ranked items.
- `Scores`: List of Item-Rating pairs.
- `Diagnostics`: List of metadata terms, including the effective options and dataset summary.

6.58.9 References

1. Elo, A. E. (1978). *The Rating of Chessplayers, Past and Present*. Arco.

6.59 events

The objects `event_registry`, `before_event_registry`, and `after_event_registry` implement convenient predicates for registering before and after events.

The code makes use of the monitoring built-in protocol, which declares the two basic event handler predicates (`before/3` and `after/3`). You will need to refer to this protocol in your objects if you want to use the super control structure (`^^/1`) with these predicates.

The monitor object implements more sophisticated event handling predicates.

6.59.1 API documentation

Open the ../apis/library_index.html#events link in a web browser.

6.59.2 Loading

To load all entities in this library, load the `loader.lgt` loader file:

```
| ?- logtalk_load(events(loader)).
```

6.60 ewma_anomaly_detector

EWMA (Exponentially Weighted Moving Average) anomaly detector for continuous sequence-like datasets. It is a statistical anomaly-detection method based on a two-sided EWMA control chart: for each known step value x_t , it computes a standardized deviation, updates the EWMA statistic E_t , and uses the maximum normalized excursion $|E_t| / (L * c_t)$ as the raw anomaly score, so a score of 1.0 corresponds exactly to reaching the chosen EWMA control limit.

The library implements the `anomaly_detector_protocol` defined in the `anomaly_detection_protocols` library. It learns a detector from a continuous dataset, computes anomaly scores for new instances, predicts normal or anomaly, and exports learned detectors as clauses or files.

Datasets are represented as objects implementing the `anomaly_dataset_protocol` protocol from the `anomaly_detection_protocols` library. Declared continuous attributes are interpreted as ordered monitoring steps in a sequence. See the `ewma_anomaly_detector/tests.lgt` file for example datasets.

6.60.1 API documentation

Open the ../apis/library_index.html#ewma_anomaly_detector link in a web browser.

6.60.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ewma_anomaly_detector(loader)).
```

6.60.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ewma_anomaly_detector(tester)).
```

6.60.4 Features

- **Statistical method:** implements anomaly detection based on a two-sided EWMA control chart, using learned per-step population means and standard deviations for continuous attributes.
- **Ordered sequence interpretation:** declared continuous attributes are treated as ordered monitoring steps. For each known step value x_t , the library computes $z_t = (x_t - \mu_t) / \sigma_t$ and updates the EWMA recurrence $E_t = \lambda * z_t + (1 - \lambda) * E_{(t-1)}$ with $E_0 = 0$.
- **EWMA control limits:** the raw anomaly score is the maximum normalized excursion $|E_t| / (L * c_t)$, where L is the learn-time `control_limit_multiplier/1` option and $c_t = \sqrt{\lambda / (2 - \lambda) * (1 - (1 - \lambda)^{(2 * t))})}$ is the classical EWMA control-limit factor after t EWMA updates. The learned detector stores a precomputed attribute schema so that query values can be reordered efficiently during scoring.
- **Continuous features only:** accepts datasets whose declared attributes are all continuous.
- **Baseline training selection:** supports learn-time `baseline_class_values(ClassValues)` and `baseline_selection_policy(Policy)` options. The default baseline class values are `[normal]`. The default reject policy throws an error if any non-baseline training example is found. The filter policy removes non-baseline examples before fitting the baseline statistics.
- **Missing-value tolerant:** ignores missing values when fitting per-step statistics. During scoring, missing step values do not update the EWMA state or advance the EWMA update count. Queries must still provide at least one known value.
- **Bounded scoring:** maps the raw EWMA excursion to $[0.0, 1.0]$ using $\text{Score} = \text{Raw} / (1 + \text{Raw})$.
- **EWMA control parameters:** supports learn-time `control_limit_multiplier/1` and `smoothing_factor/1` options. The default `control_limit_multiplier(3.0)` and `smoothing_factor(0.2)` correspond to a common EWMA monitoring setup.
- **Default threshold:** the default `anomaly_threshold(0.5)` corresponds to the chosen EWMA control limit after score normalization. Because the raw score is already normalized by the learned control limit, this threshold remains the same for any learned `control_limit_multiplier/1` value.
- **Learn-time control parameters:** `control_limit_multiplier/1` and `smoothing_factor/1` are recorded in the learned detector and reused for subsequent scoring and prediction. Passing them to `predict/4` does not override the learned values. Only `anomaly_threshold/1` can be overridden at predict time.
- **All-missing queries rejected:** scoring and prediction throw a `domain_error(non_empty_known_values, AttributeNames)` exception when every declared step is missing in the query.

- **Featureless datasets rejected:** datasets must declare at least one continuous feature; otherwise `learn/2-3` throws a `domain_error(non_empty_features, Dataset)` exception.
- **Detector export:** learned detectors can be exported as predicate clauses.
- **Explicit validation and diagnostics:** supports the shared `check_anomaly_detector/1`, `valid_anomaly_detector/1`, `diagnostics/2`, `diagnostic/2`, and `anomaly_detector_options/2` predicates.

6.60.5 Options

The following options are supported by the public API:

- `anomaly_threshold(Threshold)`: Threshold for `predict/3-4` (default: 0.5)
- `baseline_class_values(ClassValues)`: Learn-time class labels that are admissible for baseline fitting (default: [normal])
- `baseline_selection_policy(Policy)`: Learn-time handling of examples whose class is not listed in `baseline_class_values/1`. Supported values are `reject` and `filter` (default: `reject`)
- `control_limit_multiplier(ControlLimitMultiplier)`: Learn-time EWMA control-limit multiplier L (default: 3.0)
- `smoothing_factor(SmoothingFactor)`: Learn-time EWMA smoothing factor λ (default: 0.2)

6.60.6 Detector representation

The learned detector is represented by default as:

```
ewma_detector(TrainingDataset, AttributeSchema, Encoders, Diagnostics)
```

Where:

- `TrainingDataset`: training dataset object identifier
- `AttributeSchema`: precomputed attribute ordering metadata used to validate and reorder query step values efficiently during scoring
- `Encoders`: list of `ewma_encoder(Attribute, Mean, Scale)` records
- `Diagnostics`: learned metadata terms including `model/1`, `training_dataset/1`, `attribute_names/1`, `feature_count/1`, `example_count/1`, and `options/1`. The `example_count/1` value is the effective number of training examples after applying the selected baseline selection policy.

When exported using `export_to_clauses/4` or `export_to_file/4`, this detector term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.60.7 Notes

Scoring has three stages. First, the detector computes one standardized deviation $z_t = (x_t - \mu_t) / \sigma_t$ for each known monitoring step. Second, those deviations are processed sequentially using the EWMA recurrence $E_t = \lambda * z_t + (1 - \lambda) * E_{(t-1)}$ with the learned `smoothing_factor/1` value. Third, the maximum normalized excursion $|E_t| / (L * c_t)$ is mapped to the interval $[0.0, 1.0]$ using $\text{Score} = \text{Raw} / (1 + \text{Raw})$.

The control-limit factor c_t is computed from the number of actual EWMA updates, not from the declared attribute position. Accordingly, leading or intermediate missing step values neither update the EWMA state nor widen the control limits.

The `smoothing_factor/1` option changes the EWMA update rule itself. Smaller values make the detector retain longer-term history while larger values react more strongly to recent deviations. The `control_limit_multiplier/1` option scales the control limits directly in the score path. Larger values therefore make the detector less sensitive by requiring larger EWMA excursions before the raw score reaches 1.0.

The `baseline_class_values/1` option declares which dataset class labels are admissible for baseline fitting. The `baseline_selection_policy/1` option then controls what happens when other labels are present in the training data. The default reject policy raises a `domain_error(baseline_only_training_data, Dataset)` exception when any non-baseline example is found. The filter policy removes non-baseline examples before fitting and raises a `domain_error(non_empty_baseline_training_data, Dataset)` exception if no training examples remain after filtering.

Attributes with zero observed dispersion are assigned a fallback scale of 1.0. This keeps the detector well-defined for singleton datasets or constant steps while still yielding zero score for matching values and positive scores for deviating values.

6.61 `expand_library_alias_paths`

This library provides a hook object, `expand_library_alias_paths`, for expanding library alias paths in `logtalk_library_path/2` facts in source files. It is mainly used when embedding Logtalk and Logtalk applications.

6.61.1 API documentation

Open the ../apis/library_index.html#expand-library-alias-paths link in a web browser.

6.61.2 Loading

To load all entities in this library, load the `loader.lgt` utility file:

```
| ?- logtalk_load(expand_library_alias_paths(loader)).
```

6.61.3 Usage

Use the `hook/1` option when compiling a source file:

```
| ?- logtalk_load(my_source_file, [hook(expand_library_alias_paths)]).
...
```

Alternatively, assuming it is the only hook object you are using, you can set it as the default hook object:

```
| ?- set_logtalk_flag(hook, expand_library_alias_paths).
...
```

6.62 expecteds

This library provides an implementation of *expected terms* with an API that is inspired by the optional library and C++ standardization proposals for an `Expected<T>` type. An expected term is an opaque compound term that either contains an expected value or an error informing why the expected value is not present. Expected terms provide an alternative to generating an exception (or a failure) when something unexpected happens when asking for a value. This allows, e.g., separating the code that constructs expected terms from the code that processes them, which is then free to deal if necessary and at its convenience with any unexpected events.

6.62.1 API documentation

Open the ../apis/library_index.html#expecteds link in a web browser.

6.62.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(expecteds(loader)).
```

6.62.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(expecteds(tester)).
```

6.62.4 Usage

The expected object provides constructors for expected terms. For example:

```
| ?- expected::of_expected(1, Expected).
...
```

The created expected terms can then be passed as parameters to the `expected/1` parametric object. For example:

```
| ?- expected::of_expected(1, Expected), expected(Expected)::or_else(Value, 0).
Expected = expected(1),
Value = 1
yes

| ?- expected::of_unexpected(-1, Expected), expected(Expected)::or_else(Value, 0).
Expected = unexpected(-1),
Value = 0
yes
```

The `either` object provides types and predicates for extended type-checking and predicates for handling lists of expected terms, including `sequence/2` and `traverse/3`.

The `expected/1` parametric object provides `filter/3` for conditionally rejecting values, converting them to unexpected terms with a given error:

```
| ?- expected::of_expected(1, Expected),
    expected(Expected)::filter(integer, not_integer, NewExpected).
NewExpected = expected(1)
yes

| ?- expected::of_expected(a, Expected),
    expected(Expected)::filter(integer, not_integer, NewExpected).
NewExpected = unexpected(not_integer)
yes
```

The `map_or_else/3` predicate applies a closure to the value if present, returning a default value otherwise (symmetric with the `optionals` library):

```
| ?- expected::of_expected(a, Expected),
    expected(Expected)::map_or_else(char_code, 0, Value).
Value = 97
yes

| ?- expected::of_unexpected(-1, Expected),
    expected(Expected)::map_or_else(char_code, 0, Value).
Value = 0
yes
```

The `or/2` predicate chains expected terms, returning the current term if it holds a value, or calling a closure to produce an alternative otherwise:

```
| ?- expected::of_expected(1, Expected),
    expected(Expected)::or(NewExpected, expected::of_expected(2)).
NewExpected = expected(1)
yes

| ?- expected::of_unexpected(-1, Expected),
    expected(Expected)::or(NewExpected, expected::of_expected(2)).
NewExpected = expected(2)
yes
```

The `zip/3` predicate combines two expected terms using a closure when both hold values, returning the first error otherwise:

```
| ?- expected::of_expected(1, E1), expected::of_expected(3, E2),
    expected(E1)::zip([X,Y,Z]>>(Z is X+Y), E2, NewExpected).
NewExpected = expected(4)
yes

| ?- expected::of_unexpected(-1, E1), expected::of_expected(3, E2),
    expected(E1)::zip([X,Y,Z]>>(Z is X+Y), E2, NewExpected).
NewExpected = unexpected(-1)
yes
```

The map_unexpected/2 predicate transforms the error held by an expected term:

```
| ?- expected::of_unexpected(-1, Expected),
    expected(Expected)::map_unexpected([X,Y]>>(Y is abs(X)), NewExpected).
NewExpected = unexpected(1)
yes
```

The map_catching/2 predicate applies a closure that may throw an error, catching it and wrapping it as an unexpected term:

```
| ?- expected::of_expected(a, Expected),
    expected(Expected)::map_catching(char_code, NewExpected).
NewExpected = expected(97)
yes

| ?- expected::of_expected(1, Expected),
    expected(Expected)::map_catching(char_code, NewExpected),
    expected(NewExpected)::is_unexpected.
yes
```

The map_both/3 predicate is a bifunctor map that transforms both the expected value and unexpected error using separate closures:

```
| ?- expected::of_expected(1, Expected),
    expected(Expected)::map_both([X,Y]>>(Y is X+1), [X,Y]>>(Y is abs(X)), NewExpected).
NewExpected = expected(2)
yes
```

The swap/1 predicate swaps expected and unexpected terms:

```
| ?- expected::of_expected(1, Expected),
    expected(Expected)::swap(NewExpected).
NewExpected = unexpected(1)
yes

| ?- expected::of_unexpected(error, Expected),
    expected(Expected)::swap(NewExpected).
NewExpected = expected(error)
yes
```

The flatten/1 predicate unwraps a nested expected term:

```
| ?- expected::of_expected(1, Inner), expected::of_expected(Inner, Outer),
    expected(Outer)::flatten(NewExpected).
```

(continues on next page)

(continued from previous page)

```

NewExpected = expected(1)
yes

| ?- expected::of_unexpected(oops, Inner), expected::of_expected(Inner, Outer),
    expected(Outer)::flatten(NewExpected).
NewExpected = unexpected(oops)
yes

```

Conversion between expected and optional terms is provided by the `to_optional/1`, `from_optional/3`, and `optional/1::to_expected/2` predicates:

```

| ?- expected::of_expected(1, Expected),
    expected(Expected)::to_optional(Optional).
Optional = optional(1)
yes

| ?- expected::of_unexpected(error, Expected),
    expected(Expected)::to_optional(Optional).
Optional = empty
yes

| ?- optional::of(1, Optional),
    expected::from_optional(Optional, missing, Expected).
Expected = expected(1)
yes

| ?- optional::empty(Optional),
    expected::from_optional(Optional, missing, Expected).
Expected = unexpected(missing)
yes

```

Examples:

```

| ?- expected::of_expected(1, E1), expected::of_expected(2, E2),
    either::sequence([E1, E2], Expected).
Expected = expected([1,2])
yes

| ?- either::traverse({expected}/[X,E]>>expected::of_expected(X, E), [1,2], Expected).
Expected = expected([1,2])
yes

| ?- either::traverse({expected}/[X,E]>>(
    integer(X) -> expected::of_expected(X, E)
    ; expected::of_unexpected(not_integer(X), E)
), [1,a,2], Expected).
Expected = unexpected(not_integer(a))
yes

| ?- expected::of_expected(1, E1), expected::of_unexpected(e, E2),
    either::sequence([E1, E2], Expected).
Expected = unexpected(e)
yes

```

6.62.5 See also

The optionals and validations libraries.

6.63 format

The format object provides a portable abstraction over how the de facto standard format/2-3 predicates are made available by the supported backend Prolog systems. Some systems provide these predicates as built-in predicates, while others make them available using a library that must be explicitly loaded.

Calls to the library predicates are inlined when compiled with the optimize flag turned on for most of the backends. When that's the case, there is no overhead compared with calling the abstracted predicates directly.

This library provides linter checks for calls to the format/2-3 predicates. Given the differences between the implementations of these predicates among Prolog systems, the linter checks focus on detecting common errors such as missing arguments and too many arguments. The linter warnings are printed when the suspicious_calls flag is set to warning (its usual default).

6.63.1 API documentation

Open the ../apis/library_index.html#format link in a web browser.

6.63.2 Loading

To load all entities in this library, load the loader.lgt file:

```
| ?- logtalk_load(format(loader)).
```

6.63.3 Testing

Minimal tests for this library predicates can be run by loading the tester.lgt file:

```
| ?- logtalk_load(format(tester)).
```

Detailed tests for the format/2-3 predicates are available in the tests/prolog/predicates directory as part of the Prolog standards conformance test suite. Use those tests to confirm the portability of the format specifiers that you want to use.

6.63.4 Usage

Load this library from your application loader file. To call the format/2-3 predicates using implicit message-sending, add the following directive to any object or category calling the predicates:

```
:- uses(format, [
    format/2, format/3
]).
```

6.63.5 Portability

Some Prolog systems provide only a subset of the expected format specifiers. Notably, table-related format specifiers are only fully supported by a few systems. See the section below on testing.

Only some of the supported Prolog backends provide implementations of the `format/2-3` predicates that allow using not only an atom or a list of character codes for the format string (as de facto standard) but also using a list of characters. These currently include ECLiPSe, GNU Prolog, SICStus Prolog, SWI-Prolog, Trealla Prolog, XVM, and YAP. Therefore, when wide portability is sought, atoms must be used for the format specifier argument as they bypass any dependency on the `double_quotes` standard Prolog flag. Some systems, like Tau Prolog, only accept a list of characters for the format string. In this case, this library will convert the atom format string before calling these systems native implementations.

6.64 `fp_growth_pattern_miner`

FP-growth frequent itemset miner for transaction datasets. The library depends on the `frequent_pattern_mining_protocols` support library, implements the generic `pattern_miner_protocol` defined in the `pattern_mining_protocols` core library, and mines frequent itemsets using recursive conditional pattern-base projection over a compact FP-tree whose nodes store parent links directly, plus header-table node chains derived from the final tree for direct conditional-base reconstruction, without candidate generation.

Requires a dataset implementing `transaction_dataset_protocol` with transactions represented as canonical sorted lists of unique declared items.

6.64.1 API documentation

Open the `../..apis/library_index.html#fp_growth_pattern_miner` link in a web browser.

6.64.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(fp_growth_pattern_miner(loader)).
```

6.64.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(fp_growth_pattern_miner(tester)).
```

6.64.4 Features

- **FP-tree Construction:** Builds a compact prefix tree from frequent items ordered by global support.
- **Header and Parent Links:** Stores parent links directly in tree nodes and derives header-table node chains from the final tree so conditional pattern bases are reconstructed by walking parent links instead of using a separate node index.
- **Pattern Growth:** Mines frequent itemsets recursively from conditional pattern bases without candidate generation.
- **Canonical Transactions:** Validates that transactions are sorted, duplicate-free, and restricted to declared items.
- **Flexible Support Thresholds:** Supports relative minimum support and absolute minimum support count.
- **Model Export:** Mined pattern collections can be exported as predicate clauses or written to a file.

6.64.5 Options

The `mine/3` predicate accepts the following options:

- `minimum_support/1`: Relative minimum support threshold in the interval $]0.0, 1.0]$. The default is `0.5`.
- `minimum_support_count/1`: Absolute minimum support count. When both support options are provided, this option takes precedence.
- `maximum_pattern_length/1`: Maximum itemset length to mine. The default is `1000`, which is effectively capped by the longest transaction in the dataset.
- `minimum_pattern_length/1`: Minimum itemset length retained in the mined result. The default is `1`.

6.64.6 Pattern miner representation

The mined pattern miner result is represented by a compound term with the functor chosen by the implementation and arity 3. For example:

```
fp_growth_pattern_miner(ItemDomain, Patterns, Options)
```

Where:

- `ItemDomain`: Canonical sorted list of declared dataset items.
- `Patterns`: List of `itemset(Items, SupportCount)` terms ordered first by pattern length and then lexicographically.
- `Options`: Effective mining options used to mine the frequent itemsets.

6.64.7 References

1. Han, J., Pei, J., and Yin, Y. (2000) - “Mining frequent patterns without candidate generation”.

6.65 frequent_pattern_mining_protocols

This library provides support entities for frequent itemset mining algorithms. Transactional datasets are represented as objects implementing the `transaction_dataset_protocol` protocol. The generic `pattern_miner_protocol` protocol and the `pattern_miner_common` category used by concrete miners are loaded from the `pattern_mining_protocols` core library.

The `frequent_pattern_mining_common` category builds on that generic core with frequent-itemset-specific helpers for dataset validation, support accumulation, and itemset ordering/filtering.

This library also provides reusable transaction smoke-test datasets and a small smoke-test suite.

6.65.1 API documentation

Open the ../apis/library_index.html#frequent_pattern_mining_protocols link in a web browser.

6.65.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(frequent_pattern_mining_protocols(loader)).
```

6.65.3 Testing

To run the library smoke tests, load the `tester.lgt` file:

```
| ?- logtalk_load(frequent_pattern_mining_protocols(tester)).
```

6.65.4 Test datasets

Several sample transaction datasets are included in the `test_datasets` directory:

- `market_basket_basics.lgt`: A compact transaction dataset with 6 transactions and 5 items intended for frequent-itemset smoke tests.
- `layered_baskets.lgt`: A transaction dataset with overlapping co-occurrence layers intended for support-count and candidate-pruning tests.
- `deep_intersection_baskets.lgt`: A compact transaction dataset with one frequent length-4 itemset and multiple overlapping length-3 itemsets intended to stress deeper vertical tidset intersections.

The directory also includes invalid fixtures useful for validation and error-handling tests:

- `invalid_undeclared_item_baskets.lgt`: Uses an item not listed in the declared item domain.
- `invalid_unsorted_transaction_baskets.lgt`: Uses a transaction whose items are not in canonical sorted order.
- `invalid_duplicate_item_baskets.lgt`: Uses a transaction with a duplicate item.

- `invalid_empty_baskets.lgt`: Declares an item domain but no transactions.
- `invalid_item_domain_baskets.lgt`: Declares a non-canonical item domain with duplicate items.

6.66 gaussian_mixture_clusterer

Gaussian mixture model clusterer. It uses deterministic expectation-maximization with diagonal covariance matrices. Supports continuous attributes only.

The library implements the `clusterer_protocol` defined in the `clustering_protocols` library. It provides predicates for learning a clusterer from a dataset, assigning new instances to clusters, returning Gaussian-mixture posterior component probabilities for new instances, and exporting the learned clusterer as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol from the `clustering_protocols` library.

6.66.1 API documentation

Open the ../apis/library_index.html#gaussian_mixture_clusterer link in a web browser.

6.66.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(gaussian_mixture_clusterer(loader)).
```

6.66.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(gaussian_mixture_clusterer(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(gaussian_mixture_clusterer(tester_performance)).
```

6.66.4 Features

- **Expectation-Maximization**: Learns diagonal Gaussian components using deterministic EM updates.
- **Continuous Datasets**: Accepts datasets containing only continuous attributes.
- **Configurable Dead-Component Handling**: Dead components can either be preserved with zero weight or deterministically reseeded to the least-confident training row.
- **Deterministic Initialization**: Supports `first_k` and deterministic spread initialization for component means. The spread strategy uses a canonical first seed and canonical tie-breaking so equivalent row permutations produce the same initialization.
- **Optional Feature Scaling**: Continuous attributes can be standardized using z-score scaling.

- **Posterior Prediction:** New instances are assigned to the component with the highest posterior score, and `cluster_probabilities/3` can return the full posterior distribution over components.
- **Training Diagnostics:** Learned clusterers record convergence reason, iteration count, average log-likelihood, final delta, and effective options.
- **Portable Export:** Learned clusterers can be exported as clauses or files and reused later.

6.66.5 Gaussian-Mixture-Specific Prediction API

In addition to the shared `cluster/3` predicate from the clustering protocols library, this package provides a Gaussian-mixture-specific predicate:

- `cluster_probabilities(Clusterer, Instance, Probabilities)`: Returns posterior component probabilities for `Instance` as `Cluster-Probability` pairs in component-id order.

6.66.6 Options

The following options can be passed to the `learn/3` predicate:

- `k(K)`: Number of mixture components. Default is 2.
- `initialization(Initialization)`: Mean initialization strategy. Options: `spread` (default) or `first_k`.
- `feature_scaling(FeatureScaling)`: Whether to standardize continuous attributes before clustering. Options: `on` (default) or `off`.
- `maximum_iterations(MaximumIterations)`: Maximum number of EM iterations. Default is 100.
- `tolerance(Tolerance)`: Per-example average log-likelihood convergence tolerance. Default is 0.0001.
- `covariance_regularization(Regularization)`: Positive diagonal covariance regularization constant. Default is 0.001.
- `dead_component_policy(Policy)`: Handling for components whose total responsibility collapses below the dead-component threshold. Options: `zero_weight` (default) keeps the previous component with zero weight; `reseed` relocates the component to the least-confident training row and gives it one-example prior weight.

6.66.7 Clusterer representation

The learned clusterer is represented as a compound term with the functor chosen by the user when exporting the clusterer and arity 5. For example:

```
gaussian_mixture_clusterer(Encoders, Components, Weights, Options, Diagnostics)
```

Where:

- **Encoders:** List of continuous attribute encoders storing attribute name, mean, and scale.
- **Components:** List of `component(Mean, Variances)` terms in component-id order.
- **Weights:** List of mixture weights in component-id order.
- **Options:** Effective training options used to learn the clusterer.
- **Diagnostics:** Training diagnostics including convergence status, iteration count, average log-likelihood, final delta, and options.

6.67 gaussian_process_regression

Gaussian process regression regressor supporting continuous and mixed-feature datasets. Uses exact Gaussian process regression with a mixed covariance kernel: an automatic-relevance-determination squared-exponential component over continuous encoded features and a field-wise categorical overlap component over categorical attributes. Hyperparameters are selected by maximizing the log marginal likelihood using a deterministic coordinate search in log space.

The library implements the `regressor_protocol` defined in the `regression_protocols` library and learns an exact mixed gaussian process using an automatic-relevance-determination squared-exponential kernel for continuous encoded features together with a categorical overlap kernel for categorical attributes. Hyperparameters are selected by maximizing the log marginal likelihood using a deterministic coordinate search in log space.

6.67.1 API documentation

Open the ../apis/library_index.html#gaussian_process_regression link in a web browser.

6.67.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(gaussian_process_regression(loader)).
```

6.67.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(gaussian_process_regression(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(gaussian_process_regression(tester_performance)).
```

6.67.4 Features

- **Continuous and Mixed Features:** Supports continuous attributes and categorical attributes encoded using reference-level dummy coding.
- **Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Missing Values:** Missing numeric and categorical values are encoded using explicit missing-value indicator features.
- **Unknown Values:** Prediction requests containing categorical values that are not declared by the dataset raise a domain error.
- **Exact Bayesian Regression:** Uses exact Gaussian process regression with a mixed covariance kernel combining an automatic-relevance-determination squared-exponential component for continuous encoded features and a field-wise categorical overlap component for categorical attributes.

- **Automatic Hyperparameter Selection:** By default performs deterministic log-marginal-likelihood optimization of the continuous-feature length scales, categorical mismatch penalties, signal variance, and noise variance. The `length_scale/1`, `categorical_penalty/1`, `signal_variance/1`, and `noise_variance/1` options also accept `auto` when optimization is disabled.
- **Uncertainty Quantification:** Exposes posterior predictive Gaussian distributions for new instances (including observation noise variance) using the `predict_distribution/3` predicate. Small negative posterior variances caused by floating-point roundoff are clipped to zero while larger negative values raise an error.
- **Adaptive Stabilization:** Retries covariance factorization with progressively larger jitter values and records the effective retry count in the learned diagnostics.
- **Memory-Based Representation:** Stores the encoded training rows and cached Cholesky factor required for exact posterior prediction plus the dual coefficients required for exact posterior prediction.
- **Model Export:** Learned regressors can be exported as predicate clauses or written to a file.
- **Reference Benchmarks:** Includes a dedicated performance suite reporting training time, RMSE, and MAE for representative regression datasets.

6.67.5 Regressor representation

The learned regressor is represented by default as:

- `gaussian_process_regressor(Encoders, TrainingFeatures, TargetMean, Alpha, CholeskyFactor, Kernel, Diagnostics)`

The exported predicate clauses therefore use the shape:

- `Functor(Encoders, TrainingFeatures, TargetMean, Alpha, CholeskyFactor, Kernel, Diagnostics)`

In this representation, `Encoders` stores feature encoding metadata, `TrainingFeatures` stores the encoded training rows, `TargetMean` stores the centered-mean offset, `Alpha` stores the cached dual coefficients, `CholeskyFactor` stores the lower-triangular covariance factor, `Kernel` stores the learned mixed-kernel hyperparameters including one continuous length scale per encoded continuous feature and one categorical penalty per categorical attribute, and `Diagnostics` stores training metadata including the effective options and learned hyperparameters.

6.67.6 Prediction API

The standard `predict/3` predicate returns the posterior predictive mean.

The `predict_distribution/3` predicate returns a term with the form:

- `gaussian(Mean, Variance)`

where `Variance` is the posterior predictive variance for observed targets, including the learned observation noise variance. Tiny negative values caused by floating-point roundoff are clipped to zero; larger negative values raise a domain error.

6.67.7 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(gaussian_process_regression),
  target(Target),
  training_example_count(TrainingExampleCount),
  options(Options),
  kernel(squared_exponential),
  length_scales(LengthScales),
  categorical_penalties(CategoricalPenalties),
  signal_variance(SignalVariance),
  noise_variance(NoiseVariance),
  jitter(Jitter),
  continuous_feature_count(ContinuousFeatureCount),
  categorical_feature_count(CategoricalFeatureCount),
  jitter_attempts(JitterAttempts),
  log_marginal_likelihood(LogMarginalLikelihood),
  convergence(Convergence),
  iterations(Iterations),
  final_delta(FinalDelta),
  encoded_feature_count(FeatureCount)
]
```

Where:

- `model(gaussian_process_regression)` identifies the learning algorithm that produced the regressor.
- `target(Target)` stores the target attribute name declared by the training dataset.
- `training_example_count(TrainingExampleCount)` stores the number of examples used during training.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `kernel(squared_exponential)` records the covariance-kernel family used by the learned model.
- `length_scales(LengthScales)` stores the learned per-feature squared-exponential length scales for the continuous encoded feature dimensions.
- `categorical_penalties(CategoricalPenalties)` stores the learned mismatch penalties for the categorical attributes.
- `signal_variance(SignalVariance)` stores the learned latent-process marginal variance.
- `noise_variance(NoiseVariance)` stores the learned observation noise variance.
- `jitter(Jitter)` stores the effective diagonal jitter used to stabilize the covariance factorization.
- `continuous_feature_count(ContinuousFeatureCount)` stores the number of continuous encoded feature dimensions used by the squared-exponential kernel.
- `categorical_feature_count(CategoricalFeatureCount)` stores the number of categorical attributes handled by the overlap-kernel component.
- `jitter_attempts(JitterAttempts)` stores the number of adaptive jitter retries required by the final covariance factorization.
- `log_marginal_likelihood(LogMarginalLikelihood)` stores the final training log marginal likelihood.

- `convergence(Convergence)` records the hyperparameter-search stop reason.
- `iterations(Iterations)` stores the number of hyperparameter-search sweeps performed.
- `final_delta(FinalDelta)` stores the last log-marginal-likelihood improvement magnitude.
- `encoded_feature_count(FeatureCount)` stores the number of numeric features induced by the encoder list, including missing-value indicator features.

Use the `regression_protocols diagnostic/2` and `regressor_options/2` helper predicates when you only need a single metadata term or the effective options.

6.67.8 Options

The `learn/3` predicate accepts the following options:

- `kernel/1`: Kernel family. The current accepted value is `squared_exponential`. The default is `squared_exponential`.
- `feature_scaling/1`: Controls z-score standardization of continuous attributes before training and prediction. Accepted values are `true` and `false`. The default is `true`.
- `optimize_hyperparameters/1`: Controls deterministic log-marginal-likelihood optimization of the kernel hyperparameters. Accepted values are `true` and `false`. The default is `true`.
- `length_scale/1`: Initial or fixed squared-exponential length-scale specification for the continuous encoded feature dimensions. Accepted values are `auto`, a positive float that is broadcast to every continuous encoded feature, or a list of positive floats with one value per continuous encoded feature. The default is `auto`.
- `categorical_penalty/1`: Initial or fixed categorical mismatch-penalty specification. Accepted values are `auto`, a positive float that is broadcast to every categorical attribute, or a list of positive floats with one value per categorical attribute. The default is `auto`.
- `signal_variance/1`: Initial or fixed latent-process variance. Accepted values are `auto` or a positive float. The default is `auto`.
- `noise_variance/1`: Initial or fixed observation noise variance. Accepted values are `auto` or a non-negative float. The default is `auto`.
- `jitter/1`: Positive diagonal stabilization jitter added to the covariance matrix before factorization. When factorization fails, the implementation retries with progressively larger jitter values until it succeeds or exhausts the retry budget. The default is `1.0e-8`.
- `maximum_iterations/1`: Maximum number of hyperparameter-search sweeps when optimization is enabled. The default is `12`.
- `tolerance/1`: Minimum log-marginal-likelihood improvement floor required to continue hyperparameter optimization. The default floor is `1.0e-6`.
- `relative_improvement_factor/1`: Relative log-marginal-likelihood improvement factor used together with `tolerance/1` to stop hyperparameter optimization when only numerically insignificant gains remain. The default is `1.0e-4`.
- `hyperparameter_minimum/1`: Lower bound used when proposing scaled hyperparameter candidates during coordinate search. The default is `1.0e-6`.
- `maximum_continuous_length_scale/1`: Upper bound used when proposing scaled continuous length-scale candidates during coordinate search. The default is `32.0`.
- `maximum_categorical_penalty/1`: Upper bound used when proposing scaled categorical mismatch-penalty candidates during coordinate search. The default is `32.0`.

- `max_factorization_attempts/1`: Maximum number of covariance-factorization retries performed with progressively increased jitter before training raises a positive-definiteness error. The default is 32.
- `jitter_scale_factor/1`: Multiplicative factor used to increase the diagonal jitter on each covariance-factorization retry. The default is 2.0.

6.68 genint

The `genint` library implements predicates for generating positive integers in increasing order. The public predicates are declared synchronized to prevent race conditions when using a backend Prolog compiler with multi-threading support.

6.68.1 API documentation

Open the ../apis/library_index.html#genint link in a web browser.

6.68.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(genint(loader)).
```

6.68.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(genint(tester)).
```

6.68.4 Usage

The `genint_core` category implements the library predicates. This category is imported by the default `genint` object to provide application global named counters. To make the counters local and thus minimize the potential for counter name clashes, the category can be imported by one or more application objects. Use `protected` or `private` import to restrict the scope of the library predicates.

6.69 gensym

The `gensym` library implements predicates for generating unique atoms. The public predicates are declared synchronized to prevent race conditions when using a backend Prolog compiler with multi-threading support.

6.69.1 API documentation

Open the `../apis/library_index.html#gensym` link in a web browser.

6.69.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(gensym(loader)).
```

6.69.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(gensym(tester)).
```

6.69.4 Usage

The `gensym_core` category implements the library predicates. This category is imported by the default `gensym` object to provide application global generators. To make the generators local and thus minimize the potential for generator name clashes, the category can be imported by one or more application objects. Use protected or private import to restrict the scope of the library predicates. For example:

```
:- object(foo,  
    imports(private::gensym_core)).  
  
    bar :-  
        ^^gensym(p, S),  
        ...  
:- end_object.
```

6.70 geojson

The `geojson` library provides predicates for parsing, generating, and validating GeoJSON documents as specified by RFC 7946:

- <https://www.rfc-editor.org/rfc/rfc7946>

It builds on top of the `json` library for JSON parsing and generation and complements the `geospatial` library by providing a standard interchange format for geometry, features, and feature collections.

6.70.1 API documentation

Open the ../apis/library_index.html#geojson link in a web browser.

6.70.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(geojson(loader)).
```

6.70.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(geojson(tester)).
```

6.70.4 Representation

The library parses GeoJSON documents into native terms using the following representation:

- `point(Position)` or `point(Position, Options)`
- `multi_point(Positions)` or `multi_point(Positions, Options)`
- `line_string(Positions)` or `line_string(Positions, Options)`
- `multi_line_string(LineStrings)` or `multi_line_string(LineStrings, Options)`
- `polygon(Rings)` or `polygon(Rings, Options)`
- `multi_polygon(Polygons)` or `multi_polygon(Polygons, Options)`
- `geometry_collection(Geometries)` or `geometry_collection(Geometries, Options)`
- `feature(Geometry, Properties)` or `feature(Geometry, Properties, Options)`
- `feature_collection(Features)` or `feature_collection(Features, Options)`

Positions follow the GeoJSON order and are represented as lists of numbers, e.g. `[Longitude, Latitude]` or `[Longitude, Latitude, Altitude]`. Validation also checks longitude and latitude values against the RFC 7946 geographic ranges.

The Options lists may contain:

- `bbox(BBox)` for the optional bounding box member
- `id(Id)` for feature identifiers
- `foreign_members(ForeignMembers)` for additional non-reserved members, represented either as a list of pairs or as an embedded JSON object term using the selected object representation

Feature geometries and properties may be `@null` to represent GeoJSON null.

6.70.5 Examples

Parse a point document:

```
| ?- geojson::parse(atom('{ "type": "Point", "coordinates": [100.0, 0.0] }'), GeoJSON).
GeoJSON = point([100.0, 0.0])
yes
```

Generate a GeoJSON feature collection:

```
| ?- geojson::generate(atom(JSON), feature_collection([
    feature(point([102.0, 0.5]), {prop0-value0}),
    feature(@null, @null)
], [bbox([100.0, 0.0, 105.0, 1.0])])).
```

Validate a native GeoJSON term:

```
| ?- geojson::validate(polygon([[100.0, 0.0], [101.0, 0.0], [101.0, 1.0], [100.0, 1.0], [100.0, 0.0]])).
↪0]]))).
```

Inspect validation errors:

```
| ?- geojson::validate(polygon([[100.0, 0.0], [101.0, 0.0], [101.0, 1.0], [100.0, 1.0]])), Errors).
Errors = [ring_not_closed([coordinates, 0])]
yes
```

Convert between JSON terms and native GeoJSON terms:

```
| ?- geojson::json_to_geojson({type-'Point', coordinates-[1,2]}, GeoJSON).
GeoJSON = point([1,2])
yes

| ?- geojson::geojson_to_json(feature(@null, @null, [id(1)]), JSON).
JSON = {type-'Feature', geometry=@null, properties=@null, id-1}
yes
```

6.70.6 Notes

- The library validates the RFC 7946 structural rules for geometry, features, feature collections, linear rings, bounding boxes, and prohibited crs members.
- Bounding boxes are represented as plain lists of numbers with either four or six elements. Longitude values must stay within $[-180, 180]$, latitude values within $[-90, 90]$, south must not exceed north, and three-dimensional boxes also require minimum altitude not to exceed maximum altitude. Bounding box dimensionality must match the represented geometry dimensionality, so two-dimensional geometries use four-element boxes while geometries with altitude use six-element boxes. Mixed two-dimensional and three-dimensional geometry or feature collections are treated as three-dimensional for bounding-box validation. Antimeridian-crossing boxes are accepted, so west may still be greater than east.
- Parsing rejects duplicate reserved GeoJSON members instead of silently normalizing them.
- Properties and foreign member values are recursively validated as JSON values compatible with the selected embedded object representation.

- When using the chars or codes string representations, string terms are validated as proper character or character-code lists.
- Foreign members are preserved but reserved GeoJSON member names are rejected when used as foreign member keys, regardless of whether the foreign members are supplied as pair lists or embedded object terms.
- The `geojson/3` parametric object mirrors the `json/3` customization parameters for embedded object, pair, and string representations.

6.71 geospatial

This library provides a `geospatial_protocol` protocol and a `geospatial` object for common geographic computations over coordinates represented as `geographic(Latitude,Longitude)`. By default, distances are returned in kilometers.

6.71.1 API documentation

Open the ../apis/library_index.html#geospatial link in a web browser.

6.71.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(geospatial(loader)).
```

6.71.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(geospatial(tester)).
```

6.71.4 Available predicates

The library currently includes predicates for:

- Coordinate validation
- Coordinate normalization
- Point-to-point distances (haversine and vincenty)
- Rhumb-line distance, bearing, and destination predicates
- Rhumb-line interpolation and midpoint
- Generic distance dispatch (`distance/4` and `distance/5`)
- Initial and final bearings
- Midpoint and destination point computations
- Great-circle interpolation and track-distance

- Proximity checks and nearest coordinate search
- Mean center and coordinates bounding boxes
- Minimum enclosing circle
- Coordinates-to-bounding-box conversion
- Equirectangular projection and inverse
- Point-in-polygon checks
- Polygon area, polygon centroid, and polygon bounding boxes
- Polygon closure and orientation
- Polygon orientation normalization
- Polygon validity checks
- Bounding-box utilities
- Strict bounding-box overlap checks
- Bounding-box and polygon/polyline relation
- Nearest-point and point-to-polyline distance
- Polyline length and polygon perimeter
- Polyline simplification by tolerance
- Polyline split and resampling
- Polygon intersection checks
- Spherical bounding boxes
- Route distance accumulation (`route_distance/2`, `route_distance/3`, and `route_distance/4`)

For explicit units, predicates `distance/5` and `route_distance/4` support kilometers, meters, miles, and `nautical_miles`.

Generic metric dispatch accepts `rhumb` (and alias `loxodrome`) in addition to existing spherical and ellipsoidal metrics.

6.71.5 Notes

The generic vector distance predicates already available in the `types` library (notably in the `listp` and `numberlistp` protocols) remain the recommended choice for non-geographic n-dimensional data.

6.71.6 Usage

Load the library:

```
| ?- logtalk_load(geospatial(loader)).
```

Validate a coordinate:

```
| ?- geospatial::valid_coordinate(geographic(38.7223, -9.1393)).
```

Normalize coordinates and convert to/from local planar coordinates:

```
| ?- geospatial::normalize_coordinate(geographic(95.0, 10.0), Normalized).
| ?- geospatial::equirectangular_projection(geographic(38.7223, -9.1393), 38.0, X, Y).
| ?- geospatial::equirectangular_inverse(X, Y, 38.0, Coordinate).
```

Compute default distance in kilometers (Haversine):

```
| ?- geospatial::distance(geographic(38.7223, -9.1393), geographic(41.1579, -8.6291),
↳haversine, Distance).
| ?- geospatial::rhumb_distance(geographic(38.7223, -9.1393), geographic(41.1579, -8.6291),
↳Distance).
| ?- geospatial::rhumb_bearing(geographic(38.7223, -9.1393), geographic(41.1579, -8.6291),
↳Bearing).
| ?- geospatial::rhumb_destination_point(geographic(38.7223, -9.1393), 45.0, 50.0,
↳Destination).
| ?- geospatial::interpolate_rhumb(geographic(38.7223, -9.1393), geographic(41.1579, -8.
↳6291), 0.5, Intermediate).
| ?- geospatial::rhumb_midpoint(geographic(38.7223, -9.1393), geographic(41.1579, -8.6291),
↳Midpoint).
```

Compute distance with explicit unit (meters, miles, or nautical_miles):

```
| ?- geospatial::distance(geographic(38.7223, -9.1393), geographic(41.1579, -8.6291),
↳vincenty, miles, Distance).
```

Compute route distance using the default metric (haversine) in kilometers:

```
| ?- geospatial::route_distance([geographic(38.7223, -9.1393), geographic(39.7440, -8.8070),
↳geographic(41.1579, -8.6291)], Distance).
```

Compute route distance with explicit metric and unit:

```
| ?- geospatial::route_distance([geographic(38.7223, -9.1393), geographic(39.7440, -8.8070),
↳geographic(41.1579, -8.6291)], vincenty, nautical_miles, Distance).
```

Compute midpoint and destination point:

```
| ?- geospatial::midpoint(geographic(38.7223, -9.1393), geographic(41.1579, -8.6291),
↳Midpoint).
| ?- geospatial::destination_point(geographic(38.7223, -9.1393), 45.0, 50.0, Destination).
| ?- geospatial::interpolate_great_circle(geographic(38.7223, -9.1393), geographic(41.1579, -
↳8.6291), 0.5, Intermediate).
| ?- geospatial::cross_track_distance(geographic(40.0, -9.0), geographic(38.7223, -9.1393),
↳geographic(41.1579, -8.6291), CrossTrackKm).
| ?- geospatial::along_track_distance(geographic(40.0, -9.0), geographic(38.7223, -9.1393),
↳geographic(41.1579, -8.6291), AlongTrackKm).
```

Compute final bearing and proximity checks:

```
| ?- geospatial::final_bearing(geographic(38.7223, -9.1393), geographic(41.1579, -8.6291),
↳Bearing).
| ?- geospatial::within_distance(geographic(38.7223, -9.1393), geographic(41.1579, -8.6291),
↳300.0, haversine).
```

Find the nearest coordinate from a list:

```
| ?- geospatial::nearest_coordinate(geographic(38.7223, -9.1393), [geographic(37.7749, -122.
↪4194), geographic(41.1579, -8.6291), geographic(40.4168, -3.7038)], vincenty, Nearest,
↪Distance).
```

Compute center and coordinate-list bounding boxes:

```
| ?- geospatial::mean_center([geographic(38.7223, -9.1393), geographic(41.1579, -8.6291),
↪geographic(40.4168, -3.7038)], Center).
| ?- geospatial::minimum_enclosing_circle([geographic(38.7223, -9.1393), geographic(41.1579,
↪-8.6291), geographic(40.4168, -3.7038)], Center, Radius).
| ?- geospatial::coordinates_bounding_box([geographic(38.7223, -9.1393), geographic(41.1579,
↪-8.6291), geographic(40.4168, -3.7038)], BoundingBox).
| ?- geospatial::bbox_from_coordinates([geographic(38.7223, -9.1393), geographic(41.1579, -8.
↪6291), geographic(40.4168, -3.7038)], BoundingBox).
```

Work with polygons:

```
| ?- geospatial::point_in_polygon(geographic(0.5, 0.5), [geographic(0.0, 0.0), geographic(0.
↪0, 1.0), geographic(1.0, 1.0), geographic(1.0, 0.0)]).
| ?- geospatial::polygon_area([geographic(0.0, 0.0), geographic(0.0, 1.0), geographic(1.0, 1.
↪0), geographic(1.0, 0.0)], Area).
| ?- geospatial::polygon_centroid([geographic(0.0, 0.0), geographic(0.0, 1.0), geographic(1.
↪0, 1.0), geographic(1.0, 0.0)], Centroid).
| ?- geospatial::polygon_bounding_box([geographic(1.0, -2.0), geographic(0.0, 1.0),
↪geographic(-1.0, -1.0)], BoundingBox).
| ?- geospatial::close_polygon([geographic(0.0, 0.0), geographic(0.0, 1.0), geographic(1.0,
↪0.0)], ClosedPolygon).
| ?- geospatial::polygon_orientation([geographic(0.0, 0.0), geographic(1.0, 0.0),
↪geographic(1.0, 1.0), geographic(0.0, 1.0)], Orientation).
| ?- geospatial::is_clockwise_polygon([geographic(0.0, 0.0), geographic(1.0, 0.0),
↪geographic(1.0, 1.0), geographic(0.0, 1.0)]).
| ?- geospatial::normalize_polygon_orientation([geographic(0.0, 0.0), geographic(0.0, 1.0),
↪geographic(1.0, 1.0), geographic(1.0, 0.0)], clockwise, Oriented).
| ?- geospatial::clockwise_polygon([geographic(0.0, 0.0), geographic(0.0, 1.0), geographic(1.
↪0, 1.0), geographic(1.0, 0.0)], Clockwise).
| ?- geospatial::counterclockwise_polygon([geographic(0.0, 0.0), geographic(1.0, 0.0),
↪geographic(1.0, 1.0), geographic(0.0, 1.0)], Counterclockwise).
| ?- geospatial::is_valid_polygon([geographic(0.0, 0.0), geographic(0.0, 1.0), geographic(1.
↪0, 0.0)]).
```

Work with bounding boxes:

```
| ?- geospatial::bbox_contains(bbox(geographic(-1.0, -1.0), geographic(1.0, 1.0)),
↪geographic(0.5, 0.5)).
| ?- geospatial::bbox_contains(bbox(geographic(-1.0, 170.0), geographic(1.0, -170.0)),
↪bbox(geographic(-0.5, 175.0), geographic(0.5, -175.0))).
| ?- geospatial::bbox_intersects(bbox(geographic(0.0, 0.0), geographic(1.0, 1.0)),
↪bbox(geographic(0.5, 0.5), geographic(2.0, 2.0))).
| ?- geospatial::bbox_overlaps(bbox(geographic(0.0, 0.0), geographic(1.0, 1.0)),
↪bbox(geographic(0.5, 0.5), geographic(2.0, 2.0))).
| ?- geospatial::bbox_intersects_polygon(bbox(geographic(0.0, 0.0), geographic(1.0, 1.0)),
↪[geographic(0.5, 0.5), geographic(0.5, 1.5), geographic(1.5, 1.5), geographic(1.5, 0.5)]).
| ?- geospatial::bbox_contains_polygon(bbox(geographic(0.0, 0.0), geographic(2.0, 2.0)),
```

(continues on next page)

(continued from previous page)

```

→[geographic(0.5, 0.5), geographic(0.5, 1.5), geographic(1.5, 1.5), geographic(1.5, 0.5)]).
| ?- geospatial::bbox_intersects_polyline(bbox(geographic(0.0, 0.0), geographic(1.0, 1.0)),_
→[geographic(-1.0, 0.5), geographic(2.0, 0.5)]).
| ?- geospatial::bbox_union(bbox(geographic(0.0, 0.0), geographic(1.0, 1.0)),_
→bbox(geographic(-1.0, 0.5), geographic(0.5, 2.0)), BoundingBox).
| ?- geospatial::bbox_expand(bbox(geographic(0.0, 0.0), geographic(0.0, 0.0)), 111.195,_
→ExpandedBoundingBox).

```

Compute nearest points and distances to paths:

```

| ?- geospatial::nearest_point_on_segment(geographic(1.0, 1.0), geographic(0.0, 0.0),_
→geographic(0.0, 2.0), Nearest).
| ?- geospatial::nearest_point_on_polyline(geographic(1.0, 1.0), [geographic(0.0, 0.0),_
→geographic(0.0, 2.0), geographic(2.0, 2.0)], Nearest, Distance).
| ?- geospatial::point_to_polyline_distance(geographic(1.0, 1.0), [geographic(0.0, 0.0),_
→geographic(0.0, 2.0)], Distance).

```

Compute polyline and polygon lengths:

```

| ?- geospatial::polyline_length([geographic(0.0, 0.0), geographic(0.0, 1.0), geographic(0.0,
→2.0)], Length).
| ?- geospatial::polyline_length([geographic(0.0, 0.0), geographic(0.0, 1.0)], vincenty,_
→Length).
| ?- geospatial::polyline_simplify([geographic(0.0, 0.0), geographic(0.5, 0.1), geographic(1.
→0, 0.0)], 20.0, Simplified).
| ?- geospatial::polyline_split_at_distance([geographic(0.0, 0.0), geographic(0.0, 1.0),_
→geographic(0.0, 2.0)], 111.195, Left, Right).
| ?- geospatial::polyline_resample([geographic(0.0, 0.0), geographic(0.0, 1.0), geographic(0.
→0, 2.0)], 50.0, Resampled).
| ?- geospatial::polygon_perimeter([geographic(0.0, 0.0), geographic(0.0, 1.0), geographic(1.
→0, 1.0), geographic(1.0, 0.0)], Perimeter).

```

Check polygon intersections:

```

| ?- geospatial::polygons_intersect([geographic(0.0, 0.0), geographic(0.0, 2.0),_
→geographic(2.0, 2.0), geographic(2.0, 0.0)], [geographic(1.0, 1.0), geographic(1.0, 3.0),_
→geographic(3.0, 3.0), geographic(3.0, 1.0)]).

```

6.72 git

This library provides access to a git project current branch and latest commit data (e.g., commit hash). Support for using this library on Windows operating-systems is experimental and may or may not work depending on the used Prolog backend.

6.72.1 API documentation

Open the `../..apis/library_index.html#git` link in a web browser.

6.72.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(git(loader)).
```

6.72.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(git(tester)).
```

6.72.4 Usage

All predicates take as first argument a directory, which should be either a git repository directory or a sub-directory of a git repository directory. The main predicate is `commit_log/3`. It provides access to the `git log` command when called with the `--oneline -n 1 --pretty=format: options`. By passing as second argument the desired format, it returns an atom with the formatted output. For example:

```
| ?- git::commit_log('/Users/pmoura/logtalk3', '%h%n%B', Output),
    write(Output), nl.

eccaa1a2a
Update SVG diagrams

Output = 'eccaa1a2a\nUpdate SVG diagrams\n'
yes
```

See e.g. the official documentation on `git log` pretty formats for details:

<https://git-scm.com/docs/pretty-formats>

Convenient predicates are also provided for commonly used formats such as the commit author and the commit hash. For example:

```
| ?- git::commit_author('/Users/pmoura/Documents/Logtalk/logtalk3', Author).

Author = 'Paulo Moura'
yes
```

(continues on next page)

(continued from previous page)

```
| ?- git::commit_hash('/Users/pmoura/Documents/Logtalk/logtalk3', Hash).
Hash = ecca1a2a9495fef441915bbace84e0a4b0394a2
yes
```

It's also possible to get the name of the current local branch. For example:

```
| ?- git::branch('/Users/pmoura/Documents/Logtalk/logtalk3', Branch).
Branch = master
yes
```

6.73 glicko2_ranker

Glicko-2 pairwise preference ranker. It applies the standard Glicko-2 rating, rating-deviation, and volatility update equations over a single synthetic rating period built from the aggregated pairwise outcomes of the dataset.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from pairwise preferences, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The current implementation requires a well-formed connected pairwise dataset so that learned ratings remain globally comparable across all ranked items.

6.73.1 API documentation

Open the ../apis/library_index.html#glicko2_ranker link in a web browser.

6.73.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(glicko2_ranker(loader)).
```

6.73.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(glicko2_ranker(tester)).
```

6.73.4 Features

- **Pairwise Preference Learning:** Learns one deterministic rating per item from pairwise outcomes.
- **Deterministic Single-Period Glicko-2 Semantics:** Applies the standard Glicko-2 rating-period update equations over one synthetic rating period built from the aggregated pairwise results.
- **Configurable Rating Parameters:** Exposes the initial rating, initial rating deviation, initial volatility, volatility constraint parameter, and volatility-solver tolerance as user options.
- **Integer Weight Fidelity:** Preference weights must be positive integers and are interpreted as repeated unit outcomes inside the same rating period.
- **Deterministic Ranking:** Orders candidate items by learned rating with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items, self-preferences, non-positive weights, and disconnected comparison graphs.
- **Extended Diagnostics:** Preserves per-item rating deviations and volatilities in the learned ranker diagnostics.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.
- **Shared Ranking Infrastructure:** Reuses the `ranking_protocols` helpers for option processing, dataset validation, diagnostics access, export, and candidate ranking.

6.73.5 Rating semantics

This implementation uses a deterministic batch interpretation of Glicko-2. The whole dataset is treated as a single synthetic rating period because the current pairwise dataset protocol does not encode timestamps or rating-period boundaries.

The implementation converts ratings and rating deviations to the internal Glicko-2 scale, applies the standard opponent-scaling factor $g(\phi)$, expected-score function, variance term v , improvement estimate δ , and volatility update iteration, and then converts the updated ratings and deviations back to the conventional Glicko scale.

Positive integer preference weights are replayed as repeated unit outcomes. Datasets using non-integer preference weights are rejected because they do not map cleanly to standard Glicko-2 update semantics.

6.73.6 Usage

Learning a ranker

```
% Learn from a pairwise ranking dataset object
| ?- glicko2_ranker::learn(my_dataset, Ranker).
...

% Learn with custom Glicko-2 parameters
| ?- glicko2_ranker::learn(my_dataset, Ranker, [initial_rating(1400.0), initial_
↪ deviation(300.0), initial_volatility(0.07), tau(0.4)]).
...
```

Inspecting diagnostics

```
% Inspect model, options, deviations, volatilities, and dataset metadata
| ?- glicko2_ranker::learn(my_dataset, Ranker),
    glicko2_ranker::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
...
```

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- glicko2_ranker::learn(my_dataset, Ranker),
    glicko2_ranker::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
...
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- glicko2_ranker::learn(my_dataset, Ranker),
    glicko2_ranker::export_to_clauses(my_dataset, Ranker, my_ranker, Clauses).
Clauses = [my_ranker(glicko2_ranker(...))]
...

% Export to a file
| ?- glicko2_ranker::learn(my_dataset, Ranker),
    glicko2_ranker::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.pl').
...
```

6.73.7 Options

The following options can be passed to the learn/3 predicate:

- `initial_rating(Rating)`: Initial rating assigned to every item.
- `initial_deviation(Deviation)`: Initial rating deviation assigned to every item.
- `initial_volatility(Volatility)`: Initial volatility assigned to every item.
- `tau(Tau)`: Positive volatility-constraint parameter used by the Glicko-2 update.
- `volatility_tolerance(Tolerance)`: Positive stopping tolerance used by the volatility root-finding iteration.

Datasets supplied to the ranker must use positive integer preference weights. Non-integer weights are rejected.

6.73.8 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(glicko2_ranker),
  options(Options),
  rating_deviations(Deviations),
  volatilities(Volatilities),
  dataset_summary(DatasetSummary)
]
```

6.73.9 Ranker representation

The learned ranker is represented by a compound term of the form:

```
glicko2_ranker(Items, Ratings, Diagnostics)
```

Where:

- `Items`: List of ranked items.
- `Ratings`: List of Item-Rating pairs.
- `Diagnostics`: List of metadata terms, including the effective options, per-item rating deviations, per-item volatilities, and dataset summary.

6.73.10 References

1. Glickman, M. E. (2012). *Example of the Glicko-2 system*.

6.74 glicko2_periodic_ranker

Multi-period Glicko-2 ranker over temporal pairwise game datasets. Applies the standard Glicko-2 rating, rating-deviation, and volatility update equations period by period using simultaneous player updates inside each declared rating period.

This library implements the `ranker_protocol` defined in the `ranking_protocols` library. It learns one rating per item from datasets implementing the `temporal_pairwise_ranking_dataset_protocol` protocol, processing declared rating periods in order and applying simultaneous Glicko-2 player updates inside each period using the standard Glicko-2 rating, rating-deviation, and volatility update equations.

Draws are represented directly using game scores on the set $\{0.0, 0.5, 1.0\}$. Players who are inactive in a declared period keep their rating and volatility while their rating deviation is inflated for that period.

Players are initialized when they first play instead of being forced to appear in the first declared period.

The library provides predicates for learning a ranker from temporal pairwise games, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `temporal_pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples.

6.74.1 API documentation

Open the `../apis/library_index.html#glicko2_periodic_ranker` link in a web browser.

6.74.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(glicko2_periodic_ranker(loader)).
```

6.74.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(glicko2_periodic_ranker(tester)).
```

6.74.4 Features

- **Temporal Pairwise Game Learning:** Learns one deterministic rating per item from temporal pairwise game results.
- **Deterministic Multi-Period Glicko-2 Semantics:** Processes the declared rating periods in order and applies simultaneous Glicko-2 player updates within each period.
- **Direct Draw Support:** Uses explicit game scores on the set $\{0.0, 0.5, 1.0\}$, allowing wins, draws, and losses to be represented directly.
- **Inactive-Period Deviation Inflation:** Players who do not play in a declared period keep their rating and volatility while their rating deviation is inflated for that period.
- **Late Player Initialization:** Players are initialized when they first appear in the dataset instead of being forced to play in the first declared period.
- **Configurable Rating Parameters:** Exposes the initial rating, initial rating deviation, initial volatility, volatility constraint parameter, and volatility-solver tolerance as user options.
- **Deterministic Ranking:** Orders candidate items by learned rating with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects duplicate periods, unknown periods, undeclared items, self-games, illegal scores, and disconnected comparison graphs.
- **Extended Diagnostics:** Preserves per-item rating deviations, volatilities, processed period count, and final period metadata in the learned ranker diagnostics.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.
- **Shared Ranking Infrastructure:** Reuses the `ranking_protocols` helpers for option processing, dataset validation, diagnostics access, export, and candidate ranking.

6.74.5 Rating semantics

This implementation uses the standard Glicko-2 update equations over explicit rating periods supplied by the dataset. The declared periods are processed in order, and all active players in a period are updated simultaneously from the pre-period ratings, deviations, and volatilities.

Game results are represented directly using scores in the set $\{0.0, 0.5, 1.0\}$. A score of 0.5 denotes a draw. Players who are inactive in a period keep their rating and volatility, but their rating deviation is inflated for that period according to the Glicko-2 inactivity rule.

Players are initialized when they first appear in a game rather than being required to occur in the first declared period. Items declared by the dataset but never seen in a game are initialized after training so they remain present in the learned ranker with the configured initial parameters.

6.74.6 Usage

Learning a ranker

```
% Learn from a temporal pairwise ranking dataset object
| ?- glicko2_periodic_ranker::learn(my_dataset, Ranker).
...

% Learn with custom Glicko-2 parameters
| ?- glicko2_periodic_ranker::learn(my_dataset, Ranker, [initial_rating(1400.0), initial_
↪ deviation(300.0), initial_volatility(0.07), tau(0.4)]).
...
```

Inspecting diagnostics

```
% Inspect model, options, deviations, volatilities, and period metadata
| ?- glicko2_periodic_ranker::learn(my_dataset, Ranker),
    glicko2_periodic_ranker::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
...
```

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- glicko2_periodic_ranker::learn(my_dataset, Ranker),
    glicko2_periodic_ranker::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
...
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- glicko2_periodic_ranker::learn(my_dataset, Ranker),
    glicko2_periodic_ranker::export_to_clauses(my_dataset, Ranker, my_ranker, Clauses).
Clauses = [my_ranker(glicko2_periodic_ranker(...))]
...

% Export to a file
| ?- glicko2_periodic_ranker::learn(my_dataset, Ranker),
    glicko2_periodic_ranker::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.pl').
...
```

6.74.7 Options

The following options can be passed to the learn/3 predicate:

- `initial_rating(Rating)`: Initial rating assigned to a player when it is first initialized.
- `initial_deviation(Deviation)`: Initial rating deviation assigned to a player when it is first initialized.
- `initial_volatility(Volatility)`: Initial volatility assigned to a player when it is first initialized.
- `tau(Tau)`: Positive volatility-constraint parameter used by the Glicko-2 update.
- `volatility_tolerance(Tolerance)`: Positive stopping tolerance used by the volatility root-finding iteration.

Datasets supplied to the ranker must use legal game scores from the set $\{0.0, 0.5, 1.0\}$.

6.74.8 Diagnostics syntax

The diagnostics/2 predicate returns a list of metadata terms with the form:

```
[
  model(glicko2_periodic_ranker),
  options(Options),
  rating_deviations(Deviations),
  volatilities(Volatilities),
  periods_processed(PeriodsProcessed),
  final_period(FinalPeriod),
  dataset_summary(DatasetSummary)
]
```

6.74.9 Ranker representation

The learned ranker is represented by a compound term of the form:

```
glicko2_periodic_ranker(Items, Ratings, Diagnostics)
```

Where:

- *Items*: List of ranked items.
- *Ratings*: List of Item-Rating pairs.
- *Diagnostics*: List of metadata terms, including the effective options, per-item rating deviations, per-item volatilities, processed period count, final period, and dataset summary.

6.74.10 References

1. Glickman, M. E. (2012). *Example of the Glicko-2 system*.

6.75 gradient_boosting_classifier

Gradient boosting classifier for tabular datasets using multinomial additive models fitted by regression trees to softmax residuals. At each boosting stage the implementation fits one regression tree per class and updates additive class scores using the configured learning rate.

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning a classifier from a dataset, using it to make predictions, estimating class probabilities, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. Continuous, categorical, and mixed-feature datasets are supported through the reused `regression_tree` backend.

6.75.1 API documentation

Open the ../docs/library_index.html#gradient_boosting_classifier link in a web browser.

6.75.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(gradient_boosting_classifier(loader)).
```

6.75.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(gradient_boosting_classifier(tester)).
```

6.75.4 Features

- **Multiclass Boosting:** Fits one additive score function per class and predicts using the highest boosted score.
- **Regression Tree Backend:** Reuses the `regression_tree` library to fit residual models at each boosting stage.
- **Probability Estimation:** Provides class probabilities using a softmax over the final additive scores.
- **Configurable Tree Complexity:** Exposes boosting-stage tree depth, minimum leaf size, and minimum variance reduction options.
- **Tabular Datasets:** Supports continuous, categorical, and mixed-feature datasets.
- **Classifier Export:** Learned classifiers can be exported as predicate clauses or written to a file.

6.75.5 Options

The `learn/3` predicate supports these options:

- `number_of_estimators/1` - number of boosting stages to fit (default: 25)
- `learning_rate/1` - shrinkage factor applied to each stage prediction (default: 0.1)
- `maximum_depth/1` - maximum depth of each regression tree (default: 3)
- `minimum_samples_leaf/1` - minimum number of examples in a leaf (default: 1)
- `minimum_variance_reduction/1` - minimum variance reduction required for a split (default: 0.0)
- `feature_scaling/1` - whether to scale continuous features in the regression-tree backend (default: false)

6.75.6 Usage

Learning a classifier

```
| ?- gradient_boosting_classifier::learn(weather, Classifier).
| ?- gradient_boosting_classifier::learn(iris_small, Classifier, [number_of_estimators(50),
↪ learning_rate(0.05)]).
```

Making predictions

```
| ?- gradient_boosting_classifier::learn(weather, Classifier),  
    gradient_boosting_classifier::predict(Classifier, [outlook-overcast, temperature-mild,   
↪humidity-high, windy-false], Class).  
  
| ?- gradient_boosting_classifier::learn(mixed, Classifier),  
    gradient_boosting_classifier::predict_probabilities(Classifier, [age-40, income-60000,   
↪student-yes, credit_rating-fair], Probabilities).
```

Exporting the classifier

```
| ?- gradient_boosting_classifier::learn(weather, Classifier),  
    gradient_boosting_classifier::export_to_clauses(weather, Classifier, classify, Clauses).  
  
| ?- gradient_boosting_classifier::learn(weather, Classifier),  
    gradient_boosting_classifier::export_to_file(weather, Classifier, classify, 'classifier.  
↪pl').
```

6.75.7 Classifier representation

The learned classifier is represented as a compound term with the form:

```
gradient_boosting_classifier(Classes, InitialScores, StageTrees, Options)
```

Where:

- **Classes:** list of class labels
- **InitialScores:** list of initial log-prior scores, one per class
- **StageTrees:** list of `stage_trees(ClassTrees)` terms, where each `ClassTrees` value contains `class_tree(Class, LearningRate, Tree)` terms
- **Options:** merged training options used to learn the classifier

When exported using `export_to_clauses/4` or `export_to_file/4`, this classifier term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.75.8 References

1. Friedman, J.H. (2001). “Greedy Function Approximation: A Gradient Boosting Machine”.
2. Hastie, T., Tibshirani, R. and Friedman, J. (2009). “The Elements of Statistical Learning”. Chapter 10.
3. Bishop, C.M. (2006). “Pattern Recognition and Machine Learning”. Section 14.4.

6.76 gradient_boosting_regression

Gradient boosting regression supporting continuous and mixed-feature datasets. Builds an additive ensemble of regression trees by repeatedly fitting the current residuals under squared-error loss. Starts from the arithmetic mean of the training targets and then adds a scaled tree prediction at each boosting stage.

The library implements the `regressor_protocol` defined in the `regression_protocols` library and learns an additive ensemble of regression trees by repeatedly fitting the current residuals under squared-error loss.

6.76.1 API documentation

Open the ../apis/library_index.html#gradient_boosting_regression link in a web browser.

6.76.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(gradient_boosting_regression(loader)).
```

6.76.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(gradient_boosting_regression(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(gradient_boosting_regression(tester_performance)).
```

6.76.4 Features

- **Residual Fitting:** Fits each regression tree to the residuals of the current additive model under squared-error loss.
- **Shrinkage:** Supports a fixed learning rate to scale stage contributions and reduce overfitting.
- **Continuous and Mixed Features:** Supports continuous attributes and categorical attributes encoded by the underlying regression-tree learner.
- **Tree Configuration:** Exposes the underlying regression-tree depth, minimum-leaf, variance-reduction, and scaling options.
- **Diagnostics Metadata:** Learned regressors record model name, target, training example count, initial prediction, fitted stage count, and effective options, accessible using the shared regression diagnostics predicates.
- **Model Export:** Learned regressors can be exported as predicate clauses or written to a file.
- **Reference Benchmarks:** Includes a dedicated performance suite reporting training time, RMSE, and MAE for representative regression datasets.

6.76.5 Regressor representation

The learned regressor is represented by default as:

- `gradient_boosting_regressor(InitialPrediction, WeightedTrees, Diagnostics)`

The exported predicate clauses therefore use the shape:

- `Functor(InitialPrediction, WeightedTrees, Diagnostics)`

In this representation, `WeightedTrees` contains `weighted_tree(LearningRate, Tree)` terms and `Diagnostics` stores training metadata including the effective options.

6.76.6 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(gradient_boosting_regression),
  target(Target),
  training_example_count(TrainingExampleCount),
  options(Options),
  initial_prediction(InitialPrediction),
  stage_count(StageCount)
]
```

Where:

- `model(gradient_boosting_regression)` identifies the learning algorithm that produced the regressor.
- `target(Target)` stores the target attribute name declared by the training dataset.
- `training_example_count(TrainingExampleCount)` stores the number of examples used during training.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `initial_prediction(InitialPrediction)` stores the constant prediction used to initialize the additive model before fitting any trees.
- `stage_count(StageCount)` stores the number of boosting stages that were actually fitted.

Use the `regression_protocols` `diagnostic/2` and `regressor_options/2` helper predicates when you only need a single metadata term or the effective options.

6.76.7 Options

The `learn/3` predicate accepts the following options:

- `number_of_estimators/1`: Maximum number of boosting stages to fit. Each stage adds one regression tree to the ensemble. The default is 50. Training can stop before reaching this limit when the residual sum of squares becomes negligible.
- `learning_rate/1`: Shrinkage factor applied to each stage prediction before it is added to the current model. Smaller values usually require more stages but can improve generalization. The default is 0.1.

- `maximum_depth/1`: Maximum depth allowed for each regression tree used as a base learner. Lower values produce weaker, simpler trees; higher values allow each stage to model more complex residual structure. The default is 3.
- `minimum_samples_leaf/1`: Minimum number of training examples required in each leaf of a base learner tree. Increasing this value makes the fitted trees more conservative and can reduce overfitting. The default is 1.
- `minimum_variance_reduction/1`: Minimum reduction in target variance required to accept a split when fitting a base learner tree. Larger values make tree growth stricter by rejecting weak splits. The default is 0.0.
- `feature_scaling/1`: Controls continuous-feature scaling in the underlying regression-tree learner. The accepted values are `true` and `false`. The default is `false`.

6.77 grammars

This library provides Definite Clause Grammars (DCGs) for common parsing tasks. The DCGs support parsing both lists of characters (aka chars) and lists of character codes (aka codes).

Currently, four groups of DCGs are available, each defined in its own file:

- blanks (`blank_grammars.lgt`)
- numbers (`number_grammars.lgt`)
- sequences (`sequence_grammars.lgt`)
- IP addresses (`ip_grammars.lgt`; depends on `number_grammars.lgt`)

6.77.1 API documentation

Open the ../apis/library_index.html#grammars link in a web browser.

