

Lush: Reference Manual

Léon Bottou, Yann LeCun

January 7, 2003

Contents

1	Introduction to Lush	17
1.1	What is Lush	17
1.2	A tour of the features	19
1.3	A tour of the Lush Libraries	20
1.4	Why Lisp?	21
1.5	What does LUSH stand for?	22
1.6	History of Lush	22
1.7	What is Lush good for?	24
2	Getting Started with Lush: a quick tutorial	27
2.1	Installing Lush on Unix	27
2.2	starting Lush	28
2.3	Setting things up	28
2.4	Reporting Bugs and Problems	29
2.5	Getting Help and Browsing the Documentation	29
2.6	Basic Syntax	30
2.6.1	Symbol Names and Special Characters	31
2.7	Basic Types	32
2.8	defining functions	32
2.9	Loops, conditionals, local variables, and such	33
2.10	Compiling Lush functions	35
2.11	Writing Lush program files	37
2.12	On-line Documentation	38
2.13	Multiple file programs and search path	39
2.14	Calling C functions from Lush	40
2.15	Mixing Lisp and inline C code	41
2.16	Lists	41
2.17	Strings and Regular Expressions	43
2.18	Scalars, Vectors, Matrices, Tensors, and Storages	43
2.18.1	Working with Storages	44
2.18.2	Manipulating Tensors and IDX	45
2.18.3	tensor iterators	45
2.18.4	idx manipulations	46
2.18.5	simple tensor operations	48

2.18.6	operations between idx and scalars	49
2.18.7	contracting operations with scalar result	49
2.18.8	special inner and outer products	49
2.18.9	Simple Linear Algebra	50
2.19	Complex Numbers	50
2.20	Object-Oriented Lush Programming	51
2.20.1	defining a new class	51
2.20.2	inheritance	52
2.21	Input and Output	53
2.22	A quick tutorial on Lush Graphics	54
2.23	Writing Executable "Shell Scripts" with Lush	55
2.24	Writing Standalone Binary Executables with Lush	56
2.25	The Ogre GUI Toolkit	56
2.26	Using the Simple Directmedia Layer interface (SDL)	57
2.26.1	sdlidx-screen class: Mapping the content of an IDX to an SDL screen	58
2.26.2	sdl-screen class: doubled buffered animations	59
2.26.3	sdl-sprite: movable animated objects	59
2.26.4	sdlgl-screen class: OpenGL 3D graphics in SDL screens .	60
2.27	Using OpenGL	60
3	Lush Basics	61
3.1	Lush Interpreter Basics.	61
3.2	Lush Startup.	62
3.3	The Lush Reader.	63
3.4	The Lush Evaluator.	64
3.5	Errors in Lush.	64
3.6	Interruptions in Lush.	65
3.7	Leaving Lush.	66
3.8	Lush Memory Management.	67
3.9	Lush Libraries.	67
3.10	Lush Runtime.	68
4	Core Interpreter and Startup Library	71
4.1	Lists.	71
4.1.1	List Manipulation Functions	71
4.1.2	Physical List Manipulation Functions	79
4.2	Numbers.	81
4.2.1	Elementary Numerical Functions.	81
4.2.2	Integer Arithmetic Functions.	84
4.2.3	Mathematical Numerical Functions.	85
4.2.4	Random Numbers	92
4.2.5	Constant Numbers.	93
4.2.6	Hashing	93
4.3	Booleans Operators.	93
4.3.1	Boolean Arithmetic	93

4.3.2	Predicates.	97
4.4	Symbols.	100
4.4.1	(defvar name val)	101
4.4.2	(setq v1 a1 ... [vn an])	101
4.4.3	(set v a)	102
4.4.4	(incr v [n])	102
4.4.5	(named s)	102
4.4.6	(nameof s)	103
4.4.7	(namedclean s)	103
4.4.8	(lock-symbol s1 ... sn)	103
4.4.9	(unlock-symbol s1 ... sn)	103
4.4.10	(symlist)	104
4.4.11	(oblist)	104
4.4.12	(macrochp s)	104
4.4.13	(putp anything name value)	104
4.4.14	(getp anything name)	104
4.5	Namespaces and Packages	105
4.5.1	(private ..symbols..)	105
4.5.2	(package name ..expressions..)	105
4.6	Control structures.	105
4.6.1	(eval l1...ln)	106
4.6.2	(apply f l)	106
4.6.3	(quote a)	106
4.6.4	'a	107
4.6.5	'expr ,a ,@a	107
4.6.6	(progn l1...ln)	108
4.6.7	(progl l1...ln)	108
4.6.8	(let ((s1 v1) ... (sn vn)) l1 ... ln)	108
4.6.9	(let* ((s1 v1) ... (sn vn)) l1 ... ln)	109
4.6.10	(let-filter ((filter data)) l1...ln)	110
4.6.11	(if cond yes [no1...non])	110
4.6.12	(when cond yes1...yesn)	110
4.6.13	(while cond l1...ln)	111
4.6.14	(do-while cond l1...ln)	111
4.6.15	(repeat n l1...ln)	111
4.6.16	(mapwhile cond l1...ln)	112
4.6.17	(for (symb start end [step]) l1...ln)	112
4.6.18	(mapfor (symb start end [step]) l1...ln)	112
4.6.19	(cond l1...ln)	113
4.6.20	(selectq s l1...ln)	113
4.6.21	(mapc f l1...ln)	114
4.6.22	(mapcar f l1...ln)	114
4.6.23	(rmapc f l1...ln)	114
4.6.24	(rmapcar f l1...ln)	114
4.6.25	(each ((s1 v1) ... (sn vn)) l1 ... ln)	115
4.6.26	(all ((s1 v1) ... (sn vn)) l1 ... ln)	115