6.77.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(grammars(loader)).
```

6.77.3 Testing

Minimal tests for this library predicates can be run by loading the `tester.lgt` file:

```
| ?- logtalk_load(grammars(tester)).
```

6.77.4 Usage

The library uses (when necessary) parametric objects where the single parameter can be either chars or codes. The parameter must be bound when using the DCGs. For example, when using implicit message-sending:

```
:- uses(blank_grammars(chars), [
    white_spaces//0, new_lines//0
]).
```

6.78 graphs

The graphs library implements predicates for handling directed/undirected weighted/unweighted graphs. Graphs are represented using a dictionary where keys are vertices and values are sorted lists of neighbors (vertex names for unweighted graphs, Vertex-Weight pairs for weighted graphs).

The four graph types are implemented as parametric objects where the parameter specifies the dictionary implementation to use (e.g., `avltree`, `rbtree`, `bintree`, or `splaytree`). This allows for easy experimentation (e.g., for performance analysis) with different dictionary backends. Non-parametric objects are also provided for each graph type using the `avltree` dictionary implementation.

The common protocol for all graph types is defined in `graph_protocol`. There are also `directed_graph_protocol` and `undirected_graph_protocol` for protocols specific to directed and undirected graphs, respectively. Additional predicates specific to each graph type (e.g., `transpose/2`, `topological_sort/2`, `min_path/5`, `min_tree/3`) are declared in the individual graph objects.

6.78.1 API documentation

Open the ../apis/library_index.html#graphs link in a web browser.

6.78.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(graphs(loader)).
```

6.78.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(graphs(tester)).
```

This library provides also a `graph_types` category for type-checking and arbitrary generation of graph-related terms.

6.78.4 Usage

Graph types

- `unweighted_directed_graph(Dictionary)` - Unweighted directed graph
- `unweighted_undirected_graph(Dictionary)` - Unweighted undirected graph
- `weighted_directed_graph(Dictionary)` - Weighted directed graph
- `weighted_undirected_graph(Dictionary)` - Weighted undirected graph

Creating graphs

Create an empty graph:

```
| ?- unweighted_directed_graph(avltree)::new(Graph).
```

Create a graph from a list of edges:

```
| ?- unweighted_directed_graph(avltree)::new([1-2, 2-3, 3-1], Graph).
```

Create a graph from vertices and edges:

```
| ?- unweighted_directed_graph(avltree)::new([1,2,3,4], [1-2, 2-3], Graph).
```

For weighted graphs, edges use the format (V1-V2)-Weight:

```
| ?- weighted_directed_graph(avltree)::new([(a-b)-5, (b-c)-3], Graph).
```

Querying graphs

```
| ?- unweighted_directed_graph(avltree)::new([1-2, 2-3], Graph),
    unweighted_directed_graph(avltree)::vertices(Graph, Vertices),
    unweighted_directed_graph(avltree)::edges(Graph, Edges),
    unweighted_directed_graph(avltree)::neighbors(1, Graph, Neighbors).
```

Undirected graphs

Undirected edges are stored internally as two directed edges. The `edges/2` predicate returns each undirected edge only once (with `V1 @=< V2`). For undirected graphs, `neighbors/3` excludes self-loops.

6.79 gsp_pattern_miner

GSP sequential pattern miner for sequence datasets. The library depends on the `sequential_pattern_mining_protocols` support library, implements the `generic_pattern_miner_protocol` defined in the `pattern_mining_protocols` core library, and mines frequent sequential patterns using prefix-indexed candidate joins, pruning candidates whose immediate subpatterns are not all frequent, and counting singleton and candidate supports in batched horizontal dataset scans.

Requires a dataset implementing `sequence_dataset_protocol` with sequences represented as ordered lists of canonical sorted itemsets over a declared item domain.

6.79.1 API documentation

Open the [../apis/library_index.html#gsp_pattern_miner](#) link in a web browser.

6.79.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(gsp_pattern_miner(loader)).
```

6.79.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(gsp_pattern_miner(tester)).
```

6.79.4 Features

- **Level-Wise Mining:** Builds longer frequent sequences by extending the current frequent level.
- **Join-Based Candidates:** Generates candidate sequences by joining compatible frequent patterns from the previous level.
- **Support-Pruned Candidates:** Prunes candidates whose immediate subpatterns are not all frequent.
- **Batched Horizontal Support Counting:** Counts singleton and candidate supports using horizontal scans over the sequence database at each level.
- **Canonical Sequences:** Validates that itemsets are sorted, duplicate-free, non-empty, and restricted to declared items.
- **Flexible Support Thresholds:** Supports relative minimum support and absolute minimum support count.
- **Model Export:** Mined pattern collections can be exported as predicate clauses or written to a file.

6.79.5 Options

The `mine/3` predicate accepts the following options:

- `minimum_support/1`: Relative minimum support threshold in the interval `[0.0, 1.0]`. The default is `0.5`.
- `minimum_support_count/1`: Absolute minimum support count. When both support options are provided, this option takes precedence.
- `maximum_pattern_length/1`: Maximum total number of items in a mined sequential pattern. The default is `1000`, effectively capped by the longest sequence in the dataset.
- `minimum_pattern_length/1`: Minimum total number of items retained in the mined result. The default is `1`.

6.79.6 Pattern miner representation

The mined pattern miner result is represented by a compound term with the functor chosen by the implementation and arity 3. For example:

```
gsp_pattern_miner(ItemDomain, Patterns, Options)
```

Where:

- ItemDomain: Canonical sorted list of declared dataset items.
- Patterns: List of `sequence_pattern(Pattern, SupportCount)` terms ordered first by total item count and then lexicographically.
- Options: Effective mining options used to mine the frequent sequential patterns.

6.79.7 References

1. Srikant, R. and Agrawal, R. (1996) - "Mining sequential patterns: Generalizations and performance improvements".

6.80 hashes

The hashes library provides portable implementations of several commonly used hashing algorithms. All hash objects implement the `hash_protocol` protocol by providing a `hash/2` predicate that takes a list of bytes and returns the computed hash as a lowercase hexadecimal atom.

The fixed-size cryptographic hash objects that can be safely used with HMAC (`md5`, `sha1`, `sha256`, `sha3_224`, `sha3_256`, `sha3_384`, and `sha3_512`) also implement the `hash_digest_protocol` protocol. This protocol adds `digest/2`, `digest_size/1`, and `block_size/1` predicates so that libraries such as `hmac` can compute keyed digests without duplicating hash function internals.

The library implements the following hashing algorithms:

- DJB2 32-bit (`djb2_32`)
- DJB2 64-bit (`djb2_64`)
- SDBM 32-bit (`sdbm_32`)
- SDBM 64-bit (`sdbm_64`)
- FNV-1a 32-bit (`fnv1a_32`)
- FNV-1a 64-bit (`fnv1a_64`)
- SipHash (`siphash_2_4`)
- CRC-32 parametric reflected implementation (`crc32_reflected(Polynomial)`)
- CRC-32 parametric non-reflected implementation (`crc32_non_reflected(Polynomial, Initial, FinalXor, AppendLength)`)
- CRC-32/ISO-HDLC (`crc32b`)
- CRC-32C/Castagnoli (`crc32c`)
- CRC-32/POSIX (`crc32posix`)
- CRC-32/MPEG-2 (`crc32mpeg2`)

- CRC-32/BZIP2 (`crc32bzip2`)
- CRC-32Q (aka CRC-32/AIXM) (`crc32q`)
- MurmurHash3 x86 32-bit (`murmurhash3_x86_32`)
- MurmurHash3 x86 128-bit (`murmurhash3_x86_128`)
- MurmurHash3 x64 128-bit (`murmurhash3_x64_128`)
- SHA3-224 (`sha3_224`)
- SHA3-256 (`sha3_256`)
- SHA3-384 (`sha3_384`)
- SHA3-512 (`sha3_512`)
- SHAKE128 (`shake128(OutputBytes)`)
- SHAKE256 (`shake256(OutputBytes)`)
- MD5 (`md5`)
- SHA1 (`sha1`)
- SHA256 (`sha256`)

The `djb2_64`, `sdbm_64`, `fnv1a_64`, `siphash_2_4`, `murmurhash3_x86_128`, `murmurhash3_x64_128`, `sha3_224`, `sha3_256`, `sha3_384`, `sha3_512`, `shake128(OutputBytes)`, `shake256(OutputBytes)`, `sha1`, and `sha256` objects are only loaded on backend Prolog compilers supporting unbounded integer arithmetic.

The SHAKE objects are parametric extensible-output functions. Pass the number of output bytes to generate when constructing the object.

The `crc32_reflected(Polynomial)` object implements a reflected CRC-32 family using initial value `0xFFFFFFFF` and final xor value `0xFFFFFFFF`, where `Polynomial` is the reflected CRC-32 polynomial.

The `crc32_non_reflected(Polynomial, Initial, FinalXor, AppendLength)` object implements a non-reflected CRC-32 family using a canonical polynomial, configurable initial and final xor values, and an `AppendLength` flag that controls whether the message length is appended as little-endian bytes.

The `crc32b` object implements the CRC-32/ISO-HDLC variant, also widely used by Ethernet, gzip, and PKZip. It uses reflected input/output processing, the reflected polynomial `0xEDB88320` (equivalent to the canonical polynomial `0x04C11DB7`), initial value `0xFFFFFFFF`, and final xor value `0xFFFFFFFF`.

The `crc32c` object implements the CRC-32C/Castagnoli variant using reflected input/output processing, the reflected polynomial `0x82F63B78` (equivalent to the canonical polynomial `0x1EDC6F41`), initial value `0xFFFFFFFF`, and final xor value `0xFFFFFFFF`.

The `crc32posix` object implements the CRC-32/POSIX variant used by the standard `cksum` utility. It is an instance of the non-reflected `crc32_non_reflected(Polynomial, Initial, FinalXor, AppendLength)` family: it uses the canonical polynomial `0x04C11DB7`, initial value `0x00000000`, processes bits most-significant first, appends the message length as little-endian bytes, and applies a final xor value of `0xFFFFFFFF`.

The `crc32mpeg2` object implements the CRC-32/MPEG-2 variant using the canonical polynomial `0x04C11DB7`, initial value `0xFFFFFFFF`, no appended length bytes, and final xor value `0x00000000`.

The `crc32bzip2` object implements the CRC-32/BZIP2 variant using the canonical polynomial `0x04C11DB7`, initial value `0xFFFFFFFF`, no appended length bytes, and final xor value `0xFFFFFFFF`.

The `crc32q` object implements the CRC-32Q variant, also used by AIXM-style formats, using the canonical polynomial `0x814141AB`, initial value `0x00000000`, no appended length bytes, and final xor value `0x00000000`.

The `siphash_2_4` object uses the standard reference key `00 01 02 ... 0f`. For custom keys, use the parametric object `siphash_2_4(Key)` where `Key` is a list of 16 bytes.

The implementations of the hashing algorithms make no attempt to validate that the input is a list of bytes. When necessary, use the types library `type::check(list(byte), Input)` goal before calling the hash/2 predicate.

6.80.1 API documentation

Open the ../apis/library_index.html#hashes link in a web browser.

6.80.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(hashes(loader)).
```

6.80.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(hashes(tester)).
```

6.80.4 Examples

Compute the SHA-256 hash of a text message:

```
| ?- atom_codes('The quick brown fox jumps over the lazy dog', Bytes),
    sha256::hash(Bytes, Hash).
Bytes = [84,104,101,32,113,117,105,99,107,32,98,114,111,119,110,32,102,111,120,32,106,117,
↪109,112,115,32,111,118,101,114,32,116,104,101,32,108,97,122,121,32,100,111,103],
Hash = 'd7a8fbb307d7809469ca9abcb0082e4f8d5651e46d3cdb762d02d0bf37c9e592'
yes
```

Compute a 32-byte SHAKE128 digest:

```
| ?- atom_codes('The quick brown fox jumps over the lazy dog', Bytes),
    shake128(32)::hash(Bytes, Hash).
Bytes = [84,104,101,32,113,117,105,99,107,32,98,114,111,119,110,32,102,111,120,32,106,117,
↪109,112,115,32,111,118,101,114,32,116,104,101,32,108,97,122,121,32,100,111,103],
Hash = 'f4202e3c5852f9182a0430fd8144f0a74b95e7417ecae17db0f8cfeed0e3e66e'
yes
```

Compute the CRC-32/ISO-HDLC checksum for the standard 123456789 test vector:

```
| ?- atom_codes('123456789', Bytes), crc32b::hash(Bytes, Hash).
Bytes = [49,50,51,52,53,54,55,56,57],
Hash = 'cbf43926'
yes
```

Compute the CRC-32C/Castagnoli checksum for the standard 123456789 test vector:

```
| ?- atom_codes('123456789', Bytes), crc32c::hash(Bytes, Hash).
Bytes = [49,50,51,52,53,54,55,56,57],
Hash = 'e3069283'
yes
```

Compute the CRC-32/POSIX checksum for the standard 123456789 test vector:

```
| ?- atom_codes('123456789', Bytes), crc32posix::hash(Bytes, Hash).
Bytes = [49,50,51,52,53,54,55,56,57],
Hash = '377a6011'
yes
```

Compute the CRC-32/MPEG-2 checksum for the standard 123456789 test vector:

```
| ?- atom_codes('123456789', Bytes), crc32mpeg2::hash(Bytes, Hash).
Bytes = [49,50,51,52,53,54,55,56,57],
Hash = '0376e6e7'
yes
```

Compute the CRC-32/BZIP2 checksum for the standard 123456789 test vector:

```
| ?- atom_codes('123456789', Bytes), crc32bzip2::hash(Bytes, Hash).
Bytes = [49,50,51,52,53,54,55,56,57],
Hash = 'fc891918'
yes
```

Compute the CRC-32Q checksum for the standard 123456789 test vector:

```
| ?- atom_codes('123456789', Bytes), crc32q::hash(Bytes, Hash).
Bytes = [49,50,51,52,53,54,55,56,57],
Hash = '3010bf7f'
yes
```

Use SipHash-2-4 with a custom key on a backend supporting unbounded integer arithmetic:

```
| ?- siphash_2_4([0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15])::hash([0,1,2,3], Hash).
Hash = 'cf2794e0277187b7'
yes
```

6.81 hdbscan_clusterer

Simplified HDBSCAN-style clusterer. It builds the mutual-reachability graph, computes a minimum spanning tree, derives the single-linkage hierarchy, condenses the hierarchy using `minimum_cluster_size`, and selects clusters using `eom` or `leaf selection`. Supports continuous attributes only.

The library implements the `clusterer_protocol` defined in the `clustering_protocols` library. It provides predicates for learning a clusterer from a dataset, assigning new instances to clusters, and exporting the learned clusterer as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol from the `clustering_protocols` library.

6.81.1 API documentation

Open the ../apis/library_index.html#hdbscan_clusterer link in a web browser.

6.81.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(hdbscan_clusterer(loader)).
```

6.81.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(hdbscan_clusterer(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(hdbscan_clusterer(tester_performance)).
```

6.81.4 Features

- **Hierarchical Density Clustering:** Builds the mutual-reachability graph, computes a minimum spanning tree, derives the single-linkage hierarchy, condenses it using `minimum_cluster_size`, and selects clusters using `eom` or `leaf` selection.
- **Continuous Datasets:** Accepts datasets containing only continuous attributes.
- **Cluster Selection Methods:** Supports both `eom` and `leaf` cluster selection.
- **Distance Metrics:** Supports Euclidean and Manhattan distances.
- **Optional Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Reachability-Based Prediction:** New instances are assigned to the selected cluster with the nearest training point when the distance is within the learned cluster reachability threshold; otherwise the atom noise is returned.
- **Noise Detection:** Points not assigned to any extracted cluster are retained as noise.
- **Portable Export:** Learned clusterers can be exported as clauses or files and reused later.

6.81.5 Options

The following options can be passed to the `learn/3` predicate:

- `minimum_points(MinimumPoints)`: Minimum neighborhood size used when computing core distances and mutual reachability. Default is 2.
- `minimum_cluster_size(MinimumClusterSize)`: Minimum number of points required for an extracted cluster. Default is 2.
- `cluster_selection_method(Method)`: Cluster extraction policy. Options: `eom` (default) or `leaf`.
- `distance_metric(Metric)`: Distance metric to use. Options: `euclidean` (default) or `manhattan`.

- `feature_scaling`(`FeatureScaling`): Whether to standardize continuous attributes before clustering. Options: on (default) or off.

6.81.6 Clusterer representation

The learned clusterer is represented as a compound term with the functor chosen by the user when exporting the clusterer and arity 4. For example:

```
hdbscan_clusterer(Encoders, Clusters, Noise, Options)
```

Where:

- Encoders: List of continuous attribute encoders storing attribute name, mean, and scale.
 - Clusters: List of `cluster`(`Id`, `Points`, `MaxCoreDistance`, `Stability`) terms in cluster-id order.
- Noise: List of encoded training points classified as noise.
- Options: Effective training options used to learn the clusterer.

6.82 heaps

This library defines a heap protocol and provides minimum and maximum heaps using two different implementations: pairing heaps and binary heaps.

The heap representations should be regarded as opaque terms, subjected to be changed without notice, and only accessed using the library predicates.

For backward-compatibility, the `legacy.lgt` file defines the `heapp` protocol extending the `heap_protocol` protocol plus the `heap/1`, `minheap`, and `maxheap` objects extending the `binary_heap/1` object. These legacy protocol and objects are deprecated and should not be used in new code.

6.82.1 API documentation

Open the ../apis/library_index.html#heaps link in a web browser.

6.82.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(heaps(loader)).
```

6.82.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(heaps(tester)).
```

6.82.4 Usage

Choose the heap implementation that best suits your needs. The pairing heap implementation is usually the best choice, specially if insertion and merging operations dominate in your use case. You can easily benchmark and compare the two implementations using either a `uses/2` directive with a parameter variable for the heap object and calling the heap predicates using implicit message-sending goals or a `uses/1` directive defining an object alias for the heap object and using the alias with explicit message-sending goals.

The two main library objects are `binary_heap(Order)` and `pairing_heap(Order)`, where `Order` is either `<` for a minimum heap or `>` for a maximum heap. For convenience, there are also `binary_heap_min` and `pairing_heap_min` objects for minimum heaps and `binary_heap_max` and `pairing_heap_max` objects for maximum heaps.

6.82.5 Credits

Original binary heap code by Richard O’Keefe and adapted to Logtalk by Paulo Moura and Victor Lagerkvist.

6.83 hierarchical_clustering

Hierarchical clusterer. Supports continuous attributes only. It builds the full bottom-up agglomerative_clusterer hierarchy and derives the requested partition by cutting the learned dendrogram at the largest remaining merge distances.

The library implements the `clusterer_protocol` defined in the `clustering_protocols` library. It provides predicates for learning a full agglomerative_clusterer hierarchy from a dataset, deriving a k-cluster cut for prediction, and exporting the learned clusterer as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol from the `clustering_protocols` library.

6.83.1 API documentation

Open the ../apis/library_index.html#hierarchical_clustering link in a web browser.

6.83.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(hierarchical_clustering(loader)).
```

6.83.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(hierarchical_clustering(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(hierarchical_clustering(tester_performance)).
```

6.83.4 Features

- **Full Hierarchy Learning:** Builds the complete bottom-up merge hierarchy before deriving the requested cluster partition.
- **Deterministic Dendrogram Cutting:** Produces the final k clusters by repeatedly splitting the highest remaining merge.
- **Continuous Datasets:** Accepts datasets containing only continuous attributes.
- **Linkage Strategies:** Supports single, complete, and average linkage.
- **Distance Metrics:** Supports Euclidean and Manhattan distances.
- **Optional Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Linkage-Aware Prediction:** New instances are assigned to the nearest learned cluster using the selected linkage strategy and distance metric rather than a prototype shortcut.
- **Incremental Distance Updates:** Training caches singleton pair distances once and updates merge distances incrementally for single, complete, and average linkage instead of recomputing all member-pair distances after every merge.
- **Heap-Based Merge Selection:** Training uses a lazy min-heap of candidate cluster pairs and skips stale entries for merged-away nodes instead of rescanning all active pairs at every agglomeration step.
- **Reusable Hierarchy Cuts:** Learned hierarchies can be re-cut to a different k using `cut/3` without retraining.
- **Rich Diagnostics:** Learned clusterers record merge counts, dendrogram height, heap rebuilds, scan fallbacks, maximum heap size, and effective options.
- **Deterministic Tie-Breaking:** Equal merge distances and equal split heights are resolved by preferring the smallest node-id pair.
- **Portable Export:** Learned clusterers can be exported as clauses or files and reused later.

6.83.5 Options

The following options can be passed to the `learn/3` predicate:

- `k(K)`: Number of clusters to retain after cutting the learned hierarchy. Default is 2.
- `linkage(Linkage)`: Linkage strategy to use. Options: `single`, `complete`, or `average` (default).
- `distance_metric(Metric)`: Distance metric to use. Options: `euclidean` (default) or `manhattan`.
- `feature_scaling(FeatureScaling)`: Whether to standardize continuous attributes before clustering. Options: `on` (default) or `off`.

6.83.6 Clusterer representation

The learned clusterer is represented as a compound term with the functor chosen by the user when exporting the clusterer and arity 5. For example:

```
hierarchical_clustering_clusterer(Encoders, hierarchy(RootState, MergeRecords, Dendrogram),  
↳Clusters, Prototypes, Diagnostics)
```

Where:

- `Encoders`: List of continuous attribute encoders storing attribute name, mean, and scale.

- `hierarchy(RootState, MergeRecords, Dendrogram)`: Reusable hierarchy state. `RootState` and `MergeRecords` are used internally by `cut/3`, and `Dendrogram` is the learned merge tree represented with `leaf(Id)` and `merge(Left, Right, Distance, Size)` terms.
- `Clusters`: List of `cluster(Id, Points)` terms for the selected k-cluster cut.
- `Prototypes`: List of average vectors kept for display and export metadata.
- `Diagnostics`: Diagnostics metadata including the effective training options used to learn the clusterer.

6.83.7 Additional API

- `cut(+Clusterer, +K, -RecutClusterer)`: Reuses the learned hierarchy to derive a new K-cluster cut without retraining.

6.83.8 Diagnostics

The `diagnostics/2` predicate returns metadata terms including:

- `model(hierarchical_clustering)`
- `cluster_count(Count)`
- `prototype_count(Count)`
- `training_example_count(Count)`
- `merge_count(Count)`
- `dendrogram_height(Height)`
- `heap_rebuild_count(Count)`
- `scan_fallback_count(Count)`
- `maximum_heap_size(Size)`
- `tie_breaking(node_id_order)`
- `options(Options)`

6.84 hierarchies

This library provides categories implementing reflection predicates over class and prototype hierarchies. These categories can be imported by any object that requires reasoning about hierarchies.

6.84.1 API documentation

Open the ../apis/library_index.html#hierarchies link in a web browser.

6.84.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(hierarchies(loader)).
```

6.84.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(hierarchies(tester)).
```

6.85 hmac

The `hmac` library provides a portable implementation of HMAC (Keyed-Hashing for Message Authentication) as specified in RFC 2104:

<https://datatracker.ietf.org/doc/html/rfc2104.html>

The library exports a single object, `hmac`, implementing the `hmac_protocol` protocol with the predicates:

- `digest/4`
- `hex_digest/4`
- `digest/5`
- `hex_digest/5`

The first argument is a hash object implementing the `hash_digest_protocol` protocol from the `hashes` library. Currently supported hash objects are:

- `md5`
- `sha1`
- `sha256`
- `sha3_224`
- `sha3_256`
- `sha3_384`
- `sha3_512`

On backend Prolog compilers supporting only bounded integer arithmetic, only `md5` is available. On backends supporting unbounded integer arithmetic, all the listed hash objects are available.

The `digest/5` and `hex_digest/5` predicates implement the truncation rule described in RFC 2104 and RFC 2202 by returning the requested number of leftmost digest bytes.

6.85.1 API documentation

Open the `../apis/library_index.html#hmac` link in a web browser.

6.85.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(hmac(loader)).
```

6.85.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(hmac(tester)).
```

6.85.4 Examples

Compute the HMAC-SHA-256 digest for a text message and return it as a hexadecimal atom:

```
| ?- atom_codes('Jefe', Key),
    atom_codes('what do ya want for nothing?', Message),
    hmac::hex_digest(sha256, Key, Message, Digest).
Key = [74,101,102,101],
Message = [119,104,97,116,32,100,111,32,121,97,32,119,97,110,116,32,102,111,114,32,110,111,
↪116,104,105,110,103,63],
Digest = '5bdcc146bf60754e6a042426089575c75a003f089d2739839dec58b964ec3843'
yes
```

Compute a truncated 16-byte HMAC-SHA-256 digest:

```
| ?- atom_codes('Jefe', Key),
    atom_codes('what do ya want for nothing?', Message),
    hmac::hex_digest(sha256, Key, Message, 16, Digest).
Key = [74,101,102,101],
Message = [119,104,97,116,32,100,111,32,121,97,32,119,97,110,116,32,102,111,114,32,110,111,
↪116,104,105,110,103,63],
Digest = '5bdcc146bf60754e6a042426089575c7'
yes
```

6.86 hodge_rank

HodgeRank pairwise measurement ranker. It builds a weighted graph Laplacian from the pairwise measurement support graph, solves the anchored normal equations using deterministic Gaussian elimination with partial pivoting and residual validation, and computes edge residuals against the fitted score differences.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from weighted signed pairwise measurements, using it to order candidate items, inspecting the resulting edge residuals, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_measurement_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The current implementation requires a well-formed connected pairwise measurement dataset so that learned scores remain globally comparable across all ranked items.

6.86.1 API documentation

Open the `../apis/library_index.html#hodge_rank` link in a web browser.

6.86.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(hodge_rank(loader)).
```

6.86.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(hodge_rank(tester)).
```

6.86.4 Features

- **Weighted Measurement Learning:** Learns one deterministic score per item from weighted signed pairwise measurements instead of winner/loser counts.
- **HodgeRank Global Component:** Solves the weighted graph-Laplacian least-squares system with a zero-sum anchoring convention.
- **Residual Inspection:** Exposes edge residuals through the `residuals/2` predicate so the non-global residual edge flow can be inspected explicitly.
- **Numerically Hardened Solver:** Uses Gaussian elimination with partial pivoting plus residual checks before accepting the learned scores.
- **Deterministic Ranking:** Orders candidate items by learned score with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items, self-measurements, non-numeric values, non-positive weights, and disconnected support graphs.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.

6.86.5 Scoring semantics

This implementation interprets each measurement fact

```
measurement(Item1, Item2, Value, Weight)
```

as a weighted signed observation for the oriented edge `Item1 -> Item2`. The learner fits a global zero-sum score vector `s` by minimizing the weighted least-squares objective

```
sum_ij Weight_ij * (Value_ij - (s_i - s_j))^2
```

subject to $\text{sum}(\text{scores}) = 0$. The resulting normal equations are a weighted graph-Laplacian system with one anchoring row enforcing the zero-sum gauge.

The learned scores are relative rather than probabilistic: only their differences and ordering matter. Larger positive values indicate items whose fitted global potential is higher than average. The `residuals/2` predicate returns the unexplained part of each observed edge measurement after removing the fitted score difference, i.e. the non-global residual edge flow left after fitting the global component.

6.86.6 Residual semantics

Residuals capture the part of each observed edge measurement that is not explained by the fitted global score difference, exposing the non-global residual edge flow left after fitting the global component.

6.86.7 Usage

Learning a ranker

```
% Learn from a pairwise measurement dataset object
| ?- hodge_rank::learn(my_dataset, Ranker).
...

% Learn with an explicit empty options list
| ?- hodge_rank::learn(my_dataset, Ranker, []).
...
```

The current implementation accepts only the empty options list `[]`. Any non-empty options list is rejected.

Inspecting residuals

```
% Inspect edge residuals from the fitted global scores
| ?- hodge_rank::learn(my_dataset, Ranker),
    hodge_rank::residuals(Ranker, Residuals).
Residuals = [...]
...
```

6.86.8 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(hodge_rank),
  options(Options),
  residuals(Residuals),
  residual_norm(ResidualNorm),
  dataset_summary(DatasetSummary)
]
```

6.86.9 Ranker representation

The learned ranker is represented by a compound term of the form:

```
hodge_rank_ranker(Items, Scores, Diagnostics)
```

Where:

- Items: List of ranked items.
- Scores: List of Item-Score pairs.
- Diagnostics: List of metadata terms, including the residuals, residual norm, effective options, and dataset summary.

6.86.10 References

1. Jiang, X., Lim, L.-H., Yao, Y., & Ye, Y. (2011). *Statistical ranking and combinatorial Hodge theory*.

6.87 hook_flows

Hook objects (i.e., objects that define term- and goal-expansion rules) can be combined to define expansion *workflows*. While in some cases the expansions are independent and thus can be applied in any order, in other cases a specific order is required. The `hook_pipeline` and `hook_set` parametric objects in this library implement the two most common scenarios of combining multiple hook objects for the expansion of source files. These parametric hook objects can be combined to define workflows of any complexity (e.g., a pipeline where one of the steps is set with an element that is a pipeline). These two basic hook flows can also be used as examples of how to construct your own custom expansion workflows.

6.87.1 API documentation

Open the `../..apis/library_index.html#hook-flows` link in a web browser.

6.87.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(hook_flows(loader)).
```

6.87.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(hook_flows(tester)).
```

6.87.4 Usage

Pre-processing a source file with a hook object can be accomplished by either compiling the source file using the option `hook(HookObject)` or by adding to the top of the file the directive:

```
:- set_logtalk_flag(hook, HookObject).
```

Note that `set_logtalk_flag/2` directives are local to an entity or a source file.

The `hook_pipeline(Pipeline)` is a parametric object where the parameter is a list of hook objects, interpreted as a pre-processing pipeline: the results of a hook object are passed to the next hook object. Terms and goals that are not expanded by a hook object are passed as-is to the next hook object. This parametric object is used when a set of expansions must be applied in a specific order. It also allows overriding the default compiler semantics where term-expansion rules are tried in sequence only until one of them succeeds.

The `hook_set(Set)` is a parametric object where the parameter is a list of hook objects, interpreted as a set of hook objects. The hook objects are tried in sequence until one of them succeeds in expanding the current term (goal) into a different term (goal). This parametric object is used when applying multiple independent expansions (i.e., expansions that don't apply to the same terms or goals).

Both the `hook_pipeline(Pipeline)` and `hook_set(Set)` objects assume that the individual hook objects publicly implement the expanding protocol (this is the default) as they send `term_expansion/2` and `goal_expansion/2` messages to those objects.

When using a backend Prolog compiler that supports modules, it's also possible to use as parameter a list of hook modules as long as their names do not coincide with the names of loaded objects.

6.88 hook_objects

This library provides a set of convenient hook objects for defining custom expansion workflows (using, e.g., the `hook_flows` library) and for debugging. They are usable and useful as-is but should also be regarded as term- and goal-expansion examples that you can learn from, clone, and change to fit your application requirements.

6.88.1 API documentation

Open the ../apis/library_index.html#hook-objects link in a web browser.

6.88.2 Loading

To load all hook objects in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(hook_objects(loader)).
```

To load a specific hook object, e.g., the `backend_adapter_hook` object:

```
| ?- logtalk_load(hook_objects(backend_adapter_hook)).
```

6.88.3 Testing

To test this library hook objects, load the `tester.lgt` file:

```
| ?- logtalk_load(hook_objects(tester)).
```

6.88.4 Usage

The provided hook objects cover different expansion scenarios as follows.

Using the Prolog backend adapter file expansion rules

Useful when defining a custom expansion workflow. This can be accomplished by loading the `backend_adapter_hook.lgt` file, which defines a `backend_adapter_hook` hook object that can be used as a workflow step.

Restoring the default compiler expansion workflow

In this case, load the `default_workflow_hook.lgt` file, which defines a `default_workflow_hook` hook object, and use the following **goal** to set the default hook object:

```
| ?- set_logtalk_flag(hook, default_workflow_hook).
```

Preventing applying any (other) user-defined expansion rules

When compiling a source file, we sometimes want to prevent applying expansion rules. This can be accomplished by simply loading the `identity_hook.lgt` file, which defines the `identity_hook` hook object, whose expansion rules simply succeed without changing the terms and goals, and setting it as the file-specific hook object writing as the first term in the file the **directive**:

```
:- set_logtalk_flag(hook, identity_hook).
```

Note that the compiler will always convert any grammar rules defined in the file into clauses. Although this conversion can also be performed as an expansion, grammar rules are part of the Logtalk language. If you want to preserve the grammar rules, use the hook objects described below to write them to a stream.

Expanding grammar rules into clauses independently of the compiler

Load the `grammar_rules_hook.lgt` and use the term-expansion rules in the `grammar_rules_hook` object. For example:

```
| ?- grammar_rules_hook::term_expansion((a --> [b],c), Clause).  
  
Clause = (a([b|T], C) :- c(T, C))  
yes
```

Using the expansion rules defined in a Prolog module

Load the `prolog_module_hook.lgt`, which defines the parametric hook object `prolog_module_hook(Module)`. To use this hook object, you need to instantiate the parameter to the name of the module. For example:

```
:- set_logtalk_flag(hook, prolog_module_hook(user)).
```

Wrap the contents of a plain Prolog file as an object

Load the `object_wrapper_hook.lgt`, which defines the `object_wrapper_hook/0-2` hook objects. Use them to wrap the contents of a plain Prolog file as an object named after the file (optionally implementing a protocol) or an object with the given name and object relations. Can be used to apply Logtalk developer tools to plain Prolog code or when porting a Prolog application to Logtalk. For example:

```
| ?- logtalk_load('plain.pl', [hook(object_wrapper_hook)]).
...

| ?- current_object(plain).
yes
```

Or:

```
| ?- logtalk_load('world_1.pl', [hook(object_wrapper_hook(some_protocol))]).
...

| ?- current_object(world_1).
yes

| ?- implements_protocol(world_1, Protocol).
Protocol = some_protocol
yes
```

Or:

```
| ?- logtalk_load('foo.pl', [hook(object_wrapper_hook(bar,[imports(some_category))])).
...

| ?- current_object(bar).
yes

| ?- imports_category(bar, Category).
Category = some_category
yes
```

The `object_wrapper_hook` object sets the `context_switching_calls` flag to allow for the generated object. This enables calling the predicates using the `(<<)/2` context-switching control construct. But it's usually better to define a protocol for the predicates being encapsulated and use instead the `object_wrapper_hook/1-2` objects.

Outputting term-expansion results to a stream or a file

Load the `write_to_stream_hook.lgt` file and use the `write_to_stream_hook(Stream)` or `write_to_stream_hook(Stream, Options)` hook objects. Alternatively, you can load the `write_to_file_hook.lgt` file and use the `write_to_file_hook(File)` or `write_to_file_hook(File, Options)` hook objects. The terms are not modified and thus these hook objects may be used at any point in an expansion workflow. The terms are written followed by a period and a new line.

For example, assume that we want to expand all terms in a `input.pl` source file, writing the resulting terms to a `output.pl` file, using the expansion rules defined in a `expansions` hook object. Taking advantage of the `hook_flows` library `hook_pipeline/1` object, we can write:

```
| ?- logtalk_compile(  
    'input.pl',  
    [hook(hook_pipeline(  
        expansions,  
        write_to_file_hook('output.pl')  
    ))]  
).
```

Printing entity predicate goals before or after calling them

This is helpful for quick debugging. Load the `print_goal_hook.lgt` file and use the `print_goal_hook` hook object. For example, we can set this hook object as the default hook:

```
| ?- set_logtalk_flag(hook, print_goal_hook).
```

Then, edit the entity source code to print selected goals:

```
foo :-  
    - bar,    % print goal before calling it  
    + baz,    % print goal after calling it  
    * quux.   % print goal before and after calling it
```

Suppressing goals

The `suppress_goal_hook.lgt` file provides the `suppress_goal_hook` hook object that supports suppressing a goal in a clause body by prefixing it using the `--` operator. We can set this hook object as the default hook using the goal:

```
| ?- set_logtalk_flag(hook, suppress_goal_hook).
```

If the expansion is only to be used in a single file, use instead the source file directive:

```
:- set_logtalk_flag(hook, suppress_goal_hook).
```

Then, edit entity predicates to suppress goals. For example:

```
foo :-  
    bar,  
    -- baz,  
    quux.
```

The suppressed goals are replaced by calls to `true/0`.

6.89 html

This library provides predicates for generating HTML content using either HTML 5 or XHTML 1.1 formats from a term representation. The library performs minimal validation, checking only that all elements are valid. No attempt is made to generate nicely indented output.

Normal elements are represented using a compound term with one argument (the element content) or two arguments (the element attributes represented by a list of Key=Value or Key-Value pairs and the element content). The element content can be another element or a list of elements. For example:

```
ol([type=a], [li(foo), li(bar), li(baz)])
```

The two exceptions are the pre or code elements whose content is never interpreted as an element or a list of elements. For example, the fragment:

```
pre([foo,bar,baz])
```

is translated to:

```
<pre>
[foo,bar,baz]
</pre>
```

Void elements are represented using a compound term with one argument, the (possibly empty) list of attributes represented by a list of Key=Value or Key-Value pairs. For example:

```
hr([class=separator])
```

Atomic arguments of the compound terms are interpreted as element content. Non-atomic element content can be represented as a quoted atom or by using the pre or code elements as explained above.

This library is a work in progress.

6.89.1 API documentation

Open the ../apis/library_index.html#html link in a web browser.

6.89.2 Loading

To load all entities in this library, load the loader.lgt file:

```
| ?- logtalk_load(html(loader)).
```

6.89.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(html(tester)).
```

6.89.4 Generating a HTML document

HTML documents can be generated from a compound term representation and written to a file or a stream. For example, assuming we want to generate a HTML 5 file:

```
| ?- html5::generate(  
    file('hello.html'),  
    html([lang=en], [head(title('Hello world!')), body(p('Bye!'))])  
).
```

When the second argument is a `html/1` or `html/2` compound term, a *doctype* is automatically written. If we prefer instead e.g. a XHTML 1.1 document, we use the `xhtml11` object:

```
| ?- xhtml11::generate(  
    file('hello.html'),  
    html([lang=en], [head(title('Hello world!')), body(p('Bye!'))])  
).
```

6.89.5 Generating a HTML fragment

It's also possible to generate just a fragment of a (X)HTML document by using a list of compound terms or a compound term for an element other than `html`. For example:

```
| ?- current_output(Stream),  
    html5::generate(stream(Stream), ul([li(foo), li(bar), li(baz)]))).  
  
<ul>  
<li>  
foo</li>  
<li>  
bar</li>  
<li>  
baz</li>  
</ul>  
  
Stream = ...
```

6.89.6 Working with callbacks to generate content

Often we need to programmatically generate HTML content from queries. In other cases, we may have fixed content that we don't want to keep repeating (e.g., a navigation bar). The library supports a `(: :)/2` pseudo-element that sends a message to an object to retrieve content. As an example, assume the following predicate definition in `user`:

```
content(strong('Hello world!')).
```

This predicate can then be called from the HTML term representation. For example:

```
| ?- current_output(Stream),
    html5::generate(stream(Stream), span(user::content)).

<span><strong>Hello world!</strong></span>

Stream = ...
```

Note that the callback always takes the form `Object::Closure` where `Closure` is extended with a single argument (to be bound to the generated content). More complex callbacks are possible by using lambda expressions.

6.89.7 Working with custom elements

The `html5` and `xhtml11` objects recognize the same set of standard HTML 5 normal and void elements and generate an error for non-standard elements. If you need to generate HTML content containing custom elements, define a new object that extends one of the library objects. For example:

```
:- object(html5custom,
    extends(html5)).

    normal_element(foo, inline).
    normal_element(bar, block).
    normal_element(Name, Display) :-
        ^^normal_element(Name, Display).

:- end_object.
```

6.90 ica_projection

Independent Component Analysis reducer for continuous datasets (missing or non-numeric values are rejected). The library implements the `dimension_reducer_protocol` defined in the `dimension_reduction_protocols` library and learns a linear unmixing projection by centering the training data, optionally standardizing continuous attributes, whitening the covariance matrix using the shared deterministic symmetric eigen-decomposition from `linear_algebra`, and then extracting independent components using a deterministic cubic FastICA fixed-point iteration with orthogonal deflation.

6.90.1 API documentation

Open the ../apis/library_index.html#ica_projection link in a web browser.

6.90.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ica_projection(loader)).
```

6.90.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ica_projection(tester)).
```

6.90.4 Features

- **Continuous Datasets:** Accepts datasets containing only continuous attributes.
- **Shared Symmetric Whitening Plus Fixed-Point ICA:** Whitens the training covariance matrix using the shared symmetric eigendecomposition from `linear_algebra` and then extracts independent directions using deterministic cubic FastICA with orthogonal deflation.
- **Always Centered:** Training and transform inputs are always centered using the training-set means before whitening and projection.
- **Optional Scaling:** Can optionally standardize each continuous attribute before whitening.
- **Training Diagnostics:** Records whitening eigenvalues plus per-component convergence reasons, iteration counts, and final deltas.
- **Projection API:** Transforms a new instance into a list of component_N-Value pairs.
- **Model Export:** Learned reducers can be exported as predicate clauses or written to a file.

6.90.5 Options

The `learn/3` predicate accepts the following options:

- `n_components/1`: Number of independent components to extract. Because training data is always centered, requests that exceed `min(feature_count, sample_count - 1)` raise `domain_error(component_count, Requested-Maximum)`. Requests that still exceed the numerical rank of the whitened covariance matrix raise `domain_error(component_count, Requested-Extracted)`. The default is 2.
- `feature_scaling/1`: Whether to divide each continuous attribute by its training-set standard deviation before whitening. Options: `true` or `false` (default).
- `maximum_iterations/1`: Maximum iteration bound used both by the whitening eigensolver and by the FastICA fixed-point steps for each independent component. The default is 1000.
- `tolerance/1`: Positive convergence tolerance used both for whitening rank detection and for FastICA fixed-point stopping. The default is `1.0e-8`.

The learned diagnostics also include:

- `whitening_eigenvalues(Values)`: Eigenvalues used to build the whitening transform, aligned with the extracted components.
- `convergence(Statues)`: Per-component stop reasons, such as tolerance or maximum_iterations_exhausted.
- `iterations(Counts)`: Per-component iteration counts aligned with the extracted components.
- `final_delta(Deltas)`: Per-component final update magnitudes aligned with the extracted components.

6.90.6 Usage

The following examples use the sample dataset shipped with the `dimension_reduction_protocols` library:

```
| ?- logtalk_load(dimension_reduction_protocols('test_datasets/mixed_independent_sources')).
```

Learning a reducer

```
| ?- ica_projection::learn(mixed_independent_sources, DimensionReducer).
| ?- ica_projection::learn(mixed_independent_sources, DimensionReducer, [n_components(2),
↪ feature_scaling(true), maximum_iterations(200), tolerance(1.0e-7)]).
```

Transforming new instances

```
| ?- ica_projection::learn(mixed_independent_sources, DimensionReducer),
    ica_projection::transform(DimensionReducer, [x1-(-5.0), x2-(-4.0), x3-(-4.0)],
↪ ReducedInstance).
```

Exporting and reusing the reducer

```
| ?- ica_projection::learn(mixed_independent_sources, DimensionReducer, [n_components(2)]),
    ica_projection::export_to_file(mixed_independent_sources, DimensionReducer, reducer,
↪ 'ica_reducer.pl').
| ?- logtalk_load('ica_reducer.pl'),
    reducer(Reducer),
    ica_projection::transform(Reducer, [x1-(-5.0), x2-(-4.0), x3-(-4.0)], ReducedInstance).
```

6.90.7 Dimension reducer representation

The learned dimension reducer is represented by a compound term with the functor chosen by the implementation and arity 3. For example:

```
ica_reducer(Encoders, Components, Diagnostics)
```

Where:

- Encoders: List of continuous attribute encoders storing attribute name, centering offset, and scale factor.
- Components: List of learned unmixing vectors in feature space.
- Diagnostics: Learned reducer metadata including the effective training options, whitening eigenvalues, and per-component convergence information.

6.90.8 References

1. Hyvarinen, A. and Oja, E. (2000) - “Independent Component Analysis: Algorithms and Applications”.

6.91 ids

This library generates random identifiers given the number of bytes of randomness. The identifiers are Base64 encoded. By default, 20 bytes (160 bits) are used.

The generation of random identifiers uses the `/dev/urandom` random number generator when available. This includes macOS, Linux, *BSD, and other POSIX operating-systems. On Windows, a pseudo-random generator is used, but randomized using the current wall time.

Identifiers can be generated as atoms, lists of characters, or lists of character codes.

See also the `cuid2`, `ksuid`, `nanoid`, `snowflakeid`, `uuid`, and `ulid` libraries.

6.91.1 API documentation

Open the ../apis/library_index.html#ids link in a web browser.

6.91.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ids(loader)).
```

6.91.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ids(tester)).
```

6.91.4 Usage

To generate an identifier using the default 160 bits of randomness:

```
| ?- ids::generate(Identifier).
Identifier = '2gpMzqAFXB05mYFIPX1qMkHxgGE='
yes
```

To generate an identifier represented by an atom using 240 bits (30 bytes) of randomness:

```
| ?- ids(atom, 30)::generate(Identifier).
Identifier = 'ie/jYcLsqo8ZguCOF1ZNPFDvJ03Ww5Qa9e0FxRB'
yes
```

To generate an identifier represented by a list of characters using 64 bits (8 bytes) of randomness:

```
| ?- ids(chars, 8)::generate(Identifier).
Identifier = ['5','0','8','V',d,'S',c,y,n,o,'A',=]
yes
```

To generate an identifier represented by a list of character codes using 64 bits (8 bytes) of randomness:

```
| ?- ids(codes, 8)::generate(Identifier).
Identifier = [111,81,86,55,99,79,70,77,65,74,103,61]
yes
```

6.92 ieee_754

The `ieee_754` library is a support package for parsing and generating IEEE 754 floating-point encodings shared by binary interchange libraries such as `message_pack`, `avro`, `protobuf`, `cbor`, and `wkt_wkb`.

6.92.1 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ieee_754(loader)).
```

6.92.2 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ieee_754(tester)).
```

6.92.3 Object model

The library provides two parameterized implementation objects:

```
ieee_754(Precision, ByteOrder, NaNRepresentation)
```

`ieee_754_fields(Precision, ByteOrder)`

with the following parameter values:

- Precision: half, single, or double
- ByteOrder: big or little
- NaNRepresentation: canonical or payloads

The parameterization keeps format-specific policy out of the shared core. Advanced callers can use `ieee_754_fields` to inspect exact sign, exponent, mantissa, finite binary-rational decompositions, and NaN payload bits without going through backend float values.

6.92.4 Value representation

The intended term representation is:

- finite values: backend Prolog floats
- positive infinity: `@infinity`
- negative infinity: `@negative_infinity`
- canonical NaN: `@not_a_number`
- payload-preserving NaN values: `not_a_number(Bytes)`

The `not_a_number(Bytes)` representation is only intended to be accepted and produced by objects configured with `NaNRepresentation = payloads`.

Under the payloads policy, decoding a canonical quiet NaN should still yield `@not_a_number`. Non-canonical NaN encodings should decode to `not_a_number(Bytes)` using canonical byte order for the selected precision.

The high-level `ieee_754/3` object API deliberately abstracts away the IEEE 754 quiet-versus-signaling NaN distinction. In canonical mode, all NaN encodings decode to `@not_a_number`. In payloads mode, the designated canonical quiet NaN encoding decodes to `@not_a_number` while all other NaN encodings decode to `not_a_number(Bytes)` for roundtrip preservation. Callers that need exact NaN inspection should use `ieee_754_fields/2` object API.

6.92.5 Public API

The high-level value API is defined in `ieee_754_protocol.lgt` and consists of:

- `parse/2` Decodes `bytes(Bytes)` or `bits(Bits)` into a value term.
- `generate/2` Encodes a value term into `bytes(Bytes)` or `bits(Bits)`.
- `generate/3` Encodes a value term to a byte list with an open tail for difference-list based binary generators.
- `valid/1` Checks whether a term is a valid value representation for the selected object.
- `exactly_representable/1` Checks whether a term can be encoded and decoded in the selected precision without loss of value information.
- `precision/1` Returns the selected precision.
- `order/1` Returns the selected byte order.
- `nan_representation/1` Returns the selected NaN representation policy.
- `byte_count/1` Returns the number of bytes used by the selected precision.

The low-level exact field API is defined in `ieee_754_fields_protocol.lgt` and consists of:

- `classify/2` Classifies `bytes(Bytes)` or `bits(Bits)` as `zero`, `subnormal`, `normal`, `infinity`, or `not_a_number`.
- `fields/5` Extracts the exact sign bit, exponent bits, mantissa bits, and class.
- `finite_binary_rational/4` Extracts finite encodings as exact $(-1)^{\text{Sign}} * \text{Significand} * 2^{\text{Exponent}}$ terms.
- `nan_payload/2` Extracts the raw NaN mantissa payload bits.
- `nan_kind/2` Classifies a NaN encoding as `quiet` or `signaling`.
- `precision/1` Returns the selected precision.
- `order/1` Returns the selected byte order.
- `byte_count/1` Returns the number of bytes used by the selected precision.

6.92.6 Semantics

The intended shared semantics are:

- finite encoding rounds according to the selected IEEE 754 precision
- `exactly_representable/1` succeeds only when no value information is lost
- signed zero is preserved when supported by the backend Prolog compiler
- infinities and NaN values are represented independently from backend float syntax
- subnormal values are supported for all selected precisions
- payload-preserving NaN decoding is available when explicitly requested
- the `ieee_754_fields` API works from exact encodings and therefore does not depend on backend float decomposition behavior
- exact quiet/signaling NaN inspection is exposed by `ieee_754_fields/2` instead of the high-level value codec API

6.93 intervals

This library provides:

- an `interval_protocol` protocol and an `interval` object implementing the 13 Allen base interval relations plus interval pair classification using `relation/3`
- an `interval_algebra_protocol` protocol and an `interval_algebra` object implementing Allen relation algebra predicates over the 13 base relation atoms
- an `interval_relation_set_protocol` protocol and an `interval_relation_set` object implementing canonical relation-set operations over Allen base relation atoms
- an `interval_constraint_network_protocol` protocol and an `interval_constraint_network` object implementing first interval constraint-network predicates over canonical relation sets, including batch refinement, explanation lists, and network inspection/comparison helpers

The `interval` object works on interval terms represented as `i(Start, End)`. The `interval_algebra` object works on relation atoms such as `before`, `overlaps`, and `contains`. The `interval_relation_set` object works on canonical ordered duplicate-free lists of relation atoms. The `interval_constraint_network` object works on `network(Nodes, Constraints)` terms whose pair constraints are relation sets. Internally, the `interval_constraint_network` object compiles that public representation to an a dense indexed form for propagation and for most query operations, while still returning ordinary `network/2` terms at the API boundary.

The `interval_algebra::compose/3` predicate returns canonical ordered duplicate-free lists of base relation atoms.

6.93.1 Practical limits

The `interval_constraint_network` object is intended for small-to-medium symbolic networks. It uses a complete `network(Nodes, Constraints)` representation, so the number of stored pair constraints grows quadratically with the number of nodes. Propagation is based on path consistency and now uses an internal dense indexed representation with direct pair access, a reverse node-to-index table, and scheduled-pair tracking to avoid repeated linear worklist duplicate checks. This substantially reduces overhead relative to a purely list-based implementation, but the representation is still dense and updates still rebuild parts of the internal matrix before the result is decompiled back to a public `network/2` term.

In practice, the network predicates are suitable for moderate qualitative temporal reasoning workloads, but they should not be treated as a large-scale or latency-sensitive temporal CSP solver. Exact limits depend on the Prolog backend, network density, how often closure is recomputed, and how often callers invoke single-step API operations that must compile or decompile the public network representation.

6.93.2 API documentation

Open the ../apis/library_index.html#intervals link in a web browser.

6.93.3 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(intervals(loader)).
```

6.93.4 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(intervals(tester)).
```

6.93.5 Examples

Classify two valid intervals according to their unique Allen base relation:

```
| ?- interval::new(1, 3, I1), interval::new(3, 5, I2), interval::relation(I1, I2, Relation).
Relation = meets.
```

Compose two Allen base relation atoms:

```
| ?- interval_algebra::compose(meets, met_by, Relations).
Relations = [finishes, finished_by, equal].
```

Normalize an Allen relation set:

```
| ?- interval_relation_set::normalize([during, before, during], RelationSet).
RelationSet = [before, during].
```

Propagate interval constraints across a small symbolic network:

```
| ?- interval_constraint_network::new([a, b, c], Network0),
    interval_constraint_network::refine(Network0, a, b, [before], Network1),
    interval_constraint_network::refine(Network1, b, c, [before], Network2),
    interval_constraint_network::path_consistency(Network2, Network3),
    interval_constraint_network::relation(Network3, a, c, RelationSet).
RelationSet = [before].
```

Check whether a query is entailed by the current symbolic constraints:

```
| ?- interval_constraint_network::entails(Network3, a, c, [before, meets]).
true.
```

Inspect a simple explanation for an entailed relation:

```
| ?- interval_constraint_network::entails(Network3, a, c, [before], Explanation).
Explanation = propagated(b, [before], [before], [before]).
```

Inspect a simple contradiction explanation after closure:

```
| ?- interval_constraint_network::contradiction(Closure, Explanation).
Explanation = contradiction(a, c, propagated(b, [before], [before], [before])).
```

Refine and propagate in a single operation that fails on contradiction:

```
| ?- interval_constraint_network::new([a, b, c], Network0),  
    interval_constraint_network::refine_propagate(Network0, a, b, [before], Network1).  
Network1 = ...
```

Post a batch of constraints and propagate them in one step:

```
| ?- interval_constraint_network::refine_propagate(Network0, [constraint(a, b, [before]),  
↳constraint(b, c, [before])], Network1).  
Network1 = ...
```

Collect direct and propagated changes while propagating:

```
| ?- interval_constraint_network::propagate(Network2, Network3, Changes).  
Changes = [...].
```

Query which node triples caused the propagated refinements:

```
| ?- interval_constraint_network::propagation_triples(Changes, Triples).  
Triples = [triple(a, b, c)].
```

List all immediate explanations supporting an entailed query:

```
| ?- interval_constraint_network::entailment_explanations(Network3, a, c, [before, meets],  
↳Explanations).  
Explanations = [propagated(b, [before], [before], [before])].
```

List all contradiction explanations currently present in an inconsistent network:

```
| ?- interval_constraint_network::contradiction_explanations(Closure, Explanations).  
Explanations = [...].
```

Inspect a closure and compare networks by generality or equivalence:

```
| ?- interval_constraint_network::constraints(Network3, Constraints).  
Constraints = [constraint(a, b, [before]), constraint(a, c, [before]), constraint(b, c,  
↳[before])].  
  
| ?- interval_constraint_network::subsumes(Network0, Network3).  
true.  
  
| ?- interval_constraint_network::equivalent(Network3, Network5).  
true.
```

6.94 iqr_anomaly_detector

Statistical interquartile-range anomaly detector for continuous datasets. It is a statistical anomaly-detection method based on Tukey interquartile fences: for each known continuous attribute value it learns $Q1$ and $Q3$, computes the exceedance beyond $[Q1, Q3]$ in interquartile-range units normalized by the learned $fence_multiplier/1$, and then aggregates the per-attribute normalized deviations according to $score_mode/1$, so any value at or beyond $[Q1 - k*IQR, Q3 + k*IQR]$ reaches the default anomaly boundary when using $fence_multiplier(k)$.

The library implements the `anomaly_detector_protocol` defined in the `anomaly_detection_protocols` library. It learns a detector from a continuous dataset, computes anomaly scores for new instances, predicts normal or anomaly, and exports learned detectors as clauses or files.

Datasets are represented as objects implementing the `anomaly_dataset_protocol` protocol from the `anomaly_detection_protocols` library. See the `anomaly_detection_protocols/test_datasets` directory for examples.

6.94.1 API documentation

Open the ../apis/library_index.html#iqr_anomaly_detector link in a web browser.

6.94.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(iqr_anomaly_detector(loader)).
```

6.94.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(iqr_anomaly_detector(tester)).
```

6.94.4 Features

- **Statistical method:** implements anomaly detection based on interquartile-range fences, using per-attribute first and third quartiles to measure how far new observations deviate from the central baseline distribution.
- **Quartile-based scoring:** for each known attribute value x , the library computes the positive exceedance of x beyond the interval $[Q1, Q3]$ in interquartile-range units, where $Q1$ and $Q3$ are the learned sample quartiles.
- **Continuous features only:** accepts datasets whose declared attributes are all continuous.
- **Robust statistics:** reuses the statistics library `sample_object_quartiles/4` predicate to compute per-attribute quartiles.
- **Baseline training selection:** supports learn-time `baseline_class_values(ClassValues)` and `baseline_selection_policy(Policy)` options. The default baseline class values are `[normal]`. The default reject policy throws an error if non-baseline examples are present, while `filter` removes them before fitting.

- **Missing-value tolerant:** ignores missing values when fitting attribute statistics. During scoring, queries must provide at least one known value. In the default `score_mode(root_mean_square)`, the raw score is aggregated over attributes with positive normalized deviation so that neutral inlier attributes do not dilute fence anomalies. The learned detector stores a precomputed attribute schema so that scoring reuses the same attribute ordering without rebuilding it on every call.
- **Configurable scoring semantics:** supports both dense multivariate deviation scoring using `score_mode(root_mean_square)` and sparse anomaly detection using `score_mode(any_feature_extreme)`. The default root-mean-square mode reuses the numberlist library Euclidean norm predicate as part of the computation.
- **Configurable Tukey fences:** supports a learn-time `fence_multiplier/1` option. The default 1.5 corresponds to the classical Tukey inner fence cutoff, and the learned multiplier is applied directly in the score path.
- **Bounded scoring:** maps the raw multivariate IQR score to $[0.0, 1.0)$ using $\text{Score} = \text{Raw} / (1 + \text{Raw})$.
- **Default threshold:** the default `anomaly_threshold(0.5)` corresponds to the learned Tukey fence cutoff after score scaling, while remaining overrideable in `learn/3` and `predict/4`.
- **Learn-time options:** `fence_multiplier/1` and `score_mode/1` are recorded in the learned detector and reused for subsequent scoring and prediction. Passing either option to `predict/4` does not override the learned value.
- **All-missing queries rejected:** scoring and prediction throw a `domain_error(non_empty_known_values, AttributeNames)` exception when every declared feature is missing in the query.
- **Featureless datasets rejected:** datasets must declare at least one continuous feature; otherwise `learn/2-3` throws a `domain_error(non_empty_features, Dataset)` exception.
- **Detector export:** learned detectors can be exported as predicate clauses.
- **Explicit validation and diagnostics:** supports the shared `check_anomaly_detector/1`, `valid_anomaly_detector/1`, `diagnostics/2`, `diagnostic/2`, and `anomaly_detector_options/2` predicates.

6.94.5 Options

The following options are supported by the public API:

- `anomaly_threshold(Threshold)`: Threshold for `predict/3-4` (default: 0.5)
- `baseline_class_values(ClassValues)`: Learn-time class labels that are admissible for baseline fitting (default: `[normal]`)
- `baseline_selection_policy(Policy)`: Learn-time handling of examples whose class is not listed in `baseline_class_values/1`. Supported values are `filter` and `reject` (default: `reject`)
- `fence_multiplier(Multiplier)`: Learn-time Tukey fence multiplier stored in the learned detector (default: 1.5)
- `score_mode(Mode)`: Learn-time score aggregation mode for `learn/3`. Supported values are `root_mean_square` and `any_feature_extreme` (default: `root_mean_square`). If passed to `predict/4`, it is ignored and the value stored in the learned detector is used.

6.94.6 Detector representation

The learned detector is represented by default as:

```
iqr_detector(TrainingDataset, AttributeSchema, Encoders, Diagnostics)
```

Where:

- TrainingDataset: training dataset object identifier
- AttributeSchema: precomputed attribute ordering used for validation and scoring
- Encoders: list of iqr_anomaly_detector(Attribute, Q1, Q3, Scale) records
- Diagnostics: learned metadata terms including model/1, training_dataset/1, attribute_names/1, feature_count/1, example_count/1, and options/1

When exported using export_to_clauses/4 or export_to_file/4, this detector term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.94.7 Notes

Scoring has three stages. First, the detector computes one per-attribute IQR deviation score for each known attribute value using its exceedance beyond the interval [Q1, Q3] in interquartile-range units. Second, each per-attribute score is normalized by the learned fence_multiplier/1, so that it reaches 1.0 exactly at the chosen Tukey fence cutoff. Third, those normalized per-attribute scores are aggregated into a single raw deviation score according to the learned score_mode/1 option before being mapped to the interval [0.0, 1.0) using $\text{Score} = \text{Raw} / (1 + \text{Raw})$.

The score_mode/1 option does not change the per-attribute quartile formula. It only changes the aggregation step. With score_mode(root_mean_square), the raw score is the root mean square of the positive normalized per-attribute deviations, computed over the attributes with positive deviation so that inlier padding does not dilute fence-reaching anomalies. With score_mode(any_feature_extreme), the raw score is the maximum normalized per-attribute deviation.

The fence_multiplier/1 option defines the classical Tukey anomaly cutoff and directly normalizes the per-attribute deviation scores. With any learned fence_multiplier(K), a per-attribute score of 1.0 means the query has reached the chosen Tukey fence on that attribute, so the default normalized threshold 0.5 corresponds to that cutoff in both supported aggregation modes.

The baseline_class_values/1 option declares which dataset class labels are admissible for fitting the baseline quartiles and interquartile ranges. The baseline_selection_policy/1 option then controls what happens when other labels are present in the training data. The default reject policy raises a domain_error(baseline_only_training_data, Dataset) exception when any non-baseline example is found. The filter policy removes non-baseline examples before fitting.

Attributes with zero observed interquartile range are assigned a fallback scale of 1.0. This keeps the detector well-defined for singleton datasets or constant columns while still yielding zero score for matching values and positive scores for deviating values.

The root-mean-square aggregation keeps the default threshold stable while avoiding dilution from missing or neutral inlier attributes.

Use score_mode(any_feature_extreme) when a single extreme feature should be sufficient to flag an anomaly in high-dimensional data.

6.95 iso_639

This library provides ISO 639 sets 1, 2, 3, and 5 lookups generated from the Library of Congress and SIL maintenance-agency sources. It requires Unicode support from the backend Prolog compiler.

The set 1 and set 2 data are generated from the Library of Congress ISO 639-2 code list. The set 3 data are generated from the SIL tab-delimited code set download. The set 5 data are generated from the Library of Congress ISO 639-5 identifier list.

The refresh tooling lives in `library/iso_639/scripts/` together with the checked-in source snapshots used to regenerate the fact tables using shell scripts only.

6.95.1 API documentation

Open the `../..//apis/library_index.html#iso_639` link in a web browser.

6.95.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(iso_639(loader)).
```

6.95.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(iso_639(tester)).
```

6.96 iso_3166

This library provides generated ISO 3166 country and subdivision code lookups. The current worktree includes a real ISO 3166-1 fact table generated from the official United Nations M49 overview page, which republishes ISO alpha-2 and alpha-3 country codes together with the M49 numeric codes used here. It requires Unicode support from the backend Prolog compiler.

ISO 3166-2 subdivision data is bundled in the public facade and generated from the Debian `iso-codes` machine-readable JSON snapshot. Subdivision codes and country alpha-2 codes are stored as lowercase atoms to match the conventions used by the generated ISO 3166-1 country fact table, while names and categories preserve the source spelling.

To refresh the generated subdivision table, run:

```
$ library/iso_3166/scripts/update_iso_3166_2.sh
```

6.96.1 API documentation

Open the ../apis/library_index.html#iso_3166 link in a web browser.

6.96.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(iso_3166(loader)).
```

6.96.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(iso_3166(tester)).
```

6.97 iso_4217

This library provides ISO 4217 active non-fund currency code lookups generated from the authoritative SIX maintenance-agency XML list together with active fund currency code lookups. It requires Unicode support from the backend Prolog compiler.

The generated fact table keeps one entry per XML row with a currency code, therefore preserving entity-specific assignments for shared currencies such as the Euro or the US Dollar. Rows marked by SIX as fund entries are exposed using the `fund_currency/5` predicate while the `currency/5` predicate remains scoped to non-fund entries. Minor units are represented either as integers or the atom `na` when the source lists N.A..

6.97.1 API documentation

Open the ../apis/library_index.html#iso_4217 link in a web browser.

6.97.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(iso_4217(loader)).
```

6.97.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(iso_4217(tester)).
```

6.98 iso_9362

This library provides ISO 9362 Business Identifier Code (BIC) structure parsing and normalization based on the public Swift and ISO 9362 documentation. It requires Unicode support from the backend Prolog compiler.

The current public implementation focuses on the normative structure that Swift and ISO 9362 publish openly: a BIC consists of a 4-character business party prefix, a 2-character ISO 3166-1 alpha-2 country code, a 2-character business party suffix, and an optional 3-character branch identifier. Eight-character primary-office BICs are normalized to the branch code XXX.

This library does not currently ship the full ISO 9362 directory as facts. The public SwiftRef TXT directory download is available only through a secure, time-limited email link, so it is not treated as an automatable checked-in source snapshot in the same way as the ISO 4217 and ISO 639 libraries.

6.98.1 API documentation

Open the `../apis/library_index.html#iso_9362` link in a web browser.

6.98.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(iso_9362(loader)).
```

6.98.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(iso_9362(tester)).
```

6.99 iso_13616

This library provides ISO 13616 International Bank Account Number (IBAN) structure parsing, checksum validation, and normalization based on the public ISO 13616 and Swift IBAN registry documentation. It requires Unicode support from the backend Prolog compiler.

The current public implementation includes a checked-in snapshot of the public SWIFT IBAN registry for all currently registered IBAN countries. Validation therefore checks:

- the ISO 3166-1 alpha-2 country code
- the country-specific total IBAN length
- the country-specific BBAN structure pattern
- the standard MOD-97 checksum

The checked-in registry facts use a derived Prolog segment representation for BBAN patterns, e.g. `[a-4, n-6, n-8]` instead of the original SWIFT text syntax `4!a6!n8!n`. This keeps the published structure information while making validation code simpler and more direct.

The public API remains structural. The library validates the published national IBAN formats but does not attempt to ship any bank directory or account existence data beyond the public registry pattern definitions.

6.99.1 API documentation

Open the ../apis/library_index.html#iso_13616 link in a web browser.

6.99.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(iso_13616(loader)).
```

6.99.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(iso_13616(tester)).
```

6.100 isolation_forest_anomaly_detector

This library implements the Extended Isolation Forest (EIF) algorithm for anomaly detection as described by Hariri et al. (2019). The Extended Isolation Forest improves upon the original Isolation Forest algorithm (Liu et al., 2008) by using random hyperplane cuts instead of axis-aligned cuts, eliminating bias artifacts in anomaly scores along coordinate axes. Learning fits the forest from baseline training examples selected from the dataset class labels.

The algorithm builds an ensemble of isolation trees (iTrees) by recursively partitioning data using random hyperplanes. Anomalous points, being few and different from normal points, require fewer partitions (shorter path lengths) to be isolated. The anomaly score for an instance is computed based on the average path length across all trees in the forest.

Datasets are represented as objects implementing the `anomaly_dataset_protocol` protocol from the `anomaly_detection_protocols` library. See the `anomaly_detection_protocols/test_datasets` directory for examples.

6.100.1 API documentation

Open the ../apis/library_index.html#isolation-forest link in a web browser.

6.100.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(isolation_forest_anomaly_detector(loader)).
```

6.100.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(isolation_forest_anomaly_detector(tester)).
```

6.100.4 Implemented features

- Extended Isolation Forest with random hyperplane cuts: splits are defined by random normal vectors and intercept points drawn from the data range, producing $(x - p) \cdot n \leq 0$ partitions that generalize to arbitrary orientations
- Configurable extension level: level 0 corresponds to the original axis-aligned Isolation Forest; levels up to $d - 1$ (the default) use fully extended random hyperplanes where d is the number of dimensions
- Anomaly score computation following Liu et al. (2008): $s(x) = 2^{-(E(h(x)) / c(\psi))}$ where $E(h(x))$ is the average path length across all trees, $c(\psi)$ is the average path length of unsuccessful searches in a BST, and ψ is the subsample size
- Handling of both continuous (numeric) and discrete (categorical) attributes: discrete attributes are mapped to numeric indices based on their position in the attribute value list declared by the dataset
- Baseline training selection: `baseline_class_values/1` declares which class labels are admissible for fitting the forest, while `baseline_selection_policy/1` controls whether non-baseline examples are rejected (default) or filtered before training
- Handling of missing attribute values (represented using anonymous variables): during tree construction, missing values are replaced with random values drawn from the observed range of the corresponding attribute; during scoring, each internal tree node stores its own feasible per-dimension bounds so that missing dimensions can be routed using subtree-local support instead of only the global training ranges
- Scoring all dataset instances with results sorted by descending anomaly score for easy identification of top anomalies
- Pretty-printing of learned models with tree depth and node count summaries
- Learning rejects empty datasets with a `domain_error(non_empty_dataset, Dataset)` exception

6.100.5 Options

The following options can be passed to the `learn/3` and `predict/4` predicates:

- `number_of_trees(N)`: number of isolation trees to build (default: 100)
- `subsample_size(N)`: subsample size used to build each isolation tree. When omitted, the implementation uses 256 or the number of training instances if smaller

- `extension_level(N)`: controls the dimensionality of the random hyperplane cuts. 0 reproduces the original axis-aligned Isolation Forest; when omitted, the implementation uses $d - 1$, where d is the number of dimensions
- `anomaly_threshold(T)`: threshold used by `predict/3-4` (default: 0.5)
- `baseline_class_values(Classes)`: learn-time list of admissible baseline class labels (default: `[normal]`)
- `baseline_selection_policy(Policy)`: learn-time handling of non-baseline examples. Supported values are `reject` (default) and `filter`

6.100.6 Detector representation

The learned detector is represented by default as:

```
if_model(Trees, Psi, AttributeNames, Attributes, Ranges, Diagnostics)
```

Where:

- `Trees`: List of learned isolation trees
- `Psi`: Effective subsample size used to build each tree
- `AttributeNames`: List of attribute names in order
- `Attributes`: List of Attribute-Values declarations from the training dataset
- `Ranges`: Observed numeric ranges used for imputing missing values during training
- `Diagnostics`: Learned metadata terms including `model/1`, `tree_count/1`, `subsample_size/1`, `attribute_names/1`, `feature_count/1`, and `options/1`

Each internal tree node additionally stores node-local dimension bounds used when resolving missing-value routing during scoring.

When exported using `export_to_clauses/4` or `export_to_file/4`, this detector term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.100.7 Limitations

- No incremental learning (the forest must be rebuilt from scratch when new examples are added)
- No streaming or online variant

6.100.8 References

- Liu, F.T., Ting, K.M. and Zhou, Z.-H. (2008). Isolation Forest. *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining*, 413-422. <https://doi.org/10.1109/ICDM.2008.17>
- Hariri, S., Kind, M.C. and Brunner, R.J. (2019). Extended Isolation Forest. *IEEE Transactions on Knowledge and Data Engineering*, 33(4), 1479-1489. <https://doi.org/10.1109/TKDE.2019.2947676>

6.100.9 Usage

To learn an isolation forest model from a dataset with default options:

```
| ?- isolation_forest_anomaly_detector::learn(gaussian_anomalies, Model, [base-  
line_selection_policy(filter)]).
```

To learn with custom options:

```
| ?- isolation_forest_anomaly_detector::learn(gaussian_anomalies, Model, [  
    baseline_selection_policy(filter),  
    number_of_trees(200),  
    subsample_size(128),  
    extension_level(1),  
    anomaly_threshold(0.6)  
]).
```

To compute the anomaly score for a new instance:

```
| ?- isolation_forest_anomaly_detector::learn(gaussian_anomalies, Model, [base-  
line_selection_policy(filter)]), isolation_forest_anomaly_detector::score(Model, [x-0.12, y-0.34], Score).
```

To predict whether an instance is an anomaly or normal:

```
| ?- isolation_forest_anomaly_detector::learn(gaussian_anomalies, Model, [base-  
line_selection_policy(filter)]), isolation_forest_anomaly_detector::predict(Model, [x-4.50, y-4.20], Pre-  
diction).
```

To compute and rank anomaly scores for all instances in a dataset:

```
| ?- isolation_forest_anomaly_detector::learn(gaussian_anomalies, Model, [base-  
line_selection_policy(filter)]), isolation_forest_anomaly_detector::score_all(gaussian_anomalies, Model,  
Scores).
```

The Scores list contains Id-Class-Score triples sorted by descending anomaly score. This makes it easy to inspect top anomalies:

```
| ?- isolation_forest_anomaly_detector::learn(gaussian_anomalies, Model, [base-  
line_selection_policy(filter)]), isolation_forest_anomaly_detector::score_all(gaussian_anomalies, Model,  
[Top1, Top2, Top3| _]).
```

To print a summary of the learned model:

```
| ?- isolation_forest_anomaly_detector::learn(gaussian_anomalies, Model, [base-  
line_selection_policy(filter)]), isolation_forest_anomaly_detector::print_anomaly_detector(Model).
```

To use the original (non-extended) Isolation Forest, set the extension level to 0:

```
| ?- isolation_forest_anomaly_detector::learn(gaussian_anomalies, Model, [base-  
line_selection_policy(filter), extension_level(0)]).
```

6.101 java

The library Java entities define a minimal abstraction for calling Java from Logtalk. This abstraction makes use of Logtalk parametric objects and allows creating Java objects, accessing Java class fields, and calling Java class and object methods using syntax closer to Logtalk. It also gives access to some Java utility predicates.

This abstraction was developed primarily to work with XVM (with its jni plug-in installed) or the JPL library bundled with SWI-Prolog and YAP. However, it's expected to be implementable with alternative Java interfaces found in other backend Prolog compilers. Currently, a preliminary implementation is also available for JIProlog.

The main idea in this abstraction layer is to use parametric objects where the first parameter holds the Java reference (usually to a class or object) and an optional second parameter holds the return value. Together with a forward message handler, this allows the use of Java messages with the same functor and number of arguments as found in the relevant JavaDocs.

6.101.1 API documentation

Open the ../apis/library_index.html#java link in a web browser.

6.101.2 Loading

To load all entities in this library, load the loader.lgt file:

```
| ?- logtalk_load(java(loader)).
```

6.101.3 Testing

To test this library predicates, load the tester.lgt file:

```
| ?- logtalk_load(java(tester)).
```

6.101.4 Usage

The two main objects in this library are `java(Reference, ReturnValue)` and `java(Reference)`. Use the latter if you want to ignore the return value or when calling a void Java method.

The `java` object implements utility predicates. For some backend Java interfaces such as JPL (available in SWI-Prolog, XVM, and YAP) there is also a `java_hook` hook object for removing any overhead when using this library abstraction.

For usage examples and additional tests, see the `clustering`, `document_converter`, `jp1`, and `neo4j` examples.

6.101.5 Known issues

When running Java GUI examples on the macOS Terminal application, you may get a Java error saying that the AWT cannot be started. In alternative, try to run the example from within the SWI-Prolog macOS application instead of using the shell integration script. This issue is due to a macOS Java issue that's orthogonal to both Prolog backends and Logtalk.

6.102 json

The `json` library provides predicates for parsing and generating data in the JSON format based on the specification and standard found at:

- <https://www.rfc-editor.org/rfc/rfc8259>
- <https://ecma-international.org/publications-and-standards/standards/ecma-404/>

It includes parametric objects whose parameters allow selecting the representation for parsed JSON objects (curly or list), JSON text strings (atom, chars, or codes) and JSON pairs (dash, equal, or colon).

6.102.1 API documentation

Open the `../apis/library_index.html#json` link in a web browser.

6.102.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(json(loader)).
```

6.102.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(json(tester)).
```

Some of the sample JSON test files are based on examples published at:

<https://www.json.org/>

6.102.4 Representation

The following choices of syntax have been made to represent JSON elements as terms:

- By default, JSON objects are represented using curly-bracketed terms, `{Pairs}`, where each pair uses the representation `Key-Value` (see below for alternative representations).
- Arrays are represented using lists.
- Text strings can be represented as atoms, `chars(List)`, or `codes(List)`. The default when decoding is to use atoms when using the `json` object. To decode text strings into lists of chars or codes, use the `json/1` with the parameter bound to `chars` or `codes`. For example:

```

| ?- json::parse(codes([34,104,101,108,108,111,34]), Term).
Term = hello
yes

| ?- json(atom)::parse(codes([34,104,101,108,108,111,34]), Term).
Term = hello
yes

| ?- json(chars)::parse(codes([34,104,101,108,108,111,34]), Term).
Term = chars([h,e,l,l,o])
yes

| ?- json(codes)::parse(codes([34,104,101,108,108,111,34]), Term).
Term = codes([104,101,108,108,111])
yes

```

- The JSON values false, true and null are represented by, respectively, the @false, @true and @null compound terms.

The following table exemplifies the term equivalents of JSON elements using default representations for objects, pairs, and strings:

JSON	term
[1,2]	[1,2]
true	@true
false	@false
null	@null
-1	-1
[1.2345]	[1.2345]
[]	[]
[2147483647]	[2147483647]
[0]	[0]
[1234567890123456789]	[1234567890123456789]
[false]	[@false]
[-2147483648]	[-2147483648]
{"a":null,"foo":"bar"}	{a-@null, foo-bar}
[2.225073858507201e-308]	[2.225073858507201e-308]
[0,1]	[0,1]
[2.2250738585072014e-308]	[2.2250738585072014e-308]
[1.7976931348623157e+308]	[1.7976931348623157e+308]
[0.0]	[0.0]
[4294967295]	[4294967295]
[-1234567890123456789]	[-1234567890123456789]
["foo"]	[foo]
[1]	[1]
[null]	[@null]
[-1.2345]	[-1.2345]
[5.0e-324]	[5.0e-324]
[-1]	[-1]
[true]	[@true]
[9223372036854775807]	[9223372036854775807]

For JSON objects that are two possible term representations:

JSON object	term (curly)
{“a”:1, “b”:2, “c”:3}	{a-1, b-2, c-3}
{}	{}

and:

JSON object	term (list)
{“a”:1, “b”:2, “c”:3}	json([a-1, b-2, c-3])
{}	json([])

For JSON pairs that are three possible representations:

JSON object	term (dash)
{“a”:1, “b”:2, “c”:3}	{a-1, b-2, c-3}

and:

JSON object	term (equal)
{“a”:1, “b”:2, “c”:3}	{a=1, b=2, c=3}

and:

JSON object	term (colon)
{“a”:1, “b”:2, “c”:3}	{a:1, b:2, c:3}

By default, the curly-term representation and the dash pair representation are used. The json/3 parametric object allows selecting the desired representation choices. For example:

```
| ?- json(curly,dash,atom)::parse(atom('{"a":1, "b":2, "c":3}'), JSON).
JSON = {a-1, b-2, c-3}
yes

| ?- json(list,equal,atom)::parse(atom('{"a":1, "b":2, "c":3}'), JSON).
JSON = json([a=1, b=2, c=3])
yes

| ?- json(curly,colon,atom)::parse(atom('{"a":1, "b":2, "c":3}'), JSON).
JSON = {a:1, b:2, c:3}
yes
```

6.102.5 Encoding

Encoding is accomplished using the `generate/2` predicate. For example:

```
| ?- json::generate(codes(Encoding), [a,{b-c}]).
Encoding = [91,34,97,34,44,123,34,98,34,58,34,99,34,125,93]
yes
```

Alternatively:

```
| ?- json::generate(chars(Encoding), [a,{b-c}]).
Encoding = ['[','"',a,'"',',','{','"',b,'"',:',','"',c,'"','}',']']
Yes

| ?- json::generate(atom(Encoding), [a,{b-c}]).
Encoding = '["a",{b:"c"}]'
Yes
```

Notice that `generate/2` takes, as second argument, a Prolog term that corresponds to the JSON syntax in Prolog and produces the corresponding JSON output in the format specified in the first argument: (`codes(Variable)`, `stream(Stream)`, `file(File)`, `chars(Variable)` or `atom(Variable)`).

6.102.6 Decoding

Decoding is accomplished using the `parse/2` predicate. For example, to decode a given json file:

```
| ?- json::parse(file('simple/roundtrip_array_obj_array.json'), Term).
Term = [{a-[b]}]
yes
```

The `parse/2` predicate first argument must indicate the input source (`codes(Codes)`, `stream(Stream)`, `line(Stream)`, `file(Path)`, `chars(Chars)` or `atom(Atom)`) containing a JSON payload to be decoded into the Prolog term in the second argument.

6.102.7 Known issues

Some tests, notably `parse_simple_valid_files` and `roundtrip_hexadecimals`, fail on backends such as ECLiPSe and GNU Prolog that don't support Unicode.

6.103 json_ld

The `json_ld` library provides predicates for parsing, generating, expanding, and compacting JSON-LD 1.1 (JSON-based Serialization for Linked Data) documents based on the W3C Recommendation found at:

<https://www.w3.org/TR/json-ld11/>

This library builds on top of the `json` library for JSON parsing and generation, and uses the same representation choices for JSON data. It includes parametric objects whose parameters allow selecting the representation for parsed JSON objects (`curly` or `list`), JSON text strings (`atom`, `chars`, or `codes`) and JSON pairs (`dash`, `equal`, or `colon`).

6.103.1 API documentation

Open the `../apis/library_index.html#json_ld` link in a web browser.

6.103.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(json_ld(loader)).
```

6.103.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(json_ld(tester)).
```

6.103.4 Parsing

JSON-LD documents can be parsed from various sources using the `parse/2` predicate. Since JSON-LD is a superset of JSON, any valid JSON-LD document is also valid JSON:

```
| ?- json_ld::parse(file('document.jsonld'), Term).

| ?- json_ld::parse(atom('{"@context": {"name": "http://schema.org/name"}, "name": "Manu_
↳ Sporny"}'), Term).
```

6.103.5 Generating

JSON-LD documents can be generated from Logtalk terms using the `generate/2` predicate:

```
| ?- json_ld::generate(atom(Atom), {'@context'-{name-'http://schema.org/name'}, name-'Manu_
↳ Sporny'}).
```

6.103.6 Expansion

Expansion is the process of removing the context and representing all properties and types as full IRIs. This is useful for processing JSON-LD data in a context-independent way:

```
| ?- json_ld::parse(atom('{"@context": {"name": "http://schema.org/name"}, "name": "Manu_
↳ Sporny"}'), Doc),
    json_ld::expand(Doc, Expanded).
Doc = {'@context'-{name-'http://schema.org/name'}, name-'Manu Sporny'},
Expanded = [{http://schema.org/name-[{'@value'-'Manu Sporny'}]]]
yes
```

6.103.7 Compaction

Compaction is the process of applying a context to shorten IRIs to terms or compact IRIs. This is the inverse of expansion:

```
| ?- json_ld::expand({'@context'-{name-'http://schema.org/name'}, name-'Manu Sporny'}, _
↳ Expanded),
    json_ld::compact(Expanded, {'@context'-{name-'http://schema.org/name'}}, Compacted).
```

6.103.8 Flattening

Flattening collects all node objects from a document into a flat @graph array, with nested objects replaced by references. Blank node identifiers are generated for nodes that don't have an @id:

```
| ?- json_ld::parse(atom('{"@context":{"knows":"http://xmlns.com/foaf/0.1/knows"},"@id":
↳ "http://example.org/john","knows":{"@id":"http://example.org/jane"}}'), Doc),
    json_ld::expand(Doc, Expanded),
    json_ld::flatten(Expanded, Flattened).
```

6.103.9 Supported Features

The library supports the following JSON-LD 1.1 features:

Context processing:

- Inline context definitions (@context)
- Vocabulary mapping (@vocab)
- Base IRI (@base)
- Default language (@language)
- Base direction (@direction)
- Compact IRIs (prefix:suffix notation)
- Term definitions (simple and expanded)
- Type coercion (@type in term definitions)
- Context arrays (multiple contexts)

Node objects:

- Node identifiers (@id)
- Node types (@type)
- Blank node identifiers (_:name)

Value objects:

- Typed values (@value with @type)
- Language-tagged strings (@value with @language)
- Direction-tagged strings (@value with @direction)

Graph support:

- Named graphs (@graph)

- Default graph

Collections:

- Ordered lists (@list)
- Unordered sets (@set)

Other features:

- Reverse properties (@reverse)
- Included blocks (@included)
- Index preservation (@index)
- Flattening algorithm (flatten/2)

Not currently supported:

- Remote context fetching (contexts referenced by URL)
- Framing algorithm
- @import context processing

6.103.10 Representation

JSON-LD documents are represented using the same term conventions as the `json` library. The JSON-LD keywords (@context, @id, @type, etc.) are represented as atoms in the parsed terms. For example:

JSON-LD	term (default)
{"@id": "http://example.org/"}	{'@id'-'http://example.org/'}
{"@type": "Person"}	{'@type'-'Person'}
{"@value": "hello", "@language": "en"}	{'@value'-'hello', '@language'-'en'}
{"@list": [1, 2, 3]}	{'@list'-[1, 2, 3]}
{"@graph": [...]}	{'@graph'-[...]}

As with the `json` library, objects can be represented using curly terms (default) or list terms, and pairs can use dash, equal, or colon notation.

6.104 json_lines

The `json_lines` library provides predicates for parsing and generating data in the JSON Lines format based on the proposal found at:

<https://jsonlines.org>

It includes parametric objects whose parameters allow selecting the representation for parsed JSON objects (curly or list), JSON text strings (atom, chars, or codes) and JSON pairs (dash, equal, or colon).

6.104.1 API documentation

Open the `../apis/library_index.html#json_lines` link in a web browser.

6.104.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(json_lines(loader)).
```

6.104.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(json_lines(tester)).
```

Some of the sample JSON test files are based on examples published at:

<https://jsonlines.org/>

6.104.4 Representation

The following choices of syntax have been made to represent JSON elements as terms:

- By default, JSON objects are represented using curly-bracketed terms, `{Pairs}`, where each pair uses the representation Key-Value (see below for alternative representations).
- Arrays are represented using lists.
- Text strings can be represented as atoms, `chars(List)`, or `codes(List)`. The default when decoding is to use atoms when using the `json_lines` object. To decode text strings into lists of chars or codes, use the `json_lines/1` object with the parameter bound to `chars` or `codes`. For example:

```
| ?- json_lines::parse(codes([34,104,101,108,108,111,34]), Terms).
Terms = [hello]
yes

| ?- json_lines(atom)::parse(codes([34,104,101,108,108,111,34]), Terms).
Terms = [hello]
yes

| ?- json_lines(chars)::parse(codes([34,104,101,108,108,111,34]), Terms).
Terms = [chars([h,e,l,l,o])]
yes

| ?- json_lines(codes)::parse(codes([34,104,101,108,108,111,34]), Terms).
Terms = [codes([104,101,108,108,111])]
yes
```

- The JSON values `false`, `true` and `null` are represented by, respectively, the `@false`, `@true` and `@null` compound terms.

The following table exemplifies the term equivalents of JSON elements using default representations for objects, pairs, and strings:

JSON	term
[1,2]	[1,2]
true	@true
false	@false
null	@null
-1	-1
[1.2345]	[1.2345]
[]	[]
[2147483647]	[2147483647]
[0]	[0]
[1234567890123456789]	[1234567890123456789]
[false]	[@false]
[-2147483648]	[-2147483648]
{"a":null,"foo":"bar"}	{a-@null, foo-bar}
[2.225073858507201e-308]	[2.225073858507201e-308]
[0,1]	[0,1]
[2.2250738585072014e-308]	[2.2250738585072014e-308]
[1.7976931348623157e+308]	[1.7976931348623157e+308]
[0.0]	[0.0]
[4294967295]	[4294967295]
[-1234567890123456789]	[-1234567890123456789]
["foo"]	[foo]
[1]	[1]
[null]	[@null]
[-1.2345]	[-1.2345]
[5.0e-324]	[5.0e-324]
[-1]	[-1]
[true]	[@true]
[9223372036854775807]	[9223372036854775807]

For JSON objects that are two possible term representations:

JSON object	term (curly)
{"a":1, "b":2, "c":3}	{a-1, b-2, c-3}
{}	{}

and:

JSON object	term (list)
{"a":1, "b":2, "c":3}	json([a-1, b-2, c-3])
{}	json([])

For JSON pairs that are three possible representations:

JSON object	term (dash)
{"a":1, "b":2, "c":3}	{a-1, b-2, c-3}

and:

JSON object	term (equal)
{“a”:1, “b”:2, “c”:3}	{a=1, b=2, c=3}

and:

JSON object	term (colon)
{“a”:1, “b”:2, “c”:3}	{a:1, b:2, c:3}

By default, the curly-term representation and the dash pair representation are used. The json/3 parametric object allows selecting the desired representation choices. For example:

```
| ?- json_lines(curly,dash,atom)::parse(atom('{"a":1, "b":2, "c":3}'), JSONL).
JSONL = [{a-1, b-2, c-3}]
yes

| ?- json_lines(list,equal,atom)::parse(atom('{"a":1, "b":2, "c":3}'), JSONL).
JSONL = [json([a=1, b=2, c=3])]
yes

| ?- json_lines(curly,colon,atom)::parse(atom('{"a":1, "b":2, "c":3}'), JSONL).
JSONL = [{a:1, b:2, c:3}]
yes
```

6.104.5 Encoding

Encoding is accomplished using the generate/2 predicate. For example:

```
| ?- json_lines::generate(codes(Encoding), [a,{b-c}]).
Encoding = [34,97,34,10,123,34,98,34,58,34,99,34,125,10]
yes
```

Alternatively:

```
| ?- json_lines::generate(chars(Encoding), [a,{b-c}]).
Encoding = ['"',a,'"','\n',{'"',b,'"':','"',c,'"','}'},'\n']
Yes

| ?- json_lines::generate(atom(Encoding), [a,{b-c}]).
Encoding = '"a"\n{"b":"c"}\n'
Yes
```

Notice that generate/2 takes, as second argument, a Prolog term that corresponds to the JSON syntax in Prolog and produces the corresponding JSON output in the format specified in the first argument: (codes(Variable), stream(Stream), file(File), chars(Variable) or atom(Variable)).

6.104.6 Decoding

Decoding is accomplished using the `parse/2` predicate. For example, to decode a given json file:

```
| ?- json_lines::parse(file('simple/data.jsonl'), Terms).  
Term = [{a-[b]}]  
yes
```

The `parse/2` predicate first argument must indicate the input source (`codes(Codes)`, `stream(Stream)`, `line(Stream)`, `file(Path)`, `chars(Chars)` or `atom(Atom)`) containing a JSON payload to be decoded into the Prolog term in the second argument.

6.104.7 Known issues

Some tests may fail on backends such as ECLiPSe and GNU Prolog that don't support Unicode.

6.105 json_pointer

The `json_pointer` library provides predicates for parsing, generating, and evaluating JSON Pointers as specified by RFC 6901:

- <https://www.rfc-editor.org/rfc/rfc6901>

It supports both the plain string syntax and the URI fragment syntax. Reference tokens can be represented as atoms, `chars(List)`, or `codes(List)`.

6.105.1 API documentation

Open the `../apis/library_index.html#json_pointer` link in a web browser.

6.105.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(json_pointer(loader)).
```

6.105.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(json_pointer(tester)).
```

6.105.4 Representation

The library defines the `json_pointer(_Representation_)` parametric object where `_Representation_` can be one of:

- `atom` - reference tokens are represented as atoms
- `chars` - reference tokens are represented as `chars(List)`
- `codes` - reference tokens are represented as `codes(List)`

When using the default `json_pointer` object, reference tokens are represented as atoms.

6.105.5 Examples

Parse and evaluate a pointer:

```
| ?- json_pointer::parse(atom('/foo/0'), Pointer),
    json_pointer::evaluate(Pointer, {foo-[bar, baz]}, Value).
Pointer = [foo, '0']
Value = bar
yes
```

Generate a URI fragment representation:

```
| ?- json_pointer::generate_fragment(atom(Fragment), ['a b', 'c/d']).
Fragment = '#/a%20b/c~1d'
yes
```

6.106 json_rpc

JSON-RPC 2.0 protocol encoding and decoding library. Provides predicates for constructing, parsing, classifying, and inspecting JSON-RPC 2.0 messages (requests, notifications, responses, and error responses). Also provides stream-based message I/O for implementing JSON-RPC clients and servers, including Content-Length header framing as used by the Language Server Protocol (LSP) and the Model Context Protocol (MCP). Uses the `json` library for JSON parsing and generation.

6.106.1 API documentation

Open the `../apis/library_index.html#json_rpc` link in a web browser.

6.106.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(json_rpc(loader)).
```

6.106.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(json_rpc(tester)).
```

6.106.4 Usage

Constructing JSON-RPC 2.0 messages

Construct requests, notifications, responses, and error responses:

```
| ?- json_rpc::request(subtract, [42,23], 1, Request).
Request = {jsonrpc-'2.0', method-subtract, params-[42,23], id-1}.

| ?- json_rpc::notification(update, [1,2,3], Notification).
Notification = {jsonrpc-'2.0', method-update, params-[1,2,3]}.

| ?- json_rpc::response(19, 1, Response).
Response = {jsonrpc-'2.0', result-19, id-1}.

| ?- json_rpc::error_response(-32601, 'Method not found', 1, ErrorResponse).
```

Standard error constructors

```
| ?- json_rpc::parse_error(Error).
Error = ...

| ?- json_rpc::invalid_request(Error).
Error = ...

| ?- json_rpc::method_not_found(1, Error).
Error = ...

| ?- json_rpc::invalid_params(1, Error).
Error = ...

| ?- json_rpc::internal_error(1, Error).
Error = ...
```

Encoding and decoding

```
| ?- json_rpc::request(subtract, [42,23], 1, Request),
    json_rpc::encode(Request, JSON).
JSON = '{"jsonrpc":"2.0","method":"subtract","params":[42,23],"id":1}'.

| ?- json_rpc::decode('{"jsonrpc":"2.0","result":19,"id":1}', Term).
```

Message classification

```
| ?- json_rpc::is_request(Message).
...

| ?- json_rpc::is_notification(Message).
...

| ?- json_rpc::is_response(Message).
...

| ?- json_rpc::is_error_response(Message).
...

| ?- json_rpc::is_batch(Messages).
...
```

Field extraction

```
| ?- json_rpc::id(Message, Id).
...

| ?- json_rpc::method(Message, Method).
...

| ?- json_rpc::params(Message, Params).
...

| ?- json_rpc::result(Message, Result).
...

| ?- json_rpc::error_code(Message, Code).
...

| ?- json_rpc::error_message(Message, ErrorMessage).
...

| ?- json_rpc::error_data(Message, Data).
...
```

Stream I/O (client and server API)

Write and read newline-delimited JSON-RPC messages over streams:

```
| ?- json_rpc::write_message(Output, Message).
...

| ?- json_rpc::read_message(Input, Message).
...
```

These predicates can be used with any stream, including socket streams from the `sockets` library, to implement JSON-RPC clients and servers.

Content-Length framed I/O (LSP/MCP protocols)

Write and read JSON-RPC messages using Content-Length header framing as defined by the Language Server Protocol (LSP) and the Model Context Protocol (MCP). Each message is preceded by a Content-Length: N\r\n\r\n header where N is the byte length of the JSON body:

```
| ?- json_rpc::write_framed_message(Output, Message).  
...  
  
| ?- json_rpc::read_framed_message(Input, Message).  
...
```

`read_framed_message/2` fails at end of stream or if the header is missing or malformed.

6.106.5 API summary

Message construction

Predicates for building JSON-RPC 2.0 message terms. Requests include an `id` for matching responses; notifications do not.

- `request(+Method, +Params, +Id, --Request)` - Construct a request
- `request(+Method, +Id, --Request)` - Construct a request with no parameters
- `notification(+Method, +Params, --Notification)` - Construct a notification
- `notification(+Method, --Notification)` - Construct a notification with no parameters
- `response(+Result, +Id, --Response)` - Construct a successful response
- `error_response(+Code, +Message, +Id, --ErrorResponse)` - Construct an error response
- `error_response(+Code, +Message, +Data, +Id, --ErrorResponse)` - Construct an error response with data

Standard error responses

Convenience predicates for the five standard JSON-RPC 2.0 error codes.

- `parse_error(--ErrorResponse)` - Parse error (-32700)
- `invalid_request(--ErrorResponse)` - Invalid request (-32600)
- `method_not_found(+Id, --ErrorResponse)` - Method not found (-32601)
- `invalid_params(+Id, --ErrorResponse)` - Invalid params (-32602)
- `internal_error(+Id, --ErrorResponse)` - Internal error (-32603)

Encoding and decoding

Convert between JSON-RPC message terms and JSON atoms (strings).

- `encode(+Term, --JSON)` - Encode a JSON-RPC term to a JSON atom
- `decode(+JSON, --Term)` - Decode a JSON atom to a JSON-RPC term

Message classification

Test what kind of JSON-RPC message a term represents.

- `is_request(+Term)` - Test if a term is a request
- `is_notification(+Term)` - Test if a term is a notification
- `is_response(+Term)` - Test if a term is a successful response
- `is_error_response(+Term)` - Test if a term is an error response
- `is_batch(+Term)` - Test if a term is a batch (non-empty list)

Field extraction

Extract individual fields from JSON-RPC message terms.

- `id(+Message, --Id)` - Extract the id field
- `method(+Message, --Method)` - Extract the method field
- `params(+Message, --Params)` - Extract the params field
- `result(+Message, --Result)` - Extract the result field
- `error(+Message, --Error)` - Extract the error object
- `error_code(+Message, --Code)` - Extract the error code
- `error_message(+Message, --ErrorMessage)` - Extract the error message
- `error_data(+Message, --Data)` - Extract the error data

Stream I/O

Read and write JSON-RPC messages over streams. The newline-delimited variants are suitable for socket-based communication. The Content-Length framed variants implement the header-based framing used by the Language Server Protocol (LSP) and the Model Context Protocol (MCP).

- `write_message(+Output, +Message)` - Write a newline-delimited message
- `read_message(+Input, --Message)` - Read a newline-delimited message
- `write_framed_message(+Output, +Message)` - Write a message with Content-Length framing
- `read_framed_message(+Input, --Message)` - Read a message with Content-Length framing

6.107 json_schema

The `json_schema` library provides predicates for parsing JSON Schema documents and validating JSON data against schemas. It is based on the JSON Schema specification found at:

- <https://json-schema.org/>

This library builds on top of the `json` library for JSON parsing and uses the same representation choices for JSON data.

6.107.1 API documentation

Open the `../..apis/library_index.html#json_schema` link in a web browser.

6.107.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(json_schema(loader)).
```

6.107.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(json_schema(tester)).
```

6.107.4 Schema Parsing

Schemas can be parsed from various sources using the `parse/2` predicate:

```
| ?- json_schema::parse(file('schema.json'), Schema).  
  
| ?- json_schema::parse(atom('{"type": "string"}'), Schema).  
  
| ?- json_schema::parse(codes(Codes), Schema).
```

6.107.5 Validation

JSON data can be validated against a schema using the `validate/2` or `validate/3` predicates:

```
% Success/failure validation  
| ?- json_schema::parse(atom('{"type": "string"}'), Schema),  
    json_schema::validate(Schema, hello).  
yes  
  
| ?- json_schema::parse(atom('{"type": "string"}'), Schema),  
    json_schema::validate(Schema, 42).  
no
```

(continues on next page)

(continued from previous page)

```
% Validation with error collection
| ?- json_schema::parse(atom('{"type": "string"}'), Schema),
      json_schema::validate(Schema, 42, Errors).
Errors = [error([], expected_type(string))]
```

6.107.6 Supported Keywords

The library supports the following JSON Schema keywords:

Type validation:

- type (string, number, integer, boolean, array, object, null)
- Multiple types: {"type": ["string", "number"]}

Generic validation:

- enum - value must be one of the specified values
- const - value must be exactly the specified value

String validation:

- minLength - minimum string length
- maxLength - maximum string length

Numeric validation:

- minimum - minimum value (inclusive)
- maximum - maximum value (inclusive)
- exclusiveMinimum - minimum value (exclusive)
- exclusiveMaximum - maximum value (exclusive)
- multipleOf - value must be a multiple of this number

Array validation:

- items - schema for array items
- prefixItems - schemas for array items by position
- minItems - minimum number of items
- maxItems - maximum number of items
- uniqueItems - all items must be unique
- contains - at least one item must match the schema

Object validation:

- properties - schemas for object properties
- required - list of required property names
- minProperties - minimum number of properties
- maxProperties - maximum number of properties
- propertyNames - schema for property names

Composition:

- allOf - value must match all schemas
- anyOf - value must match at least one schema
- oneOf - value must match exactly one schema
- not - value must not match the schema

Conditional:

- if/then/else - conditional schema application

Schema References:

- \$defs/definitions - schema definitions
- \$ref - local JSON Pointer references and local file references (e.g., #/\$defs/name, address.json, common.json#/\$defs/name)

Format validation:

The format keyword validates string formats. Non-string values pass format validation. Unknown formats are ignored per JSON Schema specification. Supported formats:

- email - basic email format (local@domain)
- date - ISO 8601 date (YYYY-MM-DD)
- time - ISO 8601 time (HH:MM:SS) with optional fractional seconds and timezone
- date-time - ISO 8601 combined date and time
- uri - URI with scheme (e.g., http://example.com)
- uri-reference - URI or relative reference
- ipv4 - IPv4 address (dotted decimal notation)
- ipv6 - IPv6 address (hex groups with colons, supports :: compression)
- uuid - UUID format (8-4-4-4-12 hex pattern)

Boolean schemas:

- true - accepts any value
- false - rejects all values

6.107.7 Known Limitations

The following JSON Schema features are not yet supported:

- pattern - regular expression validation (requires regex support)
- patternProperties - property matching by pattern
- format - formats not listed above (e.g., hostname, regex, json-pointer)
- \$ref with remote URLs - external schema references require HTTP support

6.108 kcenters_clusterer

k-Centers clusterer. It uses a deterministic farthest-first center selection heuristic. Supports continuous attributes only.

The library implements the `clusterer_protocol` defined in the `clustering_protocols` library. It provides predicates for learning a clusterer from a dataset, assigning new instances to clusters, and exporting the learned clusterer as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol from the `clustering_protocols` library.

6.108.1 API documentation

Open the ../apis/library_index.html#kcenters_clusterer link in a web browser.

6.108.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(kcenters_clusterer(loader)).
```

6.108.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(kcenters_clusterer(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(kcenters_clusterer(tester_performance)).
```

6.108.4 Features

- **Continuous Datasets:** Accepts datasets containing only continuous attributes.
- **Farthest-First Center Selection:** Uses a deterministic farthest-first heuristic to choose centers.
- **Deterministic Initialization:** Supports `first_k` and deterministic spread initialization that repeatedly chooses the farthest example from the centers selected so far..
- **Distance Metrics:** Supports Euclidean and Manhattan distances..
- **Optional Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Rich Training Diagnostics:** Learned clusterers report training example count, selected center count, and the center-selection strategy used during learning.
- **Portable Export:** Learned clusterers can be exported as clauses or files and reused later.

6.108.5 Options

The following options can be passed to the `learn/3` predicate:

- `k(K)`: Number of clusters to learn. Default is 2.
- `initialization(Initialization)`: Center initialization strategy. Options: `spread` (default) or `first_k`.
- `distance_metric(Metric)`: Distance metric to use. Options: `euclidean` (default) or `manhattan`.
- `feature_scaling(FeatureScaling)`: Whether to standardize continuous attributes before clustering. Options: `on` (default) or `off`.

6.108.6 Diagnostics

The `diagnostics/2` predicate returns a list containing:

- `model(kcenters_clusterer)`
- `center_count(Count)`
- `training_example_count(Count)`
- `selection_strategy(Strategy)`
- `options(Options)`

6.108.7 Clusterer representation

The learned clusterer is represented as a compound term with the functor chosen by the user when exporting the clusterer and arity 4. For example:

```
kcenters_clusterer(Encoders, Centers, Options, Diagnostics)
```

Where:

- `Encoders`: List of continuous attribute encoders storing attribute name, mean, and scale.
- `Centers`: List of center vectors in cluster-id order.
- `Options`: Effective training options used to learn the clusterer.
- `Diagnostics`: Training diagnostics metadata returned by the `diagnostics/2` predicate.

6.108.8 References

1. Gonzalez (1985) - "Clustering to minimize the maximum intercluster distance". Theoretical Computer Science, 38, 293-306.

6.109 kemeny_young_ranker

Kemeny-Young pairwise preference ranker. It aggregates head-to-head preference weights, then performs an exact branch-and-bound search over linear orders to maximize the total pairwise agreement score under the Kemeny-Young criterion.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning an exact consensus ranker from pairwise preferences, using it to order candidate items, inspecting the learned consensus ranking and consensus score, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The current implementation requires a well-formed connected pairwise dataset so that all ranked items remain part of a single comparison graph.

6.109.1 API documentation

Open the ../apis/library_index.html#kemeny_young_ranker link in a web browser.

6.109.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(kemeny_young_ranker(loader)).
```

6.109.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(kemeny_young_ranker(tester)).
```

6.109.4 Features

- **Exact Consensus Ranking:** Learns one deterministic Kemeny-Young consensus order from aggregated pairwise outcomes.
- **Branch-and-Bound Search:** Uses exact search with pruning over linear orders to maximize total pairwise agreement.
- **Deterministic Optimal Tie Breaking:** Supports both standard term-order and dataset declaration-order selection when multiple optimal consensus rankings exist.
- **Consensus Access:** Exposes the learned full-item consensus ranking and its maximum agreement score using the `consensus_ranking/2` and `consensus_score/2` predicates.
- **Deterministic Ranking:** Orders candidate items by the selected consensus order with deterministic tie-breaking. When multiple optimal Kemeny orders exist, the `tie_breaking/1` option selects the deterministic search order used to choose one representative consensus ranking.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items, self-preferences, non-positive weights, and disconnected comparison graphs.

- **Training Diagnostics:** Learned rankers include dataset summary metadata, the effective tie-breaking mode, the selected consensus ranking, and the maximum agreement score.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.

6.109.5 Usage

Learning a ranker

```
% Learn from a pairwise ranking dataset object
| ?- kemeny_young_ranker::learn(my_dataset, Ranker).
...

% Learn with custom optimal-order tie breaking
| ?- kemeny_young_ranker::learn(my_dataset, Ranker, [tie_breaking(declaration_order)]).
...
```

Inspecting the learned consensus

```
% Inspect the learned full-item consensus ranking
| ?- kemeny_young_ranker::learn(my_dataset, Ranker),
    kemeny_young_ranker::consensus_ranking(Ranker, ConsensusRanking).
ConsensusRanking = [...]
...

% Inspect the maximum pairwise agreement score
| ?- kemeny_young_ranker::learn(my_dataset, Ranker),
    kemeny_young_ranker::consensus_score(Ranker, ConsensusScore).
ConsensusScore = ...
...
```

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- kemeny_young_ranker::learn(my_dataset, Ranker),
    kemeny_young_ranker::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
...
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

6.109.6 Options

The following options can be passed to the `learn/3` predicate:

- `tie_breaking(term_order)`: Select the lexicographically earliest optimal consensus ranking using the standard term order of the item identifiers.
- `tie_breaking(declaration_order)`: Select the earliest optimal consensus ranking using the dataset declaration order preserved by the pairwise dataset helpers.

The default is `tie_breaking(term_order)`.

6.109.7 Scoring semantics

The learned `scores/2` values are reverse consensus positions. For a learned consensus order with N items, the first item receives score $N-1$ and the last item receives score 0 .

This positional encoding is a deterministic surrogate used to satisfy the shared `ranker_protocol` scoring interface. The authoritative Kemeny-Young output is the full `consensus_ranking/2` order together with its `consensus_score/2` agreement value.

6.109.8 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(kemeny_young_ranker),
  options(Options),
  consensus_ranking(ConsensusRanking),
  consensus_score(ConsensusScore),
  dataset_summary(DatasetSummary)
]
```

6.109.9 Ranker representation

The learned ranker is represented by a compound term of the form:

```
kemeny_young_ranker(Items, Scores, Diagnostics)
```

Where:

- **Items**: List of ranked items.
- **Scores**: List of Item-Score pairs.
- **Diagnostics**: List of metadata terms, including the effective options, selected consensus ranking, maximum agreement score, and dataset summary.

6.109.10 Notes

Kemeny-Young optimization is NP-hard in general. This implementation uses an exact branch-and-bound search and is intended for small to moderate item sets.

6.109.11 References

1. Kemeny, J. G. (1959). *Mathematics without numbers*.
2. Young, H. P. (1988). *Condorcet's theory of voting*.

6.110 kernel_svm_classifier

Kernel support vector machine classifier using one-vs-rest dual margin models with linear, polynomial, and radial basis function kernels. The implementation encodes tabular datasets using the shared linear encoder pipeline so mixed continuous and categorical datasets are handled consistently with the existing linear classifiers.

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning a classifier from a dataset, using it to make predictions, estimating class probabilities, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. Continuous, categorical, and mixed-feature datasets are supported.

6.110.1 API documentation

Open the ../docs/library_index.html#kernel_svm_classifier link in a web browser.

6.110.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(kernel_svm_classifier(loader)).
```

6.110.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(kernel_svm_classifier(tester)).
```

6.110.4 Features

- **Binary and Multiclass Classification:** Learns one-vs-rest dual models and predicts the class with the highest decision score.
- **Multiple Kernels:** Supports linear, polynomial(Degree, Gamma, Coef0), and rbf(Gamma) kernels.
- **Mixed Features:** Reuses the shared tabular encoders for continuous and categorical attributes, including missing-value indicators.
- **Probability Estimation:** Provides class probabilities using a softmax over kernel decision scores.
- **Regularized Training:** Supports configurable learning-rate scheduling, tolerance, and L2 regularization.
- **Classifier Export:** Learned classifiers can be exported as predicate clauses or written to a file.

6.110.5 Options

The learn/3 predicate supports these options:

- kernel/1 - kernel function to use (default: linear)
- learning_rate/1 - base learning rate for the dual optimization loop (default: 0.5)
- learning_schedule/1 - learning-rate schedule, either constant or inverse_scaling(Power) (default: constant)
- maximum_iterations/1 - maximum number of optimization epochs (default: 25)
- tolerance/1 - convergence threshold for the maximum parameter update (default: 1.0e-5)
- l2_regularization/1 - L2 penalty factor applied during optimization (default: 0.001)
- feature_scaling/1 - whether to standardize continuous attributes before encoding (default: true)

6.110.6 Usage

Learning a classifier

```
| ?- kernel_svm_classifier::learn(weather, Classifier).
| ?- kernel_svm_classifier::learn(mixed, Classifier, [kernel(rbf(0.5)), maximum_
↪ iterations(50)]).
```

Making predictions

```
| ?- kernel_svm_classifier::learn(mixed, Classifier),
   kernel_svm_classifier::predict(Classifier, [age-35, income-65000, student-yes, credit_
↪ rating-fair], Class).
| ?- kernel_svm_classifier::learn(weather, Classifier),
   kernel_svm_classifier::predict_probabilities(Classifier, [outlook-sunny, temperature-
↪ hot, humidity-high, windy-false], Probabilities).
```

Exporting the classifier

```
| ?- kernel_svm_classifier::learn(weather, Classifier),  
    kernel_svm_classifier::export_to_clauses(weather, Classifier, classify, Clauses).  
  
| ?- kernel_svm_classifier::learn(weather, Classifier),  
    kernel_svm_classifier::export_to_file(weather, Classifier, classify, 'classifier.pl').
```

6.110.7 Classifier representation

The learned classifier is represented as a compound term with the form:

```
kernel_svm_classifier(Classes, Encoders, Kernel, TrainingRows, Models, Options)
```

Where:

- **Classes**: list of class labels
- **Encoders**: list of continuous scaling descriptors and categorical value encoders
- **Kernel**: selected kernel specification
- **TrainingRows**: encoded feature vectors for the training examples
- **Models**: list of `class_model(Class, Bias, Coefficients)` terms
- **Options**: merged training options used to learn the classifier

When exported using `export_to_clauses/4` or `export_to_file/4`, this classifier term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.110.8 References

1. Cortes, C. and Vapnik, V. (1995). “Support-Vector Networks”.
2. Scholkopf, B. and Smola, A.J. (2002). “Learning with Kernels”.
3. Bishop, C.M. (2006). “Pattern Recognition and Machine Learning”. Chapter 7.

6.111 kernel_pca_projection

Kernel Principal Component Analysis reducer for continuous datasets. The library implements the `dimension_reducer_protocol` defined in the `dimension_reduction_protocols` library and learns a nonlinear projection by centering the training data, optionally standardizing continuous attributes, building a centered kernel Gram matrix, and extracting deterministic principal directions in sample space using portable power iteration with deflation.

6.111.1 API documentation

Open the ../apis/library_index.html#kernel_pca_projection link in a web browser.

6.111.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(kernel_pca_projection(loader)).
```

6.111.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(kernel_pca_projection(tester)).
```

6.111.4 Features

- **Continuous Datasets:** Accepts datasets containing only continuous attributes. Missing or nonnumeric values are rejected.
- **Centering and Optional Scaling:** Centers all attributes and optionally standardizes them before evaluating kernels.
- **Configurable Shortfall Handling:** Lets callers choose whether a component-extraction shortfall raises an error or truncates the learned reducer with explicit diagnostics.
- **Supported Kernels:** Supports linear, polynomial with non-negative offset, and radial basis function kernels through the `kernel/1` option.
- **Shared Gram-Centering Helpers:** Delegates training Gram-matrix centering and out-of-sample Gram-vector centering to the shared `linear_algebra` library.
- **Portable Eigensolver:** Uses deterministic power iteration with deflation instead of backend-specific linear algebra libraries.
- **Projection API:** Transforms a new instance into a list of `component_N-Value` pairs using centered kernel evaluations against the training rows.
- **Model Export:** Learned reducers can be exported as predicate clauses or written to a file.

6.111.5 Options

The `learn/3` predicate accepts the following options:

- `n_components/1`: Number of kernel principal components to extract. Requests that exceed `SampleCount - 1` raise `domain_error(component_count, Requested-Maximum)`. The default is 2.
- `feature_scaling/1`: Whether to standardize continuous attributes before evaluating kernels. Options: `true` (default) or `false`.
- `shortfall_policy/1`: Controls what happens when the centered kernel Gram matrix yields fewer numerically significant components than requested. Options: `truncate` (default), which returns a reducer with fewer components and records a `shortfall(truncated(Requested,`

Learned, ResidualEigenvalue, Tolerance)) diagnostic, or error, which raises `domain_error(component_count, Requested-Learned)`.

- `kernel/1`: Kernel specification. Supported values are `linear` (default), `polynomial`(Degree, Gamma, Coef0) with positive Degree, positive Gamma, and non-negative Coef0, and `rbf`(Gamma) with positive Gamma.
- `maximum_iterations/1`: Maximum number of power-iteration steps used when estimating each dual principal direction. The default is 1000.
- `tolerance/1`: Positive convergence tolerance used both for power-iteration stopping and for deciding when deflated eigenvalues are negligible. The default is 1.0e-8.

6.111.6 Usage

The following examples use the sample datasets shipped with the `dimension_reduction_protocols` library:

```
| ?- logtalk_load(dimension_reduction_protocols('test_datasets/correlated_plane')).
```

Learning a reducer

```
| ?- kernel_pca_projection::learn(correlated_plane, DimensionReducer).
| ?- kernel_pca_projection::learn(correlated_plane, DimensionReducer, [n_components(1),
↪kernel(rbf(0.25))]).
```

Transforming new instances

```
| ?- kernel_pca_projection::learn(correlated_plane, DimensionReducer, [n_components(2),
↪kernel(rbf(0.25))]),
   kernel_pca_projection::transform(DimensionReducer, [x-2.0, y-4.0, z-6.0],
↪ReducedInstance).
```

Exporting and reusing the reducer

```
| ?- kernel_pca_projection::learn(correlated_plane, DimensionReducer, [n_components(1)]),
   kernel_pca_projection::export_to_file(correlated_plane, DimensionReducer, reducer,
↪'kernel_pca_reducer.pl').
| ?- logtalk_load('kernel_pca_reducer.pl'),
   reducer(Reducer),
   kernel_pca_projection::transform(Reducer, [x-1.0, y-2.0, z-3.0], ReducedInstance).
```

6.111.7 Dimension reducer representation

The learned dimension reducer is represented by a compound term with the functor chosen by the implementation and arity 7. For example:

```
kernel_pca_reducer(Encoders, TrainingRows, RowMeans, TotalMean, Components,
↳ExplainedVariances, Diagnostics)
```

Where:

- Encoders: List of continuous attribute encoders storing attribute name, mean, and scale.
- TrainingRows: Encoded training rows used when evaluating kernels for new instances.
- RowMeans: Per-training-row kernel means used for centering out-of-sample kernel vectors.
- TotalMean: Global kernel mean used for centering both the training Gram matrix and new kernel vectors.
- Components: List of normalized dual projection vectors in descending variance order.
- ExplainedVariances: List of kernel Gram matrix eigenvalues matching the extracted components.
- Diagnostics: Learned metadata including the effective training options, kernel preprocessing, sample count, explained variances, and optional truncate-mode shortfall details.

When exported using `export_to_clauses/4` or `export_to_file/4`, this reducer term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.111.8 References

1. Schölkopf, B., Smola, A., and Müller, K.-R. (1998) - “Nonlinear component analysis as a kernel eigenvalue problem”.
2. Shawe-Taylor, J. and Cristianini, N. (2004) - “Kernel Methods for Pattern Analysis”.

6.112 kmeans_clusterer

k-Means clusterer. It uses Lloyd’s algorithm with deterministic initialization. Supports continuous attributes only.

The library implements the `clusterer_protocol` defined in the `clustering_protocols` library. It provides predicates for learning a clusterer from a dataset, assigning new instances to clusters, and exporting the learned clusterer as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol from the `clustering_protocols` library.

6.112.1 API documentation

Open the `../apis/library_index.html#kmeans_clusterer` link in a web browser.

6.112.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(kmeans_clusterer(loader)).
```

6.112.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(kmeans_clusterer(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(kmeans_clusterer(tester_performance)).
```

6.112.4 Features

- **Continuous Datasets:** Accepts datasets containing only continuous attributes.
- **Deterministic Initialization:** Supports `first_k` and deterministic spread initialization that repeatedly chooses the farthest example from the centroids selected so far.
- **Optional Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Rich Training Diagnostics:** Learned clusterers report training example count, convergence status, iteration count, and final centroid shift.
- **Portable Export:** Learned clusterers can be exported as clauses or files and reused later.
- **Stable Empty-Cluster Handling:** Empty clusters keep their previous centroids instead of failing.

6.112.5 Options

The following options can be passed to the `learn/3` predicate:

- `k(K)`: Number of clusters to learn. Default is 2.
- `maximum_iterations(Iterations)`: Maximum number of Lloyd iterations. Default is 100.
- `tolerance(Tolerance)`: Maximum centroid shift threshold for convergence. Default is `1.0e-6`.
- `initialization(Initialization)`: Centroid initialization strategy. Options: `spread` (default) or `first_k`.
- `feature_scaling(FeatureScaling)`: Whether to standardize continuous attributes before clustering. Options: `on` (default) or `off`.

6.112.6 Diagnostics

The `diagnostics/2` predicate returns a list containing:

- `model(kmeans_clusterer)`
- `centroid_count(Count)`
- `training_example_count(Count)`
- `convergence(Reason)`
- `iterations(Count)`
- `final_shift(Shift)`
- `options(Options)`

6.112.7 Clusterer representation

The learned clusterer is represented as a compound term with the functor chosen by the user when exporting the clusterer and arity 4. For example:

```
kmeans_clusterer(Encoders, Centroids, Options, Diagnostics)
```

Where:

- **Encoders:** List of continuous attribute encoders storing attribute name, mean, and scale.
- **Centroids:** List of centroid vectors in cluster-id order.
- **Options:** Effective training options used to learn the clusterer.
- **Diagnostics:** Training diagnostics metadata returned by the `diagnostics/2` predicate.

6.112.8 References

1. MacQueen (1967) - "Some Methods for Classification and Analysis of Multivariate Observations". Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability.
2. Lloyd (1982) - "Least Squares Quantization in PCM". IEEE Transactions on Information Theory, 28(2), 129-137.

6.113 kmedians_clusterer

k-Medians clusterer. It uses an iterative median-update algorithm with deterministic initialization. Supports continuous attributes only.

The library implements the `clusterer_protocol` defined in the `clustering_protocols` library. It provides predicates for learning a clusterer from a dataset, assigning new instances to clusters, and exporting the learned clusterer as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol from the `clustering_protocols` library.

6.113.1 API documentation

Open the `../apis/library_index.html#kmedians_clusterer` link in a web browser.

6.113.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(kmedians_clusterer(loader)).
```

6.113.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(kmedians_clusterer(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(kmedians_clusterer(tester_performance)).
```

6.113.4 Features

- **Continuous Datasets:** Accepts datasets containing only continuous attributes.
- **Deterministic Initialization:** Supports `first_k` and deterministic spread initialization that repeatedly chooses the farthest example from the medians selected so far.
- **Optional Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Manhattan Distance:** Uses Manhattan distance for cluster assignment and convergence checks.
- **Rich Training Diagnostics:** Learned clusterers report training example count, convergence status, iteration count, and final median shift.
- **Portable Export:** Learned clusterers can be exported as clauses or files and reused later.
- **Stable Empty-Cluster Handling:** Empty clusters keep their previous medians instead of failing.

6.113.5 Options

The following options can be passed to the `learn/3` predicate:

- `k(K)`: Number of clusters to learn. Default is 2.
- `maximum_iterations(Iterations)`: Maximum number of median-update iterations. Default is 100.
- `tolerance(Tolerance)`: Maximum median shift threshold for convergence. Default is $1.0e-6$.
- `initialization(Initialization)`: Median initialization strategy. Options: `spread` (default) or `first_k`.
- `feature_scaling(FeatureScaling)`: Whether to standardize continuous attributes before clustering. Options: `on` (default) or `off`.

6.113.6 Diagnostics

The `diagnostics/2` predicate returns a list containing:

- `model(kmedians_clusterer)`
- `median_count(Count)`
- `training_example_count(Count)`
- `convergence(Reason)`
- `iterations(Count)`
- `final_shift(Shift)`
- `options(Options)`

6.113.7 Clusterer representation

The learned clusterer is represented as a compound term with the functor chosen by the user when exporting the clusterer and arity 4. For example:

```
kmedians_clusterer(Encoders, Medians, Options, Diagnostics)
```

Where:

- **Encoders:** List of continuous attribute encoders storing attribute name, mean, and scale.
- **Medians:** List of median vectors in cluster-id order.
- **Options:** Effective training options used to learn the clusterer.
- **Diagnostics:** Training diagnostics metadata returned by the `diagnostics/2` predicate.

6.113.8 References

1. MacQueen (1967) - "Some Methods for Classification and Analysis of Multivariate Observations". Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability.
2. Kaufman & Rousseeuw (1990) - "Finding Groups in Data: An Introduction to Cluster Analysis". Wiley.

6.114 kmedoids_clusterer

k-Medoids clusterer. It uses an iterative medoid-update algorithm with deterministic initialization and deterministic cluster assignments. Supports continuous attributes only.

The library implements the `clusterer_protocol` defined in the `clustering_protocols` library. It provides predicates for learning a clusterer from a dataset, assigning new instances to clusters, and exporting the learned clusterer as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol from the `clustering_protocols` library.

6.114.1 API documentation

Open the `../apis/library_index.html#kmedoids_clusterer` link in a web browser.

6.114.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(kmedoids_clusterer(loader)).
```

6.114.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(kmedoids_clusterer(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(kmedoids_clusterer(tester_performance)).
```

6.114.4 Features

- **Continuous Datasets:** Accepts datasets containing only continuous attributes.
- **Distance Metrics:** Supports Euclidean and Manhattan distances.
- **Deterministic Initialization:** Supports `first_k` and deterministic spread initialization that repeatedly chooses the farthest example from the medoids selected so far.
- **Optional Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Rich Training Diagnostics:** Learned clusterers report training example count, convergence status, iteration count, and final medoid shift.
- **Portable Export:** Learned clusterers can be exported as clauses or files and reused later.
- **Stable Empty-Cluster Handling:** Empty clusters keep their previous medoids instead of failing.

6.114.5 Options

The following options can be passed to the `learn/3` predicate:

- `k(K)`: Number of clusters to learn. Default is 2.
- `maximum_iterations(Iterations)`: Maximum number of medoid-update iterations. Default is 100.
- `tolerance(Tolerance)`: Maximum medoid shift threshold for convergence. Default is $1.0e-6$.
- `initialization(Initialization)`: Medoid initialization strategy. Options: `spread` (default) or `first_k`.
- `distance_metric(Metric)`: Distance metric to use. Options: `euclidean` (default) or `manhattan`.
- `feature_scaling(FeatureScaling)`: Whether to standardize continuous attributes before clustering. Options: `on` (default) or `off`.

6.114.6 Diagnostics

The `diagnostics/2` predicate returns a list containing:

- `model(kmedoids_clusterer)`
- `medoid_count(Count)`
- `training_example_count(Count)`
- `convergence(Reason)`
- `iterations(Count)`
- `final_shift(Shift)`
- `options(Options)`

6.114.7 Clusterer representation

The learned clusterer is represented as a compound term with the functor chosen by the user when exporting the clusterer and arity 4. For example:

```
kmedoids_clusterer(Encoders, Medoids, Options, Diagnostics)
```

Where:

- **Encoders:** List of continuous attribute encoders storing attribute name, mean, and scale.
- **Medoids:** List of medoid vectors in cluster-id order.
- **Options:** Effective training options used to learn the clusterer.
- **Diagnostics:** Training diagnostics metadata returned by the `diagnostics/2` predicate.

6.114.8 References

1. Kaufman & Rousseeuw (1987) - "Clustering by Means of Medoids". In Statistical Data Analysis Based on the L1-Norm and Related Methods.
2. Kaufman & Rousseeuw (1990) - "Finding Groups in Data: An Introduction to Cluster Analysis". Wiley.

6.115 kmodes_clusterer

k-Modes clusterer. It uses an iterative mode-update algorithm with deterministic initialization and deterministic cluster assignments. Supports discrete attributes only.

The library implements the `clusterer_protocol` defined in the `clustering_protocols` library. It provides predicates for learning a clusterer from a dataset, assigning new instances to clusters, and exporting the learned clusterer as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol from the `clustering_protocols` library.

6.115.1 API documentation

Open the `../apis/library_index.html#kmodes_clusterer` link in a web browser.

6.115.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(kmodes_clusterer(loader)).
```

6.115.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(kmodes_clusterer(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(kmodes_clusterer(tester_performance)).
```

6.115.4 Features

- **Discrete Datasets:** Accepts datasets containing only discrete attributes.
- **Deterministic Initialization:** Supports `first_k` and deterministic spread initialization that repeatedly chooses the farthest example from the modes selected so far.
- **Rich Training Diagnostics:** Learned clusterers report training example count, convergence status, iteration count, and final mode shift.
- **Portable Export:** Learned clusterers can be exported as clauses or files and reused later.
- **Stable Empty-Cluster Handling:** Empty clusters keep their previous modes instead of failing.

6.115.5 Options

The following options can be passed to the `learn/3` predicate:

- `k(K)`: Number of clusters to learn. Default is 2.
- `maximum_iterations(Iterations)`: Maximum number of mode-update iterations. Default is 100.
- `tolerance(Tolerance)`: Maximum mode shift threshold for convergence. Default is 0.0.
- `initialization(Initialization)`: Mode initialization strategy. Options: `spread` (default) or `first_k`.

6.115.6 Diagnostics

The `diagnostics/2` predicate returns a list containing:

- `model(kmodes_clusterer)`
- `mode_count(Count)`
- `training_example_count(Count)`
- `convergence(Reason)`
- `iterations(Count)`
- `final_shift(Shift)`
- `options(Options)`

6.115.7 Clusterer representation

The learned clusterer is represented as a compound term with the functor chosen by the user when exporting the clusterer and arity 4. For example:

```
kmodes_clusterer(Encoders, Modes, Options, Diagnostics)
```

Where:

- `Encoders`: List of discrete attribute encoders.
- `Modes`: List of learned categorical modes in cluster-id order.
- `Options`: Effective training options used to learn the clusterer.
- `Diagnostics`: Training diagnostics metadata returned by the `diagnostics/2` predicate.

6.115.8 References

1. Huang (1998) - "Extensions to the k-means algorithm for clustering large data sets with categorical values". *Data Mining and Knowledge Discovery*, 2, 283-304.

6.116 knn_classifier

k-Nearest Neighbors classifier supporting multiple distance metrics, weighting schemes, and both categorical and continuous features. This is a lazy learning algorithm that classifies instances based on the majority class among the k nearest training instances.

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning a classifier from a dataset, using it to make predictions, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. See `test_files` directory for examples.

6.116.1 API documentation

Open the `.././docs/library_index.html#knn_classifier` link in a web browser.

6.116.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(knn_classifier(loader)).
```

6.116.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(knn_classifier(tester)).
```

6.116.4 Features

- **Multiple Distance Metrics:** Euclidean, Manhattan, Chebyshev, Minkowski.
- **Flexible Weighting:** Uniform, distance-based, and Gaussian weighting of neighbors.
- **Mixed Features:** Automatically handles categorical and continuous features.
- **Configurable Options:** k value, distance metric, and weighting scheme via predicate options.
- **Probability Estimation:** Provides confidence scores for predictions.
- **Classifier Export:** Learned classifiers can be exported as predicate clauses.

6.116.5 Usage

Learning a Classifier

```
% Learn from a dataset object with default options (k=3, euclidean, uniform)
| ?- knn_classifier::learn(my_dataset, Classifier).
...

% Learn with custom options
| ?- knn_classifier::learn(my_dataset, Classifier, [k(5), distance_metric(manhattan)]).
...
```

Making Predictions

```
% Predict class for a new instance
| ?- Instance = [attr1-value1, attr2-value2, ...],
    knn_classifier::learn(my_dataset, Classifier),
    knn_classifier::predict(Classifier, Instance, PredictedClass).
PredictedClass = ...
...

% Predict with custom options
| ?- knn_classifier::predict(Classifier, Instance, PredictedClass, [k(5), weight_
    ↪scheme(distance)]).
...

% Get probability distribution
| ?- knn_classifier::predict_probabilities(Classifier, Instance, Probabilities).
Probabilities = [class1-0.67, class2-0.33]
...
```

Exporting the Classifier

Learned classifiers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- knn_classifier::learn(my_dataset, Classifier),
    knn_classifier::export_to_clauses(my_dataset, Classifier, my_classifier, Clauses).
Clauses = [my_classifier(...)]
...

% Export to a file
| ?- knn_classifier::learn(my_dataset, Classifier),
    knn_classifier::export_to_file(my_dataset, Classifier, my_classifier, 'classifier.pl').
...
```

Using a learned classifier

Learned and saved classifiers can later be used for predictions without needing to access the original training dataset.

```
% Later, load the file and use the classifier
| ?- consult('classifier.pl'),
    my_classifier(Classifier),
    Instance = [...],
    knn_classifier::predict(Classifier, Instance, Class).
Class = ...
...
```

6.116.6 Options

The following options can be passed to the `predict/4` and `predict_probabilities/4` predicates:

- `k(K)`: Number of neighbors to consider (default: 3)
- `distance_metric(Metric)`: Distance metric to use. Options: `euclidean` (default), `manhattan`, `chebyshev`, `minkowski`
- `weight_scheme(Scheme)`: Weighting scheme for neighbor votes. Options: `uniform` (default), `distance`, `gaussian`

6.116.7 Classifier representation

The learned classifier is represented as a compound term:

```
knn_classifier(AttributeNames, FeatureTypes, Instances)
```

Where:

- `AttributeNames`: List of attribute names in order
- `FeatureTypes`: List of types (numeric or categorical)
- `Instances`: List of Values-Class pairs (the training data in compact form)

When exported using `export_to_clauses/4` or `export_to_file/4`, this classifier term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.116.8 References

1. Cover, T. & Hart, P. (1967). “Nearest neighbor pattern classification”. IEEE Transactions on Information Theory.
2. Hastie, T., Tibshirani, R., & Friedman, J. (2009). “The Elements of Statistical Learning”. Chapter 13.
3. Mitchell, T. (1997). “Machine Learning”. Chapter 8: Instance-Based Learning.

6.117 knn_distance_anomaly_detector

k-nearest-neighbor distance anomaly detector supporting multiple distance metrics, mixed continuous and categorical features, and missing values. The detector memorizes the training instances and computes an anomaly score from normalized distances to the nearest neighbors. Larger distances indicate more isolated and therefore more anomalous instances.

The library implements the `anomaly_detector_protocol` defined in the `anomaly_detection_protocols` library. It learns a compact detector from a dataset by selecting baseline training examples from the declared class labels, computes anomaly scores for new instances, predicts normal or anomaly, and exports learned detectors as clauses or files.

Datasets are represented as objects implementing the `anomaly_dataset_protocol` protocol from the `anomaly_detection_protocols` library. See the `anomaly_detection_protocols/test_datasets` directory for examples.

6.117.1 API documentation

Open the `../apis/library_index.html#knn_distance_anomaly_detector` link in a web browser.

6.117.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(knn_distance_anomaly_detector(loader)).
```

6.117.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(knn_distance_anomaly_detector(tester)).
```

6.117.4 Features

- **Distance-based anomaly scoring:** supports both distance to the k-th neighbor and average distance to the k nearest neighbors.
- **Mixed features:** automatically handles continuous and categorical features declared by the dataset.
- **Missing values:** ignores missing dimensions while normalizing distances (distances are normalized by the number of comparable dimensions).
- **Baseline training selection:** `baseline_class_values/1` declares which class labels are admissible for fitting the detector, while `baseline_selection_policy/1` controls whether non-baseline examples are rejected (default) or filtered before training.
- **Multiple metrics:** supports Euclidean, Manhattan, Chebyshev, and Minkowski distance metrics.
- **Detector export:** learned detectors can be exported as predicate clauses.
- **Dataset validation:** learning rejects empty datasets with a `domain_error(non_empty_dataset, Dataset)` exception.

6.117.5 Options

The following options can be passed to the `learn/3` and `predict/4` predicates:

- `k(K)`: Number of neighbors to consider (default: 5)
- `distance_metric(Metric)`: Distance metric to use. Options: `euclidean` (default), `manhattan`, `chebyshev`, `minkowski`
- `score_mode(Mode)`: Score computation mode. Options: `kth_distance` (default) and `mean_distance`
- `anomaly_threshold(Threshold)`: Threshold for `predict/3-4` (default: 0.5)
- `baseline_class_values(Classes)`: Learn-time list of admissible baseline class labels (default: `[normal]`)
- `baseline_selection_policy(Policy)`: Learn-time handling of non-baseline examples. Supported values are `reject` (default) and `filter`

6.117.6 Detector representation

The learned detector is represented by default as:

```
knn_distance_detector(TrainingDataset, AttributeNames, FeatureTypes, AttributeScales,   
↪Instances, ReferenceScores, Diagnostics)
```

Where:

- AttributeNames: List of attribute names in order
- FeatureTypes: List of feature types (numeric or categorical)
- AttributeScales: Normalization scales for numeric features
- Instances: List of retained baseline training Id-Class-Values triples
- ReferenceScores: Cached leave-one-out raw training scores for the retained baseline training instances
- Diagnostics: Learned metadata terms including model/1, training_dataset/1, attribute_names/1, feature_types/1, example_count/1, reference_score_count/1, and options/1

The score/3 predicate always treats its input as a fresh query. Only score_all/3 on the original training dataset with the reject baseline selection policy reuses the cached leave-one-out ReferenceScores for all examples. With the filter policy, retained baseline training examples reuse the cached leave-one-out scores while excluded examples are scored as fresh queries against the learned baseline detector.

When exported using export_to_clauses/4 or export_to_file/4, this detector term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.117.7 References

1. Angiulli, F. and Pizzuti, C. (2002). “Fast outlier detection in high dimensional spaces”. PKDD.
2. Chandola, V., Banerjee, A., and Kumar, V. (2009). “Anomaly detection: A survey”. ACM Computing Surveys.

6.118 knn_regression

k-Nearest Neighbors regressor supporting continuous and mixed-feature datasets. Learns lazily by storing encoded training rows and predicts targets as the weighted average of the k nearest neighbors.

The library implements the regressor_protocol defined in the regression_protocols library, learns lazily by storing encoded training rows, and predicts numeric targets using the weighted average of the nearest encoded neighbors.

6.118.1 API documentation

Open the `../..apis/library_index.html#knn_regression` link in a web browser.

6.118.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(knn_regression(loader)).
```

6.118.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(knn_regression(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(knn_regression(tester_performance)).
```

6.118.4 Features

- **Distance-Based Regression:** Predicts targets using the weighted average of the nearest neighbors.
- **Multiple Metrics:** Supports Euclidean, Manhattan, Chebyshev, and Minkowski distance metrics over encoded numeric feature vectors.
- **Weighting Schemes:** Supports uniform, inverse-distance, and Gaussian weighting of neighbors.
- **Continuous and Mixed Features:** Supports continuous attributes and categorical attributes encoded using reference-level dummy coding derived from the declared dataset attribute values.
- **Optional Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Missing Values:** Missing numeric and categorical values represented using anonymous variables are encoded using explicit missing-value indicator features.
- **Diagnostics Metadata:** Learned regressors record model name, target, training example count, encoded feature count, and effective options, accessible using the shared regression diagnostics predicates.
- **Model Export:** Learned regressors can be exported as predicate clauses or written to a file.
- **Reference Benchmarks:** Includes a dedicated performance suite reporting training time, RMSE, and MAE for representative regression datasets.

6.118.5 Regressor representation

The learned regressor is represented by default as:

- `knn_regressor(Encoders, Rows, Diagnostics)`

The exported predicate clauses therefore use the shape:

- `Functor(Encoders, Rows, Diagnostics)`

In this representation, `Rows` stores encoded feature vectors paired with numeric targets and `Diagnostics` stores training metadata including the effective options.

6.118.6 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(knn_regression),
  target(Target),
  training_example_count(TrainingExampleCount),
  options(Options),
  encoded_feature_count(FeatureCount)
]
```

Where:

- `model(knn_regression)` identifies the learning algorithm that produced the regressor.
- `target(Target)` stores the target attribute name declared by the training dataset.
- `training_example_count(TrainingExampleCount)` stores the number of examples used during training.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `encoded_feature_count(FeatureCount)` stores the number of numeric features induced by the encoder list, including missing-value indicator features.

Use the `regression_protocols` `diagnostic/2` and `regressor_options/2` helper predicates when you only need a single metadata term or the effective options.

6.118.7 Options

The `learn/3` predicate accepts the following options:

- `k/1`: Number of nearest neighbors considered for each prediction. Smaller values make predictions more local; larger values smooth them by averaging over more training rows. The default is 3.
- `distance_metric/1`: Distance function used to compare encoded feature vectors. Accepted values are `euclidean`, `manhattan`, `chebyshev`, and `minkowski`. The default is `euclidean`.
- `weight_scheme/1`: Neighbor weighting policy used when averaging targets. Accepted values are `uniform`, `distance`, and `gaussian`. The default is `uniform`.
- `minkowski_power/1`: Exponent used when `distance_metric(minkowski)` is selected. Larger values increase the influence of larger coordinate differences. The default is 3.0.

- `feature_scaling/1`: Controls z-score standardization of continuous attributes before storing rows and encoding prediction requests. Accepted values are `true` and `false`. The default is `true`.

6.119 kprototypes_clusterer

k-Prototypes clusterer. It uses an iterative prototype-update algorithm with deterministic initialization and deterministic cluster assignments. Supports continuous and discrete attributes in the same dataset.

The library implements the `clusterer_protocol` defined in the `clustering_protocols` library. It provides predicates for learning a clusterer from a dataset, assigning new instances to clusters, and exporting the learned clusterer as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol from the `clustering_protocols` library.

6.119.1 API documentation

Open the ../apis/library_index.html#kprototypes_clusterer link in a web browser.

6.119.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(kprototypes_clusterer(loader)).
```

6.119.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(kprototypes_clusterer(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(kprototypes_clusterer(tester_performance)).
```

6.119.4 Features

- **Mixed Datasets**: Accepts datasets with continuous, discrete, or mixed attributes.
- **Strict Attribute Validation**: Training examples and prediction instances must contain each declared attribute exactly once and no undeclared attributes.
- **Deterministic Initialization**: Supports `first_k` and deterministic spread initialization that repeatedly chooses the farthest example from the prototypes selected so far.
- **Optional Feature Scaling**: Continuous attributes can be standardized using z-score scaling.
- **Categorical Weighting**: Uses a gamma mismatch penalty for discrete attributes in the mixed distance function.
- **Portable Export**: Learned clusterers can be exported as clauses or files and reused later.

- **Stable Empty-Cluster Handling:** Empty clusters keep their previous prototypes instead of failing.
- **Training Diagnostics:** Exposes convergence metadata including training example count, iteration count, and final prototype shift.

6.119.5 Options

The following options can be passed to the `learn/3` predicate:

- `k(K)`: Number of clusters to learn. Default is 2.
- `maximum_iterations(Iterations)`: Maximum number of prototype-update iterations. Default is 100.
- `tolerance(Tolerance)`: Maximum prototype shift threshold for convergence. Default is $1.0e-6$.
- `initialization(Initialization)`: Prototype initialization strategy. Options: `spread` (default) or `first_k`.
- `gamma(Gamma)`: Penalty added for each discrete-feature mismatch. Default is 1.0.
- `feature_scaling(FeatureScaling)`: Whether to standardize continuous attributes before clustering. Options: `on` (default) or `off`.

6.119.6 Distance Function

The mixed k-prototypes distance used for both assignment and prototype spread initialization is:

$$D(X, P) = \text{Sum}((x_i - p_i)^2) + \text{gamma} * M$$

where the sum is taken over all continuous attributes and M is the number of discrete-attribute mismatches between the instance X and the prototype P .

For discrete prototype updates, the selected value for each categorical attribute is the most frequent value among the cluster members. When two or more values are tied, the implementation deterministically keeps the first value in the declared attribute-values list.

6.119.7 Clusterer representation

The learned clusterer is represented as a compound term with the functor chosen by the user when exporting the clusterer and arity 4. For example:

```
kprototypes_clusterer(Encoders, Prototypes, Options, Diagnostics)
```

Where:

- `Encoders`: List of continuous and discrete attribute encoders.
- `Prototypes`: List of learned mixed prototypes in cluster-id order.
- `Options`: Effective training options used to learn the clusterer.
- `Diagnostics`: Training metadata including convergence reason, iteration count, and final prototype shift.

6.119.8 References

1. Huang (1997) - “Clustering large data sets with mixed numeric and categorical values”. Proceedings of the First Pacific-Asia Conference on Knowledge Discovery and Data Mining.
2. Huang (1998) - “Extensions to the k-means algorithm for clustering large data sets with categorical values”. Data Mining and Knowledge Discovery, 2, 283-304.

6.120 ksuid

This library generates random KSUID identifiers:

<https://github.com/segmentio/ksuid>

This library requires a backend supporting unbounded integer arithmetic.

By default, identifiers are represented as atoms and encoded using the canonical Base62 alphabet:

```
0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz
```

The generation of random identifiers uses the `/dev/urandom` random number generator when available. This includes macOS, Linux, *BSD, and other POSIX operating systems. On Windows, a pseudo-random generator is used, randomized using the current wall time.

Identifiers can be generated as atoms, lists of characters, or lists of character codes.

See also the `cuid2`, `ids`, `nanoid`, `snowflakeid`, `uuid`, and `ulid` libraries.

6.120.1 API documentation

Open the `../..apis/library_index.html#ksuid` link in a web browser.

6.120.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ksuid(loader)).
```

6.120.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ksuid(tester)).
```

6.120.4 Usage

To generate a KSUID using the default configuration:

```
| ?- ksuid::generate(KSUID).
KSUID = '2YBXxVf8R5A1x6Yx5Y1AL7bEmel'
yes
```

To generate a KSUID represented as a list of characters:

```
| ?- ksuid(chars, '0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz
↳'):generate(KSUID).
KSUID = ['2','Y','B','X','x','V','f','8','R','5','A','1','x','6','Y','x','5','Y','1','A','L',
↳'7','b','E','m','e','l']
yes
```

6.121 lasso_regression

Lasso regression regressor supporting continuous and mixed-feature datasets. Uses cyclic coordinate descent with soft-thresholding updates for each encoded feature in order to minimize mean squared error plus an L1 penalty on the encoded coefficient vector.

The library implements the `regressor_protocol` defined in the `regression_protocols` library and learns a linear model using cyclic coordinate descent with L1 regularization and soft-thresholding updates for each encoded feature in order to minimize mean squared error plus an L1 penalty on the encoded coefficient vector.

6.121.1 API documentation

Open the ../apis/library_index.html#lasso_regression link in a web browser.

6.121.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(lasso_regression(loader)).
```

6.121.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(lasso_regression(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(lasso_regression(tester_performance)).
```

6.121.4 Features

- **Continuous and Mixed Features:** Supports continuous attributes and categorical attributes
- **Categorical Attributes Encoding:** Uses reference-level dummy coding derived from the declared dataset attribute values, with a missing-value indicator, and the resulting encoded coefficients are regularized independently.
- **Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Missing Values:** Missing numeric and categorical values represented using anonymous variables are encoded using explicit missing-value indicator features.
- **Unknown Values:** Prediction requests containing categorical values that are not declared by the dataset raise a domain error.
- **Coefficient-wise L1 Shrinkage:** Applies soft-thresholding updates independently to every encoded feature, including categorical dummy and missing-indicator features.
- **Diagnostics Metadata:** Learned regressors record model name, target, training example count, optimization stop reason, completed iterations, final parameter delta, encoded feature count, and effective options, accessible using the shared regression diagnostics predicates.
- **Model Export:** Learned regressors can be exported as predicate clauses or written to a file.
- **Reference Benchmarks:** Includes a dedicated performance suite reporting training time, RMSE, and MAE for representative regression datasets.

6.121.5 Regressor representation

The learned regressor is represented by default as:

- `lasso_regressor(Encoders, Bias, Weights, Diagnostics)`

The exported predicate clauses therefore use the shape:

- `Functor(Encoders, Bias, Weights, Diagnostics)`

In this representation, `Encoders` stores the feature encoding metadata, `Bias` stores the intercept, `Weights` stores one coefficient per encoded feature, and `Diagnostics` stores training metadata including the effective options.

6.121.6 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(lasso_regression),
  target(Target),
  training_example_count(TrainingExampleCount),
  options(Options),
  convergence(Status),
  iterations(Iterations),
  final_delta(FinalDelta),
  encoded_feature_count(FeatureCount)
]
```

Where:

- `model(lasso_regression)` identifies the learning algorithm that produced the regressor.
- `target(Target)` stores the target attribute name declared by the training dataset.
- `training_example_count(TrainingExampleCount)` stores the number of examples used during training.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `convergence(Status)` records the optimization stop condition. The current values are `tolerance` when the maximum Karush-Kuhn-Tucker optimality violation across the intercept and all encoded features is within the configured tolerance and `maximum_iterations_exhausted` when training stops because the iteration cap is reached.
- `iterations(Iterations)` stores the number of coordinate-descent sweeps completed during training.
- `final_delta(FinalDelta)` stores the maximum Karush-Kuhn-Tucker optimality violation measured during the final optimization check.
- `encoded_feature_count(FeatureCount)` stores the number of numeric features induced by the encoder list, including missing-value indicator features.

Use the `regression_protocols diagnostic/2` and `regressor_options/2` helper predicates when you only need a single metadata term or the effective options.

6.121.7 Options

The `learn/3` predicate accepts the following options:

- `maximum_iterations/1`: Maximum number of coordinate-descent sweeps to run before stopping even if the tolerance criterion has not been met. The default is `2000`.
- `tolerance/1`: Convergence threshold for the maximum Karush-Kuhn-Tucker optimality violation in a full coordinate-descent sweep. Training stops early when both the intercept condition and all encoded-feature subgradient conditions are satisfied within this value. The default is `1.0e-7`.
- `regularization/1`: L1 penalty coefficient applied independently to every encoded feature during optimization. Higher values increase shrinkage and can reduce overfitting. The default is `0.01`.
- `feature_scaling/1`: Controls z-score standardization of continuous attributes before training and prediction. Accepted values are `true` and `false`. The default is `true`.

6.122 `lda_classifier`

Linear Discriminant Analysis classifier for continuous datasets using a shared pooled covariance estimate with diagonal regularization. The implementation learns one linear discriminant function per class and predicts the class with the highest discriminant score.

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning a classifier from a dataset, using it to make predictions, inspecting class scores, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. All dataset attributes must be declared as continuous.

6.122.1 API documentation

Open the ../docs/library_index.html#lda_classifier link in a web browser.

6.122.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(lda_classifier(loader)).
```

6.122.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(lda_classifier(tester)).
```

6.122.4 Features

- **Continuous Datasets:** Accepts only datasets whose attributes are all declared as continuous.
- **Pooled Covariance Model:** Learns a shared covariance matrix and class-specific means and priors.
- **Regularized Estimation:** Applies configurable diagonal regularization to stabilize covariance inversion.
- **Feature Scaling:** Supports optional z-score scaling of continuous features before fitting the model.
- **Score Inspection:** Provides discriminant scores for all classes using `predict_scores/3`.
- **Classifier Export:** Learned classifiers can be exported as predicate clauses or written to a file.

6.122.5 Options

The `learn/3` predicate supports these options:

- `feature_scaling/1` - whether to standardize continuous attributes before training (default: `true`)
- `regularization/1` - positive diagonal value added to the pooled covariance matrix before inversion (default: `1.0e-6`)

6.122.6 Usage

Learning a classifier

```
| ?- lda_classifier::learn(iris_small, Classifier).
| ?- lda_classifier::learn(iris_small, Classifier, [regularization(1.0e-5)]).
```

Making predictions

```
| ?- lda_classifier::learn(iris_small, Classifier),
    lda_classifier::predict(Classifier, [sepal_length-5.1, sepal_width-3.5, petal_length-1.
↪4, petal_width-0.2], Class).

| ?- lda_classifier::learn(iris_small, Classifier),
    lda_classifier::predict_scores(Classifier, [sepal_length-6.0, sepal_width-2.9, petal_
↪length-4.5, petal_width-1.5], Scores).
```

Exporting the classifier

```
| ?- lda_classifier::learn(iris_small, Classifier),
    lda_classifier::export_to_clauses(iris_small, Classifier, classify, Clauses).

| ?- lda_classifier::learn(iris_small, Classifier),
    lda_classifier::export_to_file(iris_small, Classifier, classify, 'classifier.pl').
```

6.122.7 Classifier representation

The learned classifier is represented as a compound term with the form:

```
lda_classifier(Encoders, Models, Options)
```

Where:

- Encoders: list of continuous feature encoders with learned scaling parameters
- Models: list of `class_model(Class, Prior, Mean, Weights, Offset)` terms
- Options: merged training options used to learn the classifier

When exported using `export_to_clauses/4` or `export_to_file/4`, this classifier term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.122.8 References

1. Fisher, R.A. (1936). “The use of multiple measurements in taxonomic problems”.
2. Hastie, T., Tibshirani, R. and Friedman, J. (2009). “The Elements of Statistical Learning”. Section 4.3.
3. Bishop, C.M. (2006). “Pattern Recognition and Machine Learning”. Section 4.1.

6.123 lda_projection

Linear Discriminant Analysis projection for labeled continuous datasets. Supports continuous attributes only.

The library implements the `dimension_reducer_protocol` defined in the `dimension_reduction_protocols` library and learns discriminant directions by centering the training data, optionally standardizing continuous attributes, building regularized within-class and between-class scatter matrices, whitening the Fisher criterion using a Cholesky factorization, and extracting discriminant directions using deterministic power iteration with deflation.

Requires a dataset implementing `supervised_dimension_reduction_dataset_protocol` and therefore uses class labels during training.

6.123.1 API documentation

Open the [../apis/library_index.html#lda_projection](http://logtalk.org/..../apis/library_index.html#lda_projection) link in a web browser.

6.123.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(lda_projection(loader)).
```

6.123.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(lda_projection(tester)).
```

6.123.4 Features

- **Supervised Continuous Datasets:** Accepts labeled datasets implementing `supervised_dimension_reduction_dataset_protocol` and containing only continuous attributes.
- **Centering and Optional Scaling:** Centers all attributes and optionally standardizes them before estimating class scatter matrices.
- **Portable Fisher Eigensolver:** Uses regularized within-class scatter, Cholesky whitening, and deterministic power iteration instead of backend-specific linear algebra libraries.
- **Projection API:** Transforms a new instance into a list of `component_N-Value` pairs.
- **Model Export:** Learned reducers can be exported as predicate clauses or written to a file.
- **Missing Values:** Missing or nonnumeric values are rejected.

6.123.5 Options

The `learn/3` predicate accepts the following options:

- `n_components/1`: Maximum number of discriminant components to extract. Requests that exceed the minimum of the number of features and `number_of_classes - 1` raise `domain_error(component_count, Requested-Maximum)`. The default is 2.
- `feature_scaling/1`: Whether to standardize continuous attributes before estimating scatter matrices. Options: `true` (default) or `false`.
- `maximum_iterations/1`: Maximum number of power-iteration steps used when estimating each discriminant direction. The default is 1000.
- `tolerance/1`: Positive convergence tolerance used both for power-iteration stopping and for deciding when additional discriminant directions are numerically negligible. The default is `1.0e-8`.
- `regularization/1`: Positive diagonal regularization added to the within-class scatter matrix before whitening. The default is `1.0e-6`.

6.123.6 Usage

The following examples use the sample labeled dataset shipped with the `dimension_reduction_protocols` library:

```
| ?- logtalk_load(dimension_reduction_protocols('test_datasets/labeled_measurements')).
```

Learning a reducer

```
| ?- lda_projection::learn(labeled_measurements, DimensionReducer).  
  
| ?- lda_projection::learn(labeled_measurements, DimensionReducer, [n_components(1), feature_  
↪scaling(false), maximum_iterations(250), tolerance(1.0e-7), regularization(1.0e-5)]).
```

Transforming new instances

```
| ?- lda_projection::learn(labeled_measurements, DimensionReducer),  
    lda_projection::transform(DimensionReducer, [length-5.1, width-3.5, height-1.4, weight-  
↪0.2], ReducedInstance).  
  
| ?- lda_projection::learn(labeled_measurements, DimensionReducer, [n_components(1)]),  
    lda_projection::transform(DimensionReducer, [length-6.2, width-3.4, height-5.4, weight-  
↪2.3], ReducedInstance).
```

Exporting and reusing the reducer

```
| ?- lda_projection::learn(labeled_measurements, DimensionReducer, [n_components(1)]),
    lda_projection::export_to_file(labeled_measurements, DimensionReducer, reducer, 'lda_
    ↪projection_reducer.pl').

| ?- logtalk_load('lda_projection_reducer.pl'),
    reducer(Reducer),
    lda_projection::transform(Reducer, [length-5.1, width-3.5, height-1.4, weight-0.2], _
    ↪ReducedInstance).
```

6.123.7 Dimension reducer representation

The learned dimension reducer is represented by a compound term with the functor chosen by the implementation and arity 4. For example:

```
lda_projection_reducer(Encoders, Components, ClassValues, Diagnostics)
```

Where:

- Encoders: List of continuous attribute encoders storing attribute name, mean, and scale.
- Components: List of learned discriminant vectors in component order.
- ClassValues: Ordered list of class labels used during training.
- Diagnostics: Learned reducer metadata including the effective training options and model details.

When exported using `export_to_clauses/4` or `export_to_file/4`, this reducer term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.123.8 References

1. Fisher, R. A. (1936) - “The use of multiple measurements in taxonomic problems”.
2. Rao, C. R. (1948) - “The utilization of multiple measurements in problems of biological classification”.

6.124 linda

Linda is a classic coordination model for process communication. This library provides a Linda tuple-space implementation. Requires both multi-threading and sockets support. It supports SWI-Prolog, Trealla Prolog, and XVM.

The tuple-space is a shared blackboard where processes can:

- **Write** tuples using the `out/1-2` predicates
- **Read** tuples (without removing) using the `rd/1-2` and `rd_noblock/1-2` predicates
- **Remove** tuples using the `in/1-2` and `in_noblock/1-2` predicates

The blocking predicates (`in/1-2`, `rd/1-2`) suspend the process until a matching tuple becomes available. The non-blocking variants (`in_noblock/1-2`, `rd_noblock/1-2`) fail immediately if no matching tuple is found. There are also additional predicates for list operations. Tuples are matched using unification.

6.124.1 API documentation

Open the `../apis/library_index.html#linda` link in a web browser.

6.124.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(linda(loader)).
```

6.124.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(linda(tester)).
```

6.124.4 Usage

The main library entities are the `linda_server` and `linda_client` categories. These categories provide a clean separation between server and client code. For backward compatibility with the previous version of this library, a `linda` object importing both categories is also provided. The usage examples that follow use this object. An object importing the `linda_server` category must use the `threaded/0` directive.

Starting a server

To start a Linda server that prints its address:

```
| ?- linda::linda.  
% Server started at localhost:54321  
...
```

To start a server with a callback when it starts:

```
| ?- linda::linda([(Host:Port)-format('Server at ~w:~w~n', [Host, Port])]).
```

The server runs until all clients disconnect after a shutdown request.