4.7	Functions.	115
4.7.1	Argument List.	117
4.7.2	Defining a Function.	118
4.8	Strings.	124
4.8.1	Basic String Functions	125
4.8.2	Regular Expressions (regex).	130
4.9	Storages	134
4.9.1	Storage Creation and Allocation	134
4.9.2	Storage allocation	136
4.9.3	Storage access	137
4.9.4	Miscellaneous storage functions	137
4.10	Matrices, Vectors, Tensors, Indexes and IDX	138
4.10.1	IDX Structure	139
4.10.2	IDX Creation and Allocation	139
4.10.3	IDX Literals	142
4.10.4	IDX Copying and Cloning	143
4.10.5	IDX Iterators	143
4.10.6	IDX Information	145
4.10.7	IDX Manipulations	147
4.10.8	Loading/Saving	150
4.10.9	Component-wise Unary Operations	154
4.10.10	Component-wise dyadic operations	156
4.10.11	Contracting Operations with Scalar Result.	157
4.10.12	Operations between tensors and scalars	158
4.10.13	Matrix/Vector and 4-Tensor/Matrix products	159
4.10.14	Outer Products	159
4.10.15	Backward Compatibility Functions	159
4.11	Objects.	161
4.11.1	Object Terminology.	161
4.11.2	Inheritance.	162
4.11.3	Predefined Classes	162
4.11.4	Defining a Class.	163
4.11.5	Creating Objects.	165
4.11.6	Accessing Slots.	166
4.11.7	Defining Methods.	168
4.11.8	Sending Messages.	169
4.11.9	Constructors and Destructors.	171
4.11.10	Predefined Methods.	173
4.12	Hash Tables.	174
4.12.1	Representing Associations with Hash Tables.	174
4.12.2	Representing Sets with Hash Tables.	178
4.13	Dates.	179
4.13.1	(date-to-day date).	179
4.13.2	(day-to-date daynumber)	180
4.13.3	(today)	180
4.13.4	(now)	180

4.13.5	(split-date <i>date</i>)	180
4.13.6	(date-type <i>date</i>)	181
4.13.7	(date-extend <i>date from to</i>)	181
4.13.8	(date-to-string <i>date [format]</i>)	182
4.13.9	(string-to-date <i>string [from to]</i>)	184
4.13.10	(date-add-second <i>date count</i>)	184
4.13.11	(date-add-minute <i>date count</i>)	184
4.13.12	(date-add-hour <i>date count</i>)	184
4.13.13	(date-add-day <i>date count</i>)	185
4.13.14	(date-add-month <i>date count [opt]</i>)	185
4.13.15	(date-add-year <i>date count [opt]</i>)	185
4.13.16	(date-code <i>date flags</i>)	186
4.14	Loading and Executing Lush Program Files.	186
4.14.1	(load in [<i>out [prompt]</i>])	186
4.14.2	\hat{L} <i>filename</i>	187
4.14.3	(libload <i>libname [opt]</i>)	187
4.14.4	file-being-loaded	188
4.14.5	(find-file <i>dirlist filename extlist</i>)	188
4.14.6	(save <i>f s1 ... sn</i>)	188
4.14.7	(autoload <i>f s1 ... sn</i>)	188
4.15	Lisp-Style Input / Output.	189
4.15.1	Input.	189
4.15.2	Output.	190
4.15.3	File Names.	192
4.15.4	Searched Path.	193
4.15.5	Files and Directories Management.	194
4.15.6	Files and Descriptors.	196
4.15.7	Binary I/O.	200
4.16	Miscellaneous.	201
4.16.1	Debug.	201
4.16.2	System.	207
4.16.3	Special Numerical Values (IEEE754).	217
4.17	Graphics.	219
4.17.1	Creating a Window	220
4.17.2	Drawing.	221
4.17.3	Drawing with Colors.	225
4.17.4	Drawing Text.	228
4.17.5	The Drawing Context.	229
4.17.6	Double Buffering and Synchronization.	231
4.17.7	Plotting (obsolete).	231
4.17.8	Events.	235
4.17.9	System Specific Graphic Functions	239
4.17.10	Producing Encapsulated PostScript Files.	242
4.17.11	Graphical Utilities.	243
4.17.12	Editable Graphics with comdraw/ivtools	244
4.18	Events and Timers	248