Connecting a client

To connect to a server:

```
| ?- linda::linda_client(localhost:54321).
```

Tuple operations

Write a tuple:

```
| ?- linda::out(message(hello, world)).
```

Read a tuple (blocking):

```
| ?- linda::rd(message(X, Y)).
X = hello,
Y = world
```

Remove a tuple (blocking):

```
| ?- linda::in(message(X, Y)).
X = hello,
Y = world
```

Non-blocking operations:

```
| ?- linda::rd_noblock(message(X, Y)).
no    % if no matching tuple exists

| ?- linda::in_noblock(message(X, Y)).
no    % if no matching tuple exists
```

Wait for one of several tuples:

```
| ?- linda::in([task(1, X), task(2, X), done], Tuple).
```

Alternative syntax using in_list/2:

```
| ?- linda::in_list([task(1, X), task(2, X), done], Tuple).
```

Similarly for reading:

```
| ?- linda::rd_list([status(ready), status(waiting)], Status).
```

Collect all matching tuples atomically (read without removing):

```
| ?- linda::findall_rd_noblock(N, counter(N), Counters).
```

Collect and remove all matching tuples atomically:

```
| ?- linda::findall_in_noblock(X, item(X), Items).
```

Disconnecting

Close the client connection:

```
| ?- linda::close_client.
```

Request server shutdown (server stops accepting new connections):

```
| ?- linda::shutdown_server, linda::close_client.
```

Timeout control

Set a timeout for blocking operations:

```
| ?- linda::linda_timeout(Old, 5:0). % 5 seconds timeout
```

Disable timeout (wait forever):

```
| ?- linda::linda_timeout(_, off).
```

6.124.5 Examples

Producer-Consumer

Producer (client 1):

```
producer :-  
    produce(X),  
    linda::out(item(X)),  
    producer.
```

Consumer (client 2):

```
consumer :-  
    linda::in(item(X)),  
    consume(X),  
    consumer.
```

Critical Region

```
critical_section :-  
    linda::in(mutex), % acquire lock  
    do_critical_work,  
    linda::out(mutex). % release lock
```

Initialize the mutex:

```
| ?- linda::out(mutex).
```

Synchronization

Wait for a signal:

```
| ?- linda::in(ready). % blocks until someone does out(ready)
```

Send a signal:

```
| ?- linda::out(ready).
```

6.124.6 API Summary

Server predicates

- `linda` - Start server on automatic port
- `linda(Options)` - Start server with options

Client predicates

- `linda_client(Address)` - Connect to server at Host:Port
- `close_client` - Close connection
- `shutdown_server` - Request server shutdown
- `linda_timeout(Old, New)` - Get/set timeout

Tuple predicates (single/default server)

- `out(Tuple)` - Add tuple to space
- `in(Tuple)` - Remove matching tuple (blocking)
- `in_noblock(Tuple)` - Remove matching tuple (non-blocking)
- `in(TupleList, Tuple)` - Remove one of several tuples (blocking)
- `in_list(TupleList, Tuple)` - Remove one of several tuples (blocking)
- `rd(Tuple)` - Read matching tuple (blocking)
- `rd_noblock(Tuple)` - Read matching tuple (non-blocking)
- `rd(TupleList, Tuple)` - Read one of several tuples (blocking)
- `rd_list(TupleList, Tuple)` - Read one of several tuples (blocking)
- `findall_rd_noblock(Template, Tuple, List)` - Collect all matching tuples (atomic read)
- `findall_in_noblock(Template, Tuple, List)` - Collect and remove all matching tuples (atomic remove)

Tuple predicates (multiple servers)

Variants of the predicates above with an additional first argument for the server address or the server alias (an atom; default is blackboard).

6.124.7 Known issues

Recent versions of macOS seem to disable the mapping of localhost to 127.0.0.1. This issue may prevent some functionality from working. This can be fixed either by editing the /etc/hosts file or by using '127.0.0.1' as the host argument instead of localhost.

6.124.8 See also

- SICStus Prolog Linda documentation: https://sicstus.sics.se/sicstus/docs/latest4/html/sicstus.html/lib_002dlinda.html
- Original Linda papers:
 - Carreiro, N. & Gelernter, D. (1989). Linda in Context.
 - Carreiro, N. & Gelernter, D. (1989). How to Write Parallel Programs: A Guide to the Perplexed.

6.125 linear_algebra

This library provides predicates for numeric vectors and matrices, including vector construction and scaling, general vector norms, matrix construction and transformation, matrix-matrix products, row Gram matrices and Gram-centering helpers, offset diagonal helpers, triangular-part extraction, diagonal shifts, lower- and upper-triangular matrix solves, direct row and column lookup, tolerance-aware vector normalization and sign stabilization, generic dense square-system solves, determinants, matrix inversion, thin QR decomposition, QR-backed least-squares solving, tolerance-aware rank estimation, real symmetric eigendecomposition, pseudo-inverse and null-space helpers, covariance helpers, and Cholesky-based linear solves.

Vectors are represented as lists of numbers. Matrices are represented as lists of row lists, where each row is a numeric vector.

Predicates that index into vectors or matrices use one-based indices. Matrix access predicates use one-based row and column indices.

Square-system predicates include `solve_linear_system/3`, `solve_linear_systems/3`, `determinant/2`, and `inverse_matrix/2`. Rectangular-system predicates include `qr_decomposition/3` and `least_squares/3`, `pseudo_inverse/2-3`, and `null_space/2-3`. Row-matrix structural predicates include `gram_matrix/2`, `matrix_row_means/2`, `matrix_column_means/2`, `center_gram_matrix/4`, and `center_gram_vector/4`. Diagonal and triangular helpers include `matrix_diagonal/2-3`, `diagonal_matrix/2-4`, `upper_triangular_part/2-3`, and `lower_triangular_part/2-3`. Norm helpers include `vector_norm/3` and `frobenius_norm/2`. Real symmetric spectral predicates include `symmetric_eigen/3-5`. Tolerance-sensitive numerical predicates include `normalize_vector/2-3`, `first_significant_component/2-3`, `stabilize_vector_sign/2-3`, `matrix_rank/2-3`, `symmetric_eigen/3-5`, `pseudo_inverse/2-3`, and `null_space/2-3`.

The `symmetric_eigen/5` predicate can be used when callers need to control the power-iteration budget explicitly instead of relying on the default iteration bound.

6.125.1 API documentation

Open the ../apis/library_index.html#linear_algebra link in a web browser.

6.125.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(linear_algebra(loader)).
```

6.125.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(linear_algebra(tester)).
```

6.126 linear_regression

Linear regression regressor supporting continuous and mixed-feature datasets. The library implements the `regressor_protocol` defined in the `regression_protocols` library and learns a linear model using the shared regression encoding core to build a row-oriented design matrix with an explicit intercept column before delegating least-squares solving and rank estimation to the `linear_algebra` library. The intercept is always retained and encoded feature columns that are numerically dependent on the design matrix are assigned zero coefficients.

6.126.1 API documentation

Open the ../apis/library_index.html#linear_regression link in a web browser.

6.126.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(linear_regression(loader)).
```

6.126.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(linear_regression(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(linear_regression(tester_performance)).
```

6.126.4 Features

- **Continuous and Mixed Features:** Supports continuous attributes and categorical attributes. - **Categorical Attributes Encoding:** Uses reference-level dummy coding from the declared dataset attribute values. Encoded columns with no independent signal after accounting for the intercept and previously selected features are assigned zero coefficients.
- **Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Missing Values:** Missing numeric and categorical values represented using anonymous variables are encoded using explicit missing-value indicator features.
- **Unknown Values:** Prediction requests containing categorical values that are not declared by the dataset raise a domain error.
- **Rank Handling:** Encoded columns that are numerically dependent on the intercept or on earlier selected columns are dropped from the direct solve and assigned zero coefficients.
- **Diagnostics Metadata:** Learned regressors record model name, target, training example count, solver name, residual sum of squares, effective rank, active feature count, encoded feature count, and effective options, accessible using the shared regression diagnostics predicates.
- **Model Export:** Learned regressors can be exported as predicate clauses or written to a file.
- **Reference Benchmarks:** Includes a dedicated performance suite reporting training time, RMSE, and MAE for representative regression datasets.

6.126.5 Regressor representation

The learned regressor is represented by default as:

- `linear_regressor(Encoders, Bias, Weights, Diagnostics)`

The exported predicate clauses therefore use the shape:

- `Functor(Encoders, Bias, Weights, Diagnostics)`

In this representation, Encoders stores the feature encoding metadata, Bias stores the intercept, Weights stores one coefficient per encoded feature, and Diagnostics stores training metadata including the effective options.

6.126.6 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(linear_regression),
  target(Target),
  training_example_count(TrainingExampleCount),
  options(Options),
  solver(Solver),
  residual_sum_of_squares(ResidualSumOfSquares),
  effective_rank(EffectiveRank),
  active_feature_count(ActiveFeatureCount),
  encoded_feature_count(FeatureCount)
]
```

Where:

- `model(linear_regression)` identifies the learning algorithm that produced the regressor.
- `target(Target)` stores the target attribute name declared by the training dataset.
- `training_example_count(TrainingExampleCount)` stores the number of examples used during training.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `solver(Solver)` records the direct least-squares solver family used for the fit. The current value is `modified_gram_schmidt_column_pivoting`, which is now reported by the shared regression core while delegating the actual solve to the `linear_algebra` library.
- `residual_sum_of_squares(ResidualSumOfSquares)` stores the training residual sum of squares for the fitted regressor.
- `effective_rank(EffectiveRank)` stores the rank of the fitted row-oriented design matrix, including the intercept column.
- `active_feature_count(ActiveFeatureCount)` stores the number of encoded feature columns retained after subtracting the intercept contribution from the fitted design-matrix rank.
- `encoded_feature_count(FeatureCount)` stores the number of numeric features induced by the encoder list, including missing-value indicator features.

Use the `regression_protocols diagnostic/2` and `regressor_options/2` helper predicates when you only need a single metadata term or the effective options.

6.126.7 Options

The `learn/3` predicate accepts the following options:

- `feature_scaling/1`: Controls z-score standardization of continuous attributes before training and prediction. Accepted values are `true` and `false`. The default is `true`.

6.127 linear_svm_classifier

Linear support vector machine classifier supporting both binary and multiclass classification. Multiclass classification is implemented using a one-vs-rest scheme with one linear margin model per class, and training uses batch subgradient descent to fit each hinge-loss model.

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning a classifier from a dataset object, using it to make predictions, and exporting the learned model as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. Existing datasets in `classification_protocols/test_datasets` can be used for binary, multiclass, continuous, categorical, and mixed-feature testing.

6.127.1 API documentation

Open the `../docs/library_index.html#linear_svm_classifier` link in a web browser.

6.127.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(linear_svm_classifier(loader)).
```

6.127.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(linear_svm_classifier(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(linear_svm_classifier(tester_performance)).
```

6.127.4 Features

- **Binary and Multiclass Classification:** Learns one-vs-rest linear margin models and predicts the class with the highest score.
- **Continuous Features:** Standardizes numeric attributes using z-score scaling derived from the training data.
- **Categorical Features:** Expands discrete attributes using one-hot encoding based on the declared dataset attribute values and rejects unseen values with a domain error.
- **Missing Values:** Encodes missing numeric and categorical values represented using anonymous variables using explicit missing-value indicator features.
- **Unknown values:** Prediction requests containing categorical values that are not declared by the dataset raise a domain error.
- **Classifier Export:** Learned classifiers can be exported as predicate clauses or written to a file.
- **Pretty Printing:** Prints a compact summary of classes, encoders, and learned weight vector sizes.
- **Reference Benchmarks:** Includes a dedicated performance suite covering the weather, mixed, iris_small, missing_mixed, and breast_cancer datasets with reported training time and training accuracy.

6.127.5 Options

The learn/3 predicate supports these options:

- learning_rate/1 - base gradient descent learning rate (default: 0.1)
- maximum_iterations/1 - maximum number of optimization iterations (default: 1000)
- tolerance/1 - convergence threshold for the maximum parameter update (default: 1.0e-6)
- l2_regularization/1 - L2 penalty factor applied to weights (default: 0.01)

6.127.6 Usage

Learning a classifier

```
| ?- linear_svm_classifier::learn(weather, Classifier).

| ?- linear_svm_classifier::learn(iris_small, Classifier, [learning_rate(0.05), maximum_
↪ iterations(1500)]).
```

Making predictions

```
| ?- linear_svm_classifier::learn(mixed, Classifier),
    linear_svm_classifier::predict(Classifier, [age-45, income-75000, student-no, credit_
↪ rating-fair], Class).
```

Exporting the classifier

```
| ?- linear_svm_classifier::learn(weather, Classifier),
    linear_svm_classifier::export_to_clauses(weather, Classifier, classify, Clauses).

| ?- linear_svm_classifier::learn(weather, Classifier),
    linear_svm_classifier::export_to_file(weather, Classifier, classify, 'classifier.pl').
```

6.127.7 Classifier representation

The learned classifier is represented as a compound term with the form:

```
linear_svm_classifier(Classes, Encoders, Models, Options)
```

Where:

- Classes: list of class labels
- Encoders: list of continuous scaling descriptors and categorical value lists
- Models: list of class_model(Class, Bias, Weights) terms
- Options: merged training options used to learn the model

When exported using `export_to_clauses/4` or `export_to_file/4`, this classifier term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.127.8 References

1. Cortes, C. and Vapnik, V. (1995). “Support-Vector Networks”.
2. Bishop, C.M. (2006). “Pattern Recognition and Machine Learning”. Section 7.1.
3. Hastie, T., Tibshirani, R. and Friedman, J. (2009). “The Elements of Statistical Learning”. Section 12.3.

6.128 listing

This library provides support for listing object dynamic predicate clauses. The predicates must (also) be declared using a scope directive to make them visible to the listing predicates.

6.128.1 API documentation

Open the ../apis/library_index.html#listing link in a web browser.

6.128.2 Loading

To load all entities in this library, load the `loader.lgt` utility file:

```
| ?- logtalk_load(listing(loader)).
```

6.128.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(listing(tester)).
```

6.128.4 Usage

This library provides a `listing` category that can be imported by any number of objects. The main predicates are declared public. If you want to restrict their scope, use `protected` or `private import`. For example:

```
:- object(data_store,  
    imports(private::listing)).  
  
    debug :-  
        ^^listing(data/4).  
    ...  
:- end_object.
```

The `listing` category provides a bare-bones `portray_clause/1` predicate implementation. As this predicate is called (by the `listing/0-1` predicates) using the `:::/1` control construct, the object importing the category can easily override the inherited definition with its own or with a call to the backend system native implementation of the predicate. For example, assuming a backend that provides `portray_clause/1` as a built-in predicate, we can write:

```
:- object(thing,
    imports(listing)).

    :- uses(user, [portray_clause/1]).
    ...

:- end_object.
```

The main predicate, `listing/1`, accepts as argument a predicate indicator, a non-terminal indicator, or a clause head template (to list only clauses with a matching head).

This library is often useful as a debugging helper. For example, assuming that we want to list dynamic predicate clauses for an object data compiled (or created) with the complements flag set to allow, we can hot patch it to add the `listing` category:

```
| ?- create_category(patch, [extends(listing),complements(data)], [], []).
yes

| ?- data::listing.
...

```

Another example:

```
| ?- create_category(
    patch,
    [extends(listing),complements(data)],
    [public(debug/1)],
    [(debug(Key) :- ::listing(p(Key,Datum)))])
).
yes

| ?- data::debug(k42).
...

```

Note that the semantics of complementing categories require that we use the `:::/2` control construct instead of the `(^^)/2` control construct as a *super* call would be interpreted as made from the complemented object.

6.129 lof_anomaly_detector

Local Outlier Factor anomaly detector supporting multiple distance metrics, mixed continuous and categorical features, and missing values. The detector memorizes the training instances and computes Local Outlier Factor values by comparing the local reachability density of a query to the densities of its neighbors.

The library implements the `anomaly_detector_protocol` defined in the `anomaly_detection_protocols` library. It learns a compact detector from a dataset by selecting baseline training examples from the declared class labels, computes normalized anomaly scores for new instances, predicts normal or anomaly, and exports learned detectors as clauses or files.

Datasets are represented as objects implementing the `anomaly_dataset_protocol` protocol from the `anomaly_detection_protocols` library. See the `anomaly_detection_protocols/test_datasets` directory for examples.

6.129.1 API documentation

Open the `../apis/library_index.html#lof_anomaly_detector` link in a web browser.

6.129.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(lof_anomaly_detector(loader)).
```

6.129.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(lof_anomaly_detector(tester)).
```

6.129.4 Features

- **Density-based anomaly scoring:** computes Local Outlier Factor scores from local reachability densities.
- **Normalized scores:** Raw LOF values are normalized to the interval $[0.0, 1.0]$ by mapping the ideal baseline value 1.0 to 0.0 and scaling larger values against the largest training raw score.
- **Mixed features:** automatically handles continuous and categorical features declared by the dataset.
- **Missing values:** ignores missing dimensions while normalizing distances (distances are normalized by the number of comparable dimensions).
- **Baseline training selection:** `baseline_class_values/1` declares which class labels are admissible for fitting the detector, while `baseline_selection_policy/1` controls whether non-baseline examples are rejected (default) or filtered before training.
- **Multiple metrics:** supports Euclidean, Manhattan, Chebyshev, and Minkowski distance metrics.
- **Detector export:** learned detectors can be exported as predicate clauses.
- **Dataset validation:** learning rejects empty datasets with a `domain_error(non_empty_dataset, Dataset)` exception.

6.129.5 Options

The following options can be passed to the `learn/3` and `predict/4` predicates:

- `k(K)`: Number of neighbors to consider (default: 5)
- `distance_metric(Metric)`: Distance metric to use. Options: `euclidean` (default), `manhattan`, `chebyshev`, `minkowski`
- `anomaly_threshold(Threshold)`: Threshold for `predict/3-4` (default: 0.4)
- `baseline_class_values(Classes)`: Learn-time list of admissible baseline class labels (default: `[normal]`)
- `baseline_selection_policy(Policy)`: Learn-time handling of non-baseline examples. Supported values are `reject` (default) and `filter`

6.129.6 Detector representation

The learned detector is represented by default as:

```
lof_detector(TrainingDataset, AttributeNames, FeatureTypes, AttributeScales, Instances,
↳ReferenceScores, Diagnostics)
```

Where:

- `AttributeNames`: List of attribute names in order
- `FeatureTypes`: List of feature types (numeric or categorical)
- `AttributeScales`: Normalization scales for numeric features
- `Instances`: List of retained baseline training Id-Class-Values triples
- `ReferenceScores`: Cached leave-one-out raw training scores for the retained baseline training instances
- `Diagnostics`: Learned metadata terms including `model/1`, `training_dataset/1`, `attribute_names/1`, `feature_types/1`, `example_count/1`, `reference_score_count/1`, and `options/1`

The `score/3` predicate always treats its input as a fresh query. Only `score_all/3` on the original training dataset with the `reject` baseline selection policy reuses the cached leave-one-out `ReferenceScores` for all examples. With the `filter` policy, retained baseline training examples reuse the cached leave-one-out scores while excluded examples are scored as fresh queries against the learned baseline detector.

When exported using `export_to_clauses/4` or `export_to_file/4`, this detector term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.129.7 References

1. Breunig, M. M., Kriegel, H.-P., Ng, R. T., and Sander, J. (2000). "LOF: Identifying density-based local outliers". SIGMOD.

6.130 logging

This library provides support for logging events to files.

6.130.1 API documentation

Open the ../apis/library_index.html#logging link in a web browser.

6.130.2 Loading

To load all entities in this library, load the `loader.lgt` utility file:

```
| ?- logtalk_load(logging(loader)).
```

6.131 logistic_regression_classifier

Logistic regression classifier supporting both binary and multiclass classification. Multiclass classification is implemented using batch gradient descent to train a single multiclass softmax model. Binary classification is treated as a two-class special case of the same objective.

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning a classifier from a dataset object, using it to make predictions, returning class probabilities, and exporting the learned model as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. Existing datasets in `classification_protocols/test_datasets` can be used for binary, multiclass, continuous, categorical, and mixed-feature testing.

6.131.1 API documentation

Open the ../docs/library_index.html#logistic_regression_classifier link in a web browser.

6.131.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(logistic_regression_classifier(loader)).
```

6.131.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(logistic_regression_classifier(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(logistic_regression_classifier(tester_performance)).
```

6.131.4 Features

- **Binary and Multiclass Classification:** Learns a joint softmax logistic model with one parameter vector per class.
- **Continuous Features:** Standardizes numeric attributes using z-score scaling derived from the training data.
- **Categorical Features:** Expands discrete attributes using one-hot encoding based on the declared dataset attribute values and rejects unseen values with a domain error.
- **Missing Values:** Encodes missing numeric and categorical values represented using anonymous variables using explicit missing-value indicator features instead of being conflated with baseline feature values.
- **Unknown values:** Prediction requests containing categorical values that are not declared by the dataset raise a domain error instead of being silently mapped into an existing feature bucket.
- **Probability Estimation:** Provides class probability distributions in addition to class predictions.
- **Classifier Export:** Learned classifiers can be exported as predicate clauses or written to a file.
- **Reference Benchmarks:** Includes a dedicated performance suite covering the weather, mixed, iris_small, missing_mixed, and breast_cancer datasets with reported training time, training accuracy, and mean log loss.

6.131.5 Options

The learn/3 predicate supports these options:

- learning_rate/1 - gradient descent learning rate (default: 0.1)
- maximum_iterations/1 - maximum number of optimization iterations (default: 1000)
- tolerance/1 - convergence threshold for the maximum parameter update (default: 1.0e-6)
- l2_regularization/1 - L2 penalty factor applied to weights (default: 0.0)

6.131.6 Usage

Learning a classifier

```
| ?- logistic_regression_classifier::learn(weather, Classifier).
| ?- logistic_regression_classifier::learn(iris_small, Classifier, [learning_rate(0.05),
↪maximum_iterations(1500)]).
```

Making predictions

```
| ?- logistic_regression_classifier::learn(mixed, Classifier),
    logistic_regression_classifier::predict(Classifier, [age-45, income-75000, student-no,
↳credit_rating-fair], Class).

| ?- logistic_regression_classifier::learn(iris_small, Classifier),
    logistic_regression_classifier::predict_probabilities(Classifier, [sepal_length-6.4,
↳sepal_width-3.0, petal_length-5.8, petal_width-2.2], Probabilities).
```

Exporting the classifier

```
| ?- logistic_regression_classifier::learn(weather, Classifier),
    logistic_regression_classifier::export_to_clauses(weather, Classifier, classify,
↳Clauses).

| ?- logistic_regression_classifier::learn(weather, Classifier),
    logistic_regression_classifier::export_to_file(weather, Classifier, classify,
↳'classifier.pl').
```

6.131.7 Classifier representation

The learned classifier is represented as a compound term with the form:

```
lr_classifier(Classes, Encoders, Models, Options)
```

Where:

- *Classes*: list of class labels
- *Encoders*: list of continuous scaling descriptors and categorical value lists
- *Models*: list of `class_model(Class, Bias, Weights)` terms
- *Options*: merged training options used to learn the model

When exported using `export_to_clauses/4` or `export_to_file/4`, this classifier term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.131.8 References

1. Hosmer, D.W., Lemeshow, S. and Sturdivant, R.X. (2013). “Applied Logistic Regression”.
2. Bishop, C.M. (2006). “Pattern Recognition and Machine Learning”. Chapter 4.
3. Hastie, T., Tibshirani, R. and Friedman, J. (2009). “The Elements of Statistical Learning”. Chapter 4.

6.132 loops

This library provides implementations of several kinds of loops typical of imperative languages.

6.132.1 API documentation

Open the `../../apis/library_index.html#loops` link in a web browser.

6.132.2 Loading

To load the main entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(loops(loader)).
```

6.132.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(loops(tester)).
```

6.132.4 Usage

See e.g. the searching example.

6.133 massey_ranker

Massey pairwise preference ranker. It builds the Massey coefficient matrix with diagonal entries `games_i`, off-diagonal entries `-games_ij`, and a final `sum(ratings)=0` anchoring row, then solves the linear system using deterministic Gaussian elimination with partial pivoting and residual validation.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from pairwise preferences, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The current implementation requires a well-formed connected pairwise dataset so that learned rankings remain globally comparable across all ranked items.

6.133.1 API documentation

Open the `../apis/library_index.html#massey_ranker` link in a web browser.

6.133.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(massey_ranker(loader)).
```

6.133.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(massey_ranker(tester)).
```

6.133.4 Features

- **Pairwise Preference Learning:** Learns one deterministic rating per item from aggregated pairwise outcomes.
- **Massey Matrix Fidelity:** Solves the standard Massey linear system with diagonal entries `games_i`, off-diagonal entries `-games_ij`, and a final anchoring row enforcing `sum(ratings) = 0`.
- **Numerically Hardened Solver:** Uses Gaussian elimination with partial pivoting plus residual checks before accepting the learned ratings.
- **Zero-Sum Ratings:** Produces ratings centered at `0.0`, where positive values indicate above-average aggregate performance and negative values indicate below-average aggregate performance.
- **Deterministic Ranking:** Orders candidate items by learned rating with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items, self-preferences, non-positive weights, and disconnected comparison graphs.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.
- **Shared Ranking Infrastructure:** Reuses the `ranking_protocols` helpers for dataset validation, diagnostics, export, and candidate ranking.

6.133.5 Scoring semantics

This implementation aggregates pairwise preferences into matchup totals and then solves the Massey system

```
M r = p
```

where $M_{ii} = \text{games}_i$, $M_{ij} = -\text{games}_{ij}$ for $i \neq j$, and the final row of M is replaced with ones so that the learned ratings satisfy `sum(ratings) = 0`. The right-hand-side vector p is the per-item signed point-differential total `wins_i - losses_i`.

The linear system is solved using deterministic Gaussian elimination with partial pivoting. The implementation also verifies that the recovered solution has a small residual and only clamps negligible floating-point noise around `0.0`.

The resulting ratings are relative rather than probabilistic: only their differences and ordering matter. Larger positive values indicate stronger aggregate pairwise performance against the field.

6.133.6 Usage

Learning a ranker

```
% Learn from a pairwise ranking dataset object
| ?- massey_ranker::learn(my_dataset, Ranker).
...

% Learn with an explicit empty options list
| ?- massey_ranker::learn(my_dataset, Ranker, []).
...
```

The current implementation accepts only the empty options list []. Any non-empty options list is rejected.

Inspecting diagnostics

```
% Inspect model and dataset summary metadata
| ?- massey_ranker::learn(my_dataset, Ranker),
    massey_ranker::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
...
```

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- massey_ranker::learn(my_dataset, Ranker),
    massey_ranker::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
...
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- massey_ranker::learn(my_dataset, Ranker),
    massey_ranker::export_to_clauses(my_dataset, Ranker, my_ranker, Clauses).
Clauses = [my_ranker(massey_ranker(...))]
...

% Export to a file
| ?- massey_ranker::learn(my_dataset, Ranker),
```

(continues on next page)

(continued from previous page)

```
massey_ranker::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.pl').  
...
```

6.133.7 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[  
  model(massey_ranker),  
  options(Options),  
  dataset_summary(DatasetSummary)  
]
```

6.133.8 Ranker representation

The learned ranker is represented by a compound term of the form:

```
massey_ranker(Items, Ratings, Diagnostics)
```

Where:

- `Items`: List of ranked items.
- `Ratings`: List of Item-Rating pairs.
- `Diagnostics`: List of metadata terms, including the effective options and dataset summary.

6.133.9 References

1. Massey, K. (1997). *Statistical models applied to the rating of sports teams*.

6.134 mcp_server

MCP (Model Context Protocol) server library for Logtalk applications. Makes any Logtalk application available as a local MCP server using stdio transport. Implements the MCP 2025-06-18 specification (tools, prompts, resources, elicitation, structured output, and resource links). Supports version negotiation.

The library uses the `json_rpc` library for JSON-RPC 2.0 message handling as currently required by the MCP specification.

6.134.1 API documentation

Open the `../apis/library_index.html#mcp_server` link in a web browser.

6.134.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(mcp_server(loader)).
```

6.134.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(mcp_server(tester)).
```

6.134.4 Usage

Implementing the tool protocol

To expose a Logtalk object as an MCP tool provider, implement the `mcp_tool_protocol` protocol:

```
:- object(my_tools,
    implements(mcp_tool_protocol)).

    :- public(factorial/2).
    :- mode(factorial(+integer, -integer), one).
    :- info(factorial/2, [
        comment is 'Computes the factorial of a non-negative integer.',
        argnames is ['N', 'F']
    ]).

    tools([
        tool(factorial, factorial, 2)
    ]).

    factorial(0, 1) :- !.
    factorial(N, F) :-
        N > 0,
        N1 is N - 1,
        factorial(N1, F1),
        F is N * F1.

:- end_object.
```

The `tools/1` predicate returns a list of `tool(Name, Functor, Arity)` descriptors that declare which predicates are exposed as MCP tools:

- Name — the MCP tool name (an atom)
- Functor — the predicate functor

- Arity — the predicate arity

Tool descriptions and parameter schemas are automatically derived from the `info/2` and `mode/2` directives. This must be accurate to allow correct derivation of input and output arguments and types.

A title key in the predicate's `info/2` directive provides a human-friendly display name for the tool. If omitted, the predicate functor is used as the title.

The supported predicates types and their corresponding JSON types are:

Logtalk type	JSON type
integer	integer
float	number
number	number
atom	string
boolean	boolean
list	array
list(_)	array
compound	object
nonvar	string
term	string
chars	string
codes	string
(other)	string

Auto-dispatch vs. custom `tool_call/3`

By default, when an MCP client calls a tool, the server auto-dispatches: it calls the corresponding predicate as a message to the application object, collects output-mode arguments (`-` and `--`), and returns them as text content.

For custom result formatting, implement `tool_call/3`:

```
tool_call(factorial, Arguments, Result) :-
    member('N'-N, Arguments),
    factorial(N, F),
    number_codes(F, Codes),
    atom_codes(FAtom, Codes),
    atom_concat('The factorial is: ', FAtom, Text),
    Result = text(Text).
```

The Result term can be:

- `text(Atom)` — a text result
- `error(Atom)` — a tool-level error (sets `isError: true`)
- `results(List)` — a list of content items (`text(Atom)`, `error(Atom)`, `resource_link(URI, Name)`, or `resource_link(URI, Name, Description, MimeType)`)
- `structured(StructuredContent)` — a structured result (the server auto-generates a text representation for the content field)
- `structured(Items, StructuredContent)` — a structured result with explicit content items for the content field

Output schemas (structured tool output)

Tools can declare an output schema by defining `output_schema/2` in their application object:

```
output_schema(divide, {type-object, properties-{quotient-{type-number}}, required-[quotient]}
→).
```

When an output schema is declared, the tool descriptor includes an `outputSchema` field. The tool's `tool_call/3` can then return `structured(StructuredContent)` or `structured(Items, StructuredContent)` results. The `StructuredContent` must be a curly-term matching the schema.

Starting the server for debugging

Start the MCP server from a Logtalk top-level or script:

```
| ?- mcp_server::start('my-server', my_tools).
```

With options:

```
| ?- mcp_server::start('my-server', my_tools, [version('2.0.0'), server_title('My Server')]).
```

There should either be no output or only a Prolog backend term input prompt. Other than that, any spurious output will break the connection between a MCP client and the MCP server.

MCP client configuration

To use the server with an MCP client (e.g., VSCode or Claude Desktop), configure it as a stdio server. Example `claude_desktop_config.json`:

```
{
  "mcpServers": {
    "my-server": {
      "command": "swilgt",
      "args": [
        "-q",
        "-g", "logtalk_load(my_mcp_server(loader))",
        "-t", "halt"
      ],
      "env": {
        "LOGTALKHOME": "/usr/local/share/logtalk",
        "LOGTALKUSER": "/Users/jdoe/logtalk"
      }
    }
  }
}
```

The `env` definition of the `LOGTALKHOME` and `LOGTALKUSER` environment variables may or may not be required (it's usually necessary on macOS). When required, replace the values above with the actual values on your Logtalk setup.

The actual arguments to the integration script (`swilgt` in the example above) depend on the Prolog backend. For example, XVM requires instead:

```
{
  "mcpServers": {
    "my-server": {
      "command": "xvmlgt",
      "args": [
        "-q",
        "-g", "logtalk_load(my_mcp_server(loader)), halt.",
      ],
      "env": {
        "LOGTALKHOME": "/usr/local/share/logtalk",
        "LOGTALKUSER": "/Users/jdoe/logtalk"
      }
    }
  }
}
```

Elicitation (interactive tools)

Tools that need to ask the user questions during execution can use MCP elicitation if the MCP client supports it (tested and working with VSCode Copilot). The application declares the elicitation capability and implements `tool_call/4` instead of `tool_call/3`:

```
:- object(interactive_tools,
    implements(mcp_tool_protocol)).

capabilities([elicitation]).

tools([
    tool(ask_name, ask_name, 0)
]).

:- public(ask_name/0).
:- info(ask_name/0, [
    comment is 'Asks the user for their name and greets them.'
]).

tool_call(ask_name, _Arguments, Elicit, Result) :-
    Schema = {
        type-object,
        properties-{name-{type-string}},
        required-[name]
    },
    call(Elicit, 'What is your name?', Schema, Answer),
    ( Answer = accept(Content),
      has_pair(Content, name, Name) ->
        atom_concat('Hello, ', Name, Greeting),
        atom_concat(Greeting, '!', Text),
        Result = text(Text)
      ; Result = text('No name provided.')
    ).

has_pair({Pairs}, Key, Value) :-
```

(continues on next page)

(continued from previous page)

```

    curly_member(Key-Value, Pairs).

    curly_member(Pair, (Pair, _)) :- !.
    curly_member(Pair, (_, Rest)) :-
        !, curly_member(Pair, Rest).
    curly_member(Pair, Pair).

:- end_object.

```

The Elicit closure is called as `call(Elicit, Message, Schema, Answer)` where:

- Message — an atom with the prompt text
- Schema — a curly-term JSON Schema for the requested input
- Answer — unified with `accept(Content)`, `decline`, or `cancel`

When `accept(Content)` is returned, `Content` is a curly-term with the user's response matching the requested schema.

See the `examples/birds_mcp/` example for a complete demonstration of elicitation with a bird identification expert system.

Prompts (prompt templates)

MCP prompts are templates for structured LLM interactions. They allow an application to expose reusable prompt templates that MCP clients can discover and use. To add prompts, implement `mcp_prompt_protocol` in addition to `mcp_tool_protocol`, and declare prompts in capabilities:

```

:- object(my_prompts,
    implements([mcp_tool_protocol, mcp_prompt_protocol])).

capabilities([prompts]).

tools([]).

prompts([
    prompt(code_review, 'Reviews code for potential issues', [
        argument(code, 'The code to review', true),
        argument(language, 'The programming language', false)
    ]),
    prompt(summarize, 'Summarizes a given text', [
        argument(text, 'The text to summarize', true)
    ])
]).

prompt_get(code_review, Arguments, Result) :-
    ( member(code-Code, Arguments) ->
        atom_concat('Please review the following code for potential issues:\n\n', Code, _
↪Text)
    ; Text = 'Please provide code to review.'
    ),
    Result = messages([message(user, text(Text))]).

```

(continues on next page)

(continued from previous page)

```

prompt_get(summarize, Arguments, Result) :-
  ( member(text-Text, Arguments) ->
    atom_concat('Please summarize the following text:\n\n', Text, PromptText)
  ; PromptText = 'Please provide text to summarize.'
  ),
  Result = messages([message(user, text(PromptText))]).

:- uses(list, [member/2]).

:- end_object.

```

The `prompt_get/1` predicate returns a list of prompt descriptors:

- `prompt(Name, Description, Arguments)` — without title
- `prompt(Name, Title, Description, Arguments)` — with title

Where:

- Name — the MCP prompt name (an atom)
- Title — a human-friendly display name (an atom, optional)
- Description — a human-readable description (an atom)
- Arguments — a list of `argument(ArgName, ArgDescription, Required)` terms where `Required` is `true` or `false`

The `prompt_get/3` predicate handles prompt get requests. Its result term can be:

- `messages(MessageList)` — a list of prompt messages
- `messages(Description, MessageList)` — a list of messages with a description

Each message in the list is a `message(Role, Content)` term where:

- Role — user or assistant
- Content — `text(Text)` where `Text` is an atom

Multi-turn prompts can return multiple messages:

```

prompt_get(debate, Arguments, Result) :-
  member(topic-Topic, Arguments),
  atom_concat('Let us debate: ', Topic, UserText),
  Result = messages([
    message(user, text(UserText)),
    message(assistant, text('I would be happy to debate that topic. What is your
↪position?'))
  ]).

```

Resources (data exposure)

MCP resources expose data and content from the application that MCP clients can access. To add resources, implement `mcp_resource_protocol` in addition to `mcp_tool_protocol`, and declare resources in capabilities:

```
:- object(my_resources,
    implements([mcp_tool_protocol, mcp_resource_protocol])).

    capabilities([resources]).

    tools([]).

    resources([
        resource('logtalk://my-app/config', config, 'Application configuration',
↪ 'application/json'),
        resource('logtalk://my-app/readme', readme, 'Application readme', 'text/plain')
    ]).

    resource_read('logtalk://my-app/config', _Arguments, Result) :-
        Result = contents([
            text_content('logtalk://my-app/config', 'application/json', '{"name": "my-app",
↪ "version": "1.0"}')
        ]).

    resource_read('logtalk://my-app/readme', _Arguments, Result) :-
        Result = contents([
            text_content('logtalk://my-app/readme', 'text/plain', 'Welcome to my application.
↪ ')
        ]).

:- end_object.
```

The `resources/1` predicate returns a list of resource descriptors:

- `resource(URI, Name, Description, MimeType)` — without title
- `resource(URI, Name, Title, Description, MimeType)` — with title

Where:

- `URI` — the resource identifier (an atom, typically a URI like `logtalk://my-app/data`)
- `Name` — a human-readable name (an atom)
- `Title` — a human-friendly display name (an atom, optional)
- `Description` — a human-readable description (an atom)
- `MimeType` — the MIME type of the resource content (an atom, e.g. `'text/plain'`, `'application/json'`)

The `resource_read/3` predicate handles resource read requests. Its result term must be `contents(ContentList)` where each content item is:

- `text_content(URI, MimeType, Text)` — for text resources
- `blob_content(URI, MimeType, Base64Data)` — for binary resources encoded as base64

A resource can return multiple content items:

```
resource_read('logtalk://my-app/logs', _Arguments, Result) :-  
    Result = contents([  
        text_content('logtalk://my-app/logs', 'text/plain', 'Log entry 1'),  
        text_content('logtalk://my-app/logs', 'text/plain', 'Log entry 2')  
    ]).
```

6.134.5 Error handling

- Predicate failures result in a tool-level error with `isError: true`.
- Predicate exceptions result in a tool-level error with the exception term serialized as the error text.
- Prompt execution failures result in a JSON-RPC error response.
- Resource read failures result in a JSON-RPC error response.

6.134.6 Protocol

The `mcp_tool_protocol` protocol defines the following predicates:

- `capabilities/1` — returns the list of additional capabilities needed by the application (e.g. `[elicitation]`, `[prompts]`, `[resources]`, or `[prompts, resources, elicitation]`); optional, defaults to `[]`
- `tools/1` — returns the list of tool descriptors
- `tool_call/3` — handles a tool call (optional; auto-dispatch is used when not defined)
- `tool_call/4` — handles a tool call with an elicitation closure (optional; requires `capabilities([elicitation])` or `capabilities(..., elicitation)`)
- `output_schema/2` — declares a JSON Schema for structured tool output (optional; when defined, the tool descriptor includes `outputSchema`)

The `mcp_prompt_protocol` protocol defines the following predicates:

- `prompts/1` — returns the list of prompt descriptors
- `prompt_get/3` — handles a prompt get request

The `mcp_resource_protocol` protocol defines the following predicates:

- `resources/1` — returns the list of resource descriptors
- `resource_read/3` — handles a resource read request

6.134.7 Supported MCP methods

Method	Type	Description
initialize	Request	MCP handshake
initialized	Notification	Client acknowledgment
ping	Request	Server liveness check
tools/list	Request	Lists available tools
tools/call	Request	Calls a tool
prompts/list	Request	Lists available prompts
prompts/get	Request	Gets a prompt with arguments
resources/list	Request	Lists available resources
resources/read	Request	Reads a resource by URI
elicitation/create	Request (= >)	Asks the client for user input

6.135 memcached

Portable Memcached client implementing the text (ASCII) protocol. This library uses the sockets library and supports all backend Prolog systems supported by that library: ECLiPSe, GNU Prolog, SICStus Prolog, SWI-Prolog, Trealla Prolog, and XVM.

6.135.1 API documentation

Open the ../apis/library_index.html#memcached link in a web browser.

6.135.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(memcached(loader)).
```

6.135.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(memcached(tester)).
```

Note: Tests require a running Memcached server on localhost port 11211 (the default). To install and start Memcached:

On macOS:

```
$ brew install memcached
$ memcached -d
```

On Linux (Debian/Ubuntu):

```
$ sudo apt-get install memcached
$ sudo systemctl start memcached
```

On Linux (RHEL/Fedora):

```
$ sudo dnf install memcached
$ sudo systemctl start memcached
```

6.135.4 Protocol Version

This library implements the Memcached text (ASCII) protocol as documented in the official specification:

<https://github.com/memcached/memcached/blob/master/doc/protocol.txt>

The text protocol is the standard protocol supported by all Memcached clients and servers. All commands use simple text lines terminated by `\r\n`.

6.135.5 Features

- Full storage command support: set, add, replace, append, prepend, cas
- Retrieval commands: get, gets (with CAS token), multi-get, gat, gats
- Key deletion with delete
- Atomic increment/decrement with incr/decr
- Item expiration management with touch, gat, gats
- Server management: flush_all, version, stats
- Graceful disconnect with quit
- Proper handling of all protocol response codes

6.135.6 Usage

Connecting to a Memcached server

Basic connection to localhost on the default port (11211):

```
?- memcached::connect(localhost, Connection).
```

Connection to a specific host and port:

```
?- memcached::connect('cache.example.com', 11211, Connection).
```

Storing data

Store a value unconditionally:

```
?- memcached::set(Connection, my_key, 'Hello, World!').
```

Store with flags and expiration time (in seconds):

```
?- memcached::set(Connection, my_key, 'Hello!', 0, 3600).
```

Store only if the key does not exist:

```
?- memcached::add(Connection, new_key, 'New value', 0, 60).
```

Store only if the key already exists:

```
?- memcached::replace(Connection, existing_key, 'Updated', 0, 60).
```

Append data to an existing value:

```
?- memcached::append(Connection, my_key, ' World!').
```

Prepend data to an existing value:

```
?- memcached::prepend(Connection, my_key, 'Hello ').
```

Retrieving data

Get a value by key:

```
?- memcached::get(Connection, my_key, Value).
```

Get a value with flags:

```
?- memcached::get(Connection, my_key, Value, Flags).
```

Get a value with CAS token (for optimistic locking):

```
?- memcached::gets(Connection, my_key, Value, CasUnique).
```

Get multiple keys at once:

```
?- memcached::mget(Connection, [key1, key2, key3], Items).
```

Check-and-Set (CAS)

CAS provides optimistic locking for cache updates:

```
?- memcached::gets(Connection, my_key, Value, CasUnique),
   % ... process value ...
   memcached::cas(Connection, my_key, 'New value', 0, 60, CasUnique).
```

If another client modifies the value between gets and cas, the cas command will throw `error(memcached_error(exists), _)`.

Deleting data

Delete an item by key:

```
?- memcached::delete(Connection, my_key).
```

Fails if the key does not exist.

Increment/Decrement

Atomically increment a numeric value:

```
?- memcached::set(Connection, counter, '100'),  
   memcached::incr(Connection, counter, 5, NewValue).  
% NewValue = 105
```

Atomically decrement a numeric value:

```
?- memcached::decr(Connection, counter, 3, NewValue).  
% NewValue = 102
```

Note: The item must already exist and contain a decimal representation of a 64-bit unsigned integer. Decrementing below 0 yields 0.

Touch and Get-and-Touch

Update expiration time without fetching:

```
?- memcached::touch(Connection, my_key, 3600).
```

Fetch and update expiration time in one operation:

```
?- memcached::gat(Connection, my_key, 3600, Value).
```

Fetch with CAS and update expiration:

```
?- memcached::gats(Connection, my_key, 3600, Value, CasUnique).
```

Server management

Invalidate all items:

```
?- memcached::flush_all(Connection).
```

Invalidate all items after a delay (in seconds):

```
?- memcached::flush_all(Connection, 30).
```

Get server version:

```
?- memcached::version(Connection, Version).
```

Get server statistics:

```
?- memcached::stats(Connection, Stats).
```

Get specific statistics:

```
?- memcached::stats(Connection, items, Stats).
```

Disconnecting

```
?- memcached::disconnect(Connection).
```

6.135.7 API Summary

Connection management

- `connect(+Host, +Port, -Connection)` - Connect to server
- `connect(+Host, -Connection)` - Connect to server on default port
- `disconnect(+Connection)` - Disconnect from server

Storage commands

- `set(+Connection, +Key, +Value, +Flags, +ExpTime)` - Store unconditionally
- `set(+Connection, +Key, +Value)` - Store with default flags and no expiration
- `add(+Connection, +Key, +Value, +Flags, +ExpTime)` - Store only if new
- `replace(+Connection, +Key, +Value, +Flags, +ExpTime)` - Store only if exists
- `append(+Connection, +Key, +Value)` - Append to existing value
- `prepend(+Connection, +Key, +Value)` - Prepend to existing value
- `cas(+Connection, +Key, +Value, +Flags, +ExpTime, +CasUnique)` - Check-and-set

Retrieval commands

- `get(+Connection, +Key, -Value)` - Get value
- `get(+Connection, +Key, -Value, -Flags)` - Get value and flags
- `gets(+Connection, +Key, -Value, -CasUnique)` - Get with CAS token
- `gets(+Connection, +Key, -Value, -Flags, -CasUnique)` - Get with flags and CAS
- `mget(+Connection, +Keys, -Items)` - Multi-get

Deletion

- `delete(+Connection, +Key)` - Delete item

Increment/Decrement

- `incr(+Connection, +Key, +Amount, -NewValue)` - Increment
- `decr(+Connection, +Key, +Amount, -NewValue)` - Decrement

Touch

- `touch(+Connection, +Key, +ExpTime)` - Update expiration
- `gat(+Connection, +Key, +ExpTime, -Value)` - Get and touch
- `gats(+Connection, +Key, +ExpTime, -Value, -CasUnique)` - Get-and-touch with CAS

Server commands

- `flush_all(+Connection)` - Invalidate all items
- `flush_all(+Connection, +Delay)` - Invalidate all items after delay
- `version(+Connection, -Version)` - Get server version
- `stats(+Connection, -Stats)` - Get general statistics
- `stats(+Connection, +Argument, -Stats)` - Get specific statistics

6.135.8 Error Handling

Most predicates throw `error(memcached_error(Reason), Context)` on failure. Common error reasons:

- `connection_failed` - Could not establish TCP connection
- `not_stored` - Storage condition not met (e.g. add on existing key)
- `exists` - CAS conflict (item was modified by another client)
- `not_found` - CAS on non-existent key
- `server_error(Message)` - Server returned an error

Retrieval predicates (`get`, `gets`, `gat`, `gats`) fail (rather than throw) when the requested key is not found. The `delete`, `incr`, `decr`, and `touch` predicates also fail when the key is not found.

6.135.9 Key Format

- Keys are atoms up to 250 characters
- Keys must not contain control characters or whitespace
- Keys are case-sensitive

6.135.10 Expiration Times

- 0 means the item never expires (may still be evicted by LRU)
- Values up to 2592000 (30 days) are treated as relative seconds
- Values over 2592000 are treated as absolute Unix timestamps
- Negative values cause immediate expiration

6.136 message_pack

The `message_pack` library provides predicates for parsing and generating data in the `MessagePack` format based on the current specification and reference documents found at:

- <https://github.com/msgpack/msgpack/blob/master/spec.md>
- <https://github.com/msgpack/msgpack/blob/master/spec-old.md>

It follows the representation conventions used by the existing binary and JSON libraries whenever practical.

6.136.1 API documentation

Open the `../apis/library_index.html#message_pack` link in a web browser.

6.136.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(message_pack(loader)).
```

6.136.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(message_pack(tester)).
```

6.136.4 Representation

The following choices of syntax have been made to represent `MessagePack` elements as terms:

- Maps are represented using curly-bracketed terms, `{Pairs}`, where each pair uses the representation `Key-Value`.
- Arrays are represented using lists.
- Strings can be represented as atoms, `chars(List)`, or `codes(List)`. Use the `message_pack/1` parameterized object with the parameter bound to `atom`, `chars`, or `codes` to select the representation used when decoding strings.
- `MessagePack` strings are validated using strict UTF-8 decoding and encoding. Invalid UTF-8 byte sequences and invalid Unicode scalar values are rejected.
- Binary data is represented by `bytes(List)`.

- Extension values are represented by `ext(Type, bytes(Bytes))`, where `Type` is the signed 8-bit extension type and `Bytes` is the extension payload.
- The MessagePack values `false`, `true`, and `nil` are represented by, respectively, the `@false`, `@true`, and `@null` compound terms.
- IEEE 754 infinities and NaN values are represented by the `@infinity`, `@negative_infinity`, and `@not_a_number` compound terms.
- Only some backends distinguish between positive zero and negative zero. The decoder produces the `0.0` and `-0.0` floats, but backends that normalize both values to `0.0` cannot preserve the sign of zero when re-encoding.

The following table exemplifies the term equivalents of common MessagePack elements when using `message_pack(atom)`:

MessagePack value	term
<code>nil</code>	<code>@null</code>
<code>false</code>	<code>@false</code>
<code>true</code>	<code>@true</code>
<code>"foo"</code>	<code>foo</code>
<code>[1,2,3]</code>	<code>[1,2,3]</code>
<code>{"a":1,"b":true}</code>	<code>{a-1, b-@true}</code>
bin payload	<code>bytes([...])</code>
extension payload	<code>ext(Type, bytes([...]))</code>

6.136.5 Notes

- Generated integers always use the smallest valid MessagePack integer format.
- Generated floating point values use the single-precision format when it can exactly represent the input value; otherwise, the double-precision format is used.
- Recommended string representation choice depends on the trust boundary of the input. For trusted input, `message_pack(atom)` is usually the most convenient choice when the rest of the application already works with atoms. For untrusted input, prefer `message_pack(codes)` or `message_pack(chars)` to avoid interning arbitrary incoming strings as atoms.
- For untrusted input, `message_pack(codes)` is generally the best default choice. It avoids atom creation while preserving the exact decoded Unicode code points. `message_pack(chars)` provides the same safety with a more readable but less compact representation.
- The predefined timestamp extension is currently exposed using the generic `ext(-1, bytes(Bytes))` representation.
- The library requires a backend Prolog compiler with unbounded integer support, matching the requirements of the existing `cbor`, `avro`, and `protobuf` libraries.

6.137 meta

This library provides implementations of common meta-predicates. The meta object implements common meta-predicates like `map/3` and `fold_left/4`.

See also the `meta_compiler` library, which provides optimized compilation of meta-predicate calls.

6.137.1 API documentation

Open the ../apis/library_index.html#meta link in a web browser.

6.137.2 Loading

To load the main entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(meta(loader)).
```

6.137.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(meta(tester)).
```

6.137.4 Usage

See e.g. the `metapredicates` example and unit tests.

6.138 meta_compiler

This library supports implementations optimized compilation of meta-calls for the predicates defined in the meta library.

6.138.1 API documentation

Open the ../apis/library_index.html#meta-compiler link in a web browser.

6.138.2 Loading

To load the main entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(meta_compiler(loader)).
```

6.138.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(meta_compiler(tester)).
```

6.138.4 Usage

If `meta_compiler` is the only hook object you are using, you can set it as the default hook object (but note that the optimizations are only applied to entities compiled with the `optimize` flag turned on):

```
| ?- set_logtalk_flag(hook, meta_compiler).  
...
```

Otherwise, use the `hook(meta_compiler)` and `optimize(on)` compiler options when compiling and loading the code that you want to optimize. For example:

```
| ?- logtalk_load(my_source_file, [hook(meta_compiler), optimize(on)]).  
...
```

See also the `metapredicates_compiled` example and unit tests.

6.138.5 Known issues

Goals in clauses containing parameter variables are currently not expanded by this library and thus not optimized.

6.139 mime_types

The `mime_types` library provides convenience predicates for mapping file extensions, file names, and URL-like resources to MIME media types and content encodings.

The current implementation ships with a built-in registry containing a practical set of common standard mappings plus a smaller set of de-facto common mappings. Additional mappings can be loaded at runtime from `mime.types-style` files.

Lookup is lenient by default. The convenience predicates also accept a boolean `Strict` argument to restrict lookup to strict runtime overlays plus the built-in standard registry.

The library uses the `os` library `decompose_file_name/4` predicate to split paths, and recognizes common compound suffixes such as `.tgz` and `.svgz`.

6.139.1 API documentation

Open the `../..apis/library_index.html#mime_types` link in a web browser.

6.139.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(mime_types(loader)).
```

6.139.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(mime_types(tester)).
```

6.139.4 Usage

Guess the MIME type and content encoding for a file name:

```
| ?- mime_types::guess_file_type('archive.tgz', Type, Encoding).
Type = 'application/x-tar',
Encoding = gzip
yes
```

Lookup a preferred extension for a MIME type:

```
| ?- mime_types::guess_extension('application/json', Extension).
Extension = '.json'
yes
```

Read a `mime.types`-style file without mutating the runtime registry:

```
| ?- mime_types::read_mime_types('mime.types', Pairs).
```

Load additional runtime mappings from a `mime.types`-style file:

```
| ?- mime_types::load('mime.types').
```

6.140 modified_z_score_anomaly_detector

Statistical modified Z-score anomaly detector for continuous datasets. It is a statistical anomaly-detection method based on the modified Z-score defined by Iglewicz and Hoaglin (1993): for each known continuous attribute value x , the detector computes $0.6745 * (x - \text{median}) / \text{mad}$, where `median` is the learned sample median and `mad` is the learned median absolute deviation, and then aggregates the per-attribute modified Z-scores using the selected `learn-time score_mode/1` option.

The library implements the `anomaly_detector_protocol` defined in the `anomaly_detection_protocols` library. It learns a detector from a continuous dataset, computes anomaly scores for new instances, predicts normal or anomaly, and exports learned detectors as clauses or files.

Datasets are represented as objects implementing the `anomaly_dataset_protocol` protocol from the `anomaly_detection_protocols` library. See the `anomaly_detection_protocols/test_datasets` directory for examples.

6.140.1 API documentation

Open the ../apis/library_index.html#modified-z-score link in a web browser.

6.140.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(modified_z_score_anomaly_detector(loader)).
```

6.140.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(modified_z_score_anomaly_detector(tester)).
```

6.140.4 Features

- **Statistical method:** implements robust anomaly detection based on the modified Z-score described by Iglewicz and Hoaglin (1993), using the sample median and median absolute deviation of each continuous attribute to measure how far new observations deviate from the training data distribution.
- **Classical per-attribute modified Z-score:** for each known attribute value x , the library computes the standard robust score $0.6745 * (x - \text{median}) / \text{mad}$, where `median` is the learned sample median for that attribute and `mad` is the learned median absolute deviation.
- **Continuous features only:** accepts datasets whose declared attributes are all continuous.
- **Robust statistics:** reuses the statistics library `sample` object to compute per-attribute medians and the median of absolute deviations.
- **Baseline training selection:** supports learn-time `baseline_class_values(ClassValues)` and `baseline_selection_policy(Policy)` options. The default baseline class values are `[normal]`. The default reject policy throws an error if non-baseline examples are present, while `filter` removes them before fitting.
- **Missing-value tolerant:** ignores missing values when fitting attribute statistics. During scoring, queries must provide at least one known value. In the default `score_mode(root_mean_square)`, the raw score is normalized by the number of known values so that scores remain comparable across different missing-value patterns.
- **Configurable scoring semantics:** supports both dense multivariate deviation scoring using `score_mode(root_mean_square)` and sparse anomaly detection using `score_mode(any_feature_extreme)`. The default root-mean-square mode reuses the numberlist library Euclidean norm predicate as part of the computation. The `score_mode/1` option only controls how the per-attribute modified Z-scores are aggregated into a single raw anomaly score.
- **Bounded scoring:** maps the raw multivariate modified Z-score to $[0.0, 1.0)$ using `Score = Raw / (1 + Raw)`.

- **Default threshold:** the default `anomaly_threshold(0.777777777777778)` corresponds to the classical raw modified Z-score cutoff 3.5 recommended by Iglewicz and Hoaglin (1993), while remaining overrideable in `learn/3` and `predict/4`.
- **Learn-time score mode:** `score_mode/1` is recorded in the learned detector and reused for subsequent scoring and prediction. Passing a `score_mode/1` option to `predict/4` does not override the learned mode.
- **All-missing queries rejected:** scoring and prediction throw a `domain_error(non_empty_known_values, AttributeNames)` exception when every declared feature is missing in the query.
- **Featureless datasets rejected:** datasets must declare at least one continuous feature; otherwise `learn/2-3` throws a `domain_error(non_empty_features, Dataset)` exception.
- **Detector export:** learned detectors can be exported as predicate clauses.
- **Explicit validation and diagnostics:** supports the shared `check_anomaly_detector/1`, `valid_anomaly_detector/1`, `diagnostics/2`, `diagnostic/2`, and `anomaly_detector_options/2` predicates.

6.140.5 Options

The following options are supported by the public API:

- `anomaly_threshold(Threshold)`: Threshold for `predict/3-4` (default: `0.777777777777778`)
- `baseline_class_values(ClassValues)`: Learn-time class labels that are admissible for baseline fitting (default: `[normal]`)
- `baseline_selection_policy(Policy)`: Learn-time handling of examples whose class is not listed in `baseline_class_values/1`. Supported values are `filter` and `reject` (default: `reject`)
- `score_mode(Mode)`: Learn-time score aggregation mode for `learn/3`. Supported values are `root_mean_square` and `any_feature_extreme` (default: `root_mean_square`). If passed to `predict/4`, it is ignored and the value stored in the learned detector is used.

6.140.6 Detector representation

The learned detector is represented by default as:

```
modified_z_score_detector(TrainingDataset, Encoders, Diagnostics)
```

Where:

- `TrainingDataset`: training dataset object identifier
- `Encoders`: list of `modified_zscore(Attribute, Median, Scale)` records
- `Diagnostics`: learned metadata terms including `model/1`, `training_dataset/1`, `attribute_names/1`, `feature_count/1`, `example_count/1`, and `options/1`

When exported using `export_to_clauses/4` or `export_to_file/4`, this detector term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.140.7 Notes

Scoring has three stages. First, the detector computes one classical per-attribute modified Z-score for each known attribute value using $0.6745 * (x - \text{median}) / \text{mad}$. Second, those per-attribute modified Z-scores are aggregated into a single raw anomaly score according to the learned `score_mode/1` option. Third, the raw score is mapped to the interval $[0.0, 1.0]$ using $\text{Score} = \text{Raw} / (1 + \text{Raw})$.

The `score_mode/1` option does not change the classical per-attribute formula. It only changes the aggregation step. With `score_mode(root_mean_square)`, the raw score is the root mean square of the per-attribute modified Z-scores. With `score_mode(any_feature_extreme)`, the raw score is the maximum absolute per-attribute modified Z-score.

The `baseline_class_values/1` option declares which dataset class labels are admissible for fitting the baseline medians and median absolute deviations. The `baseline_selection_policy/1` option then controls what happens when other labels are present in the training data. The default reject policy raises a `domain_error(baseline_only_training_data, Dataset)` exception when any non-baseline example is found. The filter policy removes non-baseline examples before fitting.

Attributes with zero observed median absolute deviation are assigned a fallback scale of 1.0 . This keeps the detector well-defined for singleton datasets or constant columns while still yielding zero score for matching values and positive scores for deviating values.

The root-mean-square aggregation keeps the default threshold stable as the number of observed dimensions grows and avoids penalizing partially observed queries solely for having fewer known attributes.

Use `score_mode(any_feature_extreme)` when a single extreme feature should be sufficient to flag an anomaly in high-dimensional data.

6.141 mutations

Experimental library. Should not be used in production code. Details can be changed without advance notice.

The mutations library provides support for generating random mutations of selected types. The library defines default mutation algorithms for the following basic types:

- atom
- integer
- float
- compound
- list

The user can add additional mutation algorithms for these or other types by defining objects or categories providing clauses for the `mutation/3` predicate and expanding the entity source files using the `mutations_store` object as the hook object.

This library is expected to eventually be used to support mutation-based *fuzz testing*.

By default, loading this library loads a set of default mutation algorithms. These can be overridden by defining alternative mutations and a custom loader file.

6.141.1 API documentation

Open the `../apis/library_index.html#mutations` link in a web browser.

6.141.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(mutations(loader)).
```

6.141.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(mutations(tester)).
```

6.141.4 Usage

The mutations category complements the type object and thus its predicates are accessed via this object. For example:

```
| ?- type::mutation(integer, 123, M).
M = 1293
yes

| type::mutation(integer, 123, M).
M = 5123
yes

| type::mutation(integer, 123, M).
M = -123
yes

| type::mutation(integer, 123, M).
M = 23
yes
```

When there are multiple mutation algorithms for a given type, the predicate `type::mutation/3` chooses one of them randomly. We can query the number of mutation algorithms available per type using the `mutations_store::counter/2` predicate:

```
| ?- mutations_store::counter(Type, Count).
Type = atom,
Count = 6 ;
Type = integer,
Count = 7 ;
...
```

Loading this library also loads the arbitrary library, which provides `get_seed/1` and `set_seed/1` predicates that can be used to control the pseudo-random number generator.

6.142 multisets

This library provides predicates for generating and querying multisets over lists. Multisets are unordered selections of a given length with element repetition allowed. The following categories of predicates are provided:

- **Generation operations** - Predicates for generating multisets.
- **Ordering variants** - Predicates that support an additional order argument (default or lexicographic) for controlling output order.
- **Distinct-value generation** - Predicates for generating multisets while deduplicating equal-valued results.
- **Indexed access** - Predicates for direct access to multisets at specific positions, including order-aware access.
- **Counting operations** - Predicates for counting multisets and distinct multisets.
- **Random selection** - Predicates for randomly selecting and sampling multisets and distinct multisets.
- **Lexicographic stepping** - Predicates for stepping to the next or previous multiset in lexicographic order.

Dedicated arrangements, cartesian_products, combinations, permutations, derangements, partitions, and subsequences libraries are also available for focused APIs on related operations.

6.142.1 API documentation

Open the ../apis/library_index.html#multisets link in a web browser.

6.142.2 Loading

To load all entities in this library, load the loader.lgt file:

```
| ?- logtalk_load(multisets(loader)).
```

6.142.3 Testing

To test this library predicates, load the tester.lgt file:

```
| ?- logtalk_load(multisets(tester)).
```

6.143 nanoid

This library generates random NanoID identifiers:

<https://github.com/ai/nanoid>

By default, identifiers are represented as atoms with 21 symbols and use the standard URL-safe alphabet:

```
_0123456789abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ
```

Custom size, alphabet, and representation (atoms, lists of characters, or lists of character codes) are supported using a parametric object.

The generation of random identifiers uses the `/dev/urandom` random number generator when available. This includes macOS, Linux, *BSD, and other POSIX operating systems. On Windows, a pseudo-random generator is used, randomized using the current wall time.

See also the `ids`, `uuid`, and `ulid` libraries.

6.143.1 API documentation

Open the ../apis/library_index.html#nanoid link in a web browser.

6.143.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(nanoid(loader)).
```

6.143.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(nanoid(tester)).
```

6.143.4 Usage

To generate an identifier using the default configuration:

```
| ?- nanoid::generate(NanoID).
NanoID = 'V1StGXR8_Z5jdHi6B-myT'
yes
```

To generate a 10-symbol identifier represented as a list of characters:

```
| ?- nanoid(chars, 10, 'abcdef012345')::generate(NanoID).
NanoID = ['4','b','d','f','a','0','3','2','1','e']
yes
```

To generate a 32-symbol identifier represented as a list of character codes:

```
| ?- nanoid(codes, 32, [0'a,0'b,0'c,0'd,0'e,0'f,0'0,0'1,0'2,0'3])::generate(NanoID).
NanoID = [49,97,99,100,48,102,98,101,50,51,48,101,99,102,49,50,98,100,49,97,48,51,101,98,100,
↪99,49,102,48,97,50,100]
yes
```

6.144 naive_bayes_classifier

Naive Bayes probabilistic classifier based on Bayes theorem with strong (naive) independence assumptions between features and supporting both categorical and continuous (Gaussian) features with Laplace smoothing.

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning a classifier from a dataset, using it to make predications, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. See `test_files` directory for examples.

6.144.1 API documentation

Open the ../docs/library_index.html#naive_bayes_classifier link in a web browser.

6.144.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(naive_bayes_classifier(loader)).
```

6.144.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(naive_bayes_classifier(tester)).
```

6.144.4 Features

- **Categorical Features:** Handles discrete-valued features with Laplace smoothing
- **Continuous Features:** Uses Gaussian (normal) distribution for numeric features
- **Classifier Export:** Learned classifiers can be exported as predicate clauses
- **Probability Estimation:** Provides both class predictions and probability distributions

6.144.5 Usage

Learning a Classifier

```
% Learn from a dataset object
| ?- naive_bayes_classifier::learn(my_dataset, Classifier).
...
```

Making Predictions

```
% Predict class for a new instance and the probability distribution
| ?- Instance = [...],
    naive_bayes_classifier::learn(my_dataset, Classifier),
    naive_bayes_classifier::predict(Classifier, Instance, PredictedClass),
    naive_bayes_classifier::predict_probability(Classifier, Instance, Probabilities).
PredictedClass = ...,
Probabilities = [...],
...
```

Exporting the Classifier

Learned classifiers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- naive_bayes_classifier::learn(my_dataset, Classifier),
    naive_bayes_classifier::export_to_clauses(Classifier, my_classifier, Clauses).
Clauses = [my_classifier(...)]
...

% Export to a file
| ?- naive_bayes_classifier::learn(my_dataset, Classifier),
    naive_bayes_classifier::export_to_file(Classifier, my_classifier, 'classifier.pl').
...
```

Using a learned classifier

Learned and saved classifiers can later be used for predictions without needing to access the original training dataset.

```
% Later, load the file and use the classifier
| ?- consult('classifier.pl'),
    my_classifier(Classifier),
    Instance = [...],
    naive_bayes_classifier::predict(Classifier, Instance, Class).
Class = ...
...
```

6.144.6 Classifier representation

The learned classifier is represented as a compound term:

```
nb_classifier(Classes, ClassPriors, AttributeNames, FeatureTypes, FeatureParams)
```

Where:

- **Classes:** List of class labels
- **ClassPriors:** List of Class-Prior probability pairs
- **AttributeNames:** List of attribute names in order

- FeatureTypes: List of types (categorical or continuous)
- FeatureParams: List of learned parameters for each feature

When exported using `export_to_clauses/4` or `export_to_file/4`, this classifier term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.144.7 References

1. Rish, I. (2001). “An empirical study of the naive Bayes classifier”.
2. Russell, S. & Norvig, P. (2020). “Artificial Intelligence: A Modern Approach”.
3. Mitchell, T. (1997). “Machine Learning”. Chapter 6: Bayesian Learning.

6.145 nearest_centroid_classifier

Nearest Centroid classifier. Assigns to an instance the class of the training samples whose mean (centroid) is closest to the instance.

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning a classifier from a dataset, using it to make predictions, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. See `test_files` directory for examples.

6.145.1 API documentation

Open the ../docs/library_index.html#nearest_centroid_classifier link in a web browser.

6.145.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(nearest_centroid_classifier(loader)).
```

6.145.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(nearest_centroid_classifier(tester)).
```

6.145.4 Features

- **Multiple Distance Metrics:** Euclidean, Manhattan, cosine
- **Mixed Features:** Automatically handles categorical and continuous features
- **Configurable Option:** distance metric via predicate options
- **Probability Estimation:** Provides confidence scores for predictions
- **Classifier Export:** Learned classifiers can be exported as predicate clauses

6.145.5 Usage

Learning a Classifier

```
% Learn from a dataset object with default options (euclidean distance)
| ?- nearest_centroid_classifier::learn(my_dataset, Classifier).
...

% Learn with custom options
| ?- nearest_centroid_classifier::learn(my_dataset, Classifier, [distance_
↪metric(manhattan)]).
...
```

Making Predictions

```
% Predict class for a new instance
| ?- Instance = [attr1-value1, attr2-value2, ...],
    nearest_centroid_classifier::learn(my_dataset, Classifier),
    nearest_centroid_classifier::predict(Classifier, Instance, PredictedClass).
PredictedClass = ...
...

% Predict with custom options
| ?- nearest_centroid_classifier::predict(Classifier, Instance, PredictedClass, [distance_
↪metric(cosine)]).
...

% Get probability distribution
| ?- nearest_centroid_classifier::predict_probabilities(Classifier, Instance, Probabilities).
Probabilities = [class1-0.67, class2-0.33]
...
```

Exporting the Classifier

Learned classifiers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- nearest_centroid_classifier::learn(my_dataset, Classifier),
    nearest_centroid_classifier::export_to_clauses(my_dataset, Classifier, my_classifier, _
    ↪Clauses).
Clauses = [my_classifier(...)]
...

% Export to a file
| ?- nearest_centroid_classifier::learn(my_dataset, Classifier),
    nearest_centroid_classifier::export_to_file(my_dataset, Classifier, my_classifier,
    ↪'classifier.pl').
...
```

Using a learned classifier

Learned and saved classifiers can later be used for predictions without needing to access the original training dataset.

```
% Later, load the file and use the classifier
| ?- consult('classifier.pl'),
    my_classifier(AttributeNames, FeatureTypes, Centroids),
    Instance = [...],
    nearest_centroid_classifier::predict(my_classifier(AttributeNames, FeatureTypes, _
    ↪Centroids), Instance, Class).
Class = ...
...
```

6.145.6 Options

The following options can be passed to the `predict/4` and `predict_probabilities/4` predicates:

- `distance_metric(Metric)`: Distance metric to use. Options: euclidean (default), manhattan, cosine

6.145.7 Classifier representation

The learned classifier is represented as a compound term:

```
nc_classifier(AttributeNames, FeatureTypes, Centroids)
```

Where:

- `AttributeNames`: List of attribute names in order
- `FeatureTypes`: List of types (numeric or categorical)
- `Centroids`: List of computed Class-Centroid pairs

When exported using `export_to_clauses/4` or `export_to_file/4`, this classifier term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.145.8 References

1. Manning, Raghavan & Schütze (2008) - "Introduction to Information Retrieval". Cambridge University Press.
2. Tibshirani, Hastie, Narasimhan & Chu (2002) - "Diagnosis of multiple cancer types by shrunken centroids of gene expression". Proceedings of the National Academy of Sciences, 99(10), 6567-6572.
3. Hastie, Tibshirani & Friedman (2009) - "The Elements of Statistical Learning: Data Mining, Inference, and Prediction" (2nd Edition). Springer.

6.146 nested_dictionaries

This library provides nested dictionary implementations based on private extensions to the dictionaries library objects. The representations of a nested dictionary should be regarded as opaque terms and only accessed using the library predicates.

This library is experimental, a work in progress, and future versions can introduce incompatible changes.

6.146.1 API documentation

Open the ../apis/library_index.html#nested-dictionaries link in a web browser.

6.146.2 Loading

To load all entities in this library, load the loader.lgt file:

```
| ?- logtalk_load(nested_dictionaries(loader)).
```

6.146.3 Testing

To test this library predicates, load the tester.lgt file:

```
| ?- logtalk_load(nested_dictionaries(tester)).
```

6.146.4 Usage

First, select the nested dictionary implementation that you want to use. For cases where the number of elements is relatively small and performance is not critical, nbintree can be a good choice. For other cases, navltree or nrbtree are likely better choices. If you want to compare the performance of the implementations, either define an object alias or use a uses/2 directive so that you can switch between implementations by simply changing the alias definition or the first argument of the directive. Note that you can switch between implementations at runtime without code changes by using a parameter variable in the first argument of a uses/2 directive.

To create an empty nested dictionary, you can use the new/1 predicate. For example:

```
| ?- navltree::new(Dictionary).
Dictionary = ...
yes
```

You can also create a new nested dictionary from a *curly bracketed term representation* (see below) by using the predicate `as_nested_dictionary/2`. For example:

```
| ?- navltree::as_nested_dictionary(
    {a-1, b-{c-3, d-{e-7,f-8}}},
    Dictionary
).
```

```
Dictionary = ...
yes
```

Several predicates are provided to insert, lookup, update, and delete key-value pairs given a list of keys interpreted as an *access path* to a nested dictionary. For example:

```
| ?- navltree::as_nested_dictionary(
    {a-1, b-{c-3, d-{e-7,f-8}}},
    Dictionary
),
    navltree::lookup_in([b,d,f], Value, Dictionary).
```

```
Dictionary = ...
Value = 8
yes
```

For details on these and other provided predicates, consult the library API documentation.

6.146.5 Curly term representation

To simplify importing and exporting data into a nested dictionary, the library provides `as_nested_dictionary/2` and `as_curly_bracketed/2` predicates that work with a *curly term representation*. This format is based on the JSON data interchange format.

A dictionary is represented by the `{Pairs}` term where `Pairs` is a conjunction of `Key-Value` or `Key:Value` pairs and `Value` can be a nested dictionary or lists of pairs. An empty dictionary is represented using the `{}` term.

the raw parsed term.

- `data/2` maps supported sentence types into typed semantic terms.

6.147 Raw Sentence Representation

Parsed sentences are represented as:

```
nmea_sentence(Talker, Type, Fields, checksum(Provided, Computed))
```

Where:

- `Talker` is a normalized lowercase talker atom such as `gp` or `gn`
- proprietary `$P...` sentences use the atom `proprietary`
- `Type` is the normalized lowercase sentence type atom
- `Fields` is the ordered list of field atoms after the sentence identifier
- `Provided` is either a normalized uppercase hexadecimal checksum atom or missing

- Computed is the normalized uppercase hexadecimal checksum atom computed from the sentence payload

6.148 Typed data/2 Results

The library currently provides semantic decoding for these sentence types:

- `gga(Time, Coordinate, fix(FixQuality, SatellitesUsed, HDOP), altitude(AntennaAltitudeMeters, GeoidSeparationMeters), dgps(DifferentialAgeSeconds, StationId))`
- `rmc(Time, Status, Coordinate, movement(SpeedKnots, CourseDegrees), Date, MagneticVariation, Mode)`
- `gsa(SelectionMode, FixType, SatelliteIds, dop(PDOP, HDOP, VDOP), SystemId)`
- `gsv(MessageCount, MessageNumber, SatellitesInView, Satellites)`
- `vtg(track(TrueCourseDegrees, MagneticCourseDegrees), speed(SpeedKnots, SpeedKph), Mode)`
- `gll(Coordinate, Time, Status, Mode)`

Shared semantic subterms include:

- `geographic(Latitude, Longitude)`
- `utc_time(Hour, Minute, Second, fraction(Numerator, Denominator))`
- `date(Year, Month, Day)`
- `magnetic_variation(Degrees, east|west)`
- `satellite(PRN, ElevationDegrees, AzimuthDegrees, SnrDbHz)`
- `missing`

6.149 Date Rule

Two-digit RMC years are expanded using a fixed GPS-era rule:

- 80..99 map to 1980..1999
- 00..79 map to 2000..2079

6.150 Current Scope

Implemented in this first version:

- raw parsing from text sources
- configurable checksum handling
- proprietary \$P... sentence passthrough
- semantic decoding for GGA, RMC, GSA, GSV, VTG, and GLL

Deliberately out of scope for this first version:

- sentence generation

- AIS ! sentences
- serial or socket transport helpers
- corrupted-stream resynchronization
- GSV multi-fragment aggregation
- typed proprietary sentence decoding

6.151 nmf_projection

This library implements Non-negative Matrix Factorization (NMF) for continuous datasets whose attribute values are all non-negative. It learns a non-negative basis using deterministic multiplicative updates and represents each transformed instance as a list of non-negative component weights.

6.151.1 API documentation

Open the ../apis/library_index.html#nmf_projection link in a web browser.

6.151.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(nmf_projection(loader)).
```

6.151.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(nmf_projection(tester)).
```

6.151.4 Features

- Implements deterministic multiplicative-update training for NMF.
- Requires all training and inference attribute values to be non-negative.
- Supports exporting learned reducers and reloading them for later use.
- Records diagnostics including convergence status, iteration count, final update delta, and reconstruction error.

6.151.5 Options

- `n_components/1`
Number of latent components to learn. The current implementation requires the requested count to not exceed the minimum of the sample count and the number of features. Larger requests raise `domain_error(component_count, Requested-Maximum)`.
- `center/1`
Always false. Centering is intentionally disabled because NMF requires non-negative features.
- `feature_scaling/1`
When true, each continuous attribute is divided by its maximum observed training value. When false, raw non-negative values are used.
- `maximum_iterations/1`
Maximum number of multiplicative-update iterations used during both training and per-instance coefficient inference.
- `tolerance/1`
Stop iteration when the maximum absolute update delta is less than or equal to this value.

6.151.6 Usage

Assuming a dataset object `parts_based_measurements` implementing the `dimension_reduction_dataset_protocol` protocol:

```
| ?- nmf_projection::learn(parts_based_measurements, DimensionReducer).
| ?- nmf_projection::learn(parts_based_measurements, DimensionReducer, [n_components(2),
↪ feature_scaling(true), maximum_iterations(250), tolerance(1.0e-7)]).
| ?- nmf_projection::learn(parts_based_measurements, DimensionReducer),
    nmf_projection::transform(DimensionReducer, [f1-3.0, f2-0.0, f3-1.5, f4-0.0],
↪ ReducedInstance).
| ?- nmf_projection::learn(parts_based_measurements, DimensionReducer, [n_components(2)]),
    nmf_projection::export_to_file(parts_based_measurements, DimensionReducer, reducer,
↪ 'nmf_reducer.pl').
| ?- logtalk_load('nmf_reducer.pl'),
    reducer(Reducer),
    nmf_projection::transform(Reducer, [f1-3.0, f2-0.0, f3-1.5, f4-0.0], ReducedInstance).
```

6.151.7 Dimension reducer representation

Learned NMF reducers use the representation:

```
nmf_reducer(Encoders, Components, Diagnostics)
```

where `Encoders` stores the continuous attribute encoders used to map instances into a non-negative feature vector, `Components` stores the learned non-negative basis vectors, and `Diagnostics` stores metadata about the learned reducer.

6.151.8 References

- Daniel D. Lee and H. Sebastian Seung. *Algorithms for Non-negative Matrix Factorization*. Advances in Neural Information Processing Systems 13, 2001.

6.152 optics_clusterer

OPTICS clusterer. It uses deterministic OPTICS ordering with epsilon-based cluster extraction for the fixed clusterer protocol. Supports continuous attributes only.

The library implements the `clusterer_protocol` defined in the `clustering_protocols` library. It provides predicates for learning a clusterer from a dataset, assigning new instances to clusters, and exporting the learned clusterer as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `clustering_dataset_protocol` protocol from the `clustering_protocols` library.

6.152.1 API documentation

Open the ../apis/library_index.html#optics_clusterer link in a web browser.

6.152.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(optics_clusterer(loader)).
```

6.152.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(optics_clusterer(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(optics_clusterer(tester_performance)).
```

6.152.4 Features

- **OPTICS Ordering:** Learns a deterministic ordering using density-based reachability over continuous datasets.
- **Adaptive Neighborhood Indexing:** Uses a low-dimensional epsilon-grid index when it is likely to be cheaper and otherwise falls back to a deterministic metric tree for neighborhood search during ordering construction. The search backend can also be selected explicitly.
- **Continuous Datasets:** Accepts datasets containing only continuous attributes.
- **Distance Metrics:** Supports Euclidean and Manhattan distances.
- **Optional Feature Scaling:** Continuous attributes can be standardized using z-score scaling.
- **Epsilon-Based Extraction:** Extracts clusters from the ordering using a configurable extraction epsilon threshold.
- **Noise Detection:** New instances not reachable from an extracted core cluster within the extraction threshold are assigned to noise.
- **Prediction Pruning:** Classification reuses per-cluster core-point bounds to prune clusters that cannot beat the current best reachable match.
- **Portable Export:** Learned clusterers can be exported as clauses or files and reused later.

6.152.5 Options

The following options can be passed to the `learn/3` predicate:

- `ordering_and_extraction_epsilons(MaximumOrderingEpsilon, ExtractionEpsilon)`: Pair of epsilon thresholds where `MaximumOrderingEpsilon` is the neighborhood radius used while constructing the OPTICS ordering and `ExtractionEpsilon` is the threshold used when extracting clusters from the learned ordering and when classifying new instances. Default is `ordering_and_extraction_epsilons(1.0, 1.0)`. `ExtractionEpsilon` must not be greater than `MaximumOrderingEpsilon`.
- `search_index(SearchIndex)`: Search backend selection used while constructing the OPTICS ordering. Options are `auto` (default), `grid`, and `metric_tree`.
- `minimum_points(MinimumPoints)`: Minimum neighborhood size required for a point to be considered a core point. Default is 2.
- `distance_metric(Metric)`: Distance metric to use. Options: `euclidean` (default) or `manhattan`.
- `feature_scaling(FeatureScaling)`: Whether to standardize continuous attributes before clustering. Options: `on` (default) or `off`.

6.152.6 Clusterer representation

The learned clusterer is represented as a compound term with the functor chosen by the user when exporting the clusterer and arity 5. For example:

```
optics_clusterer(Encoders, Ordering, Clusters, Noise, Options)
```

Where:

- `Encoders`: List of continuous attribute encoders storing attribute name, mean, and scale.
- `Ordering`: List of ordered points annotated with reachability and core-distance information.

- Clusters: List of extracted clusters in cluster-id order.
- Noise: List of extracted noise points.
- Options: Effective training options used to learn the clusterer.

6.153 optionals

This library provides an implementation of *optional terms* with an API modeled after the Java 8 `Optional` class (originally due to requests by users working in Logtalk/Java hybrid applications). An optional term is an opaque compound term that may or may not hold a value. Optional terms avoid forcing the user to define a representation for the absence of a value by providing an API with predicates that depend on the presence or absence of a value. Optional terms also allow separating the code that constructs optional terms from the code that processes them, which is then free to deal if necessary and at its convenience with any case where the values held by optional terms are not present.

6.153.1 API documentation

Open the `../apis/library_index.html#optionals` link in a web browser.

6.153.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(optionals(loader)).
```

6.153.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(optionals(tester)).
```

6.153.4 Usage

The optional object provides constructors for optional terms. For example:

```
| ?- optional::of(1, Optional).  
...
```

The created optional terms can then be passed as parameters to the `optional/1` parametric object. For example:

```
| ?- optional::of(1, Optional), optional(Optional)::or_else(Term, 0).  
Optional = optional(1),  
Term = 1  
yes  
  
| ?- optional::empty(Optional), optional(Optional)::or_else(Term, 0).  
Optional = empty,
```

(continues on next page)

(continued from previous page)

```
Term = 0
yes
```

The maybe object provides types and predicates for type-checking of the term held by optional terms. It also provides some predicates for handling lists of optional terms, including `sequence/2` and `traverse/3`.

The `optional/1` parametric object also provides `map_or_else/3` for applying a closure to the value held by the optional term if not empty, returning a default value otherwise:

```
| ?- optional::of(a, Optional),
    optional(Optional)::map_or_else(char_code, 0, Value).
Value = 97
yes

| ?- optional::empty(Optional),
    optional(Optional)::map_or_else(char_code, 0, Value).
Value = 0
yes
```

The `zip/3` predicate combines two optional terms using a closure when both are not empty:

```
| ?- optional::of(1, O1), optional::of(3, O2),
    optional(O1)::zip([X,Y,Z]>>(Z is X+Y), O2, NewOptional).
NewOptional = optional(4)
yes

| ?- optional::of(1, O1), optional::empty(O2),
    optional(O1)::zip([X,Y,Z]>>(Z is X+Y), O2, NewOptional).
NewOptional = empty
yes
```

The `flatten/1` predicate unwraps a nested optional term:

```
| ?- optional::of(1, Inner), optional::of(Inner, Outer),
    optional(Outer)::flatten(NewOptional).
NewOptional = optional(1)
yes

| ?- optional::empty(Inner), optional::of(Inner, Outer),
    optional(Outer)::flatten(NewOptional).
NewOptional = empty
yes
```

The `to_expected/2` predicate converts an optional to an expected term:

```
| ?- optional::of(1, Optional),
    optional(Optional)::to_expected(missing, Expected).
Expected = expected(1)
yes

| ?- optional::empty(Optional),
    optional(Optional)::to_expected(missing, Expected).
Expected = unexpected(missing)
yes
```

The `from_goal/3` and `from_goal/2` constructors silently convert goal exceptions to empty optional terms. Use `from_goal_or_throw/3` or `from_goal_or_throw/2` if exceptions should be propagated instead:

```
| ?- optional::from_goal_or_throw(Y is 1+2, Y, Optional).
Optional = optional(3)
yes

| ?- optional::from_goal_or_throw(2 is 3, _, Optional).
Optional = empty
yes

| ?- catch(
    optional::from_goal_or_throw(Y is _, Y, _),
    Error,
    true
).
Error = error(instantiation_error, ...)
yes
```

Examples:

```
| ?- optional::of(1, 01), optional::of(2, 02),
    maybe::sequence([01, 02], Optional).
Optional = optional([1,2])
yes

| ?- maybe::traverse({optional}/[X,0]>>optional::of(X, 0), [1,2], Optional).
Optional = optional([1,2])
yes

| ?- maybe::traverse({optional}/[X,0]>>(
    integer(X) -> optional::of(X, 0)
; optional::empty(0)
), [1,a,2], Optional).
Optional = empty
yes

| ?- optional::of(1, 01), optional::empty(02),
    maybe::sequence([01, 02], Optional).
Optional = empty
yes
```

6.153.5 See also

The `expecteds` and `validations` libraries.

6.154 options

This library provides useful predicates for managing developer tool and application options.

6.154.1 API documentation

Open the `../apis/library_index.html#options` file in a web browser.

6.154.2 Loading

To load all entities in this library, load the `loader.lgt` utility file:

```
| ?- logtalk_load(options(loader)).
```

6.154.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(options(tester)).
```

6.154.4 Usage

The `options` category is usually imported by the root object of the developer tool or application. The importing object should define the `default_option/1` predicate and, if option type-checking is required, the `valid_option/1` predicate must be defined for each option. This library requires options to be represented by compound terms where the functor is the option name (e.g., `trim(true)` or `box(0,2)`). The `option/2-3` can be used to get or test an option given a list of options. When an option appears multiple times in a list, the `option/2-3` predicates get or test the first (leftmost) occurrence.

The library also supports a user-defined `fix_option/2` predicate. A usage example is when an option value can be a relative file path that should be expanded before use. Another usage example would be converting from a user-friendly option to a form more suitable for internal processing. When a call to the `fix_option/2` predicate fails, the option is used as-is.

A simple example:

```
:- object(foo,
    imports(options)).

:- uses(type, [
    valid/2
]).

:- public(p/0).
p :-
```

(continues on next page)

(continued from previous page)

```

    % use default options
    p([]).

:- public(p/1).
p(UserOptions) :-
    ^^check_options(UserOptions),
    % construct the full set of options from
    % the user options and the default options
    ^^merge_options(UserOptions, Options),
    ...
    % query an option
    ^^option(baz(Boolean), Options),
    q(Boolean),
    ...

default_option(baz(true)).
...

valid_option(baz(Boolean)) :-
    valid(boolean, Boolean).
...

:- end_object.

```

Note that you can use protected or private import of the options category if you don't want to add its public predicates to the object protocol.

6.155 os

This library provides a *portable* operating-system interface for the supported backend Prolog compilers.

The `os_types` category defines some useful operating-system types for type-checking when using with the type library object.

6.155.1 API documentation

Open the `../..apis/library_index.html#os` link in a web browser.

6.155.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(os(loader)).
```

6.155.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(os(tester)).
```

6.155.4 Known issues

Some predicates may only be supported by a subset of backend Prolog compilers on a subset of operating-systems. They should be used with care and fully tested in your application domain, as some backend Prolog compilers have buggy and inconsistent interfaces, notably across operating-systems. See the remarks section in the `os` object documentation for details.

6.156 partitions

This library provides predicates for generating and querying set partitions of lists. A set partition divides a list into non-empty, non-overlapping blocks whose union is the original list. The following categories of predicates are provided:

- **Generation operations** - Predicates for generating all set partitions and set partitions with an exact number of blocks.
- **Ordering variants** - Predicates that support an additional order argument (default or lexicographic) for controlling output order.
- **Distinct-value generation** - Predicates for generating set partitions while deduplicating equal-valued partitions after canonicalizing block order.
- **Indexed access** - Predicates for retrieving partitions and distinct partitions by zero-based index and for recovering the index of a partition.
- **Random selection** - Predicates for selecting random partitions and taking random samples, with distinct-value and exact-block-count variants.
- **Lexicographic stepping** - Predicates for moving to the next or previous distinct partition value in lexicographic order.
- **Counting operations** - Predicates for counting set partitions and distinct set partitions, including exact-block counts.

Base predicates are position-sensitive, matching the existing combinatorics libraries. Inputs with repeated values can therefore yield duplicate-valued partitions. Distinct predicates collapse those equal-valued results.

The lexicographic stepping predicates operate on the distinct partition view, using the same canonical block ordering as the distinct generation predicates.

The counting predicates are connected to Bell numbers (for all partitions of an N -element list) and Stirling numbers of the second kind (for partitions with an exact number of blocks).

Dedicated arrangements, cartesian_products, combinations, derangements, multisets, permutations, and subsequences libraries are also available for focused APIs on related operations.

6.156.1 API documentation

Open the ../apis/library_index.html#partitions link in a web browser.

6.156.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(partitions(loader)).
```

6.156.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(partitions(tester)).
```

6.157 pattern_mining_protocols

This library provides the generic core entities used in the implementation of machine learning pattern-finding algorithms. Pattern-mining algorithm implementations typically import the `pattern_miner_common` category, which implements the `pattern_miner_protocol` protocol and provides shared defaults, option handling, diagnostics accessors, and export helpers.

The family-specific dataset protocols and bundled smoke-test datasets are provided by the `frequent_pattern_mining_protocols` and `sequential_pattern_mining_protocols` support libraries.

6.157.1 API documentation

Open the ../apis/library_index.html#pattern_mining_protocols link in a web browser.

6.157.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(pattern_mining_protocols(loader)).
```

6.157.3 Testing

To run the library smoke tests, load the `tester.lgt` file:

```
| ?- logtalk_load(pattern_mining_protocols(tester)).
```

6.157.4 Common options

The `pattern_miner_common` category supports the following options used by importing pattern miners to control support thresholds and pattern length filtering:

- `minimum_support(0.5)` sets the relative minimum support threshold as a proportion in the interval `]0.0, 1.0]`.
- `minimum_support_count(N)` sets the absolute minimum support count. If both support options are provided, this option takes precedence.
- `maximum_pattern_length(1000)` sets the maximum pattern length to mine. The effective value is capped by the longest transaction or sequence in the dataset.
- `minimum_pattern_length(1)` sets the minimum pattern length retained in the mined result.

The current `apriori_pattern_miner`, `eclat_pattern_miner`, `fp_growth_pattern_miner`, and `prefix_span_pattern_miner` libraries all use these shared defaults.

6.157.5 Diagnostics

The `pattern_miner_protocol` protocol also defines the `diagnostics/2`, `diagnostic/2`, and `pattern_miner_options/2` predicates. These expose representation-independent metadata about mined results.

All pattern miners now provide at least the following generic diagnostics terms:

- `model(Model)`
- `options(Options)`
- `item_domain_size(Size)`
- `pattern_count(Count)`
- `pattern_length_histogram(Histogram)`
- `support_range(MinimumSupport, MaximumSupport)`

Each miner also provides algorithm-specific diagnostics terms describing its mining strategy and support representation, such as candidate-generation style, projected-database growth, closure filtering, or vertical support layout.

6.157.6 Related support libraries

- `frequent_pattern_mining_protocols`: Provides the `transaction_dataset_protocol` protocol plus bundled transaction smoke-test datasets for frequent itemset miners.
- `sequential_pattern_mining_protocols`: Provides the `sequence_dataset_protocol` protocol plus bundled sequence smoke-test datasets for sequential pattern miners.

6.158 pca_projection

Principal Component Analysis reducer for continuous datasets. The library implements the `dimension_reducer_protocol` defined in the `dimension_reduction_protocols` library and learns a linear projection by centering the training data, optionally standardizing continuous attributes, computing the covariance matrix, and extracting principal components using deterministic power iteration with deflation.

6.158.1 API documentation

Open the ../apis/library_index.html#pca_projection link in a web browser.

6.158.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(pca_projection(loader)).
```

6.158.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(pca_projection(tester)).
```

6.158.4 Features

- **Continuous Datasets:** Accepts datasets containing only continuous attributes.
- **Centering and Optional Scaling:** Centers all attributes and optionally standardizes them before extracting principal directions.
- **Portable Eigensolver:** Uses deterministic power iteration with deflation instead of backend-specific linear algebra libraries.
- **Projection API:** Transforms a new instance into a list of `component_N-Value` pairs.
- **Model Export:** Learned reducers can be exported as predicate clauses or written to a file.
- **Missing Values:** Missing or nonnumeric values are rejected.

6.158.5 Options

The `learn/3` predicate accepts the following options:

- `n_components/1`: Number of principal components to extract. Requests that exceed the number of features raise `domain_error(component_count, Requested-Maximum)`. The default is 2.
- `feature_scaling/1`: Whether to standardize continuous attributes before extracting components. Options: `true` (default) or `false`.
- `maximum_iterations/1`: Maximum number of power-iteration steps used when estimating each principal direction. The default is 1000.

- `tolerance/1`: Positive convergence tolerance used both for power-iteration stopping and for deciding when deflated eigenvalues are negligible. The default is `1.0e-8`.

6.158.6 Usage

The following examples use the sample datasets shipped with the `dimension_reduction_protocols` library:

```
| ?- logtalk_load(dimension_reduction_protocols('test_datasets/correlated_plane')),
    logtalk_load(dimension_reduction_protocols('test_datasets/high_dimensional_measurements
↪')).
```

Learning a reducer

```
| ?- pca_projection::learn(correlated_plane, DimensionReducer).

| ?- pca_projection::learn(correlated_plane, DimensionReducer, [n_components(1), feature_
↪scaling(false), maximum_iterations(200), tolerance(1.0e-7)]).
```

Transforming new instances

```
| ?- pca_projection::learn(high_dimensional_measurements, DimensionReducer),
    pca_projection::transform(DimensionReducer, [f1-0.9, f2-1.1, f3-1.0, f4-2.0, f5-2.2, f6-
↪2.1], ReducedInstance).

| ?- pca_projection::learn(correlated_plane, DimensionReducer, [n_components(1)]),
    pca_projection::transform(DimensionReducer, [x-1.0, y-2.0, z-3.0], ReducedInstance).
```

Exporting and reusing the reducer

```
| ?- pca_projection::learn(correlated_plane, DimensionReducer, [n_components(1)]),
    pca_projection::export_to_file(correlated_plane, DimensionReducer, reducer, 'pca_
↪reducer.pl').

| ?- logtalk_load('pca_reducer.pl'),
    reducer(Reducer),
    pca_projection::transform(Reducer, [x-1.0, y-2.0, z-3.0], ReducedInstance).
```

6.158.7 Dimension reducer representation

The learned dimension reducer is represented by a compound term with the functor chosen by the implementation and arity 4. For example:

```
pca_reducer(Encoders, Components, ExplainedVariances, Diagnostics)
```

Where:

- `Encoders`: List of continuous attribute encoders storing attribute name, mean, and scale.
- `Components`: List of principal direction vectors in descending variance order.

- **ExplainedVariances:** List of eigenvalues matching the extracted components.
- **Diagnostics:** Learned reducer metadata including the effective training options and learned model details.

When exported using `export_to_clauses/4` or `export_to_file/4`, this reducer term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.158.8 References

1. Pearson, K. (1901) - “On lines and planes of closest fit to systems of points in space”.
2. Hotelling, H. (1933) - “Analysis of a complex of statistical variables into principal components”.

6.159 permutations

This library provides predicates for generating and querying permutations over lists. The following categories of predicates are provided:

- **Generation operations** - Predicates for generating permutations and k-permutations.
- **Ordering variants** - Predicates that support an additional order argument (default, lexicographic, or shortlex) for controlling output order.
- **Distinct-value generation** - Predicates for generating permutations while deduplicating equal-valued results.
- **Lexicographic stepping** - Predicates for navigating permutations in lexicographic order.
- **Indexed access** - Predicates for direct access to permutations at specific positions, including distinct permutations.
- **Counting operations** - Predicates for counting permutations and distinct permutations.
- **Random selection** - Predicates for randomly selecting and sampling permutations and distinct permutations.

Dedicated `arrangements`, `cartesian_products`, `combinations`, `derangements`, `multisets`, `partitions`, and `subsequences` libraries are also available for focused APIs on related operations. The `arrangements` library is the repetition-allowed counterpart to this library.

6.159.1 API documentation

Open the ../apis/library_index.html#permutations link in a web browser.

6.159.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(permutations(loader)).
```

6.159.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(permutations(tester)).
```

6.160 plackett_luce_ranker

Tie-aware Plackett-Luce grouped-ranking ranker. It processes each group as a sequence of top-choice selections from highest relevance to lowest relevance, using grouped tie blocks and a deterministic fixed-point update on positive item strengths.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from grouped rankings, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The training dataset must declare each group once, use only declared groups and items in relevance judgments, assign non-negative integer relevance values, and induce a strongly connected directed strict-order graph across groups so that a finite Plackett-Luce maximum-likelihood estimate exists.

6.160.1 API documentation

Open the ../apis/library_index.html#plackett_luce_ranker link in a web browser.

6.160.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(plackett_luce_ranker(loader)).
```

6.160.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(plackett_luce_ranker(tester)).
```

6.160.4 Features

- **Grouped Top-Choice Learning:** Learns positive item strengths from grouped rankings by processing each group from highest relevance to lowest relevance.
- **Tie-Aware Likelihood:** Uses grouped tie blocks so equal relevance judgments are handled as unordered top-choice blocks instead of being broken arbitrarily. Each tie block contributes a size-constrained choice likelihood term against the remaining lower-relevance items.
- **Deterministic Ranking:** Orders candidate items by learned strength with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects malformed grouped datasets, unsupported options, duplicate candidates, and invalid ranker terms.
- **Regular MLE Fidelity:** Rejects grouped datasets whose strict-order graph does not admit a finite Plackett-Luce maximum-likelihood estimate instead of masking non-identifiability with implicit regularization.
- **Missing relevance semantics:** Missing relevance facts are treated as zero by default using the `missing_relevance(zero)` option and can be rejected using the `missing_relevance(error)` option.
- **Training Diagnostics:** Learned rankers include convergence, iteration, final update delta, and dataset summary metadata accessible using the `diagnostics/2` predicate.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.
- **Shared Grouped Infrastructure:** Reuses the shared grouped tie-block representation and iterative positive-strength scaffolding from the `ranking_protocols` library.

6.160.5 Dataset requirements

This implementation requires more than grouped-dataset well-formedness. In order to admit a finite Plackett-Luce maximum-likelihood estimate, the directed strict-order graph induced by the grouped rankings must be strongly connected. Intuitively, no partition of the items may dominate all others in only one direction across the observed groups.

Unlike `borda_ranker`, this model therefore rejects grouped datasets that consist of disconnected query universes or one-way dominance chains, because those data do not identify a finite global strength scale.

6.160.6 Usage

Learning a ranker

```
% Learn from a grouped ranking dataset object
| ?- plackett_luce_ranker::learn(my_dataset, Ranker).
...

% Learn with custom iteration and missing-relevance options
| ?- plackett_luce_ranker::learn(my_dataset, Ranker, [maximum_iterations(500), tolerance(1.
→0e-7), missing_relevance(error)]).
...
```

Inspecting diagnostics

```
% Inspect convergence and dataset summary metadata
| ?- plackett_luce_ranker::learn(my_dataset, Ranker),
    plackett_luce_ranker::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
```

6.160.7 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(plackett_luce_ranker),
  options(Options),
  convergence(Status),
  iterations(Iterations),
  final_delta(FinalDelta),
  dataset_summary(DatasetSummary)
]
```

Where:

- `model(plackett_luce_ranker)` identifies the learning algorithm that produced the ranker.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `convergence(Status)` records the training stop condition. The current values are converged and maximum_iterations_exhausted.
- `iterations(Iterations)` stores the number of update iterations that were executed.
- `final_delta(FinalDelta)` stores the maximum absolute strength update in the last iteration.
- `dataset_summary(DatasetSummary)` stores a summary list describing the validated training dataset.

Use the `ranking_protocols` `diagnostic/2` and `ranker_options/2` helper predicates when you only need a single metadata term or the effective options.

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- plackett_luce_ranker::learn(my_dataset, Ranker),
    plackett_luce_ranker::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- plackett_luce_ranker::learn(my_dataset, Ranker),
    plackett_luce_ranker::export_to_clauses(my_dataset, Ranker, my_ranker, Clauses).
Clauses = [my_ranker(plackett_luce_ranker(...))]
...

% Export to a file
| ?- plackett_luce_ranker::learn(my_dataset, Ranker),
    plackett_luce_ranker::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.pl').
...
```

6.160.8 Options

The following options can be passed to the learn/3 predicate:

- `maximum_iterations(MaximumIterations)`: Positive integer iteration bound.
- `tolerance(Tolerance)`: Positive convergence tolerance.
- `missing_relevance(zero|error)`: Policy used when a declared item in a group has no explicit relevance judgment.

6.160.9 Ranker representation

The learned ranker is represented by a compound term of the form:

```
plackett_luce_ranker(Items, Strengths, Diagnostics)
```

Where:

- `Items`: List of ranked items.
- `Scores`: List of normalized Item-Strength pairs.
- `Diagnostics`: List of metadata terms, including the effective options, convergence status, iteration count, final update delta, and dataset summary.

When exported using `export_to_clauses/4` or `export_to_file/4`, this ranker term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.160.10 See also

For the complementary grouped last-choice variant over the same dataset protocol, see the `plackett_luce_last_ranker` library.

6.161 plackett_luce_last_ranker

Tie-aware Plackett-Luce-last grouped-ranking ranker. It processes each group as a sequence of last-choice eliminations from lowest relevance to highest relevance, using grouped tie blocks and a deterministic fixed-point update on positive item strengths.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from grouped rankings, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The training dataset must declare each group once, use only declared groups and items in relevance judgments, assign non-negative integer relevance values, and induce a strongly connected directed strict-order graph across groups so that a finite Plackett-Luce-last maximum-likelihood estimate exists.

6.161.1 API documentation

Open the `../apis/library_index.html#plackett_luce_last_ranker` link in a web browser.

6.161.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(plackett_luce_last_ranker(loader)).
```

6.161.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(plackett_luce_last_ranker(tester)).
```

6.161.4 Features

- **Grouped Last-Choice Learning:** Learns positive item strengths from grouped rankings by processing each group from lowest relevance to highest relevance.
- **Tie-Aware Likelihood:** Uses grouped tie blocks so equal relevance judgments are handled as unordered last-choice blocks instead of being broken arbitrarily. Each tie block contributes a size-constrained last-choice likelihood term against the remaining higher-relevance items.
- **Deterministic Ranking:** Orders candidate items by learned strength with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects malformed grouped datasets, unsupported options, duplicate candidates, and invalid ranker terms.
- **Regular MLE Fidelity:** Rejects grouped datasets whose strict-order graph does not admit a finite Plackett-Luce-last maximum-likelihood estimate instead of masking non-identifiability with implicit regularization.
- **Missing relevance semantics:** Missing relevance facts are treated as zero by default using the `missing_relevance(zero)` option and can be rejected using `missing_relevance(error)`.

- **Training Diagnostics:** Learned rankers include convergence, iteration, final update delta, and dataset summary metadata accessible using the `diagnostics/2` predicate.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.
- **Shared Grouped Infrastructure:** Reuses the shared grouped tie-block representation and iterative positive-strength scaffolding from the `ranking_protocols` library.

6.161.5 Dataset requirements

This implementation requires more than grouped-dataset well-formedness. In order to admit a finite Plackett-Luce-last maximum-likelihood estimate, the directed strict-order graph induced by the grouped rankings must be strongly connected. Intuitively, no partition of the items may dominate all others in only one direction across the observed groups.

Unlike `borda_ranker`, this model therefore rejects grouped datasets that consist of disconnected query universes or one-way dominance chains, because those data do not identify a finite global strength scale.

6.161.6 Usage

Learning a ranker

```
% Learn from a grouped ranking dataset object
| ?- plackett_luce_last_ranker::learn(my_dataset, Ranker).
...

% Learn with custom iteration and missing-relevance options
| ?- plackett_luce_last_ranker::learn(my_dataset, Ranker, [maximum_iterations(500),
↳ tolerance(1.0e-7), missing_relevance(error)]).
...
```

Inspecting diagnostics

```
% Inspect convergence and dataset summary metadata
| ?- plackett_luce_last_ranker::learn(my_dataset, Ranker),
    plackett_luce_last_ranker::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
...
```

6.161.7 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(plackett_luce_last_ranker),
  options(Options),
  convergence(Status),
  iterations(Iterations),
  final_delta(FinalDelta),
```

(continues on next page)

(continued from previous page)

```
dataset_summary(DatasetSummary)
]
```

Where:

- `model(plackett_luce_last_ranker)` identifies the learning algorithm that produced the ranker.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `convergence(Status)` records the training stop condition. The current values are converged and `maximum_iterations_exhausted`.
- `iterations(Iterations)` stores the number of update iterations that were executed.
- `final_delta(FinalDelta)` stores the maximum absolute strength update in the last iteration.
- `dataset_summary(DatasetSummary)` stores a summary list describing the validated training dataset.

Use the `ranking_protocols diagnostic/2` and `ranker_options/2` helper predicates when you only need a single metadata term or the effective options.

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- plackett_luce_last_ranker::learn(my_dataset, Ranker),
    plackett_luce_last_ranker::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- plackett_luce_last_ranker::learn(my_dataset, Ranker),
    plackett_luce_last_ranker::export_to_clauses(my_dataset, Ranker, my_ranker, Clauses).
Clauses = [my_ranker(plackett_luce_last_ranker(...))]
...

% Export to a file
| ?- plackett_luce_last_ranker::learn(my_dataset, Ranker),
    plackett_luce_last_ranker::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.pl').
...
```

6.161.8 Options

The following options can be passed to the `learn/3` predicate:

- `maximum_iterations(MaximumIterations)`: Positive integer iteration bound.
- `tolerance(Tolerance)`: Positive convergence tolerance.
- `missing_relevance(zero|error)`: Policy used when a declared item in a group has no explicit relevance judgment.

6.161.9 Ranker representation

The learned ranker is represented by a compound term of the form:

```
plackett_luce_last_ranker(Items, Strengths, Diagnostics)
```

Where:

- *Items*: List of ranked items.
- *Scores*: List of normalized Item-Strength pairs.
- *Diagnostics*: List of metadata terms, including the effective options, convergence status, iteration count, final update delta, and dataset summary.

When exported using `export_to_clauses/4` or `export_to_file/4`, this ranker term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.161.10 See also

For the complementary grouped top-choice variant over the same dataset protocol, see the `plackett_luce_ranker` library. For a deterministic non-probabilistic grouped-ranking baseline over the same dataset protocol, see the `borda_ranker` library.

6.161.11 References

1. Plackett, R. L. (1975). The analysis of permutations. *Applied Statistics*, 24(2), 193-202.
2. Luce, R. D. (1959). *Individual Choice Behavior: A Theoretical Analysis*. Wiley.
3. Hunter, D. R. (2004). MM algorithms for generalized Bradley-Terry models. *The Annals of Statistics*, 32(1), 384-406.

6.162 pls_projection

Partial Least Squares projection reducer for target-valued continuous datasets. The library implements the `dimension_reducer_protocol` defined in the `dimension_reduction_protocols` library and learns a deterministic PLS1 projection by centering the training data, optionally standardizing continuous attributes, centering the numeric target, extracting latent directions using repeated cross-covariance maximization with sequential deflation, and storing the resulting direct projection rotations for future transforms.

Requires a dataset implementing `target_supervised_dimension_reduction_dataset_protocol` and therefore uses a numeric target during training.

6.162.1 API documentation

Open the ../apis/library_index.html#pls_projection link in a web browser.

6.162.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(pls_projection(loader)).
```

6.162.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(pls_projection(tester)).
```

6.162.4 Features

- **Continuous Datasets:** Accepts datasets containing only continuous attributes and rejects datasets that also declare the target name as an input feature. Missing or nonnumeric values are also rejected.
- **Target-Supervised Learning:** Uses a numeric target attribute declared by datasets implementing `target_supervised_dimension_reduction_dataset_protocol`.
- **Centering and Optional Scaling:** Centers all attributes and optionally standardizes them before learning latent directions.
- **Deterministic Components:** Learns components deterministically from feature-target cross-covariance without stochastic initialization.
- **Configurable Shortfall Handling:** Lets callers choose whether an extraction shortfall raises an error or truncates the learned reducer with explicit diagnostics.
- **Projection API:** Transforms a new instance into a list of `component_N-Value` pairs using the learned direct PLS rotation vectors.
- **Model Export:** Learned reducers can be exported as predicate clauses or written to a file.

6.162.5 Options

The `learn/3` predicate accepts the following options:

- `n_components/1`: Number of latent PLS components to extract. Requests that exceed `min(FeatureCount, SampleCount - 1)` raise `domain_error(component_count, Requested-Maximum)`. The default is 2.
- `feature_scaling/1`: Whether to standardize continuous attributes before projection. Options: `true` (default) or `false`.
- `shortfall_policy/1`: Controls what happens when residual feature-target cross-covariance becomes negligible before all requested components are extracted. Options: `truncate` (default), which returns a reducer with fewer components and records a `shortfall(truncated(Requested, Learned, ScoreEnergy, Tolerance))` diagnostic, or `error`, which raises `domain_error(component_count, Requested-Learned)`.

- `tolerance/1`: Positive numerical tolerance used to stop extracting components when the residual feature-target cross-covariance becomes negligible. The default is `1.0e-8`.

6.162.6 Usage

The following example uses the sample target-valued dataset shipped with the `dimension_reduction_protocols` library:

```
| ?- logtalk_load(dimension_reduction_protocols('test_datasets/target_latent_measurements')).
```

Learning a reducer

```
| ?- pls_projection::learn(target_latent_measurements, DimensionReducer).
| ?- pls_projection::learn(target_latent_measurements, DimensionReducer, [n_components(1),
↪ feature_scaling(false)]).
```

Transforming new instances

```
| ?- pls_projection::learn(target_latent_measurements, DimensionReducer, [n_components(2)]),
   pls_projection::transform(DimensionReducer, [f1-4.0, f2-8.0, f3-2.0, f4-(-2.0)],
↪ ReducedInstance).
```

Exporting and reusing the reducer

```
| ?- pls_projection::learn(target_latent_measurements, DimensionReducer, [n_components(2)]),
   pls_projection::export_to_file(target_latent_measurements, DimensionReducer, reducer,
↪ 'pls_projection_reducer.pl').

| ?- logtalk_load('pls_projection_reducer.pl'),
   reducer(Reducer),
   pls_projection::transform(Reducer, [f1-4.0, f2-8.0, f3-2.0, f4-(-2.0)],
↪ ReducedInstance).
```

6.162.7 Dimension reducer representation

The learned dimension reducer is represented by a compound term with the functor chosen by the implementation and arity 3. For example:

```
pls_projection_reducer(Encoders, Rotations, Diagnostics)
```

Where:

- `Encoders`: List of continuous attribute encoders storing attribute name, mean, and scale.
- `Rotations`: List of direct PLS projection vectors derived from the learned latent weights and loadings.
- `Diagnostics`: Learned metadata including the effective training options, target information, target loadings, and optional truncate-mode shortfall details.

When exported using `export_to_clauses/4` or `export_to_file/4`, this reducer term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.162.8 References

1. Wold, H. (1975) - “Soft modelling by latent variables; the nonlinear iterative partial least squares (NIPALS) approach”.
2. Geladi, P. and Kowalski, B. R. (1986) - “Partial least-squares regression: a tutorial”.

6.163 prefix_span_pattern_miner

PrefixSpan sequential pattern miner for sequence datasets. The library depends on the `sequential_pattern_mining_protocols` support library, implements the `generic_pattern_miner_protocol` defined in the `pattern_mining_protocols` core library, and mines frequent sequential patterns using recursive projected databases with both same-event and next-event extensions.

Requires a dataset implementing `sequence_dataset_protocol` with sequences represented as ordered lists of canonical sorted itemsets over a declared item domain.

6.163.1 API documentation

Open the ../apis/library_index.html#prefix_span_pattern_miner link in a web browser.

6.163.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(prefix_span_pattern_miner(loader)).
```

6.163.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(prefix_span_pattern_miner(tester)).
```

6.163.4 Features

- **Projected Database Mining:** Mines sequential patterns by recursively projecting suffix databases.
- **Itemset and Sequence Extensions:** Supports both extending the last itemset and appending a new singleton itemset.
- **Canonical Sequences:** Validates that itemsets are sorted, duplicate-free, non-empty, and restricted to declared items.
- **Flexible Support Thresholds:** Supports minimum support specified either as a relative proportion or as an absolute count.
- **Model Export:** Mined pattern collections can be exported as predicate clauses or written to a file.

6.163.5 Options

The mine/3 predicate accepts the following options:

- `minimum_support/1`: Relative minimum support threshold in the interval `]0.0, 1.0]`. The default is `0.5`.
- `minimum_support_count/1`: Absolute minimum support count. When both support options are provided, this option takes precedence.
- `maximum_pattern_length/1`: Maximum total number of items in a mined sequential pattern. The default is `1000`, effectively capped by the longest sequence in the dataset.
- `minimum_pattern_length/1`: Minimum total number of items retained in the mined result. The default is `1`.

6.163.6 Pattern miner representation

The mined pattern miner result is represented by a compound term with the functor chosen by the implementation and arity 3. For example:

```
prefix_span_pattern_miner(ItemDomain, Patterns, Options)
```

Where:

- `ItemDomain`: Canonical sorted list of declared dataset items.
- `Patterns`: List of `sequence_pattern(Pattern, SupportCount)` terms ordered first by total item count and then lexicographically.
- `Options`: Effective mining options used to mine the frequent sequential patterns.

6.163.7 References

1. Pei, J., Han, J., Mortazavi-Asl, B., et al. (2001) - “PrefixSpan: Mining sequential patterns efficiently by prefix-projected pattern growth”.

6.164 probabilistic_pca_projection

Probabilistic Principal Component Analysis reducer for continuous datasets. The library implements the `dimension_reducer_protocol` defined in the `dimension_reduction_protocols` library and learns a linear latent-variable projection by centering the training data, optionally standardizing continuous attributes, estimating the sample covariance matrix, extracting deterministic leading eigenvectors using portable power iteration with deflation, and converting them into the closed-form maximum-likelihood PPCA loading matrix and posterior latent projection.

6.164.1 API documentation

Open the ../apis/library_index.html#probabilistic_pca_projection link in a web browser.

6.164.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(probabilistic_pca_projection(loader)).
```

6.164.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(probabilistic_pca_projection(tester)).
```

6.164.4 Features

- **Continuous Datasets:** Accepts datasets containing only continuous attributes. Missing or nonnumeric values are rejected.
- **Centering and Optional Scaling:** Centers all attributes and optionally standardizes them before fitting the covariance model.
- **Probabilistic Latent Model:** Estimates the PPCA loading matrix and isotropic observation noise variance from the learned covariance eigensystem.
- **Configurable Shortfall Handling:** Lets callers choose whether a numerical-rank shortfall raises an error or returns a truncated reducer with explicit diagnostics.
- **Projection API:** Transforms a new instance into posterior latent means returned as component_N-Value pairs.
- **Model Export:** Learned reducers can be exported as predicate clauses or written to a file.

6.164.5 Options

The `learn/3` predicate accepts the following options:

- `n_components/1`: Number of latent dimensions to extract. Requests that exceed the structural PPCA limit $\min(\text{FeatureCount}, \text{SampleCount} - 1)$ raise `domain_error(component_count, Requested-Maximum)`. The default is 2.
- `feature_scaling/1`: Whether to standardize continuous attributes before fitting the covariance model. Options: `true` (default) or `false`.
- `shortfall_policy/1`: Controls what happens when the covariance matrix yields fewer numerically significant components than requested after passing the structural PPCA bound above. Options: `truncate` (default), which returns a reducer with fewer components and records a `shortfall(truncated(Requested, Learned, ResidualEigenvalue, Tolerance))` diagnostic, or `error`, which raises `domain_error(component_count, Requested-Learned)`.
- `maximum_iterations/1`: Maximum number of power-iteration steps used when estimating each covariance eigenvector. The default is 1000.

- tolerance/1: Positive convergence tolerance used both for power-iteration stopping and for deciding when residual eigenvalues are negligible. The default is $1.0e-8$.

6.164.6 Usage

The following examples use the sample datasets shipped with the `dimension_reduction_protocols` library:

```
| ?- logtalk_load(dimension_reduction_protocols('test_datasets/correlated_plane')).
```

Learning a reducer

```
| ?- probabilistic_pca_projection::learn(correlated_plane, DimensionReducer).
| ?- probabilistic_pca_projection::learn(correlated_plane, DimensionReducer, [n_
↪components(1), feature_scaling(false), shortfall_policy(error)]).
```

Transforming new instances

```
| ?- probabilistic_pca_projection::learn(correlated_plane, DimensionReducer, [n_
↪components(2)]),
    probabilistic_pca_projection::transform(DimensionReducer, [x-2.0, y-4.0, z-6.0], ↪
↪ReducedInstance).
```

Exporting and reusing the reducer

```
| ?- probabilistic_pca_projection::learn(correlated_plane, DimensionReducer, [n_
↪components(1)]),
    probabilistic_pca_projection::export_to_file(correlated_plane, DimensionReducer, ↪
↪reducer, 'probabilistic_pca_reducer.pl').

| ?- logtalk_load('probabilistic_pca_reducer.pl'),
    reducer(Reducer),
    probabilistic_pca_projection::transform(Reducer, [x-1.0, y-2.0, z-3.0], ↪
↪ReducedInstance).
```

6.164.7 Dimension reducer representation

The learned dimension reducer is represented by a compound term with the functor chosen by the implementation and arity 6. For example:

```
probabilistic_pca_reducer(Encoders, Components, Loadings, NoiseVariance, ExplainedVariances, ↪
↪Diagnostics)
```

Where:

- Encoders: List of continuous attribute encoders storing attribute name, mean, and scale.
- Components: List of posterior latent projection vectors in descending explained-variance order.

- **Loadings:** List of maximum-likelihood PPCA loading vectors aligned with the extracted latent dimensions.
- **NoiseVariance:** Estimated isotropic observation noise variance.
- **ExplainedVariances:** List of retained covariance eigenvalues matching the extracted latent dimensions.
- **Diagnostics:** Learned metadata including the effective training options, sample count, retained explained variances, estimated noise variance, preprocessing details, and optional truncate-mode short-fall details.

When exported using `export_to_clauses/4` or `export_to_file/4`, this reducer term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.164.8 References

1. Tipping, M. E. and Bishop, C. M. (1999) - “Probabilistic Principal Component Analysis”.
2. Bishop, C. M. (2006) - “Pattern Recognition and Machine Learning”. Section 12.2.

6.165 process

This library provides a portable abstraction over process handling predicates found in some backend Prolog systems. Currently supports ECLiPSe, GNU Prolog, SICStus Prolog, SWI-Prolog, Trealla Prolog, and XVM.

6.165.1 API documentation

Open the ../docs/library_index.html#process link in a web browser.

6.165.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(process(loader)).
```

6.165.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(process(tester)).
```

6.165.4 Usage

The process object provides the following predicates for portable process handling:

- `create(Executable, Arguments, Options)` - Creates a new process
- `wait(Process, Status)` - Waits for a process to terminate
- `kill(Process, Signal)` - Terminates a process with a specific signal
- `kill(Process)` - Terminates a process using the default signal (SIGKILL)

The `create/3` predicate supports the following options:

- `process(Pid)` - Unifies `Pid` with the process identifier (an opaque ground term)
- `stdin(Stream)` - Binds `Stream` to the process standard input
- `stdout(Stream)` - Binds `Stream` to the process standard output
- `stderr(Stream)` - Binds `Stream` to the process standard error

Note: Process identifiers (PIDs) should be treated as opaque ground terms. Their internal representation varies between backend Prolog systems.

Example:

```
| ?- process::create('/bin/echo', ['hello.'], [stdout(Out), process(PID)]),
    read_term(Out, Term, []),
    close(Out).
Out = ..., PID = ..., Term = hello.
```

6.166 protobuf

The protobuf library implements predicates for reading (parsing) and writing (generating) data in the Google Protocol Buffers binary format.

This implementation is based on the Protocol Buffers Language Guide (proto3) and the Protocol Buffers Encoding specification:

- <https://protobuf.dev/programming-guides/proto3/>
- <https://protobuf.dev/programming-guides/encoding/>

This library requires a backend supporting unbounded integer arithmetic.

6.166.1 API documentation

Open the `../apis/library_index.html#avro` link in a web browser.

6.166.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(avro(loader)).
```

6.166.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(avro(tester)).
```

6.166.4 Protocol Buffers Overview

Protocol Buffers (protobuf) is a language-neutral, platform-neutral extensible mechanism for serializing structured data. It is widely used for data storage, communication protocols, and more.

6.166.5 Schema Representation

In this library, Protocol Buffers schemas (normally defined in `.proto` files) are represented as Logtalk terms. The schema representation uses a curly-bracketed syntax similar to JSON representation.

A message schema is represented as:

```
{message-MessageName, fields-FieldList}
```

Where `FieldList` is a list of field definitions, each with:

```
{number-FieldNumber, name-FieldName, type-FieldType}
```

Supported Types

Scalar Value Types

- `int32` - Uses variable-length encoding. Inefficient for negative numbers (use `sint32` instead).
- `int64` - Uses variable-length encoding. Inefficient for negative numbers (use `sint64` instead).
- `uint32` - Uses variable-length encoding.
- `uint64` - Uses variable-length encoding.
- `sint32` - Uses variable-length encoding with ZigZag encoding. More efficient for negative numbers.
- `sint64` - Uses variable-length encoding with ZigZag encoding. More efficient for negative numbers.
- `bool` - Boolean value: `true` or `false`.
- `fixed32` - Always four bytes. More efficient than `uint32` if values are often greater than 2^{28} .
- `fixed64` - Always eight bytes. More efficient than `uint64` if values are often greater than 2^{56} .
- `sfixed32` - Always four bytes. Signed fixed-width integer.
- `sfixed64` - Always eight bytes. Signed fixed-width integer.

- float - 32-bit IEEE 754 floating point.
- double - 64-bit IEEE 754 floating point.
- string - UTF-8 encoded text.
- bytes - Arbitrary byte sequence.

Complex Types

- {message-Name, fields-Fields} - Embedded message (nested structure).

Schema Examples

Simple primitive type schemas:

- int32
- string
- bool

Message schema example:

```
{message-'Person', fields-[
  {number-1, name-name, type-string},
  {number-2, name-id, type-int32},
  {number-3, name-email, type-string}
]}
```

Nested message schema:

```
{message-'AddressBook', fields-[
  {number-1, name-people, type-{message-'Person', fields-[
    {number-1, name-name, type-string},
    {number-2, name-id, type-int32}
  ]}}
]}
```

6.166.6 Data Representation

Data to be serialized is represented as a list of field name-value pairs using the - operator. For example:

```
[name-'John Doe', id-42, email-'john@example.com']
```

For primitive types, data is represented directly as Logtalk values:

- Integers: 42, -17
- Booleans: true, false
- Floats: 3.14, -273.15
- Strings: 'Hello World', hello
- Bytes: [0x48, 0x65, 0x6c, 0x6c, 0x6f]

6.166.7 Encoding

Encoding is accomplished using the `generate/3` or `generate/4` predicates.

Simple Value Encoding

For example, encoding an integer using the `int32` schema:

```
| ?- protobuf::generate(bytes(Bytes), int32, 150).
Bytes = [150, 1]
yes
```

Encoding a string:

```
| ?- protobuf::generate(bytes(Bytes), string, 'testing').
Bytes = [116, 101, 115, 116, 105, 110, 103]
yes
```

Message Encoding

To encode a complete message:

```
| ?- Schema = {message-'Person', fields-[
|     {number-1, name-name, type-string},
|     {number-2, name-id, type-int32}
|   ]},
|   Data = [name-'John', id-42],
|   protobuf::generate(bytes(Bytes), Schema, Data).

Schema = {...},
Data = [name-'John', id-42],
Bytes = [10, 4, 74, 111, 104, 110, 16, 84]
yes
```

Including Schema in Output

To include the schema in the output (as a custom wrapper message), use the `generate/4` predicate with the second argument set to `true`:

```
| ?- protobuf::generate(file('output.pb'), true, Schema, Data).
yes
```

This generates a wrapper message with:

- Field 1: schema (as a JSON string)
- Field 2: data (as encoded bytes)

6.166.8 Decoding

Decoding is accomplished using the `parse/2` or `parse/3` predicates.

Parsing with Schema Embedded

When parsing a file that includes a schema (wrapper message format), use `parse/2` which returns a Schema-Data pair:

```
| ?- protobuf::parse(file('input.pb'), Schema-Data).
```

When the schema is not present in the file, Schema is unified with `false`.

Parsing with Known Schema

When the schema is known and not embedded in the file, use `parse/3`:

```
| ?- Schema = {message-'Person', fields-[
|       {number-1, name-name, type-string},
|       {number-2, name-id, type-int32}
|     ]},
|   protobuf::parse(file('person.pb'), Schema, Data).

Schema = {...},
Data = [name-'Alice', id-123]
yes
```

Parsing from Different Sources

The library supports three types of sources:

- `bytes(List)` - Parse from a list of bytes
- `stream(Stream)` - Parse from an open binary stream
- `file(Path)` - Parse from a file

6.166.9 Wire Format Details

Protocol Buffers uses an efficient binary wire format with the following wire types:

- **0 (Varint)**: `int32`, `int64`, `uint32`, `uint64`, `sint32`, `sint64`, `bool`
- **1 (64-bit)**: `fixed64`, `sfixed64`, `double`
- **2 (Length-delimited)**: `string`, `bytes`, `embedded messages`
- **5 (32-bit)**: `fixed32`, `sfixed32`, `float`

Each field is encoded with a tag (field number and wire type) followed by the value.

6.166.10 Binary Format Compatibility

The binary format produced by this library is compatible with:

- Official Protocol Buffers implementations (C++, Java, Python, Go, etc.)
- Other third-party implementations that follow the protobuf specification

Note: This library implements the core binary encoding format. Advanced features like `oneof`, `map`, `enum`, and repeated fields are not yet supported in this initial version.

6.166.11 Examples and Test Files

The `test_files` subdirectory contains example schemas and data files that demonstrate the library's capabilities:

`person.proto`

A simple Person message schema. This is a classic Protocol Buffers example demonstrating basic field types:

- **Source:** Adapted from the official Protocol Buffers tutorial
- **URL:** <https://protobuf.dev/getting-started/>
- **License:** Apache License 2.0 (compatible with Protocol Buffers documentation)

`addressbook.proto`

An AddressBook schema demonstrating nested messages. Shows how to represent complex hierarchical data structures:

- **Source:** Adapted from the official Protocol Buffers tutorial
- **URL:** <https://protobuf.dev/getting-started/>
- **License:** Apache License 2.0

Performance Considerations

Protocol Buffers binary format is designed for:

- **Compactness:** Variable-length encoding for integers
- **Speed:** Simple encoding/decoding rules
- **Forward/backward compatibility:** Unknown fields can be skipped

For optimal performance:

- Use `sint32/sint64` for negative numbers
- Use `fixed32/fixed64` for large numbers
- Keep field numbers low (1-15 use only 1 byte for tags)

6.166.12 Further Reading

- Protocol Buffers Language Guide: <https://protobuf.dev/programming-guides/proto3/>
- Protocol Buffers Encoding Guide: <https://protobuf.dev/programming-guides/encoding/>
- Protocol Buffers Tutorials: <https://protobuf.dev/getting-started/>

6.167 qda_classifier

Quadratic Discriminant Analysis classifier for continuous datasets using class-specific covariance estimates with diagonal regularization. The implementation learns one quadratic discriminant model per class and predicts the class with the highest class-specific score.

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning a classifier from a dataset, using it to make predictions, inspecting class scores, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. All dataset attributes must be declared as continuous.

6.167.1 API documentation

Open the `../../docs/library_index.html#qda_classifier` link in a web browser.

6.167.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(qda_classifier(loader)).
```

6.167.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(qda_classifier(tester)).
```

6.167.4 Features

- **Continuous Datasets:** Accepts only datasets whose attributes are all declared as continuous.
- **Class-Specific Covariances:** Learns one covariance matrix, mean vector, and prior per class.
- **Regularized Estimation:** Applies configurable diagonal regularization to each class covariance matrix before inversion.
- **Feature Scaling:** Supports optional z-score scaling of continuous features before training.
- **Score Inspection:** Provides class scores for all classes using `predict_scores/3`.
- **Classifier Export:** Learned classifiers can be exported as predicate clauses or written to a file.

6.167.5 Options

The learn/3 predicate supports these options:

- feature_scaling/1 - whether to standardize continuous attributes before training (default: true)
- regularization/1 - positive diagonal value added to each class covariance matrix before inversion (default: 1.0e-6)

6.167.6 Usage

Learning a classifier

```
| ?- qda_classifier::learn(iris_small, Classifier).
| ?- qda_classifier::learn(iris_small, Classifier, [regularization(1.0e-5)]).
```

Making predictions

```
| ?- qda_classifier::learn(iris_small, Classifier),
    qda_classifier::predict(Classifier, [sepal_length-5.9, sepal_width-3.0, petal_length-5.
↪1, petal_width-1.8], Class).
| ?- qda_classifier::learn(iris_small, Classifier),
    qda_classifier::predict_scores(Classifier, [sepal_length-6.4, sepal_width-3.2, petal_
↪length-4.5, petal_width-1.5], Scores).
```

Exporting the classifier

```
| ?- qda_classifier::learn(iris_small, Classifier),
    qda_classifier::export_to_clauses(iris_small, Classifier, classify, Clauses).
| ?- qda_classifier::learn(iris_small, Classifier),
    qda_classifier::export_to_file(iris_small, Classifier, classify, 'classifier.pl').
```

6.167.7 Classifier representation

The learned classifier is represented as a compound term with the form:

```
qda_classifier(Encoders, Models, Options)
```

Where:

- Encoders: list of continuous feature encoders with learned scaling parameters
- Models: list of class_model(Class, Prior, Mean, Precision, LogDeterminant, Constant) terms
- Options: merged training options used to learn the classifier

When exported using export_to_clauses/4 or export_to_file/4, this classifier term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.167.8 References

1. Hastie, T., Tibshirani, R. and Friedman, J. (2009). “The Elements of Statistical Learning”. Section 4.3.
2. Bishop, C.M. (2006). “Pattern Recognition and Machine Learning”. Section 4.2.
3. Duda, R.O., Hart, P.E. and Stork, D.G. (2001). “Pattern Classification”. Chapter 2.

6.168 queues

This library implements queues. The queue representation should be regarded as an opaque term and only accessed using the library predicates.

6.168.1 API documentation

Open the `../..apis/library_index.html#queues` link in a web browser.

6.168.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(queues(loader)).
```

6.168.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(queues(tester)).
```

6.168.4 Usage

To create a new queue, use the `new/1` predicate:

```
| ?- queue::new(Queue).
Queue = ...
yes
```

Elements can be added to either the end of the queue or the front of the queue using, respectively, the `join/3` and `join_all/3` predicates or the `jump/3` and `jump_all/3`. For example:

```
| ?- queue::(new(Queue0), join_all([1,2,3], Queue0, Queue1)).
Queue0 = ...,
Queue1 = ...
yes
```

We can query the head of the queue or remove the head of the queue using, respectively, the `head/2` and `serve/3` predicates. For example:

```
| ?- queue::(new(Queue0), join(1, Queue0, Queue1), head(Queue1, Head)).
Queue0 = ...,
Queue1 = ...,
Head = 1
yes
```

For details on these and other provided predicates, consult the library API documentation.

6.169 random

This library provides portable random number generators and an abstraction over the native backend Prolog compiler random number generator if available. It includes predicates for generating random floats and random integers in a given interval; predicates for generating random sequences and sets; predicates for randomly selecting, enumerating, and swapping elements from a list; predicates that succeed, fail, or call another predicate with a given probability; and predicates for sampling common probability distributions.

6.169.1 API documentation

Open the ../apis/library_index.html#random link in a web browser.

6.169.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(random(loader)).
```

6.169.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(random(tester)).
```

6.169.4 Algorithms

The `random(Algorithm)` and `fast_random(Algorithm)` parametric objects support the following pseudo-random number generator algorithms:

- `as183` - Algorithm AS 183 from Applied Statistics. 3x32-bit PRNG with period 2^{60} .
- `splitmix64` - SplitMix64 1x64-bit PRNG primarily used for seeding other generators. Algorithm by Guy L. Steele Jr. et al.
- `xoshiro128pp` - Xoshiro128++ 4x32-bit state-of-the-art PRNG with period $2^{128}-1$. Algorithm by David Blackman and Sebastiano Vigna.
- `xoshiro128ss` - Xoshiro128** 4x32-bit state-of-the-art PRNG with period $2^{128}-1$. Algorithm by David Blackman and Sebastiano Vigna.
- `xoshiro256pp` - Xoshiro256++ 4x64-bit state-of-the-art PRNG with period $2^{256}-1$. Algorithm by David Blackman and Sebastiano Vigna.

- xoshiro256ss - Xoshiro256** 4x64-bit state-of-the-art PRNG with period $2^{256}-1$. Algorithm by David Blackman and Sebastiano Vigna.
- well512a - WELL512a 16x32-bit state-of-the-art PRNG with period $2^{512}-1$. Algorithm by François Panneton, Pierre L'Ecuyer, and Makoto Matsumoto.

Note that none of the above algorithms are cryptographically secure.

The SplitMix64, Xoshiro256++, and Xoshiro256** algorithms require a backend supporting unbound integer arithmetic.

The random and fast_random objects use the as183 algorithm and are provided for backward compatibility.

6.169.5 Usage

The random(Algorithm) object implements a portable random number generator and supports multiple random number generators, using different seeds, by defining derived objects. For example:

```
:- object(my_random_generator_1,
    extends(random(xoshiro128pp))).

    :- initialization(::reset_seed).

:- end_object.
```

The fast_random(Algorithm) object also implements a portable random number generator but does not support deriving multiple random number generators, which makes it a bit faster than the random(Algorithm) object.

The random(Algorithm), random, fast_random(Algorithm), and fast_random objects manage the random number generator seed using internal dynamic state. The predicates that update the seed are declared as synchronized (when running on Prolog backends that support threads). Still, care must be taken when using these objects from multi-threaded applications, as there is no portable solution to protect seed updates from signals and prevent inconsistent state when threads are canceled.

The random(Algorithm), random, fast_random(Algorithm), and fast_random objects always initialize the random generator seed to the same value, thus providing a pseudo random number generator. The randomize/1 predicate can be used to initialize the seed with a random value. The argument should be a large positive integer. In alternative, when using a small integer argument, discard the first dozen random values.

The backend_random object abstracts the native backend Prolog compiler random number generator for the basic random/1, get_seed/1, and set_seed/1 predicates providing a portable implementation for the remaining predicates. This makes the object stateless, which allows reliable use from multiple threads. Consult the backend Prolog compiler documentation for details on its random number generator properties. Note that several of the supported backend Prolog systems, notably B-Prolog, CxProlog, ECLiPSe, JIProlog, and Quintus Prolog, do not provide implementations for both the get_seed/1 and set_seed/1 predicates and calling these predicates simply succeed without performing any action.

All random objects implement the sampling_protocol protocol. To maximize performance, the shared implementations of the sampling predicates is defined in the sampling.lgt file that's included in the random objects. This allows these predicates to call the basic random/1 and random/3 predicates as locally defined predicates.

6.170 random_forest_classifier

Random Forest classifier using C4.5 decision trees as base learners. Builds an ensemble of decision trees trained on bootstrap samples with random feature subsets and combines their predictions through majority voting.

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning an ensemble classifier from a dataset, using it to make predictions (with class probabilities), and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. See `test_files` directory for examples.

6.170.1 API documentation

Open the ../docs/library_index.html#random_forest_classifier link in a web browser.

6.170.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(random_forest_classifier(loader)).
```

6.170.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(random_forest_classifier(tester)).
```

6.170.4 Features

- **Ensemble Learning:** Combines multiple C4.5 decision trees for robust predictions.
- **Bootstrap Sampling:** Each tree is trained on a bootstrap sample (random sample with replacement) of the training data.
- **Feature Randomization:** Random subset of features selected for each tree (default: `sqrt(TotalFeatures)`).
- **Portable Seeded Sampling:** Uses `fast_random(xoshiro128pp)` so bootstrap sampling and feature subset selection are portable and reproducible.
- **Majority Voting:** Final predictions determined by voting across all trees.
- **Probability Estimation:** Provides confidence scores based on vote proportions.
- **Configurable Options:** Number of trees, maximum features per tree, and random seed via predicate options.
- **Classifier Export:** Learned classifiers can be exported as predicate clauses.

6.170.5 Options

The following options can be passed to the `learn/3` predicate:

- `number_of_trees(N)`: Number of trees in the forest (default: 10).
- `maximum_features_per_tree(N)`: Maximum number of features to consider per tree (default: `sqrt(TotalFeatures)`).
- `random_seed(N)`: Positive integer seed used by the portable `fast_random(xoshiro128pp)` pseudo-random generator when drawing bootstrap samples and random feature subsets. Using the same seed with the same dataset and options reproduces the same learned classifier (default: 1357911).

6.170.6 Classifier representation

The learned classifier is represented as a compound term:

```
rf_classifier(Trees, ClassValues, Options)
```

Where:

- **Trees**: List of `tree(C45Tree, AttributeNames)` pairs
- **ClassValues**: List of possible class values
- **Options**: List of options used during learning

When exported using `export_to_clauses/4` or `export_to_file/4`, this classifier term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.170.7 References

1. Breiman, L. (2001). “Random Forests”. *Machine Learning*, 45(1), 5-32.
2. Ho, T.K. (1995). “Random Decision Forests”. *Proceedings of the 3rd International Conference on Document Analysis and Recognition*.
3. Quinlan, J.R. (1993). “C4.5: Programs for Machine Learning”. Morgan Kaufmann.

6.170.8 Usage

Learning a Classifier

```
% Learn a random forest with default options (10 trees)
| ?- random_forest_classifier::learn(play_tennis, Classifier).
...

% Learn with custom options
| ?- random_forest_classifier::learn(play_tennis, Classifier, [number_of_trees(20), maximum_
↪ features_per_tree(2), random_seed(17)]).
...
```

Making Predictions

```
% Predict class for a new instance
| ?- random_forest_classifier::learn(play_tennis, Classifier),
    random_forest_classifier::predict(Classifier, [outlook-sunny, temperature-hot, humidity-
↳high, wind-weak], Class).
Class = no
...

% Get probability distribution from ensemble voting
| ?- random_forest_classifier::learn(play_tennis, Classifier),
    random_forest_classifier::predict_probabilities(Classifier, [outlook-overcast,
↳temperature-mild, humidity-normal, wind-weak], Probabilities).
Probabilities = [yes-0.9, no-0.1]
...
```

Exporting the Classifier

```
% Export as predicate clauses
| ?- random_forest_classifier::learn(play_tennis, Classifier),
    random_forest_classifier::export_to_clauses(play_tennis, Classifier, my_forest,
↳Clauses).
Clauses = [my_forest(random_forest_classifier(...))]
...

% Export to a file
| ?- random_forest_classifier::learn(play_tennis, Classifier),
    random_forest_classifier::export_to_file(play_tennis, Classifier, my_forest, 'forest.pl
↳').
...
```

Using a Saved Classifier

```
% Load and use a previously saved classifier
| ?- logtalk_load('forest.pl'),
    my_forest(Classifier),
    random_forest_classifier::predict(Classifier, [outlook-sunny, temperature-cool,
↳humidity-normal, wind-weak], Class).
Class = yes
...
```

Printing the Classifier

```
% Print a summary of the random forest
| ?- random_forest_classifier::learn(play_tennis, Classifier),
    random_forest_classifier::print_classifier(Classifier).

Random Forest Classifier
=====

Number of trees: 10
Class values: [yes,no]
Options: [number_of_trees(10)]

Trees:
  Tree 1 (features: [outlook,humidity]):
    -> tree rooted at outlook
    ...
  ...
```

6.171 random_forest_regression

Random Forest regressor supporting continuous and mixed-feature datasets. The library implements the regressor_protocol defined in the regression_protocols library and learns an ensemble of regression trees trained on bootstrap samples and per-split random feature subsets, predicting with the arithmetic mean of the individual tree predictions.

6.171.1 API documentation

Open the ../apis/library_index.html#random_forest_regression link in a web browser.

6.171.2 Loading

To load this library, load the loader.lgt file:

```
| ?- logtalk_load(random_forest_regression(loader)).
```

6.171.3 Testing

To test this library predicates, load the tester.lgt file:

```
| ?- logtalk_load(random_forest_regression(tester)).
```

To run the performance benchmark suite, load the tester_performance.lgt file:

```
| ?- logtalk_load(random_forest_regression(tester_performance)).
```

6.171.4 Features

- **Bootstrap Ensembles:** Trains multiple regression trees on bootstrap samples.
- **Random Feature Subsets:** Samples a random subset of the available dataset attributes at each split of every tree.
- **Portable Seeded Sampling:** Uses `fast_random(xoshiro128pp)` so bootstrap and split-level feature sampling are portable and reproducible.
- **Tree Averaging:** Predicts numeric targets using the arithmetic mean of the tree predictions.
- **Tree Configuration:** Exposes the underlying regression-tree split-feature, depth, minimum-leaf, variance-reduction, and scaling options.
- **Categorical Features Encoding:** Uses reference-level dummy coding derived from the declared dataset attribute values, with a missing-value indicator, and the resulting encoded features are treated as ordinary numeric split features by the tree learners.
- **Diagnostics Metadata:** Learned regressors record model name, target, training example count, attribute count, tree count, and effective options, accessible using the shared regression diagnostics predicates.
- **Model Export:** Learned regressors can be exported as predicate clauses or written to a file.
- **Reference Benchmarks:** Includes a dedicated performance suite reporting training time, RMSE, and MAE for representative regression datasets.

6.171.5 Regressor representation

The learned regressor is represented by default as:

- `rf_regressor(Trees, Diagnostics)`

The exported predicate clauses therefore use the shape:

- `Functor(Trees, Diagnostics)`

6.171.6 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(random_forest_regression),
  target(Target),
  training_example_count(TrainingExampleCount),
  options(Options),
  attribute_count(AttributeCount),
  tree_count(TreeCount)
]
```

Where:

- `model(random_forest_regression)` identifies the learning algorithm that produced the regressor.
- `target(Target)` stores the target attribute name declared by the training dataset.
- `training_example_count(TrainingExampleCount)` stores the number of examples used during training.

- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `attribute_count(AttributeCount)` stores the number of dataset attributes available to the ensemble before split-level subsampling.
- `tree_count(TreeCount)` stores the number of trained regression trees in the ensemble.

Use the `regression_protocols diagnostic/2` and `regressor_options/2` helper predicates when you only need a single metadata term or the effective options.

6.171.7 Options

The `learn/3` predicate accepts the following options:

- `number_of_trees/1`: Number of regression trees to train in the ensemble. Increasing this value usually improves stability at the cost of additional training and prediction time. The default is 10.
- `maximum_features_per_split/1`: Number of dataset attributes randomly sampled at each split when searching for the best partition. Accepted values are a positive integer or `all`. When omitted, the library uses the square root of the total number of available attributes, with a minimum of one attribute. Passing `all` disables split-level attribute subsampling.
- `maximum_depth/1`: Maximum depth allowed for each regression-tree base learner. The default is 10.
- `minimum_samples_leaf/1`: Minimum number of training examples required in each leaf of a base learner tree. The default is 1.
- `minimum_variance_reduction/1`: Minimum split gain required by each base learner tree before accepting a partition. The default is 0.0.
- `feature_scaling/1`: Controls z-score standardization of continuous attributes inside each regression-tree base learner. Accepted values are `true` and `false`. The default is `false`.
- `random_seed/1`: Positive integer seed used by the portable `fast_random(xoshiro128pp)` pseudo-random generator when drawing bootstrap samples and split-level random feature subsets. Using the same seed with the same dataset and options reproduces the same learned regressor. The default is 1357911.

6.172 random_projection

Random projection reducer for continuous datasets. The library implements the `dimension_reducer_protocol` defined in the `dimension_reduction_protocols` library and learns a seeded dense Rademacher projection matrix using the portable `fast_random` pseudo-random generator after centering the training data, optionally standardizing continuous attributes, and sampling entries in $\{-1/\sqrt{k}, +1/\sqrt{k}\}$ where k is the requested reduced dimensionality.

6.172.1 API documentation

Open the ../apis/library_index.html#random_projection link in a web browser.

6.172.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(random_projection(loader)).
```

6.172.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(random_projection(tester)).
```

6.172.4 Features

- **Continuous Datasets:** Accepts datasets containing only continuous attributes. Missing or nonnumeric values are rejected.
- **Centering and Optional Scaling:** Centers all attributes and optionally standardizes them before projection.
- **Portable Seeded Sampling:** Uses `fast_random(xoshiro128pp)` so learned projection matrices are portable and reproducible.
- **Projection API:** Transforms a new instance into a list of `component_N-Value` pairs.
- **Model Export:** Learned reducers can be exported as predicate clauses or written to a file.

6.172.5 Options

The `learn/3` predicate accepts the following options:

- `n_components/1`: Number of random projection components to sample. Requests that exceed the number of features raise `domain_error(component_count, Requested-Maximum)`. The default is 2.
- `feature_scaling/1`: Whether to standardize continuous attributes before projection. Options: `true` (default) or `false`.
- `random_seed/1`: Positive integer used to seed the portable pseudo-random generator before sampling the projection matrix. The default is 1357911.

6.172.6 Usage

The following examples use the sample datasets shipped with the `dimension_reduction_protocols` library:

```
| ?- logtalk_load(dimension_reduction_protocols('test_datasets/correlated_plane')),
    logtalk_load(dimension_reduction_protocols('test_datasets/high_dimensional_measurements_
    ↪')).
```

Learning a reducer

```
| ?- random_projection::learn(correlated_plane, DimensionReducer).

| ?- random_projection::learn(correlated_plane, DimensionReducer, [n_components(1), feature_
    ↪scaling(false), random_seed(17)]).
```

Transforming new instances

```
| ?- random_projection::learn(high_dimensional_measurements, DimensionReducer, [random_
    ↪seed(11)]),
    random_projection::transform(DimensionReducer, [f1-0.9, f2-1.1, f3-1.0, f4-2.0, f5-2.2,
    ↪f6-2.1], ReducedInstance).

| ?- random_projection::learn(correlated_plane, DimensionReducer, [n_components(1), random_
    ↪seed(19)]),
    random_projection::transform(DimensionReducer, [x-1.0, y-2.0, z-3.0], ReducedInstance).
```

Exporting and reusing the reducer

```
| ?- random_projection::learn(correlated_plane, DimensionReducer, [n_components(1), random_
    ↪seed(29)]),
    random_projection::export_to_file(correlated_plane, DimensionReducer, reducer, 'random_
    ↪projection_reducer.pl').

| ?- logtalk_load('random_projection_reducer.pl'),
    reducer(Reducer),
    random_projection::transform(Reducer, [x-1.0, y-2.0, z-3.0], ReducedInstance).
```

6.172.7 Dimension reducer representation

The learned dimension reducer is represented by a compound term with the functor chosen by the implementation and arity 3. For example:

```
random_projection_reducer(Encoders, Components, Diagnostics)
```

Where:

- Encoders: List of continuous attribute encoders storing attribute name, mean, and scale.
- Components: List of sampled projection vectors in component order.

- **Diagnostics:** Learned reducer metadata including the effective training options and reproducibility details.

When exported using `export_to_clauses/4` or `export_to_file/4`, this reducer term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.172.8 References

1. Johnson, W. B. and Lindenstrauss, J. (1984) - “Extensions of Lipschitz mappings into a Hilbert space”.
2. Achlioptas, D. (2003) - “Database-friendly random projections: Johnson-Lindenstrauss with binary coins”.

6.173 rank_centrality

Rank Centrality pairwise preference ranker. It uses the Rank Centrality transition rule where each observed opponent contributes an outgoing transition proportional to the empirical probability of beating the current item, scaled by the maximum comparison degree, and estimates the stationary distribution using deterministic power iteration.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from pairwise preferences, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The training dataset must declare each ranked item once, enumerate positive-weight pairwise preferences between distinct declared items, induce a connected undirected comparison graph, and induce a strongly connected directed transition graph so that the learned stationary distribution is unique.

6.173.1 API documentation

Open the ../apis/library_index.html#rank_centrality link in a web browser.

6.173.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(rank_centrality(loader)).
```

6.173.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(rank_centrality(tester)).
```

6.173.4 Features

- **Pairwise Preference Learning:** Learns one stationary probability score per item from weighted head-to-head outcomes.
- **Original Rank Centrality Transition Rule:** Builds the Markov chain from empirical pairwise win probabilities scaled by the maximum comparison degree, then estimates the stationary distribution by deterministic power iteration.
- **Deterministic Ranking:** Orders candidate items by learned stationary probability with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items, self-preferences, non-positive weights, disconnected undirected comparison graphs, and pairwise datasets whose directed transition graph is not strongly connected.
- **Training Diagnostics:** Learned rankers include convergence, iteration, maximum-degree, and dataset summary metadata that can be accessed using the `diagnostics/2` predicate.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.
- **Sparse Transition Processing:** Training aggregates observed pairwise comparisons into sparse incoming transition lists so each power-iteration step scales with observed matchups instead of a dense item cross-product.

6.173.5 Dataset requirements

This implementation requires more than undirected connectedness. In order to guarantee a unique stationary distribution, the directed transition graph induced by the aggregated pairwise outcomes must be strongly connected. Datasets that create one-way dominance sinks or other disconnected directed transition components are rejected instead of producing ambiguous stationary scores.

6.173.6 Usage

Learning a ranker

```
% Learn from a pairwise ranking dataset object
| ?- rank_centrality::learn(my_dataset, Ranker).
...

% Learn with custom iteration and convergence options
| ?- rank_centrality::learn(my_dataset, Ranker, [maximum_iterations(500), tolerance(1.0e-
↪9)]).
...
```

Inspecting diagnostics

```
% Inspect convergence and dataset summary metadata
| ?- rank_centrality::learn(my_dataset, Ranker),
    rank_centrality::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
```

6.173.7 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(rank_centrality),
  options(Options),
  convergence(Status),
  iterations(Iterations),
  final_delta(FinalDelta),
  maximum_degree(MaximumDegree),
  dataset_summary(DatasetSummary)
]
```

Where:

- `model(rank_centrality)` identifies the learning algorithm that produced the ranker.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `convergence(Status)` records the training stop condition. The current values are converged and maximum_iterations_exhausted.
- `iterations(Iterations)` stores the number of power-iteration steps that were executed.
- `final_delta(FinalDelta)` stores the maximum absolute score update in the last iteration.
- `maximum_degree(MaximumDegree)` stores the maximum number of observed opponents for any single item in the training dataset.
- `dataset_summary(DatasetSummary)` stores a summary list describing the validated training dataset.

The current `dataset_summary/1` payload has the form:

```
[
  items(NumberOfItems),
  preferences(NumberOfPreferences),
  connected_components(NumberOfComponents),
  isolated_items(IsolatedItems)
]
```

Use the `ranking_protocols` `diagnostic/2` and `ranker_options/2` helper predicates when you only need a single metadata term or the effective options.

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- rank_centrality::learn(my_dataset, Ranker),
    rank_centrality::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
...
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- rank_centrality::learn(my_dataset, Ranker),
    rank_centrality::export_to_clauses(my_dataset, Ranker, my_ranker, Clauses).
Clauses = [my_ranker(rank_centrality_ranker(...))]
...