4.18.1	Event Queue.	248
4.18.2	Timers.	251
4.19	Documentation and Help System	251
4.19.1	Using the Online Manual	252
4.19.2	Writing Documentation	254
4.20	Obsolete functions	261
4.20.1	Operations on Lists of Numbers	261
4.20.2	Statistical Functions on Lists of Numbers.	263
4.20.3	DZ.	267
5	Ogre: Object Oriented GUI Toolkit	277
5.1	Introduction to Ogre.	277
5.2	Ogre Tutorial.	278
5.3	The Ogre Class Library.	280
5.3.1	Ogre Methods.	280
5.3.2	Ogre Class Hierarchy.	282
5.4	Ogre Utility Functions.	284
5.4.1	Initializing the Ogre library.	284
5.4.2	Error handling in the Ogre library.	284
5.4.3	Ogre Color Palette.	285
5.4.4	Ogre Fonts.	288
5.5	Visual Objects.	289
5.5.1	(new VisualObject w h)	289
5.5.2	VisualObject Request Messages.	290
5.5.3	VisualObject Implementation Methods.	291
5.5.4	Event Methods.	294
5.6	Control Objects.	297
5.6.1	(new Control w h f)	298
5.6.2	Enabling or Disabling a Control Object.	298
5.6.3	Activation of a Control Object.	299
5.6.4	Appearance of a Control Object.	300
5.6.5	State of a Control Object.	301
5.6.6	Ogre Callbacks.	301
5.7	Container Objects.	303
5.7.1	(new Container x y w h ...contents...)	303
5.7.2	Repainting.	303
5.7.3	Inserting an Removing Objects.	304
5.7.4	Geometry Management.	304
5.7.5	Control Management.	307
5.8	Container Flavors.	308
5.8.1	Forms.	308
5.8.2	Window Objects.	309
5.8.3	Structuring Containers	312
5.9	Buttons.	315
5.9.1	Push Buttons.	315
5.9.2	Check Boxes.	316

5.9.3	Exclusive Buttons.	318
5.9.4	File Requester Button.	319
5.10	Character String Objects.	320
5.10.1	Fixed Strings.	320
5.10.2	Editable Strings.	321
5.10.3	Editable Numbers.	323
5.10.4	EditSecretString	323
5.10.5	Multiline Text Editor.	323
5.11	Icons.	326
5.11.1	(new icon mat &optional [sx [sy [map]]])	327
5.12	OgrImage.	327
5.12.1	(new OgrImage m)	327
5.12.2	(==> OgrImage get-selected)	327
5.13	ImageViewer.	327
5.13.1	(new ImageViewer w h m [scroll])	327
5.13.2	(==> ImageViewer get-selected)	328
5.14	Menus.	328
5.14.1	Standard Menus.	328
5.14.2	Choice Menus.	330
5.15	Popup Requesters.	331
5.15.1	Requester.	332
5.15.2	Warning Requester.	334
5.15.3	Error Requester.	334
5.15.4	Yes/No Requester.	335
5.15.5	Print Requester.	337
5.15.6	File Requester.	338
5.16	Movable and Resizable Objects.	340
5.16.1	(new DragArea w h)	340
5.16.2	(new SizeArea w h)	340
5.16.3	(new Frame x y w h ...contents...)	341
5.16.4	(new FrameSize x y w h ...contents...)	341
5.17	Sliders and Scrollbars.	341
5.17.1	Sliders.	342
5.17.2	Scrollbars.	343
5.18	Composite Objects.	344
5.18.1	Viewers.	344
5.18.2	Selectors.	345
5.19	A Complete Example.	347
5.19.1	The Class Browser “classtool”.	347
5.19.2	The Program “classtool”.	348
6	Dynamic Loader/Linker	355
6.1	(mod-load filename)	356
6.2	(mod-unload filename)	357
6.3	(find-shared-library name [extlist])	357
6.4	(find-static-library name [extlist])	357