% Export to a file
| ?- rank_centrality::learn(my_dataset, Ranker),
    rank_centrality::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.pl').
...
```

6.173.8 Options

The following options can be passed to the `learn/3` predicate:

- `maximum_iterations(MaximumIterations)`: Positive integer iteration bound.
- `tolerance(Tolerance)`: Positive convergence tolerance.

6.173.9 Ranker representation

The learned ranker is represented by a compound term of the form:

```
rank_centrality_ranker(Items, Scores, Diagnostics)
```

Where:

- `Items`: List of ranked items.
- `Scores`: List of Item-Score pairs.
- `Diagnostics`: List of metadata terms, including the effective options, convergence status, iteration count, final update delta, maximum degree, and dataset summary.

The `Scores` payload is expected to contain positive stationary probability values summing to 1.0. The ranker validation logic enforces this invariant when consuming serialized or exported ranker terms so that malformed payloads are rejected instead of silently accepted.

When exported using `export_to_clauses/4` or `export_to_file/4`, this ranker term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.173.10 References

1. Negahban, S., Oh, S., and Shah, D. (2012). Rank Centrality: Ranking from pairwise comparisons. *Operations Research*, 65(1), 266-287.

6.174 ranked_pairs

Ranked Pairs pairwise preference ranker. It builds the direct pairwise victory graph from aggregated matchups, considers victories in descending direct-victory strength order, and locks each victory unless it would create a directed cycle in the accepted lock graph.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from pairwise preferences, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The current implementation requires a well-formed connected pairwise dataset so that all ranked items remain part of a single comparison graph.

6.174.1 API documentation

Open the ../apis/library_index.html#ranked_pairs link in a web browser.

6.174.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ranked_pairs(loader)).
```

6.174.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ranked_pairs(tester)).
```

6.174.4 Features

- **Pairwise Preference Learning:** Learns one deterministic score per item from aggregated pairwise outcomes.
- **Ranked Pairs Locking:** Processes direct victories in descending strength order and accepts each victory unless it would create a directed cycle in the current lock graph.
- **Configurable Direct Victory Semantics:** Supports both winning-votes and victory-margins victory-strength modes before locking.
- **Configurable Equal-Strength Tie Breaking:** Supports both standard term-order and dataset declaration-order resolution when direct victories have the same strength.
- **Locked-Pair Access:** Exposes the accepted lock graph in lock order using the `locked_pairs/2` predicate.
- **Deterministic Ranking:** Orders candidate items by the number of opponents reachable in the final locked relation with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items, self-preferences, non-positive weights, and disconnected comparison graphs.
- **Training Diagnostics:** Learned rankers include dataset summary metadata, the effective victory-strength mode, and the accepted lock graph.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.
- **Shared Condorcet Infrastructure:** Reuses the shared direct-victory preprocessing helpers from the `ranking_protocols` library.

6.174.5 Usage

Learning a ranker

```
% Learn from a pairwise ranking dataset object
| ?- ranked_pairs::learn(my_dataset, Ranker).
...

% Learn with custom direct-victory semantics
| ?- ranked_pairs::learn(my_dataset, Ranker, [victory_strength(margins)]).
...

% Learn with custom equal-strength tie breaking
| ?- ranked_pairs::learn(my_dataset, Ranker, [tie_breaking(declaration_order)]).
...
```

Inspecting diagnostics

```
% Inspect the effective options and accepted locks
| ?- ranked_pairs::learn(my_dataset, Ranker),
    ranked_pairs::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
```

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- ranked_pairs::learn(my_dataset, Ranker),
    ranked_pairs::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Inspecting locked pairs

```
% Inspect the accepted lock graph in lock order
| ?- ranked_pairs::learn(my_dataset, Ranker),
    ranked_pairs::locked_pairs(Ranker, LockedPairs).
LockedPairs = [...]
```

The `locked_pairs/2` predicate returns an ordered list of `lock(Winner, Loser, Strength)` terms representing the victories that were accepted during the Ranked Pairs locking phase.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- ranked_pairs::learn(my_dataset, Ranker),
    ranked_pairs::export_to_clauses(my_dataset, Ranker, my_ranker, Clauses).
Clauses = [my_ranker(ranked_pairs_ranker(...))]
...

% Export to a file
| ?- ranked_pairs::learn(my_dataset, Ranker),
    ranked_pairs::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.pl').
...
```

6.174.6 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(ranked_pairs),
  options(Options),
  locked_pairs(LockedPairs),
  dataset_summary(DatasetSummary)
]
```

Where:

- `model(ranked_pairs)` identifies the learning algorithm that produced the ranker.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `locked_pairs(LockedPairs)` stores the accepted lock graph in the order in which victories were locked.
- `dataset_summary(DatasetSummary)` stores a summary list describing the validated training dataset.

Use the `ranking_protocols` `diagnostic/2` and `ranker_options/2` helper predicates when you only need a single metadata term or the effective options.

6.174.7 Options

The following options can be passed to the `learn/3` predicate:

- `victory_strength(winning_votes)`: Use the winning side's aggregated vote total as the direct victory strength.
- `victory_strength(margins)`: Use the victory margin as the direct victory strength.
- `tie_breaking(term_order)`: Break equal-strength direct-victory ties using the standard term order of the winner and loser item identifiers.
- `tie_breaking(declaration_order)`: Break equal-strength direct-victory ties using the training item declaration order preserved by the pairwise dataset helpers.

The defaults are `victory_strength(winning_votes)` and `tie_breaking(term_order)`.

6.174.8 Scoring semantics

The learned `scores/2` values count how many opponents are reachable from each item in the final locked relation. These scores are therefore integers in the range from 0 to $N-1$, where N is the number of learned items.

This choice keeps the `rank/3` behavior aligned with the locked graph that the algorithm actually accepts: a higher score means that the item precedes more opponents after all accepted locks are closed transitively.

The effective scores can therefore change when `tie_breaking/1` changes the order in which equal-strength victories are considered for locking.

6.174.9 Ranker representation

The learned ranker is represented by a compound term of the form:

```
ranked_pairs_ranker(Items, Scores, Diagnostics)
```

Where:

- Items: List of ranked items.
- Scores: List of integer Item-Score pairs.
- Diagnostics: List of metadata terms, including the effective options, the accepted locked pairs, and the dataset summary.

When exported using `export_to_clauses/4` or `export_to_file/4`, this ranker term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.174.10 See also

For the strongest-path Condorcet-family alternative sharing the same `victory_strength(...)` option, see the `schulze_ranker` library.

6.175 ranking_protocols

This library provides protocols used in the implementation of machine learning ranking algorithms. Rankers are represented as objects implementing the `ranker_protocol` protocol. Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol`, `pairwise_measurement_dataset_protocol`, `temporal_pairwise_ranking_dataset_protocol`, or the `ranking_dataset_protocol` protocol.

This library also provides reusable test datasets and smoke tests for the shared ranking-family contracts.

6.175.1 Protocol requirements

The dataset protocols defined by this library impose the following semantic and validity expectations:

- `pairwise_ranking_dataset_protocol` datasets should declare each item once, use only declared items in preferences, assign positive weights to preferences between distinct items, and may still need stronger conditions such as a strongly connected directed win graph for algorithms like Bradley-Terry that require finite maximum-likelihood estimates.
- `pairwise_measurement_dataset_protocol` datasets should declare each item once, use only declared items in measurements, assign numeric measurement values, and assign positive weights to measurements between distinct items. Each `measurement(Item1, Item2, Value, Weight)` fact denotes a weighted signed scalar observation on the oriented edge `Item1 -> Item2`, where positive values favor `Item1`, negative values favor `Item2`, and zero denotes a neutral observation.
- `temporal_pairwise_ranking_dataset_protocol` datasets should declare each item and period once, use only declared items and periods in games, keep game participants distinct, and restrict scores to the set `{0.0, 0.5, 1.0}`. A fact `game(Period, Item1, Item2, Score)` records the observed result for `Item1` against `Item2`, with the score for `Item2` implicitly equal to `1.0 - Score`.

6.175.2 Shared categories

The library includes a small family of reusable categories intended to be imported by ranking algorithm implementations:

- `ranking_dataset_common` — dataset collection, summaries, graph connectivity, connected-component analysis, and pairwise/grouped dataset correctness checks, including pairwise measurement helpers, temporal pairwise rating-period helpers, and grouped tie-block extraction helpers for algorithms that consume tied rankings directly.
- `glicko2_common` — shared internal Glicko-2 numeric helpers for scale conversions, per-period player updates, and volatility root solving used by both the batch and periodic Glicko-2 rankers.
- `condorcet_victory_common` — shared direct-victory preprocessing helpers for Condorcet-family rankers that derive dense directed victory strengths from aggregated pairwise matchups under the `victory_strength(...)` option semantics.
- `ranker_common` — representation-independent access to learned-ranker diagnostics plus reusable helpers for exporting learned rankers.
- `grouped_strength_ranker_common` — reusable positive-strength, strong-connectivity, and iterative-update helpers for grouped ranking models that estimate one latent strength parameter per item.

These categories are designed to keep ranking implementations compact while keeping the shared protocol-facing behavior reusable.

6.175.3 Diagnostics

The `ranker_common` category provides shared accessor predicates such as `diagnostics/2`, `diagnostic/2`, and `ranker_options/2`, together with the shared `export_to_file/4` export helper. These predicates make it possible to inspect and export learned rankers without depending on the exact term representation used by a particular ranking algorithm implementation.

The detailed contents of the diagnostics data are ranking algorithm implementation dependent. For example, one ranker may report convergence status, iteration count, and dataset summaries, while another may report a different set of metadata terms or only a subset of those details. When using diagnostics in application code, rely on the shared access predicates and the documentation of the specific ranking algorithm you are using.

6.175.4 Export header format

The shared ranker exporter in the `ranker_common` category writes a header before the exported clauses in the following format:

```
% exported ranker predicate: Functor/Arity
% training dataset: Dataset
% diagnostics: Diagnostics
% Functor(Ranker)
Functor(Ranker)
```

When exporting a serialized ranker term, using a noun such as `ranker/1` or `model/1` is recommended. The `Ranker` argument can then be passed to the rank predicates.

6.175.5 API documentation

Open the ../apis/library_index.html#ranking_protocols link in a web browser.

6.175.6 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ranking_protocols(loader)).
```

6.175.7 Testing

To test this library predicates and datasets, load the `tester.lgt` file:

```
| ?- logtalk_load(ranking_protocols(tester)).
```

6.175.8 Test datasets

Several sample datasets are included in the `test_datasets` directory:

- `head_to_head.lgt` — A compact pairwise-comparison dataset with four items and weighted preferences suitable for smoke testing deterministic ranking.
- `search_results.lgt` — A grouped ranking dataset with two query groups, three items per group, and non-negative integer relevance judgments.
- `malformed_pairwise.lgt` — A negative fixture where a preference mentions an undeclared item.
- `malformed_duplicate_items.lgt` — A negative pairwise fixture where an item is declared more than once.
- `malformed_self_preference.lgt` — A negative pairwise fixture where an item is preferred over itself.
- `malformed_non_positive_weight.lgt` — A negative pairwise fixture where a preference weight is not positive.
- `disconnected_pairwise.lgt` — A pairwise fixture with more than one connected component, useful for testing algorithms that require identifiable global scores.
- `cyclic_pairwise.lgt` — A connected pairwise fixture with a preference cycle, useful for smoke testing algorithms on non-transitive data.
- `condorcet_divergence_pairwise.lgt` — A compact connected pairwise fixture where the current `schulze_ranker` and `ranked_pairs` implementations produce different rankings, useful for cross-method regression tests.
- `malformed_grouped.lgt` — A negative fixture where a grouped relevance value is not a non-negative integer.
- `sparse_preferences.lgt` — A sparse pairwise dataset with an isolated item, useful for testing dataset summaries and disconnected-graph detection.
- `two_item_measurements.lgt` — A compact pairwise measurement dataset with two items and one signed measurement, useful for smoke testing exact two-item least-squares fits.

- `regular_measurements.lgt` — A compact connected pairwise measurement dataset whose measurements are perfectly explained by a zero-sum score potential, useful for regression tests with zero residuals.
- `cyclic_measurements.lgt` — A compact connected pairwise measurement dataset with cyclic inconsistency, useful for smoke testing zero global scores with non-zero residuals.
- `disconnected_measurements.lgt` — A pairwise measurement fixture with more than one connected component, useful for testing identifiable-score checks.
- `malformed_measurement_unknown_item.lgt` — A negative measurement fixture where a measurement mentions an undeclared item.
- `malformed_measurement_duplicate_items.lgt` — A negative measurement fixture where an item is declared more than once.
- `malformed_measurement_self.lgt` — A negative measurement fixture where an item is measured against itself.
- `malformed_measurement_non_numeric.lgt` — A negative measurement fixture where a measurement value is not numeric.
- `malformed_measurement_non_positive_weight.lgt` — A negative measurement fixture where a measurement weight is not positive.
- `temporal_two_period_chain.lgt` — A compact temporal pairwise dataset with two rating periods and three players, useful for smoke testing periodic rating carryover.
- `temporal_draws.lgt` — A compact temporal pairwise dataset with a draw, useful for testing score handling in temporal game results.
- `temporal_idle_periods.lgt` — A temporal pairwise dataset with an empty period between two played periods, useful for testing inactivity handling.

6.176 reader

The reader object provides portable predicates for reading text file and text stream contents to lists of terms, characters, or character codes and for reading binary files to lists of bytes. The text file API is loosely based on the SWI-Prolog `readutil` module.

6.176.1 API documentation

Open the ../apis/library_index.html#reader link in a web browser.

6.176.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(reader(loader)).
```

6.176.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(reader(tester)).
```

6.177 recorded_database

The `recorded_database` library aims to help port Prolog code using the legacy recorded database. Ported applications should still consider migrating to more standard solutions to handle dynamic data that must survive backtracking.

6.177.1 API documentation

Open the ../apis/library_index.html#recorded_database link in a web browser.

6.177.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(recorded_database(loader)).
```

6.177.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(recorded_database(tester)).
```

6.177.4 Usage

The `recorded_database_core` category implements the library predicates. This category is imported by the default `recorded_database` object to provide application global database. To make the database local and thus minimize potential record clashes, the category can be imported by one or more application objects. Use protected or private import to restrict the scope of the library predicates. For example:

```
:- object(foo,
    imports(private::recorded_database_core)).

    bar :-
        ^^recorda(key, value(1)),
        ...

:- end_object.
```

6.177.5 Known issues

Currently, references are non-negative integers. They still must be regarded as opaque terms and subject to change without notice. But using integers can result in integer overflows when running on backends with bounded integers in applications performing a large number of database updates.

6.178 redis

Redis client library. Supports ECLiPSe, GNU Prolog, SICStus Prolog, SWI-Prolog, Trealla Prolog, and XVM.

For general information on Redis, including a list of the available commands, visit:

```
https://redis.io
```

6.178.1 API documentation

Open the `../..apis/library_index.html#redis` link in a web browser.

6.178.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(redis(loader)).
```

6.178.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(redis(tester)).
```

The tests assume a localhost Redis server running on the default port (6379) if the `REDIS_HOST` and `REDIS_PORT` environment variables are not defined. If the server is not detected, the tests are skipped.

The unit tests were originally written by Sean Charles for his GNU Prolog Redis client library:

```
https://github.com/emacstheviking/gnuprolog-redisclient
```

The Logtalk version is a straight-forward port of the original tests using the `test/1` dialect of `lgtunit`.

6.178.4 Supported Redis Features

This library provides wrapper predicates for commonly used Redis operations across multiple data types. For Redis commands not covered by a wrapper, use the generic `redis::send/3` predicate with the command as a compound term.

Connection Management

- connect/1 - Connect to localhost on default port (6379)
- connect/3 - Connect to specified host and port
- disconnect/1 - Disconnect from Redis server
- send/3 - Send any Redis command and receive reply

String Operations

- get/3 - Get the value of a key
- set/4 - Set the value of a key
- append/4 - Append a value to a key
- getrange/5 - Get substring of string stored at key
- setrange/5 - Overwrite part of a string at key starting at offset
- strlen/3 - Get the length of the value stored at key
- mget/3 - Get values of multiple keys
- mset/3 - Set multiple key-value pairs atomically
- incr/3 - Increment integer value of key by one
- decr/3 - Decrement integer value of key by one
- incrby/4 - Increment integer value of key by amount
- decrby/4 - Decrement integer value of key by amount

Key Operations

- del/3 - Delete a key
- exists/3 - Check if a key exists
- keys/3 - Find all keys matching a pattern
- ttl/3 - Get time to live for a key in seconds
- expire/4 - Set a timeout on a key in seconds
- persist/3 - Remove the expiration from a key
- rename/4 - Rename a key
- type/3 - Get the type of value stored at key

Hash Operations

Hashes are maps between string fields and string values, ideal for representing objects.

- `hset/5` - Set field in a hash
- `hget/4` - Get value of field in a hash
- `hgetall/3` - Get all fields and values in a hash
- `hdel/4` - Delete field from a hash
- `hexists/4` - Check if field exists in a hash
- `hkeys/3` - Get all field names in a hash
- `hvals/3` - Get all values in a hash
- `hlen/3` - Get number of fields in a hash

List Operations

Redis lists are ordered collections of strings, sorted by insertion order.

- `lpush/4` - Prepend value to a list
- `rpush/4` - Append value to a list
- `lpop/3` - Remove and return first element of list
- `rpop/3` - Remove and return last element of list
- `lrange/5` - Get range of elements from list
- `llen/3` - Get length of list
- `lrem/5` - Remove elements from list
- `ltrim/5` - Trim list to specified range

Set Operations

Redis sets are unordered collections of unique strings.

- `sadd/4` - Add member to a set
- `srem/4` - Remove member from a set
- `smembers/3` - Get all members in a set
- `sismember/4` - Check if value is member of set
- `scard/3` - Get number of members in a set

Sorted Set Operations

Sorted sets are collections of unique strings (members) ordered by an associated score. Members are unique, but scores may repeat.

- `zadd/5` - Add member with score to sorted set
- `zrem/4` - Remove member from sorted set
- `zrange/5` - Get range of members from sorted set by index
- `zrank/4` - Get rank (index) of member in sorted set
- `zcard/3` - Get number of members in sorted set
- `zscore/4` - Get score of member in sorted set

Usage Examples

```
% Connect and perform basic string operations
?- redis::connect(Connection),
   redis::set(Connection, mykey, 'Hello World', Status),
   redis::get(Connection, mykey, Value),
   redis::disconnect(Connection).
Status = 'OK',
Value = 'Hello World'.

% Working with hashes
?- redis::connect(Connection),
   redis::hset(Connection, user:1000, name, 'John Doe', _),
   redis::hset(Connection, user:1000, email, 'john@example.com', _),
   redis::hgetall(Connection, user:1000, Fields),
   redis::disconnect(Connection).
Fields = [bulk(name), bulk('John Doe'), bulk(email), bulk('john@example.com')].

% Using lists as queues
?- redis::connect(Connection),
   redis::rpush(Connection, queue, task1, _),
   redis::rpush(Connection, queue, task2, _),
   redis::lpop(Connection, queue, Task),
   redis::disconnect(Connection).
Task = task1.

% Batch operations with mget/mset
?- redis::connect(Connection),
   redis::mset(Connection, [key1, val1, key2, val2], Status),
   redis::mget(Connection, [key1, key2], Values),
   redis::disconnect(Connection).
Status = 'OK',
Values = [bulk(val1), bulk(val2)].

% Key expiration
?- redis::connect(Connection),
   redis::set(Connection, session:xyz, 'user_data', _),
   redis::expire(Connection, session:xyz, 3600, _), % Expire in 1 hour
```

(continues on next page)

(continued from previous page)

```
redis::ttl(Connection, session:xyz, TTL),
redis::disconnect(Connection).
TTL = 3600.

% Using send/3 for commands without wrappers
?- redis::connect(Connection),
   redis::send(Connection, info(server), Reply),
   redis::disconnect(Connection).
```

6.178.5 Credits

This library is inspired by the Sean Charles GNU Prolog Redis client library.

6.178.6 Known issues

Recent versions of macOS seem to disable the mapping of localhost to 127.0.0.1. This issue may prevent running this library unit tests and the `redis::connect/1` predicate from working. This can be fixed either by editing the `/etc/hosts` file or by using in alternative the predicate `redis::connect/3` with `'127.0.0.1'` as the first argument.

6.179 regression_protocols

This library provides protocols used in the implementation of machine learning regression algorithms. Datasets are represented as objects implementing the `regression_dataset_protocol` protocol. Regressors are represented as objects importing the `regressor_common` category. This category provides shared helpers for regressor defaults, dataset validation, diagnostics metadata, export, and pretty-printing support.

Learned regressors expose diagnostics using the shared `diagnostics/2`, `diagnostic/2`, and `regressor_options/2` predicates. Concrete regressor implementations store effective training options in the diagnostics metadata under an `options(Options)` term.

This library also provides regression test datasets under the `test_datasets` directory.

6.179.1 API documentation

Open the ../apis/library_index.html#regression_protocols link in a web browser.

6.179.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(regression_protocols(loader)).
```

6.179.3 Testing

To test this library predicates, shared categories, and datasets, load the `tester.lgt` file:

```
| ?- logtalk_load(regression_protocols(tester)).
```

6.179.4 Test datasets

The `test_datasets` directory includes the following compact regression datasets and validation fixtures:

- `duplicate_attribute_declaration.lgt`: invalid dataset fixture containing a duplicated attribute declaration in the dataset schema.
- `duplicate_attribute_example.lgt`: invalid dataset fixture containing a duplicated declared attribute binding in one example.
- `grouped_categorical_signal.lgt`: regression dataset with one relevant continuous attribute and one irrelevant categorical attribute for testing shrinkage of encoded categorical coefficients.
- `intercept_only.lgt`: constant-target dataset for intercept-learning tests.
- `invalid_target.lgt`: invalid dataset with a non-numeric target for negative tests.
- `mixed_signal.lgt`: mixed continuous and categorical regression dataset.
- `plane.lgt`: two-feature regression dataset following the plane $z = 3x_1 - 2x_2 + 5$.
- `simple_line.lgt`: single-feature regression dataset following the line $y = 2x + 1$.
- `sparse_mixed_signal.lgt`: mixed regression dataset with omitted attribute-value pairs used to exercise missing-value handling during training and prediction.
- `sparse_signal.lgt`: two-feature regression dataset where only the signal attribute contributes to the target and the noise attribute is orthogonal to it.
- `step_signal.lgt`: piecewise-constant regression dataset for tree-based and neighborhood regressor testing.
- `undeclared_attribute_example.lgt`: invalid dataset fixture containing an undeclared attribute binding in one example.
- `wide_mixed_signal.lgt`: synthetic wide mixed regression dataset with many continuous and categorical predictors for non-trivial linear-model benchmarking.

6.180 regression_tree

Regression tree regressor supporting continuous and mixed-feature datasets. The library implements the `regressor_protocol` defined in the `regression_protocols` library and learns a binary regression tree using recursive variance-reduction splits that select the encoded feature threshold maximizing variance reduction.

6.180.1 API documentation

Open the `../apis/library_index.html#regression_tree` link in a web browser.

6.180.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(regression_tree(loader)).
```

6.180.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(regression_tree(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(regression_tree(tester_performance)).
```

6.180.4 Export header format

The shared exporter in the `regressor_common` category writes a header before the exported clauses in the following format:

```
% exported regressor predicate: Functor/Arity
% training dataset: Dataset
% target: Target
% attributes: Attributes
% diagnostics: Diagnostics
% Functor(Encoders, FeatureLabels, Tree, Diagnostics)
Functor(Encoders, FeatureLabels, Tree, Diagnostics)
```

The exported clauses serialize the learned regressor state so that loading the file gives a regressor term that can be passed directly to the `predict/3` predicate.

When exporting a serialized regressor term, using a noun such as `regressor/4` or `model/4` is recommended.

6.180.5 Features

- **Variance-Reduction Splits:** Selects binary thresholds over encoded features to reduce target variance.
- **Continuous and Mixed Features:** Supports continuous attributes and categorical attributes.
- **Categorical Features Encoding:** Uses reference-level dummy coding derived from the declared dataset attribute values, with a missing-value indicator, and the resulting encoded features are treated as ordinary numeric split features.
- **Missing Values:** Missing feature values represented using anonymous variables or omitted attribute-value pairs are encoded using explicit missing-value indicator features during both training and prediction.

- **Per-Split Feature Sampling:** Optionally samples a subset of dataset attributes at each split before searching for the best partition.
- **Optional Feature Scaling:** Continuous attributes can be standardized using z-score scaling before tree induction.
- **Diagnostics Metadata:** Learned regressors record model name, target, training example count, encoded feature count, and effective options, accessible using the shared regression diagnostics predicates.
- **Model Export:** Learned regressors can be exported as predicate clauses or written to a file.
- **Readable Trees:** Includes a pretty-printer for inspecting learned tree structure.
- **Reference Benchmarks:** Includes a dedicated performance suite reporting training time, RMSE, and MAE for representative regression datasets.

6.180.6 Regressor representation

The learned regressor is represented by default as:

- `regression_tree_regressor(Encoders, FeatureLabels, Tree, Diagnostics)`

In this representation, `Tree` is built from `leaf(Prediction)` and `node(Index, Threshold, FallbackPrediction, Left, Right)` terms and `Diagnostics` stores training metadata including the effective options.

6.180.7 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(regression_tree),
  target(Target),
  training_example_count(TrainingExampleCount),
  options(Options),
  encoded_feature_count(FeatureCount)
]
```

Where:

- `model(regression_tree)` identifies the learning algorithm that produced the regressor.
- `target(Target)` stores the target attribute name declared by the training dataset.
- `training_example_count(TrainingExampleCount)` stores the number of examples used during training.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `encoded_feature_count(FeatureCount)` stores the number of numeric features induced by the encoder list, including missing-value indicator features.

Use the `regression_protocols` `diagnostic/2` and `regressor_options/2` helper predicates when you only need a single metadata term or the effective options.

6.180.8 Options

The `learn/3` predicate accepts the following options:

- `maximum_depth/1`: Maximum depth allowed for the induced regression tree. Lower values yield smaller trees; higher values allow more detailed partitioning of the training data. The default is 10.
- `minimum_samples_leaf/1`: Minimum number of training examples required in a leaf. This option also prevents candidate splits that would create child nodes smaller than the requested size. The default is 1.
- `minimum_variance_reduction/1`: Minimum variance-reduction gain required for accepting a split. Higher values make the learner more conservative by pruning weak splits during induction. The default is 0.0.
- `maximum_features_per_split/1`: Number of dataset attributes sampled at each split when searching for the best partition. Accepted values are a positive integer or `all`. The default is `all`.
- `feature_scaling/1`: Controls z-score standardization of continuous attributes before tree induction. Accepted values are `true` and `false`. The default is `false`.

6.181 `regularized_bradley_tery_ranker`

Regularized Bradley-Terry MAP pairwise preference ranker. It uses a deterministic MM-style posterior-mode update for a Bradley-Terry likelihood regularized by an explicit independent Gamma prior over item strengths.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from pairwise preferences, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The training dataset must declare each ranked item once, enumerate positive-weight pairwise preferences between distinct declared items, and induce a connected undirected comparison graph. Unlike the unregularized Bradley-Terry model, the directed win graph is not required to be strongly connected.

6.181.1 API documentation

Open the ../apis/library_index.html#regularized_bradley_tery_ranker link in a web browser.

6.181.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(regularized_bradley_tery_ranker(loader)).
```

6.181.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(regularized_bradley_terry_ranker(tester)).
```

6.181.4 Features

- **Pairwise Preference Learning:** Learns positive item strengths from weighted head-to-head outcomes.
- **Explicit Gamma Prior:** Uses a separate MAP model with an explicit independent Gamma prior over item strengths.
- **Finite Regularized Fit:** Admits connected pairwise datasets whose directed win graph is not strongly connected, instead of rejecting them as non-identifiable maximum-likelihood problems.
- **Deterministic Ranking:** Orders candidate items by learned strength with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items, self-preferences, non-positive weights, and disconnected undirected comparison graphs.
- **Training Diagnostics:** Learned rankers include the effective Gamma prior, convergence, iteration, and dataset summary metadata that can be accessed using the `diagnostics/2` predicate.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.
- **Shared Pairwise Infrastructure:** Reuses the shared pairwise preprocessing and MM-iteration scaffolding from the `ranking_protocols` library.

6.181.5 Dataset requirements

This implementation still requires the undirected comparison graph induced by the preferences to be connected, but unlike the unregularized Bradley-Terry model it does not require the directed win graph to be strongly connected. The explicit Gamma prior regularizes dominance partitions and one-way chains, yielding a finite posterior mode for connected datasets that would otherwise fail the maximum-likelihood existence check.

For the original unregularized maximum-likelihood Bradley-Terry model, which keeps the same pairwise-preference interface but requires the directed win graph to be strongly connected, see the `bradley_terry_ranker` library.

6.181.6 Usage

Learning a ranker

```
% Learn from a pairwise ranking dataset object
| ?- regularized_bradley_terry_ranker::learn(my_dataset, Ranker).
...

% Learn with custom iteration and Gamma-prior options
| ?- regularized_bradley_terry_ranker::learn(my_dataset, Ranker, [maximum_iterations(500),
  ↪tolerance(1.0e-7), gamma_prior(gamma(3.0, 4.0))]).
...
```

Inspecting diagnostics

```
% Inspect convergence, prior, and dataset summary metadata
| ?- regularized_bradley_terry_ranker::learn(my_dataset, Ranker),
      regularized_bradley_terry_ranker::diagnostics(Ranker, Diagnostics).
Diagnostics = [...]
...
```

6.181.7 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(regularized_bradley_terry_ranker),
  options(Options),
  prior(gamma(Shape, Rate)),
  convergence(Status),
  iterations(Iterations),
  final_delta(FinalDelta),
  dataset_summary(DatasetSummary)
]
```

Where:

- `model(regularized_bradley_terry_ranker)` identifies the learning algorithm that produced the ranker.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `prior(gamma(Shape, Rate))` stores the effective Gamma prior hyperparameters.
- `convergence(Status)` records the training stop condition. The current values are converged and maximum_iterations_exhausted.
- `iterations(Iterations)` stores the number of update iterations that were executed.
- `final_delta(FinalDelta)` stores the maximum absolute strength update in the last iteration.
- `dataset_summary(DatasetSummary)` stores a summary list describing the validated training dataset.

Use the `ranking_protocols` `diagnostic/2` and `ranker_options/2` helper predicates when you only need a single metadata term or the effective options.

Ranking candidate items

```
% Rank a candidate set from most preferred to least preferred
| ?- regularized_bradley_terry_ranker::learn(my_dataset, Ranker),
      regularized_bradley_terry_ranker::rank(Ranker, [item_a, item_b, item_c], Ranking).
Ranking = [...]
...
```

Candidate lists must be proper lists of unique, ground items declared by the training dataset. Invalid ranker terms, duplicate candidates, and candidates containing variables are rejected with errors instead of being silently accepted.

Exporting the ranker

Learned rankers can be exported as a list of clauses or to a file for later use.

```
% Export as predicate clauses
| ?- regularized_bradley_terry_ranker::learn(my_dataset, Ranker),
    regularized_bradley_terry_ranker::export_to_clauses(my_dataset, Ranker, my_ranker,
    ↪Clauses).
Clauses = [my_ranker(regularized_bt_ranker(...))]
...

% Export to a file
| ?- regularized_bradley_terry_ranker::learn(my_dataset, Ranker),
    regularized_bradley_terry_ranker::export_to_file(my_dataset, Ranker, my_ranker, 'ranker.
    ↪pl').
...
```

6.181.8 Options

The following options can be passed to the learn/3 predicate:

- `maximum_iterations(MaximumIterations)`: Positive integer iteration bound.
- `tolerance(Tolerance)`: Positive convergence tolerance.
- `gamma_prior(gamma(Shape, Rate))`: Gamma-prior hyperparameters. The current implementation requires `Shape > 1` and `Rate > 0`.

6.181.9 Ranker representation

The learned ranker is represented by a compound term of the form:

```
regularized_bt_ranker(Items, Strengths, Diagnostics)
```

Where:

- `Items`: List of ranked items.
- `Scores`: List of Item-Strength pairs.
- `Diagnostics`: List of metadata terms, including the effective options, Gamma prior, convergence status, iteration count, final update delta, and dataset summary.

When exported using `export_to_clauses/4` or `export_to_file/4`, this ranker term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.181.10 References

1. Bradley, R. A. and Terry, M. E. (1952). Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, 39(3/4), 324-345.
2. Hunter, D. R. (2004). MM algorithms for generalized Bradley-Terry models. *The Annals of Statistics*, 32(1), 384-406.
3. Caron, F. and Doucet, A. (2012). Efficient Bayesian inference for generalized Bradley-Terry models. *Journal of Computational and Graphical Statistics*, 21(1), 174-196.

6.182 ridge_regression

Ridge regression regressor supporting continuous and mixed-feature datasets. The library implements the `regressor_protocol` defined in the `regression_protocols` library and learns a linear model by solving the weighted ridge normal equations directly via the shared regression encoding core in `regressor_common`, leaving the intercept unpenalized while penalizing encoded feature columns using scale-aware weights that match standardizing penalized columns before applying the L2 penalty.

6.182.1 API documentation

Open the ../apis/library_index.html#ridge_regression link in a web browser.

6.182.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ridge_regression(loader)).
```

6.182.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ridge_regression(tester)).
```

To run the performance benchmark suite, load the `tester_performance.lgt` file:

```
| ?- logtalk_load(ridge_regression(tester_performance)).
```

6.182.4 Features

- **Continuous and Mixed Features:** Supports continuous attributes and categorical attributes encoded using reference-level dummy coding.
- **Feature Scaling and Penalty Scaling:** Continuous attributes can be standardized using z-score scaling. Ridge regularization uses scale-aware weights equivalent to standardizing each penalized encoded feature column before applying the L2 penalty.
- **Missing Values:** Missing numeric and categorical values represented using anonymous variables are encoded using explicit missing-value indicator features.

- **Unknown Values:** Prediction requests containing categorical values that are not declared by the dataset raise a domain error.
- **Zero-Variance Features:** Encoded columns with zero variance are excluded from the direct solve and assigned zero coefficients in the learned regressor.
- **Ridge Penalty:** Applies L2 regularization to the learned weights using the shared `regularization/1` option.
- **Diagnostics Metadata:** Learned regressors record model name, target, training example count, solver, linear-system residual, active feature count, penalty scaling strategy, encoded feature count, and effective options, accessible using the shared regression diagnostics predicates.
- **Model Export:** Learned regressors can be exported as predicate clauses or written to a file.
- **Reference Benchmarks:** Includes a dedicated performance suite reporting training time, RMSE, and MAE for representative regression datasets.

6.182.5 Regressor representation

The learned regressor is represented by default as:

- `ridge_regressor(Encoders, Bias, Weights, Diagnostics)`

The exported predicate clauses therefore use the shape:

- `Functor(Encoders, Bias, Weights, Diagnostics)`

6.182.6 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(ridge_regression),
  target(Target),
  training_example_count(TrainingExampleCount),
  options(Options),
  solver(Solver),
  linear_system_residual(Residual),
  active_feature_count(ActiveFeatureCount),
  penalty_scaling(encoded_feature_standardization),
  encoded_feature_count(FeatureCount)
]
```

Where:

- `model(ridge_regression)` identifies the learning algorithm that produced the regressor.
- `target(Target)` stores the target attribute name declared by the training dataset.
- `training_example_count(TrainingExampleCount)` stores the number of examples used during training.
- `options(Options)` stores the effective learning options after merging the user options with the library defaults.
- `solver(Solver)` records the direct linear-system solver used during training. The current value is `pivoted_gaussian_elimination`.

- `linear_system_residual(Residual)` stores the maximum absolute residual of the solved ridge linear system.
- `active_feature_count(ActiveFeatureCount)` stores the number of encoded feature columns retained for the direct solve after dropping zero-variance columns.
- `penalty_scaling(encoded_feature_standardization)` records that the ridge penalty is scaled as if each penalized encoded feature column had been standardized before applying the L2 penalty.
- `encoded_feature_count(FeatureCount)` stores the number of numeric features induced by the encoder list, including missing-value indicator features.

Use the `regression_protocols diagnostic/2` and `regressor_options/2` helper predicates when you only need a single metadata term or the effective options.

6.182.7 Options

The `learn/3` predicate accepts the following options:

- `regularization/1`: Ridge penalty coefficient applied to the weight vector during the direct solve. Higher values increase shrinkage and can reduce overfitting. The default is `0.01`.
- `feature_scaling/1`: Controls z-score standardization of continuous attributes before training and prediction. Accepted values are `true` and `false`. The default is `true`.

6.183 schulze_ranker

Schulze pairwise preference ranker. It builds the direct pairwise strength graph from aggregated matchups and applies the Schulze strongest-path dynamic program to derive the final pairwise preference relation.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from pairwise preferences, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The current implementation requires a well-formed connected pairwise dataset so that all ranked items remain comparable in the aggregated matchup graph.

6.183.1 API documentation

Open the ../apis/library_index.html#schulze_ranker link in a web browser.

6.183.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(schulze_ranker(loader)).
```

6.183.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(schulze_ranker(tester)).
```

6.183.4 Features

- **Pairwise Preference Learning:** Learns one deterministic score per item from aggregated pairwise outcomes.
- **Schulze Strongest Paths:** Computes the final ranking relation using the Schulze strongest-path dynamic program over aggregated head-to-head data.
- **Configurable Direct Edge Semantics:** Supports both winning-votes and victory-margins victory-strength modes. The `victory_strength/1` option selects whether direct edges use winning votes or victory margins before strongest-path propagation.
- **Strongest-Path Access:** Exposes the labeled strongest-path relation for ordered item pairs using the `strongest_paths/2` predicate.
- **Deterministic Ranking:** Orders candidate items by final Schulze relation win counts with deterministic tie-breaking.
- **Tie-breaking:** Items tied in the final Schulze relation receive the same learned score, and ranking ties are then broken deterministically using the standard term order of the item identifiers.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items, self-preferences, non-positive weights, and disconnected comparison graphs.
- **Training Diagnostics:** Learned rankers include dataset summary metadata and the effective victory-strength mode.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.

6.183.5 Options

The following options can be passed to the `learn/3` predicate:

- `victory_strength(winning_votes)`: Use the winning side's aggregated vote total as the direct edge strength.
- `victory_strength(margins)`: Use the victory margin as the direct edge strength.

The default is `victory_strength(winning_votes)`.

6.183.6 Strongest paths

The `strongest_paths/2` predicate returns the learned strongest-path matrix as an ordered list of `path(Item1, Item2, Strength)` terms for all ordered pairs of distinct learned items.

This predicate complements `scores/2` and `rank/3` without changing their semantics. The `scores/2` predicate still returns the number of opponents that each item beats in the final Schulze relation, while `strongest_paths/2` exposes the underlying pairwise path strengths used to derive that relation.

6.183.7 Diagnostics syntax

The `diagnostics/2` predicate returns a list of metadata terms with the form:

```
[
  model(schulze_ranker),
  options(Options),
  strongest_paths(StrongestPaths),
  dataset_summary(DatasetSummary)
]
```

6.183.8 Ranker representation

The learned ranker is represented by a compound term of the form:

```
schulze_ranker(Items, Scores, Diagnostics)
```

Where:

- Items: List of ranked items.
- Scores: List of Item-Score pairs.
- Diagnostics: List of metadata terms, including the effective `victory_strength/1` option, the labeled strongest paths, and the dataset summary.

6.184 sequential_pattern_mining_protocols

This library provides support entities for sequential pattern mining algorithms. Ordered sequence datasets are represented as objects implementing the `sequence_dataset_protocol` protocol. The generic `pattern_miner_protocol` protocol and the `pattern_miner_common` category used by concrete miners are loaded from the `pattern_mining_protocols` core library.

The `sequential_pattern_mining_common` category builds on that generic core with sequential-specific helpers for dataset validation, support count accumulation, and sequential pattern ordering/filtering.

This library also provides reusable sequence smoke-test datasets and a small smoke-test suite.

6.184.1 API documentation

Open the ../apis/library_index.html#sequential_pattern_mining_protocols link in a web browser.

6.184.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(sequential_pattern_mining_protocols(loader)).
```

6.184.3 Testing

To run the library smoke tests, load the `tester.lgt` file:

```
| ?- logtalk_load(sequential_pattern_mining_protocols(tester)).
```

6.184.4 Test datasets

The `test_datasets` directory includes the following sample sequence datasets:

- `clickstream_sequences.lgt`: A compact sequence dataset with repeated prefixes intended for sequential-pattern smoke tests.
- `prefix_ladder_sequences.lgt`: A small ladder-shaped dataset of singleton events intended for exact baseline checks across sequential mining algorithms.
- `same_event_vs_next_event_sequences.lgt`: A compact dataset intended to distinguish same-event extensions from next-event extensions.
- `repeated_embedding_sequences.lgt`: A dataset where the same subsequence admits multiple embeddings inside a single sequence, intended for support-count semantics checks.
- `border_threshold_sequences.lgt`: A compact dataset with patterns just above and below typical support thresholds, intended for pruning and threshold regression tests.
- `closure_sequences.lgt`: A compact dataset intended for closed-pattern tests where some frequent patterns share the same support as one of their supersequences.
- `dense_overlap_sequences.lgt`: A denser dataset with overlapping subsequences and mixed singleton and multi-item events, intended for overlap-heavy mining scenarios.
- `branching_sequences.lgt`: A dataset with a common prefix and several competing branches, intended for candidate-generation and branching coverage.

The directory also includes invalid fixtures useful for validation and error-handling tests:

- `invalid_undeclared_item_sequences.lgt`: Uses an item not listed in the declared item domain.
- `invalid_unsorted_itemset_sequences.lgt`: Uses an event with items not in canonical sorted order.
- `invalid_duplicate_item_in_event_sequences.lgt`: Uses an event with a duplicate item.
- `invalid_empty_event_sequences.lgt`: Uses an empty event inside a sequence.

6.185 sets

This library provides a set protocol, two implementations of this protocol using *ordered lists* (one of them a parametric object that takes the type of the set elements as a parameter), and an implementation using *treaps* (tree heaps).

The set representations should be regarded as *opaque terms* and only constructed, accessed, and updated them using the library predicates.

For small sets, the ordered list implementations are likely to provide the best performance. For larger sets, the treap implementation likely provides better performance, notably for the `memberchk/2`, `insert/3`, and `delete/3` operations. Benchmark both implementations to select the best one for your application.

The current implementations use `==/2` for element comparison and standard term ordering. This allows non-ground set elements. But requires caution with later unifications with output arguments and when using the

member/2 and select/3 predicates, which can break the ordered representation. Note also that, per the ISO Prolog Core Standard, variable ordering is implementation dependent. This can result in unexpected results and portability issues.

6.185.1 API documentation

Open the ../apis/library_index.html#sets link in a web browser.

6.185.2 Loading

To load all entities in this library, load the loader.lgt file:

```
| ?- logtalk_load(sets(loader)).
```

6.185.3 Testing

To test this library predicates, load the tester.lgt file:

```
| ?- logtalk_load(sets(tester)).
```

6.185.4 Usage

First, select a set implementation. Use the set(Type) object if you want to type-check the set elements. Otherwise, use the set object.

To create a new set, you can use the new/1 predicate. For example:

```
| ?- set::new(Set).
Set = []
yes
```

You can also create a new set with all unique elements from a list of terms by using the as_set/2 predicate. For example:

```
| ?- set::as_set([1,3,2,1,2], Set).
Set = [1, 2, 3]
yes
```

Predicates are provided for the most common set operations. For example:

```
| ?- set::(
    as_set([1,3,2,1,2], Set1),
    as_set([7,4,2,5,1], Set2),
    intersection(Set1, Set2, Intersection),
    symdiff(Set1, Set2, Difference)
).
Set1 = [1, 2, 3],
Set2 = [1, 2, 4, 5, 7],
Intersection = [1, 2],
Difference = [3, 4, 5, 7]
yes
```

When working with a custom type of set elements and the ordered list representation, the corresponding object must implement the `comparingp` protocol. For example:

```
:- object(rainbow_colors,
    implements(comparingp)).

    order(red,    1).
    order(orange, 2).
    order(yellow, 3).
    order(green,  4).
    order(blue,   5).
    order(indigo, 6).
    order(violet, 7).

    Color1 < Color2 :-
        order(Color1, N1),
        order(Color2, N2),
        {N1 < N2}.

    Color1 =< Color2 :-
        order(Color1, N1),
        order(Color2, N2),
        {N1 =< N2}.

    ...

:- end_object.
```

We can then use this object with the `set/1` parametric object. For example:

```
| ?- set(rainbow_colors)::as_set([blue, yellow, violet], Set).
Set = [yellow, blue, violet]
yes
```

For details on these and other provided predicates, consult the library API documentation.

6.185.5 Credits

Some predicates adapted from code authored by Richard O’Keefe.

6.186 `sgd_classifier`

Linear stochastic-gradient classifier for tabular datasets using a one-vs-rest scheme with configurable losses. The implementation reuses the shared linear encoder pipeline so continuous and categorical features, including missing values, are represented consistently with the existing linear classifiers.

The library implements the `classifier_protocol` defined in the `classification_protocols` library. It provides predicates for learning a classifier from a dataset, using it to make predictions, estimating class probabilities, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `dataset_protocol` protocol from the `classification_protocols` library. Continuous, categorical, and mixed-feature datasets are supported.

6.186.1 API documentation

Open the [.././docs/library_index.html#sgd_classifier](https://docs.logtalk.org/docs/library_index.html#sgd_classifier) link in a web browser.

6.186.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(sgd_classifier(loader)).
```

6.186.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(sgd_classifier(tester)).
```

6.186.4 Features

- **Binary and Multiclass Classification:** Learns one-vs-rest linear models and predicts the class with the highest decision score.
- **Multiple Losses:** Supports `log_loss`, `hinge`, `squared_hinge`, `modified_huber`, and `perceptron` losses.
- **Mixed Features:** Reuses the shared tabular encoders for continuous and categorical attributes, including missing-value indicators.
- **Probability Estimation:** Provides class probabilities using a softmax over linear decision scores.
- **Configurable Optimization:** Exposes learning-rate scheduling, convergence tolerance, and L2 regularization options.
- **Classifier Export:** Learned classifiers can be exported as predicate clauses or written to a file.

6.186.5 Options

The `learn/3` predicate supports these options:

- `loss/1` - optimization loss, one of `log_loss`, `hinge`, `squared_hinge`, `modified_huber`, or `perceptron` (default: `log_loss`)
- `learning_rate/1` - base learning rate for optimization (default: `0.05`)
- `learning_schedule/1` - learning-rate schedule, either `constant` or `inverse_scaling(Power)` (default: `constant`)
- `maximum_iterations/1` - maximum number of optimization epochs (default: `100`)
- `tolerance/1` - convergence threshold for the maximum parameter update (default: `1.0e-5`)
- `l2_regularization/1` - L2 penalty factor applied to the weight vectors (default: `0.0001`)
- `feature_scaling/1` - whether to standardize continuous attributes before encoding (default: `true`)

6.186.6 Usage

Learning a classifier

```
| ?- sgd_classifier::learn(weather, Classifier).
| ?- sgd_classifier::learn(mixed, Classifier, [loss(hinge), maximum_iterations(250)]).
```

Making predictions

```
| ?- sgd_classifier::learn(weather, Classifier),
    sgd_classifier::predict(Classifier, [outlook-rainy, temperature-mild, humidity-normal,
    ↪windy-false], Class).
| ?- sgd_classifier::learn(missing_mixed, Classifier),
    sgd_classifier::predict_probabilities(Classifier, [age-38, income-72000, student-yes,
    ↪credit_rating-fair], Probabilities).
```

Exporting the classifier

```
| ?- sgd_classifier::learn(weather, Classifier),
    sgd_classifier::export_to_clauses(weather, Classifier, classify, Clauses).
| ?- sgd_classifier::learn(weather, Classifier),
    sgd_classifier::export_to_file(weather, Classifier, classify, 'classifier.pl').
```

6.186.7 Classifier representation

The learned classifier is represented as a compound term with the form:

```
sgd_classifier(Classes, Encoders, Loss, Models, Options)
```

Where:

- **Classes:** list of class labels
- **Encoders:** list of continuous scaling descriptors and categorical value encoders
- **Loss:** selected optimization loss
- **Models:** list of `class_model(Class, Bias, Weights)` terms
- **Options:** merged training options used to learn the classifier

When exported using `export_to_clauses/4` or `export_to_file/4`, this classifier term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.186.8 References

1. Bottou, L. (2010). “Large-Scale Machine Learning with Stochastic Gradient Descent”.
2. Shalev-Shwartz, S. and Ben-David, S. (2014). “Understanding Machine Learning”. Chapter 15.
3. Hastie, T., Tibshirani, R. and Friedman, J. (2009). “The Elements of Statistical Learning”. Chapter 12.

6.187 snowflakeid

This library generates Snowflake-style identifiers using a generic generator object and predefined profile objects.

This library requires a backend supporting unbounded integer arithmetic.

See also the `cuid2`, `ids`, `nanoid`, `ksuid`, `uuid`, and `ulid` libraries.

6.187.1 API documentation

Open the ../apis/library_index.html#snowflakeid link in a web browser.

6.187.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(snowflakeid(loader)).
```

6.187.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(snowflakeid(tester)).
```

6.187.4 Usage

The generic generator is the parametric object:

```
snowflakeid(Representation, EpochMilliseconds, TimeUnitMilliseconds, TimestampBits, NodeBits,  
↳ SequenceBits, Node)
```

where `Representation` can be `integer`, `atom`, `chars`, or `codes`.

Predefined profiles are provided as objects extending the generic object:

- `snowflakeid_twitter(_Representation_)`
- `snowflakeid_sonyflake(_Representation_)`
- `snowflakeid_instagram(_Representation_)`

The default object `snowflakeid` uses the Twitter-style profile with `atom` representation.

To generate a Snowflake ID using the default Twitter-style profile:

```
| ?- snowflakeid::generate(ID).
ID = '1917401915609202688'
yes
```

To generate a Twitter-style Snowflake ID represented as an integer:

```
| ?- snowflakeid_twitter(integer)::generate(ID).
ID = 1917401915609202688
yes
```

To generate a Sonyflake-style Snowflake ID represented as chars:

```
| ?- snowflakeid_sonyflake(chars)::generate(ID).
ID = ['9','2','8','5','6','7','2','8','0','1','4','8','9','5','5','7','5','9']
yes
```

To define a custom profile:

```
| ?- snowflakeid(atom, 1704067200000, 1, 41, 10, 12, 7)::generate(ID).
ID = '28842191400263680'
yes
```

6.188 sockets

Portable abstraction over TCP sockets. Provides a high-level API for client and server socket operations that works across all supported backend Prolog systems: ECLiPSe, GNU Prolog, SICStus Prolog, SWI-Prolog, Trealla Prolog, and XVM.

6.188.1 Design rationale

Different Prolog systems provide socket functionality at different abstraction levels. Some backends (notably SICStus Prolog and Trealla Prolog) do not provide low-level socket creation predicates that can be separated from binding or connecting. This library therefore provides a higher-level API with predicates `client_open/4-5` and `server_open/2-3` that abstracts over these differences.

6.188.2 API documentation

Open the ../apis/library_index.html#sockets link in a web browser.

6.188.3 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(sockets(loader)).
```

6.188.4 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(sockets(tester)).
```

6.188.5 Usage

Creating a client connection

To connect to a server using default options:

```
?- socket::client_open(localhost, 8080, Input, Output).
```

The predicates `client_open/4-5` and `server_accept/4-5` return separate input and output streams. For backends where the same stream is used for bidirectional communication (SICStus Prolog and ECLiPSe), the same stream handle is returned in both arguments. Use `socket::close/2` to close both streams when done.

Creating a server

To create a server that accepts connections using default options on all interfaces:

```
?- socket::server_open(8080, ServerSocket),  
   socket::server_accept(ServerSocket, Input, Output, ClientInfo),  
   % ... communicate with client using Input and Output ...  
   socket::close(Input, Output),  
   socket::server_close(ServerSocket).
```

If the port is passed as a variable to the `server_open/2-3` predicates, an available port will be selected and unified with the variable.

The `server_open/4` predicate can be used to accept a connection only on the given interface.

Getting the current host name

```
?- socket::current_host(Host).
```

6.188.6 API Summary

- `client_open(+Host, +Port, -InputStream, -OutputStream, +Options)` - Connect to a server
- `client_open(+Host, +Port, -InputStream, -OutputStream)` - Connect to a server using default options
- `server_open(+Host, ?Port, -ServerSocket, +Options)` - Create a server socket on the given interface
- `server_open(?Port, -ServerSocket, +Options)` - Create a server socket on all interfaces
- `server_open(?Port, -ServerSocket)` - Create a server socket using default options
- `server_accept(+ServerSocket, -InputStream, -OutputStream, -ClientInfo, +Options)` - Accept connection

- `server_accept(+ServerSocket, -InputStream, -OutputStream, -ClientInfo)` - Accept connection using default options
- `server_close(+ServerSocket)` - Close a server socket
- `close(+InputStream, +OutputStream)` - Close a client or accepted connection
- `current_host(-Host)` - Get the current machine's hostname

6.188.7 Backend-specific notes

Stream representation

The library provides separate input and output stream arguments in `client_open/4-5` and `server_accept/4-5`. For backends where the same stream is used for bidirectional communication (ECLiPSe, SICStus Prolog, and Trealla Prolog), the same stream handle is returned in both the input and output arguments. For backends that use separate streams (GNU Prolog and SWI-Prolog), separate stream handles are returned.

Binary and text modes

The library automatically sets streams by default to binary mode (e.g., for sending and receiving raw bytes). An option is supported for opening streams in text mode.

6.188.8 Known issues

Recent versions of macOS seem to disable the mapping of `localhost` to `127.0.0.1`. This issue may prevent some functionality from working. This can be fixed either by editing the `/etc/hosts` file or by using `'127.0.0.1'` as the host argument instead of `localhost`.

6.189 spade_pattern_miner

SPADE sequential pattern miner for sequence datasets. The library depends on the `sequential_pattern_mining_protocols` support library, implements the generic `pattern_miner_protocol` defined in the `pattern_mining_protocols` core library, and mines frequent sequential patterns using Zaki's equivalence-class decomposition with temporal joins over vertical occurrence lists keyed by sequence and event identifiers.

Requires a dataset implementing `sequence_dataset_protocol` with sequences represented as ordered lists of canonical sorted itemsets over a declared item domain.

6.189.1 API documentation

Open the ../apis/library_index.html#spade_pattern_miner link in a web browser.

6.189.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(spade_pattern_miner(loader)).
```

6.189.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(spade_pattern_miner(tester)).
```

6.189.4 Features

- **Equivalence-Class Decomposition:** Mines frequent sequences using Zaki's equivalence-class decomposition over prefix-sharing classes.
- **Vertical Occurrence Lists:** Represents frequent patterns using sequence and event occurrence lists.
- **Active Class Candidate Pruning:** Restricts class joins to members that occur in the supporting sequences of the current prefix, using per-sequence member-id indexes instead of storing full pattern terms in the sequence index.
- **Same-Event and Sequence Extensions:** Supports both itemset growth and next-event sequence growth.
- **Canonical Sequences:** Validates that itemsets are sorted, duplicate-free, non-empty, and restricted to declared items.
- **Flexible Support Thresholds:** Supports relative minimum support and absolute minimum support count.
- **Model Export:** Mined pattern collections can be exported as predicate clauses or written to a file.

6.189.5 Options

The `mine/3` predicate accepts the following options:

- `minimum_support/1`: Relative minimum support threshold in the interval $]0.0, 1.0]$. The default is 0.5.
- `minimum_support_count/1`: Absolute minimum support count. When both support options are provided, this option takes precedence.
- `maximum_pattern_length/1`: Maximum total number of items in a mined sequential pattern. The default is 1000, effectively capped by the longest sequence in the dataset.
- `minimum_pattern_length/1`: Minimum total number of items retained in the mined result. The default is 1.

6.189.6 Pattern miner representation

The mined pattern miner result is represented by a compound term with the functor chosen by the implementation and arity 3. For example:

```
spade_pattern_miner(ItemDomain, Patterns, Options)
```

Where:

- ItemDomain: Canonical sorted list of declared dataset items.
- Patterns: List of `sequence_pattern(Pattern, SupportCount)` terms ordered first by total item count and then lexicographically.
- Options: Effective mining options used to mine the frequent sequential patterns.

6.189.7 References

1. Zaki, M. J. (2001) - "SPADE: An efficient algorithm for mining frequent sequences".

6.190 simulated_annealing

Simulated annealing is a probabilistic meta-heuristic for approximating the global optimum of a given function. It is particularly useful for combinatorial optimization problems such as the Traveling Salesman Problem (TSP), graph coloring, and scheduling.

The library provides the parametric object `simulated_annealing(Problem, RandomAlgorithm)` where `Problem` is an object implementing the `simulated_annealing_protocol` protocol and `RandomAlgorithm` is one of the algorithms supported by the `fast_random` library. The algorithm minimizes the energy (cost) function defined by the problem.

A convenience object `simulated_annealing(Problem)` is also provided, using the `Xoshiro128++` random number generator (`xoshiro128pp`) as the default.

6.190.1 API documentation

Open the ../docs/library_index.html#simulated-annealing link in a web browser.

6.190.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(simulated_annealing(loader)).
```

6.190.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(simulated_annealing(tester)).
```

6.190.4 Features

- **Configurable random number generator** — the algorithm is parameterized by a `fast_random` algorithm. Available algorithms include `xoshiro128pp`, `xoshiro128ss`, `xoshiro256pp`, `xoshiro256ss`, `well512a`, `splitmix64`, and `as183`. The convenience object `simulated_annealing(Problem)` defaults to `xoshiro128pp`.
- **Boltzmann acceptance criterion** — a worse neighbor is accepted with probability $\exp(-\Delta E / \text{Temperature})$, allowing the search to escape local minima early in the run while converging as the temperature cools.
- **Configurable cooling schedule** — default geometric cooling (`NewTemp` is `Temp * 0.995`), overridable by defining `cooling_schedule/3` in the problem object.
- **Custom stop conditions** — the search stops when the maximum number of steps is reached or the temperature drops below the minimum temperature. A problem object can define `stop_condition/3` to terminate early (e.g. when a good-enough solution is found).
- **Delta-energy optimization** — when the problem object defines `neighbor_state/3`, the algorithm uses the returned energy delta directly instead of recomputing the full energy. This is useful when computing the change is cheaper than evaluating the full cost (e.g. a 2-opt swap in TSP).
- **Best state tracking** — the algorithm tracks the best state found across all iterations and across all restart cycles, not just the final state.
- **Progress reporting** — if the problem object defines `progress/5`, it is called periodically with the current step, temperature, best energy, acceptance rate, and improvement rate. A final report is always produced when the loop terminates.
- **Run statistics** — the `run/4` predicate returns a list of statistics including the number of steps, acceptances, improvements, and the final temperature.
- **Seed control** — the `seed(S)` option initializes the random number generator for reproducible runs.
- **Reheating restarts** — the `restarts(N)` option runs `N` additional SA cycles after the first. Each restart reheats the temperature to the initial value and begins from the best state found so far, allowing the search to escape deep local minima. Statistics accumulate across all cycles.
- **Auto-temperature estimation** — the `estimate_temperature/1-2` predicates sample random neighbor transitions and compute an initial temperature that would produce a target acceptance rate, avoiding manual tuning.

6.190.5 Defining a problem

A problem object must implement the `simulated_annealing_protocol` protocol by defining (at least) the following four predicates:

- `initial_state(-State)` — returns the starting state.
- `neighbor_state(+State, -Neighbor)` — generates a neighboring state.
- `state_energy(+State, -Energy)` — computes the cost of a state (to be minimized).
- `initial_temperature(-Temperature)` — returns the starting temperature.

Optionally, the problem object may also define:

- `neighbor_state(+State, -Neighbor, -DeltaEnergy)` — generates a neighboring state and returns the energy change directly, avoiding a full energy recomputation. Useful when computing the delta is cheaper than recomputing the full energy (e.g. TSP with a 2-opt swap).
- `cooling_schedule(+Temperature, +Step, -NewTemperature)` — computes the next temperature. Default: geometric cooling (NewTemp is $\text{Temp} * 0.995$).
- `stop_condition(+Step, +Temperature, +BestEnergy)` — succeeds when the search should terminate early (e.g. when a good-enough solution is found). Default: the search runs until the maximum number of steps is reached or the temperature drops below the minimum temperature.
- `progress(+Step, +Temperature, +BestEnergy, +AcceptanceRate, +ImprovementRate)` — called periodically during the optimization to report progress. A final report is always produced when the loop terminates. The acceptance and improvement rates are floats between 0.0 and 1.0.

6.190.6 Options

Options for the `run/3-4` predicates:

- `max_steps(N)` — maximum number of iterations per cycle (default: 10000).
- `min_temperature(T)` — minimum temperature floor; the search stops when the temperature drops below this value (default: 0.001).
- `updates(N)` — number of progress reports during the run. Progress is reported by calling `progress/5` on the problem object. Set to 0 to disable (default: 0).
- `seed(S)` — positive integer seed for the random number generator, enabling reproducible runs (default: none).
- `restarts(N)` — number of additional SA cycles after the first. Each restart reheats the temperature to the initial value and begins from the best state found so far, allowing the search to escape local minima (default: 0).

Options for the `estimate_temperature/2` predicate:

- `samples(N)` — number of random neighbor transitions to sample (default: 200).
- `acceptance_rate(P)` — target initial acceptance rate as an integer percentage between 1 and 99 (default: 80).

6.190.7 Run statistics

The run/4 predicate returns a list of statistics about the completed run:

- steps(N) — total number of steps executed.
- acceptances(A) — number of accepted moves (both improving and uphill).
- improvements(I) — number of moves that strictly improved the best energy found.
- final_temperature(T) — temperature at termination.

6.190.8 Usage

Defining a problem

Define an object implementing the simulated_annealing_protocol protocol. For example, a simple quadratic minimization problem:

```
:- object(quadratic,
    implements(simulated_annealing_protocol)).

    initial_state(50.0).
    neighbor_state(X, Y) :-
        random::random(-5.0, 5.0, Delta),
        Y is X + Delta.
    state_energy(X, E) :-
        E is (X - 3.0) * (X - 3.0).
    initial_temperature(100.0).

:- end_object.
```

Running the algorithm

```
| ?- simulated_annealing(quadratic)::run(State, Energy).
State = 3.001..., Energy = 0.000...
```

Running with custom options

```
| ?- simulated_annealing(quadratic)::run(State, Energy, [max_steps(50000)]).
State = 3.000..., Energy = 0.000...
```

Running with statistics

```
| ?- simulated_annealing(quadratic)::run(State, Energy, Stats, []).
State = 3.001..., Energy = 0.000...,
Stats = [steps(10000), acceptances(...), improvements(...), final_temperature(...)]
```

Reproducible runs with seed

```
| ?- simulated_annealing(quadratic)::run(S1, E1, [seed(42)]),
    simulated_annealing(quadratic)::run(S2, E2, [seed(42)]).
S1 = S2, E1 = E2.
```

Reheating restarts

Run 3 SA cycles (1 initial + 2 restarts). Each restart reheats the temperature and begins from the best state found so far:

```
| ?- simulated_annealing(quadratic)::run(State, Energy, [restarts(2)]).
State = 3.000..., Energy = 0.000...
```

Auto-temperature estimation

Estimate a good initial temperature for the problem by sampling random neighbor transitions:

```
| ?- simulated_annealing(quadratic)::estimate_temperature(T).
T = ...
```

With custom parameters (500 samples, 90% target acceptance rate):

```
| ?- simulated_annealing(quadratic)::estimate_temperature(T, [samples(500), acceptance_
    ↪rate(90)]).
T = ...
```

Using a custom random number generator

Use the two-parameter version to select a specific fast_random algorithm:

```
| ?- simulated_annealing(quadratic, well512a)::run(State, Energy).
State = 3.001..., Energy = 0.000...

| ?- simulated_annealing(quadratic, xoshiro256ss)::run(State, Energy, [seed(42)]).
State = 3.000..., Energy = 0.000...
```

6.191 statistics

The entities in this library define some useful predicates for descriptive statistics. Data is represented as a list of numbers (integers or floats). Use the object `sample` if your data represents a sample. Use the object `population` if your data represents a population.

The `variance/2`, `standard_deviation/2`, `skewness/2`, `kurtosis/2`, `covariance/3`, `standard_error/2`, `correlation/3`, and `rank_correlation/3` predicates use different formulas depending on whether the data represents a sample (dividing by $N-1$) or a population (dividing by N). All other predicates share the same implementation.

6.191.1 API documentation

Open the ../apis/library_index.html#statistics link in a web browser.

6.191.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(statistics(loader)).
```

6.191.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(statistics(tester)).
```

6.191.4 API overview

Aggregation

Predicate	Description
<code>product/2</code>	Product of all list numbers
<code>sum/2</code>	Sum of all list numbers

Extremes and Range

Predicate	Description
<code>min/2</code>	Minimum value
<code>max/2</code>	Maximum value
<code>min_max/3</code>	Minimum and maximum values
<code>range/2</code>	Range (max - min)

Central Tendency

Predicate	Description
arithmetic_mean/2	Arithmetic mean
geometric_mean/2	Geometric mean
harmonic_mean/2	Harmonic mean
weighted_mean/3	Weighted mean
trimmed_mean/3	Trimmed mean (removing a fraction of extreme values)
median/2	Median
modes/2	Modes (in ascending order)

Measures of Position / Quantiles

Predicate	Description
fractile/3	Fractile (quantile given a fraction in (0.0, 1.0))
percentile/3	Percentile (quantile given a value in (0, 100))
quartiles/4	Quartiles (Q1, Q2, Q3)
interquartile_range/2	Interquartile range (Q3 - Q1)

Measures of Dispersion

Predicate	Description
variance/2	Variance (sample or population)
standard_deviation/2	Standard deviation (sample or population)
mean_deviation/2	Mean absolute deviation
median_deviation/2	Median absolute deviation
average_deviation/3	Average absolute deviation from a given central tendency
coefficient_of_variation/2	Coefficient of variation
relative_standard_deviation/2	Relative standard deviation (percentage)
sum_of_squares/2	Sum of squared deviations from the mean
standard_error/2	Standard error of the mean

Measures of Shape

Predicate	Description
skewness/2	Moment skewness (sample or population)
kurtosis/2	Excess kurtosis (sample or population)
central_moment/3	K-th central moment

Measures of Association

Predicate	Description
covariance/3	Covariance (sample or population)
correlation/3	Pearson correlation coefficient
rank_correlation/3	Spearman rank correlation coefficient

Error Metrics

Predicate	Description
mean_squared_error/3	Mean squared error between two lists
root_mean_squared_error/3	Root mean squared error between two lists

Normalization

Predicate	Description
z_normalization/2	Z-score normalization (mean \sim 0, std \sim 1)
min_max_normalization/2	Min-max normalization (rescale to [0, 1])

Frequency / Counting

Predicate	Description
frequency_distribution/2	Frequency distribution (Value-Count pairs)

Validation

Predicate	Description
valid/1	Term is a closed list of numbers

6.192 stemming

This library provides word stemming predicates for English text, with support for different word representations: atoms, character lists, or character code lists.

The library includes implementations of two well-known stemming algorithms:

- **Porter Stemmer** - The Porter stemming algorithm (Porter, 1980) is a widely used algorithm for reducing English words to their root form by applying a series of rules that remove common suffixes.

- **Lovins Stemmer** - The Lovins stemming algorithm (Lovins, 1968) removes the longest suffix from a word using a list of endings, each associated with a condition for removal. It then applies transformation rules to fix spelling.

6.192.1 API documentation

Open the ../apis/library_index.html#stemming link in a web browser.

6.192.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(stemming(loader)).
```

6.192.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(stemming(tester)).
```

6.192.4 Usage

The stemming predicates are defined in parametric objects where the parameter specifies the word representation:

- `atom` - words are represented as atoms
- `chars` - words are represented as lists of characters
- `codes` - words are represented as lists of character codes

The parameter must be bound when sending messages to the objects.

Porter Stemmer

To stem a single word using atoms:

```
| ?- porter_stemmer(atom)::stem(running, Stem).
Stem = run
yes
```

To stem a list of words:

```
| ?- porter_stemmer(atom)::stems([running, walks, easily], Stems).
Stems = [run, walk, easili]
yes
```

Using character lists:

```
| ?- porter_stemmer(chars)::stem([r,u,n,n,i,n,g], Stem).
Stem = [r,u,n]
yes
```

Lovins Stemmer

To stem a single word using atoms:

```
| ?- lovins_stemmer(atom)::stem(running, Stem).  
Stem = run  
yes
```

To stem a list of words:

```
| ?- lovins_stemmer(atom)::stems([running, walks, easily], Stems).  
Stems = [run, walk, eas]  
yes
```

6.192.5 Algorithms

Porter Stemmer

The Porter stemming algorithm, developed by Martin Porter in 1980, is one of the most widely used stemming algorithms for the English language. It operates through a series of steps that progressively remove suffixes from words:

1. **Step 1a:** Handle plurals (e.g., “caresses” → “caress”, “ponies” → “poni”)
2. **Step 1b:** Handle past tense and progressive forms (e.g., “agreed” → “agree”)
3. **Step 1c:** Replace terminal “y” with “i” when preceded by a vowel
4. **Steps 2-4:** Remove various suffixes based on the “measure” of the stem
5. **Step 5:** Clean up final “e” and double consonants

The algorithm uses the concept of “measure” (m), which counts vowel-consonant sequences in the stem, to determine when suffixes can be safely removed.

Reference: Porter, M.F. (1980). An algorithm for suffix stripping. *Program*, 14(3), 130-137.

Lovins Stemmer

The Lovins stemming algorithm, developed by Julie Beth Lovins in 1968, was one of the earliest stemming algorithms. It takes a different approach from Porter:

1. **Ending removal:** The algorithm maintains a list of 294 possible endings, ordered by length. It removes the longest matching ending that satisfies its associated condition (e.g., minimum stem length).
2. **Transformation rules:** After removing the ending, spelling transformations are applied to fix common irregularities (e.g., “iev” → “ief”, “uct” → “uc”).

The Lovins algorithm tends to be more aggressive than Porter, sometimes producing stems that are not actual words but are consistent across related word forms.

Reference: Lovins, J.B. (1968). Development of a stemming algorithm. *Mechanical Translation and Computational Linguistics*, 11(1-2), 22-31.

6.192.6 Choosing an Algorithm

- **Porter:** More conservative, produces more readable stems, widely used in information retrieval and search applications. Good choice for most applications.
- **Lovins:** More aggressive, may conflate more word forms together. Can be useful when broader matching is desired, but may over-stem in some cases.

Both algorithms are designed for English text only.

6.192.7 Known Limitations

- Both algorithms work only with English words.
- Stemming is not lemmatization - stems may not be valid dictionary words.
- Proper nouns and abbreviations may not be handled correctly.
- Very short words (1-2 characters) are returned unchanged.

6.193 stomp

Portable STOMP 1.2 (Simple Text Orientated Messaging Protocol) client implementation. This library uses the sockets library and supports all backend Prolog systems supported by that library: ECLiPSe, GNU Prolog, SICStus Prolog, SWI-Prolog, Trealla Prolog, and XVM.

6.193.1 Protocol Version

This library implements the STOMP 1.2 protocol specification:

<https://stomp.github.io/stomp-specification-1.2.html>

STOMP 1.2 is backwards compatible with STOMP 1.1 with minor differences in frame line endings and message acknowledgment.

6.193.2 Features

- Full STOMP 1.2 protocol support
- Heartbeat negotiation and automatic heartbeat handling
- Multiple concurrent subscriptions with unique IDs
- All acknowledgment modes: auto, client, client-individual
- Transaction support: BEGIN, COMMIT, ABORT
- Graceful disconnect with receipt confirmation
- Proper frame encoding/decoding with header escaping
- User-defined headers for messages

6.193.3 API documentation

Open the `../apis/library_index.html#stomp` link in a web browser.

6.193.4 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(stomp(loader)).
```

6.193.5 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(stomp(tester)).
```

Note: Integration tests require a running STOMP server (e.g., RabbitMQ with the STOMP plugin, ActiveMQ, or Apache Apollo). They are usually run using RabbitMQ with the STOMP plugin. As RabbitMQ uses `/` as its default virtual host, the tests assume this and use the `host('/')` option when calling the `connect/4` predicate (otherwise the connection would fail with a "Virtual host 'localhost' access denied" error message). To enable the RabbitMQ's STOMP plugin:

```
$ rabbitmq-plugins enable rabbitmq_stomp
```

RabbitMQ STOMP listens on port 61613 by default.

6.193.6 Usage

Connecting to a STOMP server

Basic connection assuming RabbitMQ, which uses `host('/')` as shown:

```
?- stomp::connect(localhost, 61613, Connection, [host('/')]).
```

Connection with authentication:

```
?- stomp::connect(localhost, 61613, Connection, [
    login('user'),
    passcode('password')
]).
```

Connection with heartbeat:

```
?- stomp::connect(localhost, 61613, Connection, [
    heartbeat(10000, 10000)
]).
```

Connection with virtual host:

```
?- stomp::connect(localhost, 61613, Connection, [
    host('/my-vhost')
]).
```

Sending messages

Send a simple text message:

```
?- stomp::send(Connection, '/queue/test', 'Hello, World!', []).
```

Send a message with content type:

```
?- stomp::send(Connection, '/queue/test', '{"key":"value"}', [
    content_type('application/json')
]).
```

Send with custom headers:

```
?- stomp::send(Connection, '/queue/test', 'Message', [
    header('priority', '9'),
    header('persistent', 'true')
]).
```

Send within a transaction:

```
?- stomp::begin_transaction(Connection, 'tx-001', []),
   stomp::send(Connection, '/queue/test', 'Message 1', [transaction('tx-001')]),
   stomp::send(Connection, '/queue/test', 'Message 2', [transaction('tx-001')]),
   stomp::commit_transaction(Connection, 'tx-001', []).
```

Subscribing to destinations

Subscribe with auto acknowledgment (default):

```
?- stomp::subscribe(Connection, '/queue/test', 'sub-0', []).
```

Subscribe with client acknowledgment:

```
?- stomp::subscribe(Connection, '/queue/test', 'sub-0', [
    ack(client)
]).
```

Subscribe with client-individual acknowledgment:

```
?- stomp::subscribe(Connection, '/topic/events', 'sub-1', [
    ack(client_individual)
]).
```

Receiving messages

Poll for a message (non-blocking):

```
?- stomp::receive(Connection, Frame, [timeout(0)]).
```

Wait for a message with timeout (milliseconds):

```
?- stomp::receive(Connection, Frame, [timeout(5000)]).
```

Process received MESSAGE frame:

```
?- stomp::receive(Connection, Frame, []),  
   stomp::frame_command(Frame, 'MESSAGE'),  
   stomp::frame_header(Frame, 'destination', Destination),  
   stomp::frame_header(Frame, 'message-id', MessageId),  
   stomp::frame_body(Frame, Body).
```

Acknowledging messages

For client or client-individual acknowledgment modes:

```
?- stomp::receive(Connection, Frame, []),  
   stomp::frame_header(Frame, 'ack', AckId),  
   stomp::ack(Connection, AckId, []).
```

Negative acknowledgment:

```
?- stomp::nack(Connection, AckId, []).
```

Acknowledgment within a transaction:

```
?- stomp::ack(Connection, AckId, [transaction('tx-001')]).
```

Unsubscribing

```
?- stomp::unsubscribe(Connection, 'sub-0', []).
```

Transactions

Begin a transaction:

```
?- stomp::begin_transaction(Connection, 'tx-001', []).
```

Commit a transaction:

```
?- stomp::commit_transaction(Connection, 'tx-001', []).
```

Abort a transaction:

```
?- stomp::abort_transaction(Connection, 'tx-001', []).
```

Disconnecting

Graceful disconnect with receipt confirmation:

```
?- stomp::disconnect(Connection, []).
```

Heartbeat handling

Send a heartbeat (if needed manually):

```
?- stomp::send_heartbeat(Connection).
```

Check connection health:

```
?- stomp::connection_alive(Connection).
```

Working with frames

Get frame command:

```
?- stomp::frame_command(Frame, Command).
```

Get frame header value:

```
?- stomp::frame_header(Frame, HeaderName, Value).
```

Get frame body:

```
?- stomp::frame_body(Frame, Body).
```

Get all frame headers:

```
?- stomp::frame_headers(Frame, Headers).
```

6.193.7 API Summary

Connection management

- `connect(+Host, +Port, -Connection, +Options)` - Connect to server
- `disconnect(+Connection, +Options)` - Disconnect from server
- `connection_alive(+Connection)` - Check if connection is alive

Messaging

- `send(+Connection, +Destination, +Body, +Options)` - Send message
- `subscribe(+Connection, +Destination, +Id, +Options)` - Subscribe
- `unsubscribe(+Connection, +Id, +Options)` - Unsubscribe
- `receive(+Connection, -Frame, +Options)` - Receive frame

Acknowledgment

- `ack(+Connection, +Id, +Options)` - Acknowledge message
- `nack(+Connection, +Id, +Options)` - Negative acknowledge

Transactions

- `begin_transaction(+Connection, +TransactionId, +Options)` - Begin
- `commit_transaction(+Connection, +TransactionId, +Options)` - Commit
- `abort_transaction(+Connection, +TransactionId, +Options)` - Abort

Heartbeat

- `send_heartbeat(+Connection)` - Send heartbeat EOL

Frame inspection

- `frame_command(+Frame, -Command)` - Get frame command
- `frame_header(+Frame, +Name, -Value)` - Get header value
- `frame_headers(+Frame, -Headers)` - Get all headers
- `frame_body(+Frame, -Body)` - Get frame body

6.193.8 Connection Options

- `login(Login)` - Username for authentication
- `passcode(Passcode)` - Password for authentication
- `host(VirtualHost)` - Virtual host name (default: from Host parameter)
- `heartbeat(ClientMs, ServerMs)` - Heartbeat timing in milliseconds

6.193.9 Send Options

- `content_type(MimeType)` - MIME type of the body
- `content_length(Length)` - Byte length of body (auto-calculated if omitted)
- `transaction(TransactionId)` - Include in transaction
- `receipt(ReceiptId)` - Request receipt from server
- `header(Name, Value)` - Add custom header

6.193.10 Subscribe Options

- `ack(Mode)` - Acknowledgment mode: `auto`, `client`, or `client_individual`

6.193.11 Receive Options

- `timeout(Milliseconds)` - Wait timeout; 0 for non-blocking, -1 for infinite

6.193.12 Error Handling

Most predicates throw `error(stomp_error(Reason), Context)` on failure. Common error reasons:

- `connection_failed` - Could not establish TCP connection
- `protocol_error(Message)` - Server sent ERROR frame
- `not_connected` - Connection was closed
- `invalid_frame` - Malformed frame received
- `timeout` - Operation timed out

6.194 string_distance

This library provides string distance predicates with support for different string representations: atoms, character lists, or character code lists.

The predicates are defined in the `string_distance(_Representation_)` parametric object where `_Representation_` can be one of:

- `atom` - strings are represented as atoms
- `chars` - strings are represented as lists of characters
- `codes` - strings are represented as lists of character codes

The parameter must be bound when sending messages to the object.

6.194.1 API documentation

Open the `../apis/library_index.html#string_distance` link in a web browser.

6.194.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(string_distance(loader)).
```

6.194.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(string_distance(tester)).
```

6.194.4 Algorithms

This library implements the following string distance algorithms:

- **Levenshtein** distance is probably the most well-known. It counts the minimum number of single-character edits — insertions, deletions, or substitutions — needed to transform one string into another. For example, the distance between “kitten” and “sitting” is 3.
- **Damerau-Levenshtein** distance extends Levenshtein by also allowing transpositions of two adjacent characters as a single edit operation. This makes it more practical for catching common typos like “ab” -> “ba”.
- **Hamming** distance only works on strings of equal length and counts the number of positions where the corresponding characters differ. It’s very fast but limited by that equal-length constraint, making it useful for things like comparing binary codes or fixed-format identifiers.
- **Jaro and Jaro-Winkler** distance produce a similarity score between 0 and 1 rather than a raw edit count. Jaro considers matching characters and transpositions, while Jaro-Winkler adds a prefix-weighting boost — the idea being that strings matching at the start are more likely to be the same (useful for name matching).
- **Edit similarity** is a normalized version of the edit distance, defined as $1 - (\text{edit distance} / \text{max length of the two strings})$. It produces a similarity score between 0 and 1. Can be computed from Levenshtein, Damerau-Levenshtein, Hamming, or Longest Common Subsequence distances.
- **Longest Common Subsequence** (LCS) finds the length of the longest sequence of characters that appears in both strings in the same relative order (not necessarily contiguous). It’s heavily used in diff tools for comparing files.
- **Longest Common Substring** is similar to LCS but requires the shared characters to be contiguous. Useful when caring about shared blocks of text rather than scattered matches.
- **Cosine Similarity** shifts the approach from character edits to vector-based comparison. Strings are converted into vectors (often using character n-grams or word tokens), and then the cosine of the angle between them is computed. It’s widely used in NLP and information retrieval.
- **Jaccard Index** compares two sets of tokens (often words or n-grams) by dividing the size of their intersection by the size of their union. Simple and intuitive, though it ignores term frequency.

- **Soundex/Metaphone** are phonetic algorithms rather than strict string comparisons. They encode strings based on how they sound, so “Smith” and “Smythe” would match. Useful for name deduplication where spelling varies but pronunciation is similar.

General advice: use Levenshtein or Damerau-Levenshtein for general-purpose edit distance, Jaro-Winkler for short strings like names, LCS for diff-style comparisons, and cosine/Jaccard when working at the word or document level rather than character level.

6.195 strings

This library provides string manipulation predicates with support for different string representations: atoms, character lists, or character code lists. Its API is partially based on work and libraries found in ECLiPSe and SWI-Prolog.

The predicates are defined in the `string(_Representation_)` parametric object where `_Representation_` can be one of:

- `atom` - strings are represented as atoms
- `chars` - strings are represented as lists of characters
- `codes` - strings are represented as lists of character codes

The parameter must be bound when sending messages to the object.

6.195.1 API documentation

Open the ../apis/library_index.html#strings link in a web browser.

6.195.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(strings(loader)).
```

6.195.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(strings(tester)).
```

6.195.4 Predicates

The library provides the following compatibility predicates:

- `atom_string/2` - converts between atoms and strings
- `number_string/2` - converts between numbers and strings
- `string_chars/2` - converts between strings and character lists
- `string_codes/2` - converts between strings and character code lists
- `string_concat/3` - concatenates two strings

- `string_length/2` - returns the length of a string
- `sub_string/5` - extracts substrings
- `string_upper/2` - converts a string to uppercase
- `string_lower/2` - converts a string to lowercase
- `split_string/4` - splits a string into substrings using separators and padding
- `atomics_to_string/2` - concatenates a list of atomic terms into a string
- `atomics_to_string/3` - concatenates a list of atomic terms into a string with separator

It also provides the following string trimming predicates:

- `trim/2-3` - trims leading and trailing characters from a string
- `trim_left/2-3` - trims leading characters from a string
- `trim_right/2-3` - trims trailing characters from a string

For converting between terms and strings, see the `term_io` library.

6.196 subsequences

This library provides predicates for working with subsequences represented using lists, including generation, search, contiguous subsequence operations, and random selection. The following categories of predicates are provided:

- **Generation operations** - Predicates for generating subsequences and variants thereof.
- **Ordering variants** - Predicates that support an additional order argument (default, lexicographic, or shortlex) for controlling output order.
- **Filtered generation** - Predicates for generating specific types of subsequences (combinations, permutations).
- **Indexed access** - Predicates for direct access to subsequences at specific positions.
- **Searching and matching** - Predicates for finding specific subsequences with desired properties.
- **Prefix and suffix operations** - Predicates for checking and finding prefixes and suffixes.
- **Contiguous subsequences** - Predicates for working with contiguous subsequences (subslices, sliding windows).
- **Random selection** - Predicates for randomly selecting subsequences.
- **Constrained operations** - Predicates for generating subsequences with specific constraints.
- **Utility predicates** - Predicates that support subsequence operations.

Dedicated arrangements, cartesian_products, derangements, combinations, multisets, and permutations libraries are also available for focused APIs on related operations.

6.196.1 API documentation

Open the ../apis/library_index.html#subsequences link in a web browser.

6.196.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(subsequences(loader)).
```

6.196.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(subsequences(tester)).
```

6.197 term_io

This library implements predicates for reading/writing terms from/to atoms, chars (lists of characters), and codes (lists of character codes). These predicates are implemented using a single temporary file created when the library is loaded. This temporary file is unique per Logtalk process. The predicates can be safely used in multi-threaded applications.

6.197.1 API documentation

Open the ../apis/library_index.html#term-io link in a web browser.

6.197.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(term_io(loader)).
```

6.197.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(term_io(tester)).
```

6.198 thurstone_mosteller_ranker

Thurstone-Mosteller Case V pairwise preference ranker. It aggregates pairwise outcomes into matchup probabilities, applies a fixed continuity correction to avoid infinite probit values, transforms them with the inverse standard normal CDF, and fits Case V latent utilities using a deterministic weighted least-squares linear solve.

The library implements the `ranker_protocol` defined in the `ranking_protocols` library. It provides predicates for learning a ranker from pairwise preferences, using it to order candidate items, and exporting it as a list of predicate clauses or to a file.

Datasets are represented as objects implementing the `pairwise_ranking_dataset_protocol` protocol from the `ranking_protocols` library. See the `test_datasets` directory for examples. The current implementation requires a well-formed connected pairwise dataset so that all learned utilities remain comparable across the ranked items.

6.198.1 API documentation

Open the ../apis/library_index.html#thurstone_mosteller_ranker link in a web browser.

6.198.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(thurstone_mosteller_ranker(loader)).
```

6.198.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(thurstone_mosteller_ranker(tester)).
```

6.198.4 Features

- **Pairwise Preference Learning:** Learns one deterministic latent utility per item from aggregated pairwise outcomes.
- **Case V Thurstone-Mosteller Fit:** Uses continuity-corrected matchup win probabilities transformed by the inverse standard normal CDF and fitted with a deterministic weighted least-squares solve.
- **Centered Utility Scale:** Learned scores are centered so that only utility differences encode preference strength while the absolute location remains arbitrary.
- **Deterministic Ranking:** Orders candidate items by learned utility with deterministic tie-breaking.
- **Strict Dataset Validation:** Rejects duplicate items, undeclared items, self-preferences, non-positive weights, and disconnected comparison graphs.
- **Training Diagnostics:** Learned rankers include metadata describing the fitting method, continuity correction, and validated dataset summary.
- **Ranker Export:** Learned rankers can be exported as self-contained terms.

6.198.5 Scoring semantics

This implementation aggregates pairwise preferences into matchup totals and then estimates latent utilities under the Thurstone-Mosteller Case V model. For each observed unordered pair, it computes a continuity-corrected empirical win probability

```
(wins + 0.5) / (total + 1.0)
```

maps that probability through the inverse standard normal CDF, and fits utility differences using weighted least squares with matchup totals as the weights.

The resulting utilities are centered, real-valued scores on an arbitrary additive scale. Only differences between utilities are meaningful when interpreting implied paired-comparison probabilities and ranking order.

6.198.6 Options

The current learn/3 implementation does not define any user options beyond the default empty list. Non-empty options lists are rejected.

6.198.7 Diagnostics syntax

The diagnostics/2 predicate returns a list of metadata terms with the form:

```
[
  model(thurstone_mosteller_ranker),
  options(Options),
  fitting(weighted_least_squares_case_v),
  continuity_correction(0.5),
  dataset_summary(DatasetSummary)
]
```

6.198.8 Ranker representation

The learned ranker is represented by a compound term of the form:

```
thurstone_mosteller_ranker(Items, Scores, Diagnostics)
```

Where:

- Items: List of ranked items.
- Scores: List of Item-Score pairs.
- Diagnostics: List of metadata terms, including the fitting method, continuity correction, and dataset summary.

6.199 time_scales

The `time_scales` library provides predicates for converting instants between UTC, TAI, TT, UT1, TDB, GPS, GST, TCG, and TCB using bundled reference data and optional user-provided override files.

The library is designed to complement (not replace) the `dates` and `iso8601` libraries:

- Use `dates` for civil date-time arithmetic and Unix epoch conversions.
- Use `iso8601` for string parsing and formatting.
- Use `time_scales` for physical time-scale conversions.

6.199.1 API documentation

Open the ../apis/library_index.html#time_scales link in a web browser.

6.199.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(time_scales(loader)).
```

6.199.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(time_scales(tester)).
```

6.199.4 Features

Current feature set:

- Supported scales: `utc`, `tai`, `tt`, `ut1`, `tdb`, `gps`, `gst`, `tcg`, and `tcb`.
- UTC support starts at 1972-01-01T00:00:00Z.
- Leap seconds are provided by a bundled static table with optional override file.
- DUT1 (UT1-UTC) is provided by bundled data with optional override file.
- Active leap-second and DUT1 tables can be queried as ordered term lists for reproducibility.
- Active leap-second and DUT1 tables can be saved to deterministic term files and reloaded as overrides.
- Fail- and error-based validation predicates for validating and type-checking instant terms and conversion requests.
- TDB conversions use a practical TT/TDB approximation suitable for application-level use.
- Real-valued approximation offsets converted to rationals using high-resolution nanosecond scaling.

6.199.5 Limitations

Current non-implemented features:

- UTC coverage before 1972-01-01T00:00:00Z is not implemented.
- network-driven leap-second and DUT1 data updates are not implemented.
- high-precision ephemeris-based relativistic modeling (beyond the current practical approximation formulas) is not implemented.

For full functionality (including all high-precision conversion paths and their corresponding tests), a Prolog backend with support for unbounded integer arithmetic is required.

6.199.6 Representation

Instants are represented as:

```
instant(Scale, Seconds, fraction(Numerator, Denominator))
```

Where:

- Scale is one of utc, tai, tt, ut1, tdb, gps, gst, tcg, or tcb;
- Seconds is integer epoch seconds;
- fraction(Numerator, Denominator) is a normalized fraction in the $[0,1[$ interval.

Offsets are returned as:

```
rational(Numerator, Denominator)
```

6.199.7 Examples

Convert UTC to TAI:

```
| ?- time_scales::convert(instant(utc, 1483228800, fraction(0,1)), utc, tai, TAI).
```

Strictly validate an instant term (throws on invalid input):

```
| ?- time_scales::check_instant(instant(utc, 1483228800, fraction(0,1))).
```

Strict conversion (throws on invalid scale or mismatched instant scale):

```
| ?- time_scales::check_convert(instant(utc, 1483228800, fraction(0,1)), utc, tai, TAI).
```

Convert TAI to TT:

```
| ?- time_scales::convert(instant(tai, 1483228837, fraction(0,1)), tai, tt, TT).
```

Convert UTC to UT1:

```
| ?- time_scales::convert(instant(utc, 1483228800, fraction(0,1)), utc, ut1, UT1).
```

Convert TT to TDB:

```
| ?- time_scales::convert(instant(tt, 1483228869, fraction(0,1)), tt, tdb, TDB).
```

Convert UTC to GPS:

```
| ?- time_scales::convert(instant(utc, 1483228800, fraction(0,1)), utc, gps, GPS).
```

Convert UTC to GST:

```
| ?- time_scales::convert(instant(utc, 1483228800, fraction(0,1)), utc, gst, GST).
```

Convert UTC to TCG:

```
| ?- time_scales::convert(instant(utc, 1483228800, fraction(0,1)), utc, tcg, TCG).
```

Convert UTC to TCB:

```
| ?- time_scales::convert(instant(utc, 1483228800, fraction(0,1)), utc, tcb, TCB).
```

Lookup an effective leap-second date:

```
| ?- time_scales::leap_second_date(Date, Offset).
```

Load leap-second and DUT1 override files:

```
| ?- time_scales::load_leap_seconds_override('test_files/leap_seconds_override.pl').
```

```
| ?- time_scales::load_dut1_override('test_files/dut1_override.pl').
```

Query active leap-second and DUT1 data tables:

```
| ?- time_scales::leap_seconds_entries(LeapEntries).
```

```
| ?- time_scales::dut1_entries(DUT1Entries).
```

Save active leap-second and DUT1 data tables:

```
| ?- time_scales::save_leap_seconds_entries('/tmp/logtalk_time_scales_leap_snapshot.pl').
```

```
| ?- time_scales::save_dut1_entries('/tmp/logtalk_time_scales_dut1_snapshot.pl').
```

Validate an instant or conversion request:

```
| ?- time_scales::valid_instant(instant(utc, 1483228800, fraction(0,1))).
```

```
| ?- time_scales::valid_conversion(instant(tai, 1483228837, fraction(0,1)), tai, tdb).
```

6.200 timeout

The `timeout` object provides a portable abstraction over calling a goal deterministically with a time limit as made available in some form by some of the supported backend Prolog systems (B-Prolog, ECLiPSe, SICStus Prolog, SWI-Prolog, Trealla Prolog, XSB, XVM, and YAP).

For better performance, compile calls to this library meta-predicates with the `optimize` flag turned on so that the meta-arguments, i.e. the goals that you are timing, are also compiled.

6.200.1 API documentation

Open the ../apis/library_index.html#timeout link in a web browser.

6.200.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(timeout(loader)).
```

6.200.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(timeout(tester)).
```

6.200.4 Known issues

Two tests are currently skipped when using the SWI-Prolog backend as they cannot be interrupted and generate the expected timeout exceptions due to tests being run from an `initialization/1` directive goal that, in the SWI-Prolog implementation of this directive, ignores signals. The test goals do generate the expected timeout exception when not called from an `initialization/1` directive.

6.201 tle_orbits

The `tle_orbits` library parses Two-Line Element (TLE) records and provides a small set of helpers for portable approximate orbit propagation and ground-track generation.

The current version defaults to an approximate automatic propagation mode that dispatches between dedicated near-earth and deep-space variants. The near-earth variant applies low-order B^* drag damping together with J2 secular and short-period position corrections. The deep-space variant uses its own long-period and resonance-aware corrections so high-period orbits no longer share the near-earth path. The two-body Keplerian approximation is also available as an explicit baseline model for comparison and simple fallback use. Propagated states can also be queried together with direct velocity outputs in inertial, Earth-fixed, or local ENU frames.

6.201.1 API documentation

Open the `../apis/library_index.html#tle_orbits` link in a web browser.

6.201.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(tle_orbits(loader)).
```

6.201.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(tle_orbits(tester)).
```

6.201.4 Representation

Parsed TLE records are represented as:

```
~ ~ ~
tle(
  Name,
  satellite(CatalogNumber, Classification, designator(LaunchYear, LaunchNumber, _
↳ Piece)),
  epoch_julian_date(EpochJulianDate),
  drag(MeanMotionDot, MeanMotionDdot, BStar),
  orbit(InclinationDegrees, RightAscensionDegrees, Eccentricity,
        ArgumentOfPerigeeDegrees, MeanAnomalyDegrees, MeanMotionRevolutionsPerDay),
  ephemeris_type(EphemerisType),
  element_set(ElementSetNumber),
  revolution(RevolutionNumber)
)
~ ~ ~
```

If the optional title line is absent, Name is the atom none.

6.201.5 Public API

- `parse/2` parses one or more TLE records from `atom(Atom)`, `chars(List)`, `codes(List)`, `stream(Stream)`, or `file(Path)` sources.
- `parse_lines/4` parses a single TLE record from an optional name and the two fixed-width element lines.
- `propagate/3` propagates to a default geographic(Latitude,Longitude,Height) coordinate in WGS84 using the default approximate model.
- `propagate/4` supports the frames `eci`, `ecef`, and `wgs84_3d` using the default approximate model.
- `propagate/5` additionally allows choosing one of the `approximate`, `approximate_near_earth`, `approximate_deep_space`, or `two_body` models.

- `propagate_state/4` returns `state(Position,Velocity)` in a requested frame using the default approximate model.
- `propagate_state/5` additionally allows choosing one of the `approximate`, `approximate_near_earth`, `approximate_deep_space`, or `two_body` models. Velocity is derived directly from the propagated orbital elements in ECI and then analytically transformed to ECEF or local ENU when requested.
- `ground_track/5` samples `sample(DateTime, geographic(Latitude,Longitude,Height))` points over a time range using a fixed step in seconds and the default approximate model.
- `ground_track/6` additionally allows choosing one of the `approximate`, `approximate_near_earth`, `approximate_deep_space`, or `two_body` models.

Supported propagation time specifications are:

- `date_time(Year, Month, Day, Hours, Minutes, Seconds)`
- `julian_date(JulianDate)`
- `offset_seconds(SecondsSinceEpoch)`

6.201.6 Examples

Parse one named TLE record:

```
| ?- tle_orbits::parse(atom('ISS (ZARYA)\n1 25544U 98067A 24120.51782528 .00012051 00000-
→0 21940-3 0 9995\n2 25544 51.6393 184.4452 0003580 32.9443 327.1663 15.50957687452123\n
→'), TLEs).
```

Propagate to a WGS84 3D geodetic coordinate using the TLE epoch:

```
| ?- TLEs = [TLE], tle_orbits::propagate(TLE, offset_seconds(0.0), Coordinate).
```

Propagate to ECI or ECEF coordinates:

```
| ?- tle_orbits::propagate(TLE, offset_seconds(600.0), eci, ECI).
| ?- tle_orbits::propagate(TLE, offset_seconds(600.0), ecef, ECEF).
```

Select the propagation model explicitly:

```
| ?- tle_orbits::propagate(TLE, offset_seconds(600.0), eci, two_body, ECI).
| ?- tle_orbits::propagate(TLE, offset_seconds(600.0), eci, approximate_near_earth, ECI).
| ?- tle_orbits::propagate(TLE, offset_seconds(600.0), eci, approximate, ECI).
| ?- tle_orbits::propagate(TLE, offset_seconds(43200.0), eci, approximate_deep_space, ECI).
```

Request propagated state plus velocity:

```
| ?- tle_orbits::propagate_state(TLE, offset_seconds(600.0), eci, State).
| ?- tle_orbits::propagate_state(TLE, offset_seconds(600.0), wgs84_3d, approximate, State).
```

Generate a ground track sampled every 10 minutes:

```
| ?- tle_orbits::ground_track(
    TLE,
    date_time(2024, 4, 29, 12, 25, 40.0),
    date_time(2024, 4, 29, 13, 25, 40.0),
    600.0,
```

(continues on next page)

(continued from previous page)

Samples

).

6.201.7 Notes

- Parsing validates the TLE checksums and rejects malformed records.
- The approximate family includes low-order B* drag handling, near-earth short-period corrections, and a separate deep-space approximation with independent solar, lunar, and resonance-aware perturbation terms. It is a portable approximation, not an implementation of standards-grade SGP4/SDP4.
- State velocity output is no longer estimated using centered finite differences. The library now computes ECI velocity directly from the propagated orbital elements and transforms that velocity analytically to ECEF and local ENU.
- Position and velocity frame conversion now share the same Earth-rotation basis to keep propagated state components internally consistent across frames.
- The geodetic output path reuses the `crs_projections` library for converting propagated `ecef(X,Y,Z)` coordinates to `geographic/3` coordinates.

6.202 toml

The `toml` library provides predicates for parsing and generating data in the TOML format:

<https://toml.io>

It includes parametric objects whose parameters allow selecting the representation for parsed TOML tables (compound or curly), TOML text strings (atom, chars, or codes) and TOML pairs (dash, equal, or colon).

6.202.1 API documentation

Open the `../apis/library_index.html#toml` link in a web browser.

6.202.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(toml(loader)).
```

6.202.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(toml(tester)).
```

6.202.4 Status

The full TOML 1.0.0 data model is covered:

- key/value assignments
- standard tables using [table] headers
- dotted keys
- arrays
- inline tables
- array-of-tables
- booleans
- decimal, hexadecimal, octal, and binary integers
- decimal floats plus inf and nan
- single-line and multiline basic and literal strings
- local dates, local times, local date-times, and offset date-times
- parsing and generation from files, streams, chars, codes, and atoms

6.202.5 Representation

The following choices of syntax have been made to represent TOML elements as terms:

- By default, TOML tables are represented using the compound term `toml(Pairs)`, where `Pairs` is a list using the selected pair representation.
- Alternatively, TOML tables can be represented using curly-bracketed terms, `{Pairs}`, where `Pairs` uses the selected pair representation.
- Arrays are represented using lists.
- Text strings can be represented as atoms, `chars(List)`, or `codes(List)`. The default when decoding is to use atoms when using the `toml` object. To decode text strings into lists of chars or codes, use the `toml/1` object with the parameter bound to `chars` or `codes`.
- TOML booleans are represented by `@true` and `@false`.
- TOML special floats are represented by `@(inf)`, `@(-inf)`, and `@(nan)`.

Examples using the default `toml` object:

TOML	term
<code>title = "Example"</code>	<code>toml([title-'Example'])</code>
<code>point = {x = 1, y = 2}</code>	<code>toml([point-toml([x-1, y-2])])</code>
<code>[owner] name = "Tom"</code>	<code>toml([owner-toml([name-'Tom'])])</code>
<code>enabled = true</code>	<code>toml([enabled-@true])</code>
<code>limits = [1, 2, 3]</code>	<code>toml([limits-[1,2,3]])</code>

6.202.6 Examples

Parsing a TOML String

```
| ?- toml::parse(atom('title = "Example"\n[owner]\nname = "Tom"\n'), TOML).  
TOML = toml([title-'Example', owner-toml([name-'Tom'])]).
```

Parsing with Curly Table Representation

```
| ?- toml(curly,dash,atom)::parse(atom('title = "Example"\n[owner]\nname = "Tom"\n'), TOML).  
TOML = {title-'Example', owner-{name-'Tom'}}.
```

Generating TOML

```
| ?- toml::generate(atom(Atom), toml([title-'Example', owner-toml([name-'Tom'])])).  
Atom = 'title = "Example"\n\n[owner]\nname = "Tom"\n'.
```

6.202.7 Notes

Generation is canonical and semantic. The current implementation does not preserve comments, original quoting style, or whether a table was originally written as an inline table or a standard table.

6.203 toon

The toon library provides predicates for parsing and generating data in the TOON (Token-Oriented Object Notation) format. TOON is a compact, human-readable, line-oriented format that encodes the JSON data model while minimizing tokens. For more information on the TOON format, see:

<https://github.com/toon-format/toon>

6.203.1 API documentation

Open the `../..apis/library_index.html#toon` link in a web browser.

6.203.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(toon(loader)).
```

6.203.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(toon(tester)).
```

6.203.4 Representation

TOON objects are represented using the same conventions as the `json` library. The representation of TOON objects and TOON pairs can be controlled using the `toon/3` parametric object. The three parameters are:

- `ObjectRepresentation`: curly (default) or list
- `PairRepresentation`: dash (default), equal, or colon
- `StringRepresentation`: atom (default), chars, or codes

The `toon(StringRepresentation)` and `toon` objects use the default representations.

Objects

TOON objects are represented using curly terms by default:

```
{Key1-Value1, Key2-Value2, ...}
```

Or using lists when using `list` for the `ObjectRepresentation` parameter:

```
toon([Key1-Value1, Key2-Value2, ...])
```

Arrays

TOON arrays are represented as lists:

```
[Element1, Element2, ...]
```

Primitives

TOON primitive values are represented as follows:

- Strings: atoms (default), `chars(ListOfChars)`, or `codes(ListOfCodes)`
- Numbers: Prolog numbers
- Booleans: `@true` and `@false`
- Null: `@null`

6.203.5 TOON format features

TOON is a line-oriented format with the following key features:

- Uses indentation (2 spaces) instead of braces for objects
- Arrays declare their length: [N]: for inline or [N]{fields}: for tabular
- Tabular form for uniform arrays of objects with primitive values
- Minimal quoting rules for strings
- Three delimiters: comma (default), tab, and pipe
- UTF-8 encoding with LF line endings
- File extension: .toon, media type: text/toon

6.203.6 Usage examples

Parsing TOON from an atom:

```
| ?- toon::parse(atom('name: John\nage: 30'), Term).
Term = {name-'John', age-30}
yes
```

Generating TOON to an atom:

```
| ?- toon::generate(atom(Atom), {name-'John', age-30}).
Atom = 'name: John\nage: 30'
yes
```

Parsing TOON from a file:

```
| ?- toon::parse(file('data.toon'), Term).
...
```

Generating TOON to a file:

```
| ?- toon::generate(file('output.toon'), {name-'John', age-30}).
...
```

Using different representations:

```
| ?- toon(list, dash, atom)::parse(atom('name: John'), Term).
Term = toon([name-'John'])
yes
```

6.204 truncated_svd_projection

Truncated singular value decomposition reducer for continuous datasets. The library implements the `dimension_reducer_protocol` defined in the `dimension_reduction_protocols` library and learns a low-rank linear projection by building a preprocessed data matrix using optional centering and scaling, extracting singular triplets using deterministic two-sided power iteration, and applying rank-one deflation directly to the data matrix.

6.204.1 API documentation

Open the ../apis/library_index.html#truncated_svd_projection link in a web browser.

6.204.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(truncated_svd_projection(loader)).
```

6.204.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(truncated_svd_projection(tester)).
```

6.204.4 Features

- **Continuous Datasets:** Accepts datasets containing only continuous attributes. Missing or nonnumeric values are rejected.
- **Low-Rank Projection:** Extracts right singular vectors and singular values in descending order.
- **Portable SVD Solver:** Uses deterministic two-sided power iteration with rank-one deflation over the data matrix instead of backend-specific linear algebra libraries.
- **Optional Centering and Scaling:** Supports explicit control over whether training and transform inputs are centered and/or scaled.
- **Training Diagnostics:** Records per-component convergence reasons, iteration counts, and final deltas alongside the learned singular values.
- **Projection API:** Transforms a new instance into a list of `component_N-Value` pairs.
- **Model Export:** Learned reducers can be exported as predicate clauses or written to a file.

6.204.5 Options

The `learn/3` predicate accepts the following options:

- `n_components/1`: Number of singular vectors to extract. The default is 2. Requests that exceed the numerical rank after shape capping raise `domain_error(component_count, Requested-Extracted)` instead of silently returning fewer components.
- `center/1`: Whether to center the training and transform inputs using the training-set means. Options: `true` or `false` (default).
- `feature_scaling/1`: Whether to divide each continuous attribute by its training-set standard deviation. Options: `true` or `false` (default).
- `maximum_iterations/1`: Maximum number of power-iteration steps used when estimating each singular vector. The default is 1000.
- `tolerance/1`: Positive convergence tolerance used both for power-iteration stopping and for deciding when the residual matrix no longer supports the requested number of components. The default is `1.0e-8`.

The learned diagnostics also include:

- `convergence(Statuses)`: Per-component stop reasons, such as `tolerance` or `maximum_iterations_exhausted`.
- `iterations(Counts)`: Per-component iteration counts aligned with the extracted singular vectors.
- `final_delta(Deltas)`: Per-component final update magnitudes aligned with the extracted singular vectors.

6.204.6 Usage

The following examples use the sample datasets shipped with the `dimension_reduction_protocols` library:

```
| ?- logtalk_load(dimension_reduction_protocols('test_datasets/low_rank_rectangular')).
```

Learning a reducer

```
| ?- truncated_svd_projection::learn(low_rank_rectangular, DimensionReducer).
| ?- truncated_svd_projection::learn(low_rank_rectangular, DimensionReducer, [n_
↪ components(1), center(false), feature_scaling(false), maximum_iterations(200), tolerance(1.
↪ 0e-7)]).
```

Transforming new instances

```
| ?- truncated_svd_projection::learn(low_rank_rectangular, DimensionReducer),
   truncated_svd_projection::transform(DimensionReducer, [f1-1.0, f2-1.0, f3-2.0], _
↪ ReducedInstance).
```

Exporting and reusing the reducer

```
| ?- truncated_svd_projection::learn(low_rank_rectangular, DimensionReducer, [n_
↳ components(1)]),
    truncated_svd_projection::export_to_file(low_rank_rectangular, DimensionReducer, _
↳ reducer, 'truncated_svd_reducer.pl').

| ?- logtalk_load('truncated_svd_reducer.pl'),
    reducer(Reducer),
    truncated_svd_projection::transform(Reducer, [f1-1.0, f2-1.0, f3-2.0], ReducedInstance).
```

6.204.7 Dimension reducer representation

The learned dimension reducer is represented by a compound term with the functor chosen by the implementation and arity 4. For example:

```
truncated_svd_reducer(Encoders, Components, SingularValues, Diagnostics)
```

Where:

- Encoders: List of continuous attribute encoders storing attribute name, centering offset, and scale factor.
- Components: List of right singular vectors in descending singular-value order.
- SingularValues: List of singular values matching the extracted components.
- Diagnostics: Learned reducer metadata including the effective training options, singular values, and per-component convergence information.

6.204.8 References

1. Eckart, C. and Young, G. (1936) - “The approximation of one matrix by another of lower rank”.

6.205 tsv

The tsv library provides predicates for reading and writing TSV files and streams:

<https://www.iana.org/assignments/media-types/text/tab-separated-values>

The main object, tsv/2, is a parametric object allowing passing two options: header handling (keep or skip) and comment handling (false or true). When comment handling is true, lines starting with the # character are skipped when reading files and streams.

The tsv/1 parametric object is kept for backward compatibility and extends tsv/2 by setting the comments option to false. The tsv object extends the tsv/2 parametric object using the default options keep and false.

Files and streams can be read into a list of rows (with each row being represented by a list of fields) or asserted using a user-defined dynamic predicate. Reading can be done by first loading the whole file (using the read_file/2-3 predicates) into memory or line by line (using the read_file_by_line/2-3 predicates). Reading line by line is usually the best option for parsing large TSV files.

Data can be saved to a TSV file or stream by providing the object and predicate for accessing the data plus the name of the destination file or the stream handle or alias.

6.205.1 API documentation

Open the `../apis/library_index.html#tsv` link in a web browser.

6.205.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(tsv(loader)).
```

6.205.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(tsv(tester)).
```

6.205.4 Usage

A TSV file can be read as a list of rows:

```
| ?- tsv::read_file('test_files/data.tsv', Rows).

Rows = [['Name', 'Age', 'Address'], ['Paul', 23, '1115 W Franklin'], ['Bessy the Cow', 5, 'Big_
↪ Farm Way'], ['Zeke', 45, 'W Main St']]
yes
```

Alternatively, The TSV data can be saved using a public and dynamic object predicate (that must be previously declared). For example:

```
| ?- assertz(p(_,_,_)), retractall(p(_,_,_)).
yes

| ?- tsv(skip)::read_file('test_files/data.tsv', user, p/3).
yes

| ?- tsv(skip, true)::read_file('test_files/data_with_comments.tsv', user, p/3).
yes

| ?- p(A,B,C).

A = 'Paul', B = 23, C = '1115 W Franklin' ? ;
...
```

Given a predicate representing a table, the predicate data can be written to a file or stream. For example:

```
| ?- tsv::write_file('output.tsv', user, p/3).
yes
```

6.206 types

This library implements predicates over standard Prolog term types and also terms representing common data structures such as lists and pairs.

It also includes a user-extensible type object defining type checking predicates over common Logtalk and Prolog term types. The types define a hierarchy with the Prolog type `term` at the root (i.e., type-checking a predicate argument of type `term` trivially succeeds). Some types are only meaningful for backend Prolog systems supporting non-universal features (e.g., `cyclic` or `char(CharSet)` with a Unicode character set). See the API documentation for a full list of the types defined by default.

6.206.1 API documentation

Open the ../apis/library_index.html#types link in a web browser.

6.206.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(types(loader)).
```

If your code only requires the most basic types, you can load in alternative the file:

```
| ?- logtalk_load(basic_types(loader)).
```

See the notes on the `basic_types` virtual library for details.

6.206.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(types(tester)).
```

6.206.4 Type-checking

This library type object can be used to type-check common Logtalk and Prolog term types (see the object documentation for a listing of all the pre-defined types). The `valid/2` predicate succeeds or fails if a term is of a given type. For example:

```
| ?- type::valid(positive_integer, 42).
yes

| ?- type::valid(positive_integer, -13).
no
```

The `check/2` and `check/3` predicates throw an exception if a term is not of a given type. For example:

```
| ?- catch(type::check(integer, abc), Error, true).
Error = type_error(integer, abc)
yes
```

If we require a standard error/2 exception term, the `check/3` predicate takes a *context* argument. For example:

```
| ?- catch(type::check(integer, abc, foo/3), Error, true).
Error = error(type_error(integer, abc), foo/3)
yes
```

Typically, the context is provided by calling the built-in `context/1` method.

6.206.5 Defining new types

To define a custom type, define clauses for the multifile predicates `type::type/1` (to declare the type) and `type::check/2` (to type-check values). For example:

```
:- multifile(type::type/1).
type::type(age).

:- multifile(type::check/2).
type::check(age, Term) :-
    type::check(between(non_negative_integer, 0, 150), Term).
```

Be careful to ensure that the new type definitions don't introduce spurious choice-points for these predicates. The unit tests of the `types` library perform this check for pre-defined and loaded user-defined ground types.

When defining a meta-type (i.e., a type with arguments that are also types), add also a clause for the `type::meta_type/3` multifile predicate. For example:

```
:- multifile(type::meta_type/3).
type::meta_type(tuple(Type1, Type2, Type3), [Type1, Type2, Type3], []).
```

This predicate is called when checking if a type is a defined type. For meta-types, that check must extend to the sub-types.

6.206.6 Examples

See e.g. the `os` library implementation of custom types for files and directories. Or the `expecteds` and `optionals` libraries custom types. See also the `my_types` programming example.

6.207 tzif

The `tzif` library loads TZif v1, v2, and v3 files from the IANA time zone database into inspectable terms and answers UTC-based lookup queries over one or more zones.

Requires a backend Prolog compiler with unbounded integer arithmetic.

Current feature set:

- Loads sources given as `file(Path, ZoneId)`, `files(Root, Paths)`, `directory(Root)`, `stream(Stream, ZoneId)`, `bytes(Bytes, ZoneId)`, or `snapshot(File)` using `load/1` or `load/2`.
- Supports TZif v1, v2, and v3, including validated skipping of the v1 compatibility block in v2/v3 files.
- Parses POSIX footers, including signed-hour and signed-minute offsets.

- Exposes zone-aware lookup predicates over loaded `tzif(...)` terms, plus cached convenience variants for zone and single-zone queries.
- Exposes strict local civil-time lookup predicates that fail cleanly for ambiguous or nonexistent times.
- Exposes `*_with_resolution` local civil-time query predicates for explicit resolution modes (strict, first, second, and all).
- Exposes `*_reified` local civil-time query predicates that return `unique(...)`, `ambiguous(...)`, or `nonexistent` results.
- Caches successful `load/1` calls automatically, while `load/2` returns a loaded term without changing the cache.
- Provides `cache/1` for making an already loaded `tzif(...)` term the active cached term.
- Saves and reloads `tzif(...)` terms as plain Prolog snapshot files, with `save/1` providing a convenience predicate for saving the cached term.
- Returns the source term recorded in the active cached term via `cache_source/1`.
- Exposes the active cached term using `cached_tzif/1`.
- Exposes the cached zone list using `zones/1`.
- Accepts UTC queries as Unix seconds or `date_time/6` terms.
- Validates zone identifiers against bundled IANA TZDB 2026a canonical names plus backward-compatible aliases.

6.207.1 API documentation

Open the ../apis/library_index.html#tzif link in a web browser.

6.207.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(tzif(loader)).
```

6.207.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(tzif(tester)).
```

6.207.4 Representation

Loaded sources are returned by `load/2` as lists of `tzif(Zone, Source, ZoneData)` terms, one term per loaded zone.

Each nested `ZoneData` term contains the parsed block data, leap-second records, and parsed footer for a single zone. These per-zone `tzif(...)` terms are stable enough to be serialized directly using `save/2` and restored using `load(snapshot(File), TZifs)`.

For directory and file-set loads, zone identifiers are the relative paths used to locate the TZif files inside the root directory, and each loaded term records its own `file(File, ZoneId)` source.

When loading a directory(`Root`), regular files whose relative paths are not recognized zone identifiers are ignored. This allows loading system zoneinfo trees that also contain metadata files such as leapseconds.

For single-zone file, stream, and byte-list loads, the caller must provide the zone identifier explicitly using `file(Path, ZoneId)`, `stream(Stream, ZoneId)`, or `bytes(Bytes, ZoneId)`.

The cache stores one `tzif(...)` term per zone id. `load/1` and `cache/1` replace only cached entries whose zone ids match the newly loaded terms.

6.207.5 Zone identifier validation

The library validates zone identifiers using bundled data derived from the official IANA TZDB 2026a release. Accepted identifiers include both canonical zone names such as `America/New_York` and backward-compatible aliases such as `US/Eastern`.

The bundled data is generated from the IANA `zone1970.tab` and backward files and does not consult the host operating system's installed zoneinfo tree.

To refresh the bundled table for a newer IANA release, run:

```
$ ./update_zone_ids.sh 2026a
```

6.207.6 Local Queries

The library provides three sets of predicates for local civil-time queries.

Strict predicates:

- `local_time_type/2-4`
- `local_offset/2-4`
- `local_daylight_saving_time/2-4`
- `local_abbreviation/2-4`

These predicates are intended for local civil times with a unique interpretation. They fail for ambiguous fold times and nonexistent gap times.

Explicit-resolution predicates:

- `local_time_type_with_resolution/3-5`
- `local_offset_with_resolution/3-5`
- `local_daylight_saving_time_with_resolution/3-5`
- `local_abbreviation_with_resolution/3-5`

These predicates accept an explicit resolution mode argument. Supported modes are `strict`, `first`, `second`, and `all`.

- `strict`: succeed only when the local civil time has a unique interpretation
- `first`: select the earliest valid interpretation
- `second`: select the latest valid interpretation
- `all`: enumerate all valid interpretations in chronological order

Reified predicates:

- `local_time_type_reified/2-4`
- `local_offset_reified/2-4`
- `local_daylight_saving_time_reified/2-4`
- `local_abbreviation_reified/2-4`

These predicates never use failure to distinguish ambiguity from absence. Instead, they return one of the following results:

- `unique(...)`
- `ambiguous(...)`
- `nonexistent`

6.207.7 Examples

Load a single TZif payload without caching it:

```
| ?- tzif::load(bytes(Bytes, 'America/New_York'), TZifs).
```

Cache an already loaded term explicitly:

```
| ?- tzif::load(bytes(Bytes, 'America/New_York'), TZifs), tzif::cache(TZifs).
```

Load a directory tree of TZif files:

```
| ?- tzif::load(directory('/usr/share/zoneinfo'), TZifs).
```

Query a specific zone in a loaded term:

```
| ?- tzif::load(bytes(Bytes, 'America/New_York'), [TZif]).
| ?- tzif::offset(TZif, 'America/New_York', date_time(2024, 7, 1, 12, 0, 0), Offset).
```

Query a specific zone in the cached term:

```
| ?- tzif::load(directory('/usr/share/zoneinfo')).
| ?- tzif::offset('America/New_York', date_time(2024, 7, 1, 12, 0, 0), Offset).
```

Persist and reload a snapshot:

```
| ?- tzif::save(TZifs, 'snapshot.pl').
| ?- tzif::load(snapshot('snapshot.pl'), ReloadedTZifs).
```

Save the cached terms directly:

```
| ?- tzif::load(directory('/usr/share/zoneinfo')).
| ?- tzif::save('snapshot.pl').
```

Load using a backward-compatible alias:

```
| ?- tzif::load(bytes(Bytes, 'US/Eastern'), TZifs).
```

Use the automatic cache populated by `load/1`:

```
| ?- tzif::load(snapshot('snapshot.pl')).  
| ?- tzif::abbreviation(date_time(2024, 7, 1, 12, 0, 0), Abbreviation).
```

Resolve an ambiguous local civil time explicitly:

```
| ?- tzif::local_offset_with_resolution('America/New_York', date_time(2024, 11, 3, 1, 30, 0),  
↪ first, Offset).
```

Inspect whether a local civil time is unique, ambiguous, or nonexistent:

```
| ?- tzif::local_time_type_reified('America/New_York', date_time(2024, 11, 3, 1, 30, 0),  
↪ Result).
```

List cached zones directly:

```
| ?- tzif::load(directory('/usr/share/zoneinfo')).  
| ?- tzif::zones(Zones).
```

6.208 ulid

This library implements a Universally Unique Lexicographically Sortable Identifier (ULID) generator.

<https://github.com/ulid/spec>

Note that most backends provide time stamps with lower granularity than required (i.e., seconds but not milliseconds). Also note that, per spec, within the same millisecond, monotonic sort order is not guaranteed.

The generation of ULIDs uses the `/dev/urandom` random number generator when available. This includes macOS, Linux, *BSD, and other POSIX operating-systems. On Windows, a pseudo-random generator is used, but randomized using the current wall time.

ULIDs can be generated as atoms, lists of characters, or lists of character codes.

See also the `cuid2`, `ids`, `ksuid`, `snowflakeid`, `nanoid`, and `uuid` libraries.

6.208.1 API documentation

Open the `../..apis/library_index.html#ulid` link in a web browser.

6.208.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(ulid(loader)).
```

6.208.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(ulid(tester)).
```

6.208.4 Generating ULIDs

By default, ULIDs are generated as atoms. For example:

```
| ?- ulid::generate(ULID).
ULID = '01H0J31SYQXHJZWPRAKHQ6YVYH'
yes
```

To generate a ULID using a list of characters representation, use instead the `ulid/1` parametric object:

```
| ?- ulid(chars)::generate(ULID).
ULID = ['0','1','H','0','J','3','2','Y','V','5','V','S','P','K','5','P','4','5','G','G','0',
↪ '9','8','8','M','2']
yes
```

Similarly, to get a ULID using a list of character codes representation:

```
| ?- ulid(codes)::generate(ULID).
ULID = [48,49,72,48,74,51,52,66,54,48,55,57,54,49,67,82,70,65,67,51,67,67,86,82,48,66]
yes
```

It's also possible to generate ULIDs from a given timestamp, i.e. the number of milliseconds since the Unix epoch (00:00:00 UTC on January 1, 1970), using the `generate/2` predicate. The timestamp must be an integer. For example, assuming a backend Prolog system providing a `time_stamp/1` predicate returning the Unix epoch in milliseconds:

```
| ?- time_stamp(Milliseconds), ulid(atom)::generate(Seconds, ULID).
Seconds = 1684245175344, ULID = '01H0JDBQ1GAWJF35C44Y5S97DX'
yes
```

You can also use timestamp discrete components using the `generate/8` predicate. For example:

```
| ?- ulid(atom)::generate(2023, 5, 17, 16, 23, 38, 591, ULID).
ULID = '01H0N8CDAZK75C5H3BJSGS4VCQ'
yes
```

To extract the timestamp from a given ULID, use the `timestamp/2` and `timestamp/8` predicates. For example:

```
| ?- ulid(atom)::timestamp('01H0JDBQ1GAWJF35C44Y5S97DX', Milliseconds).
Milliseconds = 1684245175344
yes

| ?- ulid(atom)::timestamp('01H0N8CDAZK75C5H3BJSGS4VCQ', Year, Month, Day, Hours, Minutes, ↪
↪ Seconds, Milliseconds).
Year = 2023, Month = 5, Day = 17, Hours = 16, Minutes = 23, Seconds = 38, Milliseconds = 591
yes
```

6.208.5 Type-checking ULIDs

This library also defines a `ulid(Representation)` type for type-checking ULIDs. For example, with an atom containing invalid characters:

```
| ?- type::check(ulid(atom), '01BX5ZIKBKALTAV90EVGEMMVRY').  
uncaught exception: domain_error(ulid,'01BX5ZIKBKALTAV90EVGEMMVRY')
```

6.209 unicode_data

VivoMind Prolog Unicode Resources

6.209.1 Authors

Arun Majumdar @ VivoMind LLC

Paulo Moura @ VivoMind LLC

6.209.2 License

Creative Commons CC0 1.0 Universal (CC0 1.0) - Public Domain Dedication:

```
https://creativecommons.org/publicdomain/zero/1.0/
```

We do appreciate acknowledgment if you use these resources, however, and we also welcome contributions to improve them.

6.209.3 Website

The latest release of the VivoMind Prolog Unicode Resources is available at the URL:

```
https://github.com/VivoMind
```

At this address you can also find additional information about the VivoMind Prolog Unicode Resources and submit your bug reports and contributions.

6.209.4 Description

The VivoMind Prolog Unicode Resources are a set of files resulting from the conversion of most (but not all) official UCD 6.1 files and updated for the few changes in the 6.2 standard. The original files can be downloaded from:

```
http://www.unicode.org
```

The conversion of the UCD files resulted in a large number of Prolog tables and also a set of auxiliary predicates (described below) for accessing these tables. Other than the obvious conversion in the provided predicate names, no attempt was made to convert the identifiers used for properties and other data.

6.209.5 Requirements

Most of the auxiliary predicates assume that the de facto Prolog standard predicate `between/3` is available. Unicode code point values are represented using the ISO Prolog standard notation for hexadecimal integers. In addition, the ISO Prolog standard directives `include/1` and `ensure_loaded/1` are used in some of the files to load auxiliary files.

6.209.6 Usage

Most applications only require some of the tables present in these resources. Most of these tables define properties for ranges of code points and not for single code points, but the provided auxiliary predicates allow access for a single code point. When increased performance is required, consider using the existing tables and auxiliary predicates to generate derived tables more fit for your specific application.

6.209.7 Known issues

In the file `unicode_unihan_variant.pl`, when there's more than one variant for a code point, only the first one (as listed in the original UCD file) is returned.

The `include/1` and `ensure_loaded/1` directives are specified in the ISO Prolog standard published in 1995. But some Prolog compilers either don't implement one or both directives or have flawed implementations. Thus, you may need to change how some of the files are loaded depending on the chosen Prolog compiler. Using conditional compilation directives would help in some cases, but it would also raise portability issues on its own.

6.209.8 Acknowledgements

We thank Richard O'Keefe for helpful suggestions to improve the usability of these resources.

6.209.9 Files and API Summary

The Prolog file names are derived from the original file names by prefixing them with the `unicode_` string, converting to lower case, and replacing the camel case spelling with underscores. There are, however, two exceptions: the files and directories holding the code point categories and names.

There's also a utility file, `unicode_data.pl`, that can be used to load all the files in these resources. It is mostly used to test the portability of the code across Prolog compilers. Also included is a Logtalk version of this file, `unicode_data.lgt`, which uses Logtalk's own implementation of the `include/1` directive and the `logtak_load/1` predicate to load all files. This file can be used to workaround Prolog systems with buggy or missing implementations of the `ensure_loaded/1` and `include/1` directives.

An overview of the original file names and the code point properties can be found at:

http://www.unicode.org/reports/tr44/#Directory_Structure

http://www.unicode.org/reports/tr44/#Property_Definitions

`unicode_arabic_shaping.pl`

- Provides: `unicode_arabic_shaping/4`
- Dependencies: (none)

`unicode_bidi_mirroring.pl`

- Provides: `unicode_bidi_mirroring/2`
- Dependencies: (none)

`unicode_blocks.pl`

- Provides: `unicode_block/2-3`
- Dependencies: (none)

`unicode_case_folding.pl`

- Provides: `unicode_case_folding/3`
- Dependencies: (none)

`unicode_categories.pl`

- Provides: `unicode_category/2`
- Dependencies: files in the `unicode_categories` directory

`unicode_cjk_radicals.pl`

- Provides: `unicode_cjk_radical/3`
- Dependencies: (none)

`unicode_composition_exclusions.pl`

- Provides: `unicode_composition_exclusion/1`
- Dependencies: (none)

`unicode_core_properties.pl`

- Provides: `unicode_math/1-2` `unicode_alphabetic/1-2` `unicode_range_alphabetic/2` `unicode_lowercase/1-2` `unicode_uppercase/1-2` `unicode_cased/1-2` `unicode_case_ignorable/1-2` `unicode_changes_when_lowercased/1-2` `unicode_changes_when_uppercased/1-2` `unicode_changes_when_titlecased/1-2` `unicode_changes_when_casefolded/1-2` `unicode_changes_when_casemapped/1-2` `unicode_id_start/1-2` `unicode_id_continue/1-2` `unicode_xid_start/1-2` `unicode_xid_continue/1-2` `unicode_default_ignorable/1-2` `unicode_grapheme_extend/1-2` `unicode_grapheme_base/1-2` `unicode_grapheme_link/1-2`
- Dependencies: files in the `unicode_core_properties` directory

unicode_decomposition_type.pl

- Provides: unicode_canonical/1-2 unicode_compat/1-2 unicode_font/1-2 unicode_nobreak/1-2 unicode_initial/1-2 unicode_medial/1-2 unicode_final/1-2 unicode_isolated/1-2 unicode_circle/1-2 unicode_super/1-2 unicode_sub/1-2 unicode_vertical/1-2 unicode_wide/1-2 unicode_narrow/1-2 unicode_small/1-2 unicode_square/1-2 unicode_fraction/1-2
- Dependencies: files in the unicode_decomposition_type directory

unicode_derived_age.pl

- Provides: unicode_age/2-3
- Dependencies: (none)

unicode_derived_bidi_class.pl

- Provides: unicode_bidi_class/2-3
- Dependencies: (none)

unicode_derived_combining_class.pl

- Provides: unicode_combining_class/2-3
- Dependencies: (none)

unicode_derived_core_properties.pl

- Provides: unicode_core_property/2-3
- Dependencies: (none)

unicode_derived_decomposition_type.pl

- Provides: unicode_decomposition_type/2-3
- Dependencies: (none)

unicode_derived_east_asian_width.pl

- Provides: unicode_east_asian_width/2-3
- Dependencies: (none)

`unicode_derived_joining_group.pl`

- Provides: `unicode_joining_group/2-3`
- Dependencies: (none)

`unicode_derived_joining_type.pl`

- Provides: `unicode_joining_type/2-3`
- Dependencies: (none)

`unicode_derived_line_break.pl`

- Provides: `unicode_line_break/2-3`
- Dependencies: (none)

`unicode_derived_normalization_props.pl`

- Provides: `unicode_fc_nfkc/2` `unicode_nfkc_cf/2` `unicode_full_composition_exclusion/1-2` `unicode_nfd_qc_no/1-2` `unicode_nfc_qc_no/1-2` `unicode_nfc_qc_maybe/1-2` `unicode_nfkd_qc_no/1-2` `unicode_nfkc_qc_no/1-2` `unicode_nfkc_qc_maybe/1-2` `unicode_expands_on_nfd/1-2` `unicode_expands_on_nfc/1-2` `unicode_expands_on_nfkd/1-2` `unicode_expands_on_nfkc/1-2` `unicode_changes_when_nfkc_casefolded/1-2`
- Dependencies: files in the `unicode_derived_normalization_props` directory

`unicode_derived_numeric_type.pl`

- Provides: `unicode_numeric_type/2-3`
- Dependencies: (none)

`unicode_derived_numeric_values.pl`

- Provides: `unicode_numerical_value/3`
- Dependencies: (none)

`unicode_hangul_syllable_type.pl`

- Provides: `unicode_hangul_syllable_type/2-3`
- Dependencies: (none)

unicode_indic_matra_category.pl

- Provides: unicode_indic_matra_category/2-3
- Dependencies: (none)

unicode_indic_syllabic_category.pl

- Provides: unicode_indic_syllabic_category/2-3
- Dependencies: (none)

unicode_jamo.pl

- Provides: unicode_jamo/2
- Dependencies: (none)

unicode_name_aliases.pl

- Provides: unicode_name_alias/3
- Dependencies: (none)

unicode_names.pl

- Provides: unicode_name/2
- Dependencies: files in the unicode_names directory

unicode_prop_list.pl

- Provides: unicode_white_space/1-2 unicode_bidi_control/1-2 unicode_join_control/
1-2 unicode_dash/1-2 unicode_hyphen/1-2 unicode_quotation_mark/1-2
unicode_terminal_punctuation/1-2 unicode_other_math/1-2 unicode_hex_digit/1-2
unicode_ascii_hex_digit/1-2 unicode_other_alphabetic/1-2 unicode_ideographic/
1-2 unicode_diacritic/1-2 unicode_extender/1-2 unicode_other_lowercase/
1-2 unicode_other_uppercase/1-2 unicode_noncharacter_code_point/
1-2 unicode_other_grapheme_extend/1-2 unicode_ids_binary_operator/1-2
unicode_ids_trinary_operator/1-2 unicode_radical/1-2 unicode_unified_ideograph/1-2
unicode_other_default_ignorable/1-2 unicode_deprecated/1-2 unicode_soft_dotted/1-2
unicode_logical_order_exception/1-2 unicode_other_id_start/1-2 unicode_other_id_continue/
1-2 unicode_sterm/1-2 unicode_variation_selector/1-2 unicode_pattern_white_space/1-2
unicode_pattern_syntax/1-2
- Dependencies: files in the unicode_prop_list directory

`unicode_range_scripts.pl`

- Provides: `unicode_range_script/3` `unicode_script/2`
- Dependencies: (none)

`unicode_script_extensions.pl`

- Provides: `unicode_script_extension/2-3`
- Dependencies: `unicode_scripts.pl`

`unicode_scripts.pl`

- Provides: `unicode_script/6` `unicode_script_category/3`
- Dependencies: (none)

`unicode_special_casing.pl`

- Provides: `unicode_special_casing/5`
- Dependencies: (none)

`unicode_unihan_variants.pl`

- Provides: `unicode_unihan_variant/2-3`
- Dependencies: (none)

`unicode_version.pl`

- Provides: `unicode_version/3`
- Dependencies: (none)

6.210 `union_find`

This library implements a union-find data structure. This structure tracks a set of elements partitioned into a number of disjoint (non-overlapping) subsets. It provides fast operations to add new sets, to merge existing sets, and to determine whether elements are in the same set. This implementation of the union-find algorithm provides the following features:

- Path compression: Path compression flattens the structure of the tree by making every node point to the root whenever a find predicate is used on it.
- Union by rank: Union predicates always attach the shorter tree to the root of the taller tree. Thus, the resulting tree is no taller than the original unless they were of equal height, in which case the resulting tree is taller by one node.

For a general and extended discussion on this data structure, see e.g.

https://en.wikipedia.org/wiki/Disjoint-set_data_structure

6.210.1 API documentation

Open the `../..apis/library_index.html#union-find` link in a web browser.

6.210.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(union_find(loader)).
```

6.210.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(union_find(tester)).
```

6.210.4 Usage

An usage example is Kruskal's algorithm, a minimum-spanning-tree algorithm which finds an edge of the least possible weight that connects any two trees in the forest. It is a greedy algorithm in graph theory as it finds a minimum spanning tree for a connected weighted graph by adding increasing cost arcs at each step.

```
:- object(kruskal).

    :- public(kruskal/2).

    :- uses(union_find, [
        new/2, find/4, union/4
    ]).

    kruskal(g(Vertices-Edges), g(Vertices-Tree)) :-
        new(Vertices, UnionFind),
        keysort(Edges, Sorted),
        kruskal(UnionFind, Sorted, Tree).

    kruskal(_, [], []).
    kruskal(UnionFind0, [Edge|Edges], [Edge|Tree]) :-
        Edge = _-(Vertex1, Vertex2),
        find(UnionFind0, Vertex1, Root1, UnionFind1),
        find(UnionFind1, Vertex2, Root2, UnionFind2),
        Root1 \== Root2,
        !,
        union(UnionFind2, Vertex1, Vertex2, UnionFind3),
        kruskal(UnionFind3, Edges, Tree).
    kruskal(UnionFind, [_|Edges], Tree) :-
        kruskal(UnionFind, Edges, Tree).

:- end_object.
```

Sample query:

```
| ?- kruskal::kruskal(g([a,b,c,d,e,f,g]-[7-(a,b), 5-(a,d), 8-(b,c), 7-(b,e), 9-(b,d), 5-(c,
↪e), 15-(d,e), 6-(d,f), 8-(e,f), 9-(e,g), 11-(f,g)]), Tree).

Tree = g([a,b,c,d,e,f,g]-[5-(a,d),5-(c,e),6-(d,f),7-(a,b),7-(b,e),9-(e,g)])
yes
```

6.211 url

This library implements validation, parsing, generating, and normalization of URLs, which can be represented as atoms, character lists, or code lists. It currently supports the following URL schemes:

Web protocols:

- http
- https
- ws
- wss
- gopher

File transfer and version control:

- ftp
- ftps
- sftp
- git

File access:

- file

Databases:

- jdbc
- mongodb
- mysql
- postgresql

Email and news:

- mailto
- news
- nntp

Media streaming:

- mms
- rtmp
- rtsp

Shell access:

- ssh
- telnet

Directory services:

- ldap
- ldaps

Other protocols:

- tel
- urn

The library predicates are defined in the `url(_Representation_)` parametric object where `_Representation_` can be one of:

- atom - strings are represented as atoms
- chars - strings are represented as lists of characters
- codes - strings are represented as lists of character codes

The parameter must be bound when sending messages to the object.

The library also provides an utility predicate, `file_path_components/2`, for converting a file-system path into file URL components. This predicate also handles Windows drive-letter and UNC paths.

6.211.1 API documentation

Open the `../apis/library_index.html#url` link in a web browser.

6.211.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(url(loader)).
```

6.211.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(url(tester)).
```

6.212 uuid

This library implements a Universally Unique Identifier (UUID) generator. Currently version 1, version 3, version 4, version 5, and version 7 UUIDs are supported. Version 5 support requires a Prolog backend with support for unbounded integer arithmetic. For reference material, see e.g.

https://en.wikipedia.org/wiki/Universally_unique_identifier

Some backends provide time stamps with low granularity (e.g., seconds but not milliseconds or nanoseconds). To compensate, the generation of version 1 UUIDs uses 14 random bits for the clock sequence.

The generation of version 4 and version 7 UUIDs uses the `/dev/urandom` random number generator when available. This includes macOS, Linux, *BSD, and other POSIX operating-systems. On Windows, a pseudo-random generator is used, but randomized using the current wall time.

Version 3 and version 5 UUIDs are namespace-name based UUIDs using the MD5 and SHA-1 hash functions, respectively.

Version 7 UUIDs are time-ordered using a Unix Epoch timestamp in milliseconds, as specified in RFC 9562. They are recommended over version 1 UUIDs for new applications due to improved entropy and sortability characteristics.

UUIDs can be generated as atoms, lists of characters, or lists of character codes.

See also the `cuid2`, `ksuid`, `ids`, `nanoid`, `snowflakeid`, and `ulid` libraries.

6.212.1 API documentation

Open the ../apis/library_index.html#uuid link in a web browser.

6.212.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(uuid(loader)).
```

6.212.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(uuid(tester)).
```

6.212.4 Generating version 1 UUIDs

By default, version 1 UUIDs are generated as atoms. For example:

```
| ?- uuid::uuid_v1([0xf2,0xd1,0x90,0x94,0xdc,0x4b], UUID).
UUID = '00a66fc0-82cf-11eb-bc83-f2d19094dc4b'
yes
```

To generate a UUID using a list of characters representation, use instead the `uuid/1` parametric object:

```
| ?- uuid(chars)::uuid_v1([0xf2,0xd1,0x90,0x94,0xdc,0x4b], UUID).
UUID = ['0','0',d,e,'9','0',c,'0',-,'8','2',c,f,-,'1','1',e,b,-,
        a,'9','8','5',-,f,'2',d,'1','9','0','9','4',d,c,'4',b]
yes
```

Similarly, to get a UUID using a list of character codes representation:

```
| ?- uuid(codes)::uuid_v1([0xf2,0xd1,0x90,0x94,0xdc,0x4b], UUID).
UUID = [48,48,52,99,99,54,99,48,45,56,50,99,102,45,49,49,101,98,45,
        98,57,102,52,45,102,50,100,49,57,48,57,52,100,99,52,98]
yes
```

6.212.5 Generating version 4 UUIDs

By default, version 4 UUIDs are generated as atoms. For example:

```
| ?- uuid::uuid_v4(UUID).
UUID = '1c652782-69c5-4252-88c8-09e576a44db5'
yes
```

To generate a UUID using a list of characters representation, use instead the uuid/1 parametric object:

```
| ?- uuid(chars)::uuid_v4(UUID).
UUID = [d,'3',d,'3','3','5','1','3',-, '8','1',e,c,-, '4',d,'2','6',-,
        '9',f,'2','2',-,e,d,'9','5',e,'0','0',e,'1','5','7','0']
yes
```

Similar to get a UUID using a list of character codes representation:

```
| ?- uuid(codes)::uuid_v4(UUID).
UUID = [102,97,52,54,57,98,100,50,45,51,57,54,51,45,52,97,100,55,45,
        98,50,50,55,45,101,100,52,99,56,55,99,54,53,55,102,98]
yes
```

6.212.6 Generating version 3 UUIDs

Version 3 UUIDs are namespace-name based UUIDs using the MD5 hash function. For example:

```
| ?- uuid::uuid_v3('6ba7b810-9dad-11d1-80b4-00c04fd430c8', 'www.widgets.com', UUID).
UUID = '3d813cbb-47fb-32ba-91df-831e1593ac29'
yes
```

To generate a UUID using a list of characters representation, use instead the uuid/1 parametric object:

```
| ?- uuid(chars)::uuid_v3('6ba7b810-9dad-11d1-80b4-00c04fd430c8', 'www.widgets.com', UUID).
UUID = ['3',d,'8','1','3',c,b,b,-, '4','7',f,b,-, '3','2',b,a,-,
        '9','1',d,f,-, '8','3','1',e,'1','5','9','3',a,c,'2','9']
yes
```

Similarly, to get a UUID using a list of character codes representation:

```
| ?- uuid(codes)::uuid_v3('6ba7b810-9dad-11d1-80b4-00c04fd430c8', 'www.widgets.com', UUID).
UUID = [51,100,56,49,51,99,98,98,45,52,55,102,98,45,51,50,98,97,45,
        57,49,100,102,45,56,51,49,101,49,53,57,51,97,99,50,57]
yes
```

6.212.7 Generating version 5 UUIDs

Version 5 UUIDs are namespace-name based UUIDs using the SHA-1 hash function. This predicate is only available on Prolog backends with support for unbounded integer arithmetic. For example:

```
| ?- uuid::uuid_v5('6ba7b810-9dad-11d1-80b4-00c04fd430c8', 'www.widgets.com', UUID).
UUID = '21f7f8de-8051-5b89-8680-0195ef798b6a'
yes
```

To generate a UUID using a list of characters representation, use instead the `uuid/1` parametric object:

```
| ?- uuid(chars)::uuid_v5('6ba7b810-9dad-11d1-80b4-00c04fd430c8', 'www.widgets.com', UUID).
UUID = ['2','1',f,'7',f,'8',d,e,-,'8','0','5','1',-,'5',b,'8','9',-,'8',
        '6','8','0',-,'0','1','9','5',e,f,'7','9','8',b,'6',a]
yes
```

Similarly, to get a UUID using a list of character codes representation:

```
| ?- uuid(codes)::uuid_v5('6ba7b810-9dad-11d1-80b4-00c04fd430c8', 'www.widgets.com', UUID).
UUID = [50,49,102,55,102,56,100,101,45,56,48,53,49,45,53,98,56,57,45,
        56,54,56,48,45,48,49,57,53,101,102,55,57,56,98,54,97]
yes
```

6.212.8 Generating version 7 UUIDs

Version 7 UUIDs are time-ordered using the Unix Epoch timestamp in milliseconds (as specified in RFC 9562). By default, version 7 UUIDs are generated as atoms. For example:

```
| ?- uuid::uuid_v7(UUID).
UUID = '018d5f3c-9b5a-7c4e-8f2a-1b3c4d5e6f70'
yes
```

To generate a UUID using a list of characters representation, use instead the `uuid/1` parametric object:

```
| ?- uuid(chars)::uuid_v7(UUID).
UUID = ['0','1','8',d,'5',f,'3',c,-,'9',b,'5',a,-,'7',c,'4',e,-,'8',
        'f',f,'2',a,-,'1',b,'3',c,'4',d,'5',e,'6',f,'7','0']
yes
```

Similar to get a UUID using a list of character codes representation:

```
| ?- uuid(codes)::uuid_v7(UUID).
UUID = [48,49,56,100,53,102,51,99,45,57,98,53,97,45,55,99,52,101,45,
        56,102,50,97,45,49,98,51,99,52,100,53,101,54,102,55,48]
yes
```

6.212.9 Generating the Nil and Max UUIDs

Predicates are also provided that return the Nil and Max UUIDs:

```
| ?- uuid::uuid_nil(UUID).
UUID = '00000000-0000-0000-0000-000000000000'
yes

| ?- uuid::uuid_max(UUID).
UUID = 'FFFFFFFF-FFFF-FFFF-FFFF-FFFFFFFFFFFF'
yes
```

6.213 validations

This library provides an implementation of *validation terms* with an API for applicative-style error accumulation. A validation term is an opaque compound term that either contains a valid value (`valid(Value)`) or a list of errors (`invalid(Errors)`).

The library follows the same design pattern as the optionals and expecteds libraries:

- `validation` — factory object for constructing validation terms
- `validation/1` — parametric object wrapping a validation term
- `validated` — companion helper object for list operations, type-checking, and QuickCheck support

The naming follows the convention used in Scala Cats and Kotlin Arrow, where the companion type is called `Validated`.

6.213.1 API documentation

Open the ../apis/library_index.html#validations link in a web browser.

6.213.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(validations(loader)).
```

6.213.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(validations(tester)).
```

6.213.4 Usage

The validation object provides constructors for validation terms. For example:

```
| ?- validation::of_valid(1, Validation).
...
```

The created validation terms can then be passed as parameters to the validation/1 parametric object. For example:

```
| ?- validation::of_valid(1, V1), validation::of_valid(2, V2),
    validation::sequence([V1, V2], Validation).
Validation = valid([1,2])
yes

| ?- validation::of_invalid(e1, V1), validation::of_invalid(e2, V2),
    validation::sequence([V1, V2], Validation).
Validation = invalid([e1,e2])
yes

| ?- validation(valid(1))::flat_map([X,V]>>(Y is X+1, V = valid(Y)), Validation).
Validation = valid(2)
yes
```

Validation terms can be constructed from goals, optionals, and expecteds:

```
| ?- validation::from_goal(succ(1, Value), Value, Validation).
Validation = valid(2)
yes

| ?- validation::from_optional(empty, no_value, Validation).
Validation = invalid([no_value])
yes

| ?- validation::from_expected(unexpected(e), Validation).
Validation = invalid([e])
yes
```

Validation terms can be converted to optionals and expecteds:

```
| ?- validation(valid(42))::to_optional(Optional).
Optional = optional(42)
yes

| ?- validation(invalid([e1,e2]))::to_expected(Expected).
Expected = unexpected([e1,e2])
yes
```

The key feature of validation is error accumulation via zip/3:

```
| ?- validation::of_valid(1, V1), validation::of_invalid(e1, V2),
    validation(V1)::zip([_,_,_]>>true, V2, Result).
Result = invalid([e1])
yes
```

(continues on next page)

(continued from previous page)

```
| ?- validation::of_invalid(e1, V1), validation::of_invalid(e2, V2),
    validation(V1)::zip([_,_,_]>>true, V2, Result).
Result = invalid([e1,e2])
yes
```

The companion validated object provides predicates for handling lists of validation terms, including `valids/2`, `invalids/2`, `partition/3`, and `map/3-4`.

```
| ?- validation::of_valid(1, V1), validation::of_invalids([e1,e2], V2), validation::of_
    ↪valid(2, V3),
    validated::partition([V1, V2, V3], Values, Errors).
Values = [1,2],
Errors = [e1,e2]
yes

| ?- validated::map([Term,Validation]>>(
    integer(Term) -> validation::of_valid(Term, Validation)
    ; validation::of_invalid(not_integer(Term), Validation)
    ), [1,a,2,b], ValuesErrors).
ValuesErrors = [1,2]-[not_integer(a),not_integer(b)]
yes

| ?- validated::map([Term,Validation]>>(
    integer(Term) -> validation::of_valid(Term, Validation)
    ; validation::of_invalid(not_integer(Term), Validation)
    ), [1,a,2,b], Values, Errors).
Values = [1,2],
Errors = [not_integer(a),not_integer(b)]
yes
```

6.213.5 Comparison with the expecteds library

Both this library and the `expecteds` library wrap computations that may succeed or fail. The fundamental difference is the **error model**:

- `expecteds` carries a *single* error and **short-circuits** on the first failure (monadic error handling). Once an unexpected term is produced, subsequent operations are skipped.
- `validations` carries a *list* of errors and **accumulates all failures** (applicative error handling). Operations such as `zip/3`, `sequence/2`, and `traverse/3` collect every error.

Use `expecteds` when you only care about the first error. Use `validations` when you want to **report all problems at once** (e.g. form validation, configuration checking, or batch input processing).

Example: form validation with error accumulation

Consider validating a form with a name (must be an atom) and an age (must be a positive integer). Each field is validated independently and the results are combined with `validated::sequence/2`, which accumulates all errors:

```
| ?- Name = 123, Age = young,
    validation::from_goal(atom(Name), Name, invalid_name, V1),
    validation::from_goal(integer(Age), Age > 0, Age, invalid_age, V2),
    validated::sequence([V1, V2], Result),
    validation(Result)::or_else_throw(_).
uncaught exception: [invalid_name,invalid_age]
```

Both errors are reported. When all fields are valid, `or_else_throw/1` returns the list of valid values:

```
| ?- Name = john, Age = 25,
    validation::from_goal(atom(Name), Name, invalid_name, V1),
    validation::from_goal(integer(Age), Age > 0, Age, invalid_age, V2),
    validated::sequence([V1, V2], Result),
    validation(Result)::or_else_throw(Values).
Values = [john,25]
yes
```

Compare with the `expecteds` library, where `either::sequence/2` short-circuits at the first failure:

```
| ?- Name = 123, Age = young,
    expected::from_goal(atom(Name), Name, invalid_name, E1),
    expected::from_goal(integer(Age), Age > 0, Age, invalid_age, E2),
    either::sequence([E1, E2], Result),
    expected(Result)::or_else_throw(_).
uncaught exception: invalid_name
```

Only `invalid_name` is reported — `either::sequence/2` stops at the first unexpected term.

The two libraries are complementary. Use `expecteds` for sequential pipelines where only the first error matters. Use `validations` with `sequence/2`, `traverse/3`, or `zip/3` for independent validations where all errors should be collected.

6.213.6 See also

The `expecteds` and `optionals` libraries.

6.214 wkt_wkb

The `wkt_wkb` library provides predicates for parsing, generating, and validating geometries represented using the Well-Known Text (WKT) and Well-Known Binary (WKB) interchange formats.

It complements the `geojson` and `geospatial` libraries by reusing the same geometry constructors while adding support for the standard text and binary geometry encodings widely used by spatial databases and GIS tooling.

6.214.1 API documentation

Open the ../apis/library_index.html#wkt_wkb link in a web browser.

6.214.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(wkt_wkb(loader)).
```

6.214.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(wkt_wkb(tester)).
```

6.214.4 Representation

The library uses the same native geometry constructors as the `geojson` library for geometry values:

- `point(Position)` or `point(Position, Options)`
- `multi_point(Positions)` or `multi_point(Positions, Options)`
- `line_string(Positions)` or `line_string(Positions, Options)`
- `multi_line_string(LineStrings)` or `multi_line_string(LineStrings, Options)`
- `polygon(Rings)` or `polygon(Rings, Options)`
- `multi_polygon(Polygons)` or `multi_polygon(Polygons, Options)`
- `geometry_collection(Geometries)` or `geometry_collection(Geometries, Options)`

Positions are represented as lists of numeric coordinates. Empty geometries use empty lists, e.g. `point([])` or `polygon([])`.

The optional `Options` lists currently support:

- `dimensions(xy)`
- `dimensions(z)`
- `dimensions(m)`
- `dimensions(zm)`

The `dimensions/1` option is mainly required to distinguish `Z` and `M` three-coordinate geometries and to preserve dimensionality for empty geometries. When the option is omitted, dimensionality is inferred from the coordinates and defaults to `xy` for empty geometries.

6.214.5 Sources and sinks

WKT sources and sinks are wrapped as `wkt(...)` terms:

- `wkt(file(Path))`
- `wkt(stream(Stream))`
- `wkt(atom(Atom))`
- `wkt(chars(List))`
- `wkt(codes(List))`

WKB sources and sinks are wrapped as `wkb(...)` terms:

- `wkb(file(Path))`
- `wkb(stream(Stream))`
- `wkb(bytes(List))`
- `wkb(hex(atom(Atom)))`
- `wkb(hex(chars(List)))`
- `wkb(hex(codes(List)))`

When generating WKB, byte order defaults to little-endian. A specific byte order can be selected using `wkb(Sink, little)` or `wkb(Sink, big)`.

6.214.6 Examples

Parse a WKT point:

```
| ?- wkt_wkb::parse(wkt(atom('POINT Z (1 2 3)')), Geometry).
Geometry = point([1, 2, 3], [dimensions(z)])
yes
```

Generate canonical WKT for a polygon:

```
| ?- wkt_wkb::generate(wkt(atom(WKT)), polygon([[0,0],[1,0],[1,1],[0,1],[0,0]])).
WKT = 'POLYGON ((0 0, 1 0, 1 1, 0 1, 0 0))'
yes
```

Generate WKB hex for a point:

```
| ?- wkt_wkb::generate(wkb(hex(atom(Hex))), point([1, 2])).
Hex = '010100000000000000000000f03f0000000000000040'
yes
```

Round-trip through a big-endian WKB byte sequence:

```
| ?- wkt_wkb::generate(wkb(bytes(Bytes), big), line_string([[0,0],[1,1]])),
   wkt_wkb::parse(wkb(bytes(Bytes)), Geometry).
```

Validate a geometry term:

```
| ?- wkt_wkb::validate(multi_point([[], [1, 2], [3, 4]])).
yes

| ?- wkt_wkb::validate(polygon([[[0,0],[1,0],[1,1],[0,1]]]), Errors).
Errors = [ring_not_closed([coordinates,0])]
yes
```

6.214.7 Notes

- WKT parsing accepts the standard geometry keywords POINT, LINESTRING, POLYGON, MULTIPOINT, MULTILINESTRING, MULTIPOLYGON, and GEOMETRYCOLLECTION, together with the Z, M, and ZM dimension tags.
- Canonical WKT generation always uses uppercase geometry keywords.
- WKB generation uses the ISO/OGC type-code offsets 1000, 2000, and 3000 for Z, M, and ZM dimensionality.
- Empty WKB points are encoded using quiet NaN ordinates as required by the format, while empty non-point geometries are encoded using zero counts.
- Geometry validation checks coordinate dimensionality consistency, minimum cardinalities for line strings and rings, and polygon ring closure.

6.215 yaml

The yaml library provides predicates for parsing and generating data in the YAML format:

<https://yaml.org>

It includes parsing and generation of simple scalar values, lists, mappings, and nested structures.

6.215.1 API documentation

Open the [../apis/library_index.html#yaml](#) link in a web browser.

6.215.2 Loading

To load all entities in this library, load the loader.lgt file:

```
| ?- logtalk_load(yaml(loader)).
```

6.215.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(yaml(tester)).
```

6.215.4 API

Parsing

The `parse/2` predicate parses YAML content from various sources into a structured Logtalk term representation.

```
parse(+Source, -YAML)
```

Sources can be:

- `file(Path)` - parse YAML from a file
- `stream(Stream)` - parse YAML from an open stream
- `codes(Codes)` - parse YAML from a list of character codes
- `chars(Chars)` - parse YAML from a list of characters
- `atom(Atom)` - parse YAML from an atom

The parsed YAML is returned as a compound term with the functor `yaml/1`.