6.5	(find-shared-or-static-library <i>name</i> [<i>extlist</i>])	358
6.6	(mod-list)	358
6.7	(mod-undefined)	358
6.8	(mod-status)	358
6.9	(mod-inquire <i>filename</i>)	358
6.10	(mod-create-reference <i>string1</i> ... <i>stringN</i>)	359
6.11	(mod-compatibility-flag <i>boolean</i>)	359
6.12	Low level module functions.	359
6.12.1	(module-list)	359
6.12.2	(module-filename <i>m</i>)	359
6.12.3	(module-executable-p <i>m</i>)	359
6.12.4	(module-unloadable-p <i>m</i>)	359
6.12.5	(module-initname <i>m</i>)	360
6.12.6	(module-depends <i>m</i>)	360
6.12.7	(module-never-unload <i>m</i>)	360
6.12.8	(module-defs <i>m</i>)	360
6.12.9	(module-load <i>filename</i> [<i>hookfunction</i>])	360
6.12.10	(module-unload <i>m</i>)	361
6.13	Extending the Interpreter	361
6.14	Debugging modules with gdb.	362
7	CLush: Compiled Lush	365
7.1	A Simple Example	365
7.2	Using libload and dhc-make.	365
7.3	Compilation Functions	366
7.3.1	High-Level Compilation Functions	366
7.3.2	Customizing the Behavior of dhc-make.	368
7.3.3	Querying File Dependencies	370
7.3.4	Low-Level Compilation Functions.	370
7.4	Design Philosophy	371
7.5	What is Compilable, and What isn't	372
7.5.1	(compilablep <i>function</i>)	373
7.6	CLush-specific functions and constructs	373
7.6.1	Numerical Type Declarations	373
7.6.2	Class Declaration and Compilation	375
7.6.3	Type Conversions	376
7.6.4	inline C code	377
7.6.5	Compiler Directives	379
7.6.6	Dynamic Allocation in Compiled Code	380
7.7	Interfacing Existing C/C++/FORTRAN Libraries to Lush	380
7.7.1	Symbolic constants.	380
7.7.2	Interfacing C code	381
7.7.3	Interfacing C++ code	381
7.7.4	Interfacing FORTRAN code	381
7.8	More on the Lisp-C Interface	381
7.8.1	DH: the Compiled Function Class	381

7.8.2	Classinfo	381
7.8.3	Controlling the Lisp/C interface.	382
7.9	Compiler Internals	383
7.9.1	(dhc-generate-c filename '([func1 [funcn]]))	383
7.9.2	(compilablep function)	383
7.9.3	(dhc-substitute-env str [htable])	383
7.9.4	varlushdir	384
7.9.5	(dhc-make-cdir cdir [create])	384
7.9.6	(dhc-make-c-filename src)	384
7.9.7	(dhc-make-o-filename src)	384
7.9.8	(dhc-make-get-dependencies [sname])	384
7.9.9	dhc-make-lushflags	384
7.9.10	dhc-make-command	385
7.9.11	dhc-make-overrides	385
7.9.12	dhc-make-force	385
7.9.13	dhc-make-essential-libs	385
7.9.14	(dhc-make-o src-file [obj-file [lushflags]])	385
7.9.15	(dhc-make-o-maybe src-file [obj-file [cflags]])	385
7.9.16	(dhc-make-c fname fsymblist)	386
7.9.17	(dhc-make-c-maybe sname fname fsymblist)	386
7.9.18	(dhc-make-test fname)	386
7.9.19	(dhc-make-all fname fsymblist liblist)	386
7.9.20	(dhc-make fname f1 [f2 ...[fn]])	386
7.9.21	(dhc-make-with-libs fname library-list f1 [f2 ...[fn]])	386
7.9.22	(dhc-make-with-c++ fname library-list f1 [f2 ...[fn]])	386
7.9.23	(dhc-insert-op 1 op)	387
7.9.24	(dhc-nolast 1)	387
7.9.25	(dhc-remove-dup 1)	387
7.9.26	(dhc-remove-eqdup 1)	387
7.9.27	(dhc-remove-nth n 1)	387
7.9.28	(dhc-split-list 1 test)	387
7.9.29	(dhc-postincr symb)	387
7.9.30	dhc-debug-flag	388
7.9.31	—#@—	388
7.9.32	(dhc-add-c-declarations str)	388
7.9.33	(dhc-add-c-statements str)	388
7.9.34	(dhc-add-c-epilog str)	388
7.9.35	(dhc-add-c-externs str)	388
7.9.36	(dhc-add-c-metaexterns str)	388
7.9.37	(dhc-add-c-header str)	388
7.9.38	(dhc-add-c-pheader str)	389
7.9.39	(dhc-class-to-struct-decl type)	389
7.9.40	(dhc-class-to-vtable-decl type)	389
7.9.41	(dhc-type-to-c-decl type)	389
7.9.42	(dhc-declare-var name type)	389
7.9.43	(dhc-declare-temp-var type [clue])	389

7.9.44	dhc-debug-stack	389
7.9.45	(dhc-error string [arg