Generating

The `generate/2` predicate generates YAML output from a structured term.

```
generate(+Sink, +YAML)
```

Sinks can be:

- `file(Path)` - write YAML to a file
- `stream(Stream)` - write YAML to an open stream
- `codes(Codes)` - get YAML as a list of character codes
- `chars(Chars)` - get YAML as a list of characters
- `atom(Atom)` - get YAML as an atom

The input must be a compound term representing YAML data (e.g., `yaml([...])` for mappings, `[...]` for sequences).

Multi-Document Parsing

The `parse_all/2` predicate parses all YAML documents from a source into a list.

```
parse_all(+Source, -YAMLS)
```

Documents are separated by `---` markers and optionally terminated by `...` markers. Returns a list of YAML terms, even for single-document sources.

Multi-Document Generation

The `generate_all/2` predicate generates YAML output with multiple documents.

```
generate_all(+Sink, +YAMLS)
```

Takes a list of YAML terms and generates output with `---` separators between documents. For a single document, no separator is added.

6.215.5 Data Representation

YAML data is represented using the following mappings:

YAML Value	Logtalk Representation
null	'@'(null)
true	'@'(true)
false	'@'(false)
Number (42)	42
String ("text")	text (as atom)
Array/List	[item1, item2, ...]
Mapping/Object	yaml([key1-val1, ...])

6.215.6 Examples

Parsing a Simple YAML String

```
?- parse(atom('name: John\nage: 30'), YAML).
YAML = yaml([name-'John', age-30]).
```

Parsing a YAML List

```
?- parse(codes([0'-, 0' , 0'a, ...]), YAML).
YAML = yaml([apple, banana, orange]).
```

Parsing a YAML File

```
?- parse(file('config.yaml'), YAML).  
% Reads and parses config.yaml
```

Generating YAML from a Term

```
?- generate(atom(Result), yaml([title-'Project', version-1])).  
Result = '{title:Project,version:1}'.
```

Generating YAML to a File

```
?- Data = yaml([name-'Alice', score-95]),  
   generate(file('output.yaml'), Data).  
% Writes YAML data to output.yaml
```

Working with Nested Structures

```
?- YAML = yaml([  
    person-yaml([  
        name-'Bob',  
        age-25,  
        skills-[logtalk, prolog, python]  
    ])  
]),  
   generate(atom(Result), YAML).  
% Generates YAML representation of nested structure
```

6.215.7 Features

- **Scalar values:** Supports atoms, numbers, booleans, and null
- **Collections:** Supports lists (sequences) and mappings (objects)
- **Nested structures:** Fully supports nested mappings and sequences
- **Comments:** Handles YAML comments (lines starting with #)
- **Multiple source formats:** Parse from files, streams, codes, chars, or atoms
- **Multiple output formats:** Generate to files, streams, codes, chars, or atoms
- **Error handling:** Provides appropriate error messages for invalid input

6.215.8 Supported YAML Features

The library currently supports a subset of YAML 1.2 features suitable for common configuration and data serialization tasks:

Scalars

- Strings: unquoted, single-quoted ('...'), and double-quoted ("...")
- Multi-line flow scalars: quoted strings spanning multiple lines with proper line folding
 - Single newline becomes a space
 - Blank lines (consecutive newlines) become literal newlines
 - Leading/trailing whitespace around newlines is trimmed
- Multi-word keys: keys containing spaces (e.g., Mark McGwire: 65)
- Numbers:
 - Integers: decimal (42), signed positive (+12345), octal (0o14), hexadecimal (0xC)
 - Floats: decimal (3.14), scientific notation (1.23e+3)
 - Special values: .inf, -.inf, .nan (represented as '@'(inf), '@'(-inf), '@'(nan))
- Booleans: true and false
- Null: null or empty values
- Timestamps: ISO 8601 dates and times (parsed as strings)

Collections

- Block sequences: using - indicator with proper indentation
- Flow sequences: using [] syntax with comma separation
- Block mappings: using key: value syntax with proper indentation
- Flow mappings: using { key: value } syntax
- Nested structures: arbitrary nesting of sequences and mappings
- Trailing commas: supported in flow collections

Block Scalars

- Literal block scalars: | style preserves newlines exactly
- Folded block scalars: > style folds newlines into spaces
- Chomping indicators: - (strip), + (keep), default (clip) for trailing newlines
- More-indented lines: preserved with original indentation in folded scalars
- Blank lines: preserved in both literal and folded block scalars

Aliases and Aliases

- Anchor definitions: &name to define a reusable node
- Alias references: *name to reference a previously defined anchor
- Merge key: << to merge mappings from aliases
- Works with scalars, sequences, and mappings

Document Structure

- Document start marker: `---` is recognized and skipped
- Document end marker: `...` is recognized and handled
- Multi-document streams: multiple `---` or `...` separated documents (via `parse_all/2`)
- Comments: lines starting with `#` and inline comments after values
- Tags: `!tag`, `!!type`, and `!<uri>` tags are recognized and skipped

Escape Sequences (in double-quoted strings)

- `\\` - backslash
- `\"` - double quote
- `\/` - forward slash
- `\n` - newline
- `\t` - tab
- `\r` - carriage return
- `\b` - backspace
- `\f` - form feed
- `\0` - null character
- `\xXX` - hex escape (2 hex digits)
- `\uXXXX` - Unicode escape (4 hex digits)

6.215.9 Limitations

The following YAML features are **not** currently supported:

- Tag interpretation: Tags are recognized and skipped, but not interpreted
- Complex keys: `?` indicator for multi-line or complex keys
- Multi-line plain scalars: plain scalars spanning multiple lines
- Aliases in flow sequences: `[*alias1, *alias2]` (aliases work in block context)
- Merge with list of aliases: `<<: [*a, *b]` (single alias merge works)
- Directives: `%YAML` and `%TAG` directives

6.215.10 Error Handling

The library raises the following errors:

- `instantiation_error` - when a required argument is a variable
- `domain_error(yaml_source, Source)` - when an invalid source is provided
- `domain_error(yaml_sink, Sink)` - when an invalid sink is provided
- `domain_error(yaml_term, Term)` - when an invalid YAML term is provided

6.215.11 File Organization

- `yaml_protocol.lgt` - Defines the protocol for YAML parser/generator
- `yaml.lgt` - Main implementation object
- `loader.lgt` - Library loader
- `tester.lgt` - Test runner
- `test_files/` - Directory containing test files and sample YAML files for testing

6.215.12 Test Files

The library includes several sample YAML files for testing:

- `simple.yaml` - Simple key-value pairs
- `list.yaml` - A YAML list
- `nested.yaml` - Nested structures
- `complex.yaml` - Lists of mappings
- `config.yaml` - Configuration file example
- `metadata.yaml` - Metadata example with quoted strings

6.215.13 Performance Considerations

- Parsing is done recursively, which may affect performance with deeply nested structures
- Large YAML files are processed in memory
- For very large files, consider splitting them into smaller chunks

6.215.14 Integration with Other Libraries

The YAML library can be integrated with other Logtalk libraries:

- Use with `os` library for file path handling
- Use with `reader` library for stream management
- Compatible with `term_io` library for term serialization

6.215.15 Future Enhancements

Potential future enhancements may include:

- Support for tag interpretation (e.g., `!!binary` for base64 decoding)
- Support for complex keys using `?` indicator
- Support for multi-line plain scalars
- Aliases in flow sequences
- Merge key with list of aliases
- Parametric objects for custom representation choices

6.215.16 See Also

- `json_lines` library for JSON Lines format parsing and generation
- Logtalk documentation on DCGs for grammar rules
- YAML specification: <https://yaml.org>

6.216 `z_score_anomaly_detector`

Statistical Z-score anomaly detector for continuous datasets. It is a statistical anomaly-detection method based on standard scores: the detector estimates a population mean and standard deviation for each continuous attribute and supports two learn-time score modes, `root_mean_square` for dense multivariate deviation scores and `any_feature_extreme` for the maximum absolute Z-score when sparse single-feature anomalies are more informative.

The library implements the `anomaly_detector_protocol` defined in the `anomaly_detection_protocols` library. It learns a detector from a continuous dataset, computes anomaly scores for new instances, predicts normal or anomaly, and exports learned detectors as clauses or files.

Datasets are represented as objects implementing the `anomaly_dataset_protocol` protocol from the `anomaly_detection_protocols` library. See the `anomaly_detection_protocols/test_datasets` directory for examples.

6.216.1 API documentation

Open the `../..apis/library_index.html#z-score` link in a web browser.

6.216.2 Loading

To load this library, load the `loader.lgt` file:

```
| ?- logtalk_load(z_score_anomaly_detector(loader)).
```

6.216.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(z_score_anomaly_detector(tester)).
```

6.216.4 Features

- **Statistical method:** implements anomaly detection based on standard scores, using the population mean and standard deviation of each continuous attribute to measure how far new observations deviate from the training data distribution.
- **Classical per-attribute Z-score:** for each known attribute value x , the library computes the standard score $z = (x - \mu) / \sigma$, where μ is the learned population mean for that attribute and σ is the learned population standard deviation.
- **Continuous features only:** accepts datasets whose declared attributes are all continuous.

- **Population statistics:** reuses the statistics library population object to compute per-attribute arithmetic means and standard deviations.
- **Baseline training selection:** supports learn-time `baseline_class_values(ClassValues)` and `baseline_selection_policy(Policy)` options. The default baseline class values are `[normal]`. The default reject policy throws an error if non-baseline examples are present, while `filter` removes them before fitting.
- **Missing-value tolerant:** ignores missing values when fitting attribute statistics. During scoring, queries must provide at least one known value. In the default `score_mode(root_mean_square)`, the raw score is normalized by the number of known values so that scores remain comparable across different missing-value patterns.
- **Configurable scoring semantics:** supports both dense multivariate deviation scoring using `score_mode(root_mean_square)` and sparse anomaly detection using `score_mode(any_feature_extreme)`. The default root-mean-square mode reuses the numberlist library Euclidean norm predicate as part of the computation. The `score_mode/1` option only controls how the per-attribute Z-scores are aggregated into a single raw anomaly score.
- **Bounded scoring:** maps the raw multivariate Z-score to `[0.0, 1.0)` using `Score = Raw / (1 + Raw)`.
- **Default threshold:** the default `anomaly_threshold(0.70)` provides a practical out-of-the-box cutoff for the built-in anomaly fixtures while remaining overrideable in `learn/3` and `predict/4`.
- **Learn-time score mode:** `score_mode/1` is recorded in the learned detector and reused for subsequent scoring and prediction. Passing a `score_mode/1` option to `predict/4` does not override the learned mode.
- **All-missing queries rejected:** scoring and prediction throw a `domain_error(non_empty_known_values, AttributeNames)` exception when every declared feature is missing in the query.
- **Featureless datasets rejected:** datasets must declare at least one continuous feature; otherwise `learn/2-3` throws a `domain_error(non_empty_features, Dataset)` exception.
- **Detector export:** learned detectors can be exported as predicate clauses.
- **Explicit validation and diagnostics:** supports the shared `check_anomaly_detector/1`, `valid_anomaly_detector/1`, `diagnostics/2`, `diagnostic/2`, and `anomaly_detector_options/2` predicates.

6.216.5 Options

The following options are supported by the public API:

- `anomaly_threshold(Threshold)`: Threshold for `predict/3-4` (default: `0.70`)
- `baseline_class_values(ClassValues)`: Learn-time class labels that are admissible for baseline fitting (default: `[normal]`)
- `baseline_selection_policy(Policy)`: Learn-time handling of examples whose class is not listed in `baseline_class_values/1`. Supported values are `filter` and `reject` (default: `reject`)
- `score_mode(Mode)`: Learn-time score aggregation mode for `learn/3`. Supported values are `root_mean_square` and `any_feature_extreme` (default: `root_mean_square`). If passed to `predict/4`, it is ignored and the value stored in the learned detector is used.

6.216.6 Detector representation

The learned detector is represented by default as:

```
z_score_detector(TrainingDataset, Encoders, Diagnostics)
```

Where:

- TrainingDataset: training dataset object identifier
- Encoders: list of zscore(Attribute, Mean, Scale) records
- Diagnostics: learned metadata terms including model/1, training_dataset/1, attribute_names/1, feature_count/1, example_count/1, and options/1

When exported using export_to_clauses/4 or export_to_file/4, this detector term is serialized directly as the single argument of the generated predicate clause so that the exported model can be loaded and reused as-is.

6.216.7 Notes

Scoring has three stages. First, the detector computes one classical per-attribute Z-score for each known attribute value using $z = (x - \mu) / \sigma$. Second, those per-attribute Z-scores are aggregated into a single raw anomaly score according to the learned score_mode/1 option. Third, the raw score is mapped to the interval $[0.0, 1.0]$ using $\text{Score} = \text{Raw} / (1 + \text{Raw})$.

With this normalization, a raw score of 3.0 maps to 0.75.

The score_mode/1 option does not change the classical per-attribute formula. It only changes the aggregation step. With score_mode(root_mean_square), the raw score is the root mean square of the per-attribute Z-scores. With score_mode(any_feature_extreme), the raw score is the maximum absolute per-attribute Z-score.

The baseline_class_values/1 option declares which dataset class labels are admissible for fitting the baseline means and standard deviations. The baseline_selection_policy/1 option then controls what happens when other labels are present in the training data. The default reject policy raises a domain_error(baseline_only_training_data, Dataset) exception when any non-baseline example is found. The filter policy removes non-baseline examples before fitting.

Attributes with zero observed dispersion are assigned a fallback scale of 1.0. This keeps the detector well-defined for singleton datasets or constant columns while still yielding zero score for matching values and positive scores for deviating values.

The root-mean-square aggregation keeps the default threshold stable as the number of observed dimensions grows and avoids penalizing partially observed queries solely for having fewer known attributes.

Use score_mode(any_feature_extreme) when a single extreme feature should be sufficient to flag an anomaly in high-dimensional data.

6.217 zippers

This library implements zippers over lists.

6.217.1 API documentation

Open the ../apis/library_index.html#zippers link in a web browser.

6.217.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(zippers(loader)).
```

6.217.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(zippers(tester)).
```


The documentation of each port can also be found in the port directory in the `NOTES.md` file.

7.1 `fcube`

This folder contains a Logtalk port of FCube: An Efficient Prover for Intuitionistic Propositional Logic available from:

```
https://www.vidal-rosset.net/fCube/
```

The port includes portability changes (notably, operator names) plus changes to use an ordered representation for sets. Also some code formatting changes for Logtalk coding guidelines.

The port tests are adapted from the examples available from the web page above.

To load this port and for sample queries, please see the `SCRIPT.txt` file.

For more information about FCube, see the following paper:

```
@InProceedings{10.1007/978-3-642-16242-8_21,  
  author="Ferrari, Mauro and Fiorentini, Camillo and Fiorino, Guido",  
  editor="Fermüller, Christian G. and Voronkov, Andrei",  
  title="fCube: An Efficient Prover for Intuitionistic Propositional Logic",  
  booktitle="Logic for Programming, Artificial Intelligence, and Reasoning",  
  year="2010",  
  publisher="Springer Berlin Heidelberg",  
  address="Berlin, Heidelberg",  
  pages="294--301",  
  isbn="978-3-642-16242-8"  
}
```

For sample queries, please see the `SCRIPT.txt` file.

7.1.1 API documentation

Open the `../..apis/library_index.html#fcube` link in a web browser.

7.1.2 Loading

To load all entities in this port, load the `loader.lgt` file:

```
| ?- logtalk_load(fcube(loader)).
```

7.1.3 Testing

To test this port predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(fcube(tester)).
```

7.2 metagol

This folder contains a Logtalk port of `metagol`, an inductive logic programming (ILP) system based on meta-interpretive learning available from:

```
https://github.com/metagol/metagol
```

See the original code git repo for details and bibliography on `Metagol` and ILP.

The port allows any number of datasets to be loaded simultaneously with per-dataset learning options. A dataset is simply wrapped in an object that extends and is expanded by the `metagol` object as illustrated by the ported examples.

Both the original code and the port requires the coroutining `when/2` predicate, which is only available in some backend Prolog systems. The port currently supports ECLiPSe, XVM, SICStus Prolog, SWI-Prolog, and YAP. It can be used on both POSIX and Windows operating-systems.

The examples are ported from the original `Metagol` distribution. Some of the examples are taken from the following paper (with the original Prolog examples source code files made available by `MystikNinja`):

```
@article{DBLP:journals/jair/EvansG18,
  author    = {Richard Evans and Edward Grefenstette},
  title     = {Learning Explanatory Rules from Noisy Data},
  journal   = {J. Artif. Intell. Res.},
  volume    = {61},
  pages     = {1--64},
  year      = {2018},
  url       = {https://doi.org/10.1613/jair.5714},
  doi       = {10.1613/jair.5714},
  timestamp = {Mon, 21 Jan 2019 15:01:17 +0100},
  biburl    = {https://dblp.org/rec/bib/journals/jair/EvansG18},
  bibsource = {dblp computer science bibliography, https://dblp.org}
}
```

The paper can be downloaded at

```
https://arxiv.org/pdf/1711.04574.pdf
```

For sample queries, please see the `SCRIPT.txt` file.

7.2.1 API documentation

Open the `../apis/library_index.html#metagol` link in a web browser.

7.2.2 Loading

To load all entities in this port, load the `loader.lgt` file:

```
| ?- logtalk_load(metagol(loader)).
```

7.2.3 Testing

To test this port predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(metagol(tester)).
```

There are three lengthy tests that only run when the tests are being run manually instead of automatically.

7.3 toychr

This folder contains a Logtalk port of ToyCHR, a reference implementation of Constraint Handling Rules (CHR) available from:

```
https://www.comp.nus.edu.sg/~gregory/toychr/
```

The port is work in progress and includes significant modifications to the original code:

- Instead of compiling `.chr` files, it uses the term-expansion mechanism, by defining `toychrdb` as a hook object, to support writing rules inside objects and categories. As a consequence, the original `chr_compile/1` is not available.
- The port is portable and should run on all supported backends.

The port also includes examples ported from the SWI-Prolog CHR package examples and documentation. These examples are ported using the same license of the original code (BSD-2-Clause).

For sample queries, please see the `SCRIPT.txt` file.

7.3.1 API documentation

Open the `../..apis/library_index.html#toychr` link in a web browser.

7.3.2 Loading

To load all entities in this port, load the `loader.lgt` file:

```
| ?- logtalk_load(toychr(loader)).
```

7.3.3 Testing

To test this port predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(toychr(tester)).
```

CONTRIBUTIONS

The documentation of contribution can also be found in the contribution directory in the `NOTES.md` file.

8.1 flags

Flags was developed at Katholieke Universiteit Leuven

SPDX-FileCopyrightText: 2010 Katholieke Universiteit Leuven
SPDX-License-Identifier: Artistic-2.0

Contributions to this file:

Author: Theofrastos Mantadelis

Suggestions: Paulo Moura

Version: 1

Date: 27/11/2010

8.1.1 API documentation

Open the `../..apis/library_index.html#flags` link in a web browser.

8.1.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
?- logtalk_load(flags(loader)).
```

8.1.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(flags(tester)).
```

8.2 iso8601

ISO 8601 (and European civil calendar) compliant library of date and time (clock) related predicates. That is, an ISO 8601 handler.

The library supports date and time-of-day representations plus ISO 8601 duration and interval string conversions using `time_string/3`, `date_time_string/3`, `duration_string/2`, and `interval_string/2`.

Supported **time-of-day strings** include basic and extended forms with the required `T` prefix, e.g. `T143000`, `T14:30:00`, `T143000.125`, `T14:30:00.125`, `T143000,125`, and `T14:30:00,125`.

Supported **combined date-time strings** include calendar-date, ordinal-date, and complete week-date basic and extended forms with optional UTC `Z` or numeric offsets, using either `.` or `,` as fractional separator, e.g. `20260407T143000`, `2026-04-07T14:30:00`, `2026097T143000Z`, `2026-097T14:30:00+05:45`, `2026W152T143000Z`, `2026-W15-2T14:30:00+05:45`, `20260407T143000.125+0545`, and `2026-04-07T14:30:00,125Z`.

Normalized term shapes are:

- `time(Hours,Minutes,Seconds)` for time-of-day strings
- `date_time(Year,Month,Day,Hours,Minutes,Seconds)` for local date-time strings
- `date_time(Year,Month,Day,Hours,Minutes,Seconds,OffsetSeconds)` for UTC or offset date-time strings, using an offset in seconds from UTC

An **ISO 8601 duration string** represents a time amount (e.g. `P3D` for three days or `P1Y2M3DT4H5M6S` for a mixed date/time duration).

An **ISO 8601 interval string** represents a time interval using two parts separated by `/`, where each part can be a date, a date-time, or a duration (e.g. `2026-02-25/2026-03-01`, `2026-04-07T14:30:00Z/2026-04-07T15:00:00Z`, `2026-097T14:30:00Z/2026-097T15:00:00Z`, `2026-W15-2T14:30:00Z/PT30M`, or `2026-02-25/P3D`).

Author: Daniel L. Dudley Created: 2004-02-18

Modified: 2014-09-26 (to use the `os` library object to get the current date)

Modified: 2026-02-25 (to add `duration_string/2` and `interval_string/2` predicates) Modified: 2026-04-07 (to add `time_string/3` and `date_time_string/3` predicates and extend `interval_string/2`, including ordinal-date and week-date date-time forms) Modified: 2026-04-08 (to add support for comma as fractional separator in `time_string/3` and `date_time_string/3` and for corresponding interval parsing)

8.2.1 API documentation

Open the ../apis/library_index.html#iso8601 link in a web browser.

8.2.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(iso8601(loader)).
```

8.2.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(iso8601(tester)).
```

8.3 pddl_parser

PDDL 3.0 parser
Release 1.0

Copyright (c) 2011 Robert Sasak. All Rights Reserved. This is free software. You can redistribute it and/or modify it under the terms of the “Artistic License 2.0” as published by The Perl Foundation. Consult the “LICENSE.txt” file for details.

This PDDL 3.0 file parser converts PDDL files to Logtalk/Prolog friendly syntax. For example:

PDDL	Prolog
(on ?x ?y)	on(?x, ?y)

Syntax sugar: `op(200, fy, ?)`.

For whole example check the “`pddl.lgt`” file for usage and example output.

8.3.1 API documentation

Open the ../apis/library_index.html#pddl_parser link in a web browser.

8.3.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
?- logtalk_load(pddl_parser(loader)).
```

8.3.3 Testing

The provided unit tests are based on a collection of problem set files from International Planning Competition 2008. In order to run all unit tests from the parser’s directory type:

```
?- logtalk_load(tester).
```

Or from any directory by typing:

```
?- logtalk_load(pddl_parser(tester)).
```

Some of the unit tests fail in some Prolog compilers due to limitations to the maximum arity of a term.

8.4 verdi_neruda

Verdi Neruda - Meta-interpreter collection for Prolog.
Release 1.0

Copyright (c) 2010 Victor Lagerkvist. All Rights Reserved. Verdi Neruda is free software. You can redistribute it and/or modify it under the terms of the simplified BSD license.

CONTENTS

1. License
2. About
3. Verdi Neruda web site
4. Installation and running
5. Examples
6. Authors
7. LICENSE

Copyright 2010 Victor Lagerkvist. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The views and conclusions contained in the software and documentation are those of the authors and should not be interpreted as representing official policies, either expressed or implied, of the copyright holders.

2. ABOUT

Verdi Neruda is a meta-interpreter collection for Prolog. Or, to be more precise, for a Prolog like language. Or, to be pedantically precise to the point that you are annoying people, it's not really a meta-interpreter collection at all since the interpreters themselves aren't interpreting the language that they are written in. Let's just say that it is a collection of interpreters for a logic programming language very much like pure Prolog with negation as finite failure.

Verdi Neruda is written entirely in Logtalk and compatible with most major Prolog systems. The name is sadly not a subtle wordplay or an acronym, but was generated by a computer with the help of a soundex algorithm. The purpose of the interpreter suite was to compare top-down methods to bottom-up methods and how resolution tree search rules affected performance and completeness. In the top-down family we find such interpreters as the long-time standing champion depth-first, its slow but orderly brother breadth-first

and the youngster iterative deepening. A best-first framework can be found a stone throw away. With it it's possible to define interpreters that use greedy best-first search as well as A* search.

In the bottom-up camp we find an interpreter that uses a semi naive fixpoint construction. Since bottom-up interpreters by their very nature are not goal oriented a transformation technique called magic transformation is used on logic programs before any inferences are made. This technique allows the interpreter to only generate the facts that a top-down interpreter would have used on the same logic program.

A shell akin to a Prolog top loop is also included. It has commands both for proving goals with an interpreter of choice and for benchmarking logical inferences. If Verdi Neruda is run with a Prolog system that supports statistics/2 it's possible to obtain statistics such as CPU-time as well.

3. VERDI NERUDA WEB SITE

Visit the Verdi Neruda GitHub [www-page](https://github.com/Joelbyte/verdi-neruda) at:

<https://github.com/Joelbyte/verdi-neruda>

4. INSTALLATION AND RUNNING

Verdi Neruda requires Logtalk 2.40.0 or a later version.

To use the snapshot of Verdi Neruda bundled with Logtalk:

- Start Logtalk.
- Type `{verdi_neruda(loader)}`. (Including `.`).

To use the latest version of Verdi Neruda, fetch the latest source code, either as an archive or from the git repository, extract it to a directory of your choice, and:

- Start Logtalk from that directory.
- Type `{loader}`. (Including `.`). If everything went according to the plan you should be greeted by the welcoming message. If you replace the bundled version with the new one, you can use in alternative the steps above.

5. EXAMPLES

Follow the previous instructions to get everything up and running. First we're going to run some predefined programs in the included databases. Begin by typing `databases`. from the shell - this should print a list of the currently loaded databases. The demo database `demodb` should be included in the list. Next type `listing(demodb)`. to print the contents of the database. The output should look something like:

```
append([],A,A) if
    true.
append([A|B],C,[A|D]) if
    append(B,C,D).
.
.
.
```

Which means that the `append/3` program is loaded and ready for action. Next we need to decide which interpreter to use. Fortunately the shell does not leave much to the imagination - as might be expected, the `interpreters`. command prints the currently loaded interpreters. The list should look like:

```
dfs_interpreter
bfs_interpreter
iddfs_interpreter(A)
bup_interpreter
a_star_interpreter(A)
```

The variables means that the interpreters are parametric objects and that additional information is needed in order to run them. The iddfs-interpreter needs to know the increment and the A*-interpreter needs to know what weight should be used when calculating the cost of nodes. To start with let's use the dfs-interpreter and do something exciting, namely appending two lists!

```
prove(dfs_interpreter, append([a,b], [c,d], Xs), demodb).
```

The prove command takes three arguments. The first is a interpreter, the second the goal that shall be proved and the last the database that the clauses are derived from.

To accomplish the same thing with the iddfs-interpreter with an increment of 1 we need only type

```
prove(iddfs_interpreter(1), append([a,b], [c,d], Xs), demodb).
```

The shell also has support for counting logical inferences. To compare the dfs- and iddfs-interpreter with the append program we could write:

```
benchmark(dfs_interpreter, append([a,b,c,d],[e,f], Xs), demodb). ->
dfs_interpreter inferences: 5

benchmark(iddfs_interpreter(1), append([a,b,c,d],[e,f], Xs), demodb).
-> iddfs_interpreter(1) inferences: 15
```

For more information regarding the built in shell commands consult the 'help.' command.

6. AUTHORS

The bulk of Verdi Neruda was written by Victor Lagerkvist during his bachelor thesis at Linköping university in the spring of 2010. Paulo Moura also helped a great deal during the later stages of development, especially with regards to compatibility between various Prolog systems.

8.5 xml_parser

This folder contains a Logtalk version of John Fletcher's Prolog XML parser:

https://binding-time.co.uk/index.php/Parsing_XML_with_Prolog

For a detailed description of this XML parser, please see the comments in the `xml.lgt` source file or convert the automatically generated documentation to HTML or PDF. For sample queries, please see the `SCRIPT.txt` file.

See the copyright and license information on the contributed files for usage and distributions conditions.

8.5.1 API documentation

Open the ../apis/library_index.html#xml_parser link in a web browser.

8.5.2 Loading

To load all entities in this library, load the `loader.lgt` file:

```
| ?- logtalk_load(xml_parser(loader)).
```

8.5.3 Testing

To test this library predicates, load the `tester.lgt` file:

```
| ?- logtalk_load(xml_parser(tester)).
```

8.5.4 Known issues

When using GNU Prolog as the backend compiler, you may need to use a larger default global stack size (see the GNU Prolog documentation on the environment variable `GLOBALSZ`).

GLOSSARY

abstract class

A *class* that cannot be instantiated by sending it a message. Usually used to contain common predicates that are inherited by other classes.

abstract method

A *method* implementing an algorithm whose step corresponds to calls to methods defined in the descendants of the object (or *category*) containing it.

adapter file

A Prolog source file defining a minimal abstraction layer between the Logtalk compiler/runtime and a specific *backend Prolog compiler*.

after method

The *predicate definition* used to answer an after event.

ancestor

A *class* or a parent *prototype* that contributes (via inheritance) to the definition of an object. For class-based hierarchies, the ancestors of an instance are its class(es) and all the superclasses of its class(es). For prototype-based hierarchies, the ancestors of a prototype are its parent(s) and the ancestors of its parent(s).

around method

An overriding definition for a predicate that calls additional goals before and after a call to the overridden predicate definition.

backend Prolog compiler

The Prolog compiler that is used to host and run Logtalk and that is called for compiling the intermediate Prolog code generated by the Logtalk compiler when compiling source files.

before method

The *predicate definition* used to answer a before event.

built-in method

A predefined *method* that can be called from within any object or *category*. I.e. built-in methods are built-in object and category predicates. Built-in methods cannot be redefined.

built-in predicate

A predefined predicate that can be called from anywhere. Built-in predicates can be redefined within *objects* and *categories*.

category

A set of *predicates directives* and clauses that can be (virtually) imported by any object. Categories support composing objects using fine-grained units of code reuse and also *hot patching* of existing objects. A category should be functionally-cohesive, defining a single functionality.

class

An *object* that specializes another object, interpreted as its *superclass*. A class defines the common predicates of a set of objects that instantiate it. An object can also be interpreted as a class when it instantiates itself.

clause reference

An opaque term that uniquely identifies a clause. Provided by some backends via alternative database built-in predicates.

closed-world assumption

The assumption that what cannot be proved true is false. Therefore, sending a *message* corresponding to a *declared* but not *defined* predicate, or calling a declared predicate with no clauses, fails. But messages or calls to undeclared predicates generate an error.

closure

A callable term (i.e., an atom or a compound term) passed to a *meta-predicate* call where it is extended with additional arguments to form a goal called by the meta-predicate.

coinductive predicate

A predicate whose calls are proved using greatest fixed point semantics. Coinductive predicates allows reasoning about infinite rational entities such as cyclic terms and ω -automata.

complementing category

A category used for *hot patching* an existing object (or a set of objects).

component

A unique atom or compound term template identifying a library, tool, application, or application sub-system. Component names are notably used by the message printing and question asking mechanisms. Compound terms are used instead of atoms when parameterization is required.

debug handler

An object or category that defines clauses for the multifile predicates that handle the trace and debug events generated by the runtime.

directive

A source file term that affects the interpretation of source code. Directives use the `(:-)/1` prefix operator as functor.

discontiguous predicate

A predicate whose clauses are not contiguous in a *source file*. I.e. a predicate whose clauses are mixed with clauses for other predicates.

doclet file

A *source file* whose main purpose is to generate documentation for a *library* or an application.

doclet object

An object specifying the steps necessary to (re)generate the API documentation for a project. See the *doclet* and *lgtdoc* tools for details.

dynamic binding

Runtime lookup of a *predicate declaration* and *predicate definition* to verify the validity of a *message* (or a *super call*) and find the predicate definition that will be used to answer the message (or the super call). Also known as *late binding*. See also *static binding*.

dynamic entity

See *entity*.

dynamic predicate

A predicate whose clauses can be dynamically added or retracted at runtime.

early binding

See *static binding*.

encapsulation

The hiding of an object implementation. This promotes software reuse by isolating the object clients from its implementation details. Encapsulation is enforced in Logtalk by using *predicate scope directives*.

entity

Generic name for Logtalk compilation units: *objects*, *categories*, and *protocols*. Entities share a single namespace (i.e., entity *identifiers* must be unique) and can be static (the default) or dynamic. Static entities are defined in source files. Dynamic entities can also be defined in source files but are usually created and abolished at runtime using the language built-in predicates.

entity alias

An alternative name for an entity. Entity aliases can be defined using the *uses/1* and *use_module/1* directives. Entity aliases can be used to improve code clarity by using alternative names that are more meaningful in the calling context, to bind parametric entity parameters, and to simplify experimenting with alternative entities implementing the same protocol.

entity directive

A directive that affects how Logtalk *entities* are used or compiled.

event

The sending of a *message* to an object. An event can be expressed as an ordered tuple: (Event, Object, Message, Sender). Logtalk distinguishes between the sending of a message — before event — and the return of control to the sender — after event.

expansion workflow

A sequence of *term-expansion* or *goal-expansion* steps where each step is usually defined using a *hook object* or a combination of hook objects.

flaky test

A test that can unpredictably succeed, fail, or throw an error and thus can yield different results when repeated without any code changes. Can result from, e.g., race conditions or dependencies on external resources outside user control.

goal-expansion

The transformation of a goal into another goal, defined by clauses for the *goal_expansion/2* predicate that is declared by the *expanding* built-in protocol. Goal-expansion is usually recursively applied at compile time until a fixed-point is reached (i.e., until the result of an expansion cannot be further expanded). See also *term-expansion* and *expansion workflow*.

grammar rule

An alternative notation for predicates used to parse or generate sentences in some language. This notation hides the arguments used to pass the sequences of tokens being processed, thus simplifying the representation of grammars. Grammar rules are represented using as functor the infix operator (*-->*)/2 instead of the (*:-*)/2 operator used with predicate clauses.

grammar rule non-terminal

A syntactic category of words or phrases. A non-terminal is identified by its *non-terminal indicator*, i.e. by its name and number of arguments using the notation *Name//Arity*.

grammar rule terminal

A word or basic symbol of a language.

homoiconic

A property of programming languages where *code* and *data* use the same representation. Logtalk (and Prolog) are examples of homoiconic programming languages. A Logtalk program is a set of terms (clauses and directives) that can be handled as *data* by e.g. the *term-expansion* mechanism.

hook object

An object, implementing the [expanding](#) built-in protocol, defining [term-expansion](#) and [goal-expansion](#) predicates, used in the compilation of Logtalk or Prolog source files. A hook object can be specified using the [hook](#) flag. It can also be specified using a [set_logtalk_flag/2](#) directive in the source files to be expanded.

hook predicate

A predicate, usually declared [multifile](#), that allows the user to customize another predicate or provide alternative definitions for a default predicate definition.

hot patching

The act of fixing [entity directives](#) and predicates or adding new entity directives and predicates to loaded [entities](#) in a running application without requiring access to the entities source code or restarting the application. Achieved using [complementing categories](#).

identity

Property of an [entity](#) that distinguishes it from every other entity. The identifier of an entity is its functor (i.e., its name and arity), which must be unique. Object and [category](#) identifiers can be atoms or compound terms. Protocol identities must be atoms. All Logtalk entities (objects, protocols, and categories) share the same namespace.

inheritance

An [entity](#) inherits [predicates directives](#) and clauses from related entities. In the particular case of objects, when an object extends other object, we have prototype-based inheritance. When an object specializes or instantiates another object, we have class-based inheritance. See also [public inheritance](#), [protected inheritance](#), and [private inheritance](#).

instance

An [object](#) that instantiates another object, interpreted as its [class](#). An object may instantiate multiple objects (also known as multiple instantiation).

instantiation

The process of creating a new [class instance](#). In Logtalk, this does not necessarily imply dynamic creation of an object at runtime; an instance may also be defined as a static object in a source file.

interface

See [protocol](#).

lambda expression

A compound term that can be used in place of a goal or [closure](#) meta-argument and that abstracts a [predicate definition](#) by listing its variables and a callable term that implements the definition. Lambda expressions help avoid the need of naming and defining auxiliary predicates.

lambda free variable

A variable that is global to a [lambda expression](#). All used global variables must be explicitly listed in a lambda expression for well-defined semantics.

lambda parameter

A term (usually a variable or a non-ground compound term) that is local to a [lambda expression](#). All lambda parameters must be explicitly enumerated in a lambda expression for well-defined semantics.

late binding

See [dynamic binding](#).

library

A directory containing source files. See also [library alias](#) and [library notation](#).

library alias

An atom that can be used as an alias for a [library](#) full path. Library aliases and their corresponding paths can be defined using the [logtalk_library_path/2](#) predicate. See also [library notation](#).

library notation

A compound term where the name is a *library alias* and the single argument is a *source file* relative path. Use of library notation simplifies compiling and loading source files and can make an application easily relocatable by defining an alias for the root directory of the application files.

loader file

A *source file* whose main purpose is to load a set of source files (possibly with specific compiler flags) and any library dependencies.

local predicate

A predicate that is defined in an object (or in a *category*) but that is not listed in a *scope directive*. These predicates behave like private predicates but are invisible to the reflection *built-in methods*. Local predicates are usually auxiliary predicates and only relevant to the entity where they are defined.

message

A query sent to an object. In logical terms, a message can be interpreted as a request for proof construction using an object database and the databases of related entities.

message lookup

Sending a message to an object requires a lookup for the *predicate declaration*, to check if the message is within the scope of the sender, and a lookup for the *predicate definition* that is going to be called to answer the message. Message lookup can occur at *compile* time or at *runtime*.

message to self

A message sent to the object that received the original message under processing. Messages to self require *dynamic binding* as the value of self is only known at runtime.

meta-argument

A predicate argument that is called as a goal, used as a *closure* to construct a goal that will be called, or that is handled in a way that requires awareness of the predicate calling context.

meta-interpreter

A program capable of running other programs written in the same language.

meta-predicate

A predicate with one or more *meta-arguments*. For example, *call/1-N* and *findall/3* are built-in meta-predicates.

meta-variable

A variable in a *meta-argument* position that is expected to be unified with a goal or a closure at runtime.

metaclass

The *class* of a class, when interpreted as an *instance*. Metaclass instances are themselves classes. Metaclasses are optional, except for the root class, and can be shared by several classes.

method

The *predicate definition* used to answer a *message* sent to an object. Logtalk supports both *static binding* and *dynamic binding* to find which method to run to answer a message. See also *built-in method*.

mocking

Techniques used to replace dependencies in the code being tested with controlled substitutes, simplifying testing. Can be accomplished using, e.g., *term-expansion*, hot patching, or message interception.

module

A Prolog entity characterized by an identity and a set of *predicates directives* and clauses. Prolog modules are usually static although some Prolog systems allow the creation of dynamic modules at runtime. Prolog modules can be interpreted as *prototypes*.

monitor

Any object, implementing the *monitoring* built-in protocol, that is notified by the runtime when a spied event occurs. The spied *events* can be set by the monitor itself or by any other object.

multifile predicate

A predicate whose clauses can be defined in multiple *entities* and *source files*. The object or category holding the directive without an entity prefix qualifying the predicate holds the multifile predicate *primary declaration*, which consists of both a *scope directive* and a *multifile/1* directive for the predicate.

naked meta-variable

A *meta-variable* used as the body of a predicate clause or grammar rule or used in a cut-transparent argument of a control construct. The “naked” designation highlights that the meta-variable is not wrapped by *call/1* or *phrase/1* goals.

object

An entity characterized by an *identity* and a set of *predicates directives* and clauses. Logtalk objects can be either static or dynamic. Logtalk objects can play the *role* of classes, instances, or prototypes. The role or roles an object plays are a function of its relations with other objects.

object database

The set of predicates locally defined inside an object.

parameter

An argument of a parametric object or a parametric category identifier. Parameters are *logical variables* implicitly shared by all the entity directives and predicate clauses.

parameter variable

A variable used as parameter in a parametric object or a parametric category using the syntax `_ParameterName_` (i.e., a variable whose name starts and ends with an underscore). Parameter variables are *logical variables* shared by all entity terms. Occurrences of parameter variables in *entity directives* and clauses are implicitly unified with the corresponding entity parameters.

parametric category

See *parametric entity*.

parametric entity

An *object* or *category* whose *identifier* is a compound term possibly containing free variables that can be used to parameterize the entity predicates. Parameters are *logical variables* implicitly shared by all the entity clauses. Note that the identifier of a parametric entity is its functor, irrespective of the possible values of its arguments (e.g., `foo(bar)` and `foo(baz)` are different parameterizations of the same parametric entity, `foo/1`).

parametric object

See *parametric entity*.

parametric object proxy

A compound term (usually represented as a plain Prolog fact) with the same name and number of arguments as the identifier of a parametric object.

parent

A *prototype* that is extended by another prototype.

polymorphism

Different objects (and categories) can provide different implementations of the same predicate. The predicate declaration can be inherited from a common ancestor, also known as *subtype polymorphism*. Logtalk implements *single dispatch* on the receiver of a message, which can be described as *single-argument polymorphism*. As *message lookup* only uses the predicate functor, multiple predicate implementations for different types of arguments are possible, also known as *ad hoc polymorphism*. *Parametric objects and categories* enable implementation of *parametric polymorphism* by using one or more parameters to pass object identifiers that can be used to parameterize generic predicate definitions.

predicate

Predicates describe what is true about the application domain. A predicate is identified by its *predicate*

indicator, i.e. by its name and number of arguments using the notation *Name/Arity*. See also *built-in predicate* and *method*.

predicate alias

An alternative functor (*Name/Arity*) for a predicate. Predicate aliases can be defined for any inherited predicate using the *alias/2* directive and for predicates listed in *uses/2* and *use_module/2* directives. Predicate aliases can be used to solve inheritance conflicts, to improve code clarity by using alternative names that are more meaningful in the calling context, and to use a different order of the predicate arguments.

predicate calling context

The object or category from within a predicate is called (either directly or using a control construct such as a message-sending control construct).

predicate declaration

A predicate declaration is composed by a set of *predicates directives*, which must include at least a *scope directive*.

predicate definition

The set of clauses for a predicate, contained in an object or category. Predicate definitions can be overridden or specialized in descendant entities.

predicate definition context

The object or category that contains the definition (i.e., clauses) for a predicate.

predicate directive

A directive that specifies a predicate property that affects how predicates are called or compiled.

predicate execution context

The implicit arguments (including *sender*, *self*, and *this*) required for the correct execution of a predicate call.

predicate scope container

The object that inherits a *predicate declaration* from an imported *category* or an implemented *protocol*.

predicate scope directive

A directive that declares a predicate by specifying its visibility as *public*, *protected*, or *private*.

predicate shorthand

A *predicate alias* that defines a call template, possibly using a different name, with a reduced number of arguments by hard-coding the value of the omitted arguments in the original call template. Predicate shorthands can be defined using *uses/2* and *use_module/2* directives. They can be used to simplify predicate calls and to ensure consistent call patterns when some of the arguments always use the same fixed values in the calling context.

primary predicate declaration

See *multifile predicate*.

private inheritance

All public and protected predicates are inherited as private predicates. See also *public inheritance* and *protected inheritance*.

private predicate

A predicate that can only be called from the object that contains its *scope directive*.

profiler

A program that collects data about other program performance.

protected inheritance

All public predicates are inherited as protected. No scope change for protected or private predicates. See also *public inheritance* and *private inheritance*.

protected predicate

A predicate that can only be called from the object containing its *scope directive* or from an object that inherits the predicate.

protocol

An entity that contains *predicate declarations*. A predicate is declared using a *scope directive*. It may be further specified by additional predicate directives. Protocols support the separation between interface and implementation, can be implemented by both objects and categories, and can be extended by other protocols. A protocol should be functionally-cohesive, specifying a single functionality. Also known as *interface*.

prototype

A self-describing *object* that may extend or be extended by other objects. A prototype typically describes a concrete object instead of an abstraction of a set of objects. An object with no instantiation or specialization relations with other objects is always interpreted as a prototype.

public inheritance

All inherited predicates maintain their declared scope. See also *protected inheritance* and *private inheritance*.

public predicate

A predicate that can be called from any object.

scratch directory

The directory used to save the intermediate Prolog files generated by the compiler when compiling *source files*.

self

The object that received the *message* under processing.

sender

An object that sends a *message* to other object. When a message is sent from within a *category*, the *sender* is the object importing the category on which behalf the message was sent.

settings file

A *source file*, compiled and loaded automatically by default at Logtalk startup, mainly defining default values for compiler flags that override the defaults found on the backend Prolog compiler *adapter files*.

singleton method

A *method* defined in an *instance* itself. Singleton methods are supported in Logtalk and can also be found in other object-oriented programming languages.

source file

A text file defining Logtalk and/or Prolog code. Multiple Logtalk entities may be defined in a single source file. Plain Prolog code may be intermixed with Logtalk entity definitions. Depending on the used *backend Prolog compiler*, the text encoding may be specified using an *encoding/1* directive as the first term in the first line in the file.

source file directive

A *directive* that affects how a *source file* is compiled.

specialization

A *class* is specialized by defining a new class that inherits its predicates and possibly adds new ones.

static binding

Compile time lookup of a *predicate declaration* and *predicate definition* when compiling a *message* sending call (or a *super call*). Dynamic binding is used whenever static binding is not possible (e.g., due to the predicate being dynamic or due to lack of enough information at compilation time). Also known as *early binding*. See also *dynamic binding*.

static entity

See *entity*.

steadfastness

A predicate definition is *steadfast* when it still generates only correct answers when called with unexpectedly bound arguments (notably, bound output arguments). Typically, a predicate may not be steadfast when output argument unifications can occur before a cut in a predicate clause body.

subclass

A *class* that is a specialization, direct or indirectly, of another class. A class may have multiple subclasses.

super call

Call of an inherited (or imported) *predicate definition*. Mainly used when redefining an inherited (or imported) predicate to call the overridden definition while making additional calls. Super calls preserve *self* and may require *dynamic binding* if the predicate is dynamic.

superclass

A *class* from which another class is a specialization (directly or indirectly via another class). A class may have multiple superclasses.

synchronized predicate

A synchronized predicate is protected by a mutex ensuring that, in a multi-threaded application, it can only be called by a single thread at a time.

template method

See *abstract method*.

term-expansion

The transformation of a term (usually a directive or a clause) into another term or a list of terms, defined by clauses for the *term_expansion/2* predicate that is declared by the *expanding* built-in protocol. Term-expansion is usually applied at compile time once (i.e., the result of an expansion is not further expanded). See also *goal-expansion* and *expansion workflow*.

test dialect

Predicate syntax for defining a test. I.e., the predicate name; the number, syntax, and semantics of the predicate arguments; and the syntax and semantics of any options.

test outcome

The expected result of a test goal. E.g., success, failure, an error, deterministic success, success with bindings that subsume or are a variant of specific terms.

tester file

A *source file* whose main purpose is to load and run a set of unit tests.

this

The object that contains the predicate clause under execution. When the predicate clause is contained in a *category*, *this* is a reference to the object importing the category on which behalf the predicate clause is being used to prove the current goal.

threaded engine

A computing thread running a goal whose solutions can be lazily and concurrently computed and retrieved. A threaded engine also supports a term queue that allows passing arbitrary terms to the engine. This queue can be used to pass e.g. data and new goals to the engine.

top-level interpreter shorthand

Aliases for frequently used built-in predicates such as *logtalk_load/1* and *logtalk_make/1*. These short-hands are **not** part of the Logtalk language and must only be used at the top-level interpreter.

visible predicate

A predicate that is within scope, a locally defined predicate, a *built-in method*, a Logtalk built-in predicate, or a Prolog built-in predicate.

BIBLIOGRAPHY

- [Alexiev93] Mutable Object State for Object-Oriented Logic Programming: A Survey Alexiev, V. Technical Report TR 93-15, Department of Computing Science, University of Alberta, Canada
- [Belli_et_al_92] Object-oriented programming in Prolog: rationale and a case study Belli, F., Jack, O., Naish, L. Technical Report 92/2, Department of Electrical and Electronics Engineering, University of Paderborn, Germany URL: <https://lee-naish.github.io/papers/oopl/index.html>
- [Block89] An Extended Frame Language Block, F. P., Chan, N. C. Proceedings OOPLSLA 89(10):151-157, ACM
- [Bobrow_et_al_88] Common Lisp Object System Specification Bobrow, D. G., Michiel, L. G., Gabriel, R. P., Keene, S. E., Kiczales, G., Moon, D. A. ACM SIGPLAN Notices(23)
- [Bratko90] Prolog Programming for Artificial Intelligence Bratko, I. Addison Wesley, 2^o edition, 1990
- [Champaux92] A comparative Study of Object-Oriented Analysis Methods Champaux, D., Faure, P. Journal of Object-Oriented Programming, Vol. 5, N.1, 1992
- [Clocksin87] Programming in Prolog Clocksin, W.F., Mellish, C.S. Springer-Verlag, New York, 1987
- [Cointe87] Metaclasses are First Class: the ObjVlisp Model Cointe, P. Proceedings OOPLSLA 87(10):156-167, ACM
- [Cordes91] The Literate Programming Paradigm Cordes, D., Brown, M. IEEE Computer, June 1991:52-61
- [Covington94] ISO Prolog: A Summary of the Draft Proposed Standard Covington, M. A. URL: <ftp://ai.uga.edu/pub/prolog.standard/>
- [Cox86] Object-Oriented Programming: An Evolutionary Approach Cox, Brad J. Addison-Wesley Publishing Company, Don Mills, Ontario
- [Davison89] Polka: A Parlog Object oriented language Davison, A. Ph.D. Thesis, Imperial College, London, 1989
- [Davison92] A survey of logic programming-based object oriented languages Davison, A. Tech Report 92/3, Dept. of Computer Science, University of Melbourne, Australia URL: <https://catalogue.nla.gov.au/catalog/772526>
- [Davison93] The deductive and object oriented features of BeBOP Davison, A. Tech Report 93/6, Dept. of Computer Science, University of Melbourne, Australia URL: <https://catalogue.nla.gov.au/catalog/1273317>
- [Delzanno97] Logic and Object-Oriented Programming in Linear Logic Delzanno, G. Ph.D. Thesis, University of Pisa, Italy URL: <https://opac.bncf.firenze.sbn.it/Record/BVE0136144?uri=BVE0136144>
- [Dony90] Exception Handling and Object-Oriented Programming: Towards a Synthesis Dony, C. Proceedings OOPLSLA 90:322-330, ACM

- [Fornarino_et_al_89] An Original Object-Oriented Approach for Relation Management Fornarino, M., Pinna, A.-M., Trousse, B. Proceedings of the 4th Portuguese Conference on Artificial Intelligence Lecture Notes in Artificial Intelligence, Springer-Verlag (390):13-26
- [Fromherz93] OL(P): Object Layer for Prolog Fromherz, M. URL: <ftp://parcftp.xerox.com/ftp/pub/ol/>
- [Fukunaga86] An Experience with a Prolog-based Object-Oriented Language Fukunaga, K., Hirose, S. Proceedings OOPLSLA 86, 21(11):224-231, ACM
- [Goldberg83] Smalltalk-80 The language and its implementation Goldberg, A., Robson, D. Addison-Wesley Series in Computer Science
- [Joy_et_al_00] The Java Language Specification, Second Edition Joy, B., Steele, G., Gosling, J., Bracha, G. Addison-Wesley, 2000
- [ISO95] ISO/IEC DIS 13211-1 - Programming Language Prolog Part 1: General Core Joint Technical Committee ISO/IEC JTC 1 URL: <https://www.iso.org/standard/21413.html>
- [Knuth84] Literate Programming Knuth, D. E. Computer Journal, May 84, 27(2):97-111
- [Lieberman86] Using Prototypical Objects to Implement Shared Behaviour in Object Oriented Systems Lieberman, H. Proceedings OOPLSLA 86:189-214, ACM
- [Maes87] Concepts and Experiments in Computational Reflection Maes, P. Proceedings OOPLSLA 87, ACM
- [McCabe92] Logic and Objects McCabe, F. G. Prentice Hall Series in Computer Science
- [Moon86] Object-Oriented Programming in Flavors Moon, D. Proceedings OOPLSLA 86:1-8, ACM
- [Moss94] Prolog++ The Power of Object-Oriented and Logic Programming Moss, C. Addison-Wesley International Series in Logic Programming, 1994
- [Moura94] Logtalk: Programação Orientada para Objectos em Prolog Moura, P., Costa, E. 2ª Conferência e Exposição Portuguesa de Tecnologia Orientada por Objectos 3i Consultores, Lisboa
- [Moura99] Porting Prolog: Notes on porting a Prolog program to 22 Prolog compilers or the relevance of the ISO Prolog standard Moura, P. ALP Newsletter, Vol. 12/2, May 1999
- [Moura00] Logtalk 2.6 Documentation Moura, P. Technical Report DMI 2000/1 University of Beira Interior, Portugal
- [Razek92] Combining Objects and Relations Razek, G. Communications of the ACM, 27(12):66-70
- [Rumbaugh87] Relations as Semantic Constructs in an Object-Oriented Language Rumbaugh, J. Proceedings OOPLSLA 87:466-481, ACM
- [Rumbaugh88] Controlling Propagation of Operations using Attributes on Relations Rumbaugh, J. Proceedings OOPLSLA 88:285-296, ACM
- [Schachte95] Efficient Object-Oriented Programming in Prolog Schachte, P., Saab, G. Logic Programming: Formal Methods and Practical Applications Studies in Computer Science and Artificial Intelligence, 11 Elsevier Science B.V. North-Holland, Amsterdam, 1995
- [SICStus95] SICStus Prolog Manual SICStus URL: <https://sicstus.sics.se>
- [Shan_et_al_93] Is Multiple Inheritance Essential to OOP? (Panel) Shan, Y., Cargill, T., Cox, B., Cook, W., Loomis, M., Snyder, A. Proceedings OOPLSLA 93:360-363
- [Stefik_et_al_86] Integrating Access-Oriented Programming into a Multiparadigm Environment Stefik, M. J., Bobrow, D. G., Kahn, K. M. IEEE Software, January 1986:10-18
- [Stroustrup86] The C++ Programming Language Stroustrup, B. Addison-Wesley Series in Computer Science

- [Taenzer89] Problems in Object-Oriented Software Reuse Taenzer, D., Ganti, M., Podar, S. Proceedings of ECOOP 89 British Computer Society Workshop Series, Cambridge University Press
- [Tanzer95] Remarks on Object-Oriented Modeling of Associations Tanzer, C. Journal of Object-Oriented Programming, February 1995, SIGS Publications
- [Tanenbaum87] Operating Systems - Design and Implementation Tanenbaum, A. Prentice-Hall Software Series, 1987
- [Welsch89] Reasoning Objects with Dynamic Knowledge Bases Welsch, C., Barth, G. Proceedings of the 4th Portuguese Conference on Artificial Intelligence(390):257-268 Lecture Notes in Artificial Intelligence, Springer-Verlag, 1989

Symbols

!/0
 Built-in method, 311
 (::)/1
 Control construct, 195
 (::)/2
 Control construct, 194
 (@)/1
 Control construct, 199
 (^)/1
 Control construct, 198
 (\+)/1
 Built-in method, 335
 (<<)/2
 Control construct, 202
 {}/1
 Control construct, 200
 []/1
 Control construct, 197

A

abolish/1
 Built-in method, 323
 abolish_category/1
 Built-in predicate, 257
 abolish_events/5
 Built-in predicate, 269
 abolish_object/1
 Built-in predicate, 257
 abolish_protocol/1
 Built-in predicate, 258
 abstract class, 1035
 abstract method, 1035
 adapter file, 1035
 after method, 1035
 after/3
 Built-in method, 360
 alias/2
 Directive, 227
 always_true_or_false_goals
 Flag, 115
 ancestor, 1035

arithmetic_expressions
 Flag, 115
 around method, 1035
 ask_question/5
 Built-in method, 380
 asserta/1
 Built-in method, 324
 assertz/1
 Built-in method, 327

B

backend Prolog compiler, 1035
 bagof/3
 Built-in method, 354
 before method, 1035
 before/3
 Built-in method, 360
 begin_of_file, 132
 behavioral reflection, 103
 black-box view, 103
 built_in/0
 Directive, 214
 built-in method, 1035
 built-in predicate, 1035
 Built-in method
 !/0, 311
 (\+)/1, 335
 abolish/1, 323
 after/3, 360
 ask_question/5, 380
 asserta/1, 324
 assertz/1, 327
 bagof/3, 354
 before/3, 360
 call//1-N, 362
 call/1-N, 332
 catch/3, 337
 clause/2, 329
 coinductive_success_hook/1-2, 372
 consistency_error/3, 344
 context/1, 314
 current_op/3, 319

- current_predicate/1, 320
- domain_error/2, 342
- eos//0, 364
- evaluation_error/1, 350
- existence_error/2, 345
- expand_goal/2, 370
- expand_term/2, 368
- fail/0, 312
- false/0, 312
- findall/3, 355
- findall/4, 356
- forall/2, 357
- forward/1, 361
- goal_expansion/2, 371
- ignore/1, 333
- instantiation_error/0, 339
- message_hook/4, 375
- message_prefix_file/6, 377
- message_prefix_stream/4, 376
- message_tokens//2, 374
- once/1, 334
- parameter/2, 315
- permission_error/3, 347
- phrase//1, 365
- phrase/2, 366
- phrase/3, 367
- predicate_property/2, 322
- print_message/3, 374
- print_message_token/4, 379
- print_message_tokens/3, 378
- question_hook/6, 380
- question_prompt_stream/4, 381
- repeat/0, 313
- representation_error/1, 349
- resource_error/1, 351
- retract/1, 330
- retractall/1, 331
- self/1, 317
- sender/1, 317
- setof/3, 358
- syntax_error/1, 352
- system_error/0, 353
- term_expansion/2, 369
- this/1, 318
- throw/1, 338
- true/0, 311
- type_error/2, 341
- unin instantiation_error/1, 340
- Built-in predicate
 - abolish_category/1, 257
 - abolish_events/5, 269
 - abolish_object/1, 257
 - abolish_protocol/1, 258
 - category_property/2, 248
 - complements_object/2, 265
 - conforms_to_protocol/2-3, 263
 - create_category/4, 252
 - create_logtalk_flag/3, 308
 - create_object/4, 253
 - create_protocol/3, 255
 - current_category/1, 246
 - current_event/5, 269
 - current_logtalk_flag/2, 307
 - current_object/1, 247
 - current_protocol/1, 247
 - define_events/5, 270
 - extends_category/2-3, 261
 - extends_object/2-3, 259
 - extends_protocol/2-3, 260
 - implements_protocol/2-3, 262
 - imports_category/2-3, 265
 - instantiates_class/2-3, 266
 - logtalk_compile/1, 292
 - logtalk_compile/2, 293
 - logtalk_library_path/2, 302
 - logtalk_linter_hook/7, 310
 - logtalk_load/1, 295
 - logtalk_load/2, 297
 - logtalk_load_context/2, 304
 - logtalk_make/0, 298
 - logtalk_make/1, 299
 - logtalk_make_target_action/1, 301
 - object_property/2, 249
 - protocol_property/2, 250
 - set_logtalk_flag/2, 307
 - specializes_class/2-3, 267
 - threaded/1, 272
 - threaded_call/1-2, 273
 - threaded_cancel/1, 280
 - threaded_engine/1, 285
 - threaded_engine_create/3, 283
 - threaded_engine_destroy/1, 284
 - threaded_engine_fetch/1, 291
 - threaded_engine_next/2, 287
 - threaded_engine_next_reified/2, 288
 - threaded_engine_post/2, 290
 - threaded_engine_self/1, 286
 - threaded_engine_yield/1, 289
 - threaded_exit/1-2, 277
 - threaded_ignore/1, 276
 - threaded_notify/1, 282
 - threaded_once/1-2, 274
 - threaded_peek/1-2, 278
 - threaded_wait/1, 281
- C
 - call//1-N
 - Built-in method, 362

- call/1-N
 - Built-in method, 332
- catch/3
 - Built-in method, 337
- catchall_catch
 - Flag, 116
- category, 1035
- category/1-4
 - Directive, 214
- category_property/2
 - Built-in predicate, 248
- class, 1036
- clause reference, 1036
- clause/2
 - Built-in method, 329
- clean
 - Flag, 118
- closed-world assumption, 1036
- closure, 1036
- code_prefix
 - Flag, 117
- coinduction
 - Flag, 113
- coinductive predicate, 1036
- coinductive/1
 - Directive, 228
- coinductive_success_hook/1-2
 - Built-in method, 372
- complementing category, 1036
- complements
 - Flag, 116
- complements_object/2
 - Built-in predicate, 265
- component, 1036
- conditionals
 - Flag, 115
- conforms_to_protocol/2-3
 - Built-in predicate, 263
- consistency_error/3
 - Built-in method, 344
- context/1
 - Built-in method, 314
- context_switching_calls
 - Flag, 116
- Control construct
 - (:)/1, 195
 - (:)/2, 194
 - (@)/1, 199
 - (^)/1, 198
 - (<<)/2, 202
 - {}/1, 200
 - []/1, 197
- create_category/4
 - Built-in predicate, 252

- create_logtalk_flag/3
 - Built-in predicate, 308
- create_object/4
 - Built-in predicate, 253
- create_protocol/3
 - Built-in predicate, 255
- current_category/1
 - Built-in predicate, 246
- current_event/5
 - Built-in predicate, 269
- current_logtalk_flag/2
 - Built-in predicate, 307
- current_object/1
 - Built-in predicate, 247
- current_op/3
 - Built-in method, 319
- current_predicate/1
 - Built-in method, 320
- current_protocol/1
 - Built-in predicate, 247

D

- debug
 - Flag, 118
- debug handler, 1036
- define_events/5
 - Built-in predicate, 270
- deprecated
 - Flag, 114
- Directive
 - alias/2, 227
 - built_in/0, 214
 - category/1-4, 214
 - coinductive/1, 228
 - discontiguous/1, 229
 - dynamic/0, 216
 - dynamic/1, 230
 - elif/1, 211
 - else/0, 212
 - encoding/1, 204
 - end_category/0, 217
 - end_object/0, 217
 - end_protocol/0, 218
 - endif/0, 213
 - if/1, 210
 - include/1, 205
 - info/1, 218
 - info/2, 231
 - initialization/1, 206
 - meta_non_terminal/1, 234
 - meta_predicate/1, 232
 - mode/2, 235
 - mode_non_terminal/2, 235
 - multifile/1, 236

- object/1-5, 219
- op/3, 207
- private/1, 237
- protected/1, 238
- protocol/1-2, 223
- public/1, 239
- set_logtalk_flag/2, 209
- synchronized/1, 240
- threaded/0, 224
- use_module/1, 226
- use_module/2, 243
- uses/1, 225
- uses/2, 241
- directive, **1036**
- discontiguous predicate, **1036**
- discontiguous/1
 - Directive, 229
- disjunctions
 - Flag, 115
- doclet file, **1036**
- doclet object, **1036**
- domain_error/2
 - Built-in method, 342
- duplicated_clauses
 - Flag, 115
- duplicated_directives
 - Flag, 114
- dynamic binding, **1036**
- dynamic entity, **1036**
- dynamic predicate, **1036**
- dynamic/0
 - Directive, 216
- dynamic/1
 - Directive, 230
- dynamic_declarations
 - Flag, 116

E

- early binding, **1037**
- elif/1
 - Directive, 211
- else/0
 - Directive, 212
- encapsulation, **1037**
- encoding/1
 - Directive, 204
- encoding_directive
 - Flag, 113
- encodings
 - Flag, 116
- end_category/0
 - Directive, 217
- end_object/0
 - Directive, 217

- end_of_file, 132
- end_protocol/0
 - Directive, 218
- endif/0
 - Directive, 213
- engines
 - Flag, 113
- entity, **1037**
- entity alias, **1037**
- entity directive, **1037**
- eos//0
 - Built-in method, 364
- evaluation_error/1
 - Built-in method, 350
- event, **1037**
- events
 - Flag, 116
- existence_error/2
 - Built-in method, 345
- expand_goal/2
 - Built-in method, 370
- expand_term/2
 - Built-in method, 368
- expansion workflow, **1037**
- extends_category/2-3
 - Built-in predicate, 261
- extends_object/2-3
 - Built-in predicate, 259
- extends_protocol/2-3
 - Built-in predicate, 260

F

- fail/0
 - Built-in method, 312
- false/0
 - Built-in method, 312
- findall/3
 - Built-in method, 355
- findall/4
 - Built-in method, 356
- Flag
 - always_true_or_false_goals, 115
 - arithmetic_expressions, 115
 - catchall_catch, 116
 - clean, 118
 - code_prefix, 117
 - coinduction, 113
 - complements, 116
 - conditionals, 115
 - context_switching_calls, 116
 - debug, 118
 - deprecated, 114
 - disjunctions, 115
 - duplicated_clauses, 115

- duplicated_directives, 114
- dynamic_declarations, 116
- encoding_directive, 113
- encodings, 116
- engines, 113
- events, 116
- general, 116
- grammar_rules, 115
- hook, 118
- lambda_variables, 115
- left_recursion, 116
- linter, 114
- missing_directives, 114
- modules, 113
- naming, 115
- optimize, 117
- portability, 114
- prolog_compatible_version, 113
- prolog_compiler, 117
- prolog_dialect, 113
- prolog_loader, 117
- prolog_version, 113
- redefined_built_ins, 115
- redefined_operators, 115
- relative_to, 118
- reload, 118
- report, 117
- scratch_directory, 117
- settings_file, 113
- singleton_variables, 115
- source_data, 118
- steadfastness, 114
- suspicious_calls, 115
- tabling, 113
- tail_recursive, 116
- threads, 113
- trivial_goal_fails, 115
- undefined_predicates, 114
- underscore_variables, 117
- unicode, 113
- unknown_entities, 114
- unknown_predicates, 114
- version_data, 114
- flaky test, 1037
- forall/2
 - Built-in method, 357
- forward/1
 - Built-in method, 361

G

- general
 - Flag, 116
- goal_expansion/2
 - Built-in method, 371

- goal_expansion, 1037
- grammar rule, 1037
- grammar rule non-terminal, 1037
- grammar rule terminal, 1037
- grammar_rules
 - Flag, 115

H

- homoiconic, 1037
- hook
 - Flag, 118
- hook object, 1038
- hook predicate, 1038
- hot patching, 1038

I

- identity, 1038
- if/1
 - Directive, 210
- ignore/1
 - Built-in method, 333
- implements_protocol/2-3
 - Built-in predicate, 262
- imports_category/2-3
 - Built-in predicate, 265
- include/1
 - Directive, 205
- info/1
 - Directive, 218
- info/2
 - Directive, 231
- inheritance, 1038
- initialization/1
 - Directive, 206
- instance, 1038
- instantiates_class/2-3
 - Built-in predicate, 266
- instantiation, 1038
- instantiation_error/0
 - Built-in method, 339
- interface, 1038

L

- lambda expression, 1038
- lambda free variable, 1038
- lambda parameter, 1038
- lambda_variables
 - Flag, 115
- late binding, 1038
- left_recursion
 - Flag, 116
- library, 1038
- library alias, 1038
- library notation, 1039

- linter
 - Flag, [114](#)
- loader file, [1039](#)
- local predicate, [1039](#)
- logtalk_compile/1
 - Built-in predicate, [292](#)
- logtalk_compile/2
 - Built-in predicate, [293](#)
- logtalk_library_path/2
 - Built-in predicate, [302](#)
- logtalk_linter_hook/7
 - Built-in predicate, [310](#)
- logtalk_load/1
 - Built-in predicate, [295](#)
- logtalk_load/2
 - Built-in predicate, [297](#)
- logtalk_load_context/2
 - Built-in predicate, [304](#)
- logtalk_make/0
 - Built-in predicate, [298](#)
- logtalk_make/1
 - Built-in predicate, [299](#)
- logtalk_make_target_action/1
 - Built-in predicate, [301](#)

M

- message, [1039](#)
- message lookup, [1039](#)
- message to self, [1039](#)
- message_hook/4
 - Built-in method, [375](#)
- message_prefix_file/6
 - Built-in method, [377](#)
- message_prefix_stream/4
 - Built-in method, [376](#)
- message_tokens//2
 - Built-in method, [374](#)
- meta_non_terminal/1
 - Directive, [234](#)
- meta_predicate/1
 - Directive, [232](#)
- meta-argument, [1039](#)
- meta-interpreter, [1039](#)
- meta-predicate, [1039](#)
- meta-variable, [1039](#)
- metaclass, [1039](#)
- method, [1039](#)
- missing_directives
 - Flag, [114](#)
- mocking, [1039](#)
- mode/2
 - Directive, [235](#)
- mode_non_terminal/2
 - Directive, [235](#)

- module, [1039](#)
- modules
 - Flag, [113](#)
- monitor, [1039](#)
- multifile predicate, [1040](#)
- multifile/1
 - Directive, [236](#)

N

- naked meta-variable, [1040](#)
- naming
 - Flag, [115](#)

O

- object, [1040](#)
- object database, [1040](#)
- object/1-5
 - Directive, [219](#)
- object_property/2
 - Built-in predicate, [249](#)
- once/1
 - Built-in method, [334](#)
- op/3
 - Directive, [207](#)
- optimize
 - Flag, [117](#)

P

- parameter, [1040](#)
- parameter variable, [1040](#)
- parameter/2
 - Built-in method, [315](#)
- parametric category, [1040](#)
- parametric entity, [1040](#)
- parametric object, [1040](#)
- parametric object proxy, [1040](#)
- parent, [1040](#)
- permission_error/3
 - Built-in method, [347](#)
- phrase//1
 - Built-in method, [365](#)
- phrase/2
 - Built-in method, [366](#)
- phrase/3
 - Built-in method, [367](#)
- polymorphism, [1040](#)
- portability
 - Flag, [114](#)
- predicate, [1040](#)
- predicate alias, [1041](#)
- predicate calling context, [1041](#)
- predicate declaration, [1041](#)
- predicate definition, [1041](#)
- predicate definition context, [1041](#)

predicate directive, [1041](#)
 predicate execution context, [1041](#)
 predicate scope container, [1041](#)
 predicate scope directive, [1041](#)
 predicate shorthand, [1041](#)
 predicate_property/2
 Built-in method, [322](#)
 primary predicate declaration, [1041](#)
 print_message/3
 Built-in method, [374](#)
 print_message_token/4
 Built-in method, [379](#)
 print_message_tokens/3
 Built-in method, [378](#)
 private inheritance, [1041](#)
 private predicate, [1041](#)
 private/1
 Directive, [237](#)
 profiler, [1041](#)
 prolog_compatible_version
 Flag, [113](#)
 prolog_compiler
 Flag, [117](#)
 prolog_dialect
 Flag, [113](#)
 prolog_loader
 Flag, [117](#)
 prolog_version
 Flag, [113](#)
 protected inheritance, [1041](#)
 protected predicate, [1042](#)
 protected/1
 Directive, [238](#)
 protocol, [1042](#)
 protocol/1-2
 Directive, [223](#)
 protocol_property/2
 Built-in predicate, [250](#)
 prototype, [1042](#)
 public inheritance, [1042](#)
 public predicate, [1042](#)
 public/1
 Directive, [239](#)

Q

question_hook/6
 Built-in method, [380](#)
 question_prompt_stream/4
 Built-in method, [381](#)

R

redefined_built_ins
 Flag, [115](#)
 redefined_operators

 Flag, [115](#)
 reflection, [102](#)
 relative_to
 Flag, [118](#)
 reload
 Flag, [118](#)
 repeat/0
 Built-in method, [313](#)
 report
 Flag, [117](#)
 representation_error/1
 Built-in method, [349](#)
 resource_error/1
 Built-in method, [351](#)
 retract/1
 Built-in method, [330](#)
 retractall/1
 Built-in method, [331](#)

S

scratch directory, [1042](#)
 scratch_directory
 Flag, [117](#)
 self, [1042](#)
 self/1
 Built-in method, [317](#)
 sender, [1042](#)
 sender/1
 Built-in method, [317](#)
 set_logtalk_flag/2
 Built-in predicate, [307](#)
 Directive, [209](#)
 setof/3
 Built-in method, [358](#)
 settings file, [1042](#)
 settings_file
 Flag, [113](#)
 singleton method, [1042](#)
 singleton_variables
 Flag, [115](#)
 source file, [1042](#)
 source file directive, [1042](#)
 source_data
 Flag, [118](#)
 specialization, [1042](#)
 specializes_class/2-3
 Built-in predicate, [267](#)
 static binding, [1042](#)
 static entity, [1043](#)
 steadfastness, [1043](#)
 Flag, [114](#)
 structural reflection, [102](#)
 subclass, [1043](#)
 super call, [1043](#)

- superclass, [1043](#)
- suspicious_calls
 - Flag, [115](#)
- synchronized predicate, [1043](#)
- synchronized/1
 - Directive, [240](#)
- syntax_error/1
 - Built-in method, [352](#)
- system_error/0
 - Built-in method, [353](#)

T

- tabling
 - Flag, [113](#)
- tail_recursive
 - Flag, [116](#)
- template method, [1043](#)
- term_expansion/2
 - Built-in method, [369](#)
- term-expansion, [1043](#)
- test dialect, [1043](#)
- test outcome, [1043](#)
- tester file, [1043](#)
- this, [1043](#)
- this/1
 - Built-in method, [318](#)
- threaded engine, [1043](#)
- threaded/0
 - Directive, [224](#)
- threaded/1
 - Built-in predicate, [272](#)
- threaded_call/1-2
 - Built-in predicate, [273](#)
- threaded_cancel/1
 - Built-in predicate, [280](#)
- threaded_engine/1
 - Built-in predicate, [285](#)
- threaded_engine_create/3
 - Built-in predicate, [283](#)
- threaded_engine_destroy/1
 - Built-in predicate, [284](#)
- threaded_engine_fetch/1
 - Built-in predicate, [291](#)
- threaded_engine_next/2
 - Built-in predicate, [287](#)
- threaded_engine_next_reified/2
 - Built-in predicate, [288](#)
- threaded_engine_post/2
 - Built-in predicate, [290](#)
- threaded_engine_self/1
 - Built-in predicate, [286](#)
- threaded_engine_yield/1
 - Built-in predicate, [289](#)
- threaded_exit/1-2

- Built-in predicate, [277](#)
- threaded_ignore/1
 - Built-in predicate, [276](#)
- threaded_notify/1
 - Built-in predicate, [282](#)
- threaded_once/1-2
 - Built-in predicate, [274](#)
- threaded_peek/1-2
 - Built-in predicate, [278](#)
- threaded_wait/1
 - Built-in predicate, [281](#)
- threads
 - Flag, [113](#)
- throw/1
 - Built-in method, [338](#)
- top-level interpreter shorthand, [1043](#)
- transparent-box view, [103](#)
- trivial_goal_fails
 - Flag, [115](#)
- true/0
 - Built-in method, [311](#)
- type_error/2
 - Built-in method, [341](#)

U

- undefined_predicates
 - Flag, [114](#)
- underscore_variables
 - Flag, [117](#)
- unicode
 - Flag, [113](#)
- unknown_entities
 - Flag, [114](#)
- unknown_predicates
 - Flag, [114](#)
- ununinstantiation_error/1
 - Built-in method, [340](#)
- use_module/1
 - Directive, [226](#)
- use_module/2
 - Directive, [243](#)
- uses/1
 - Directive, [225](#)
- uses/2
 - Directive, [241](#)

V

- version_data
 - Flag, [114](#)
- visible predicate, [1043](#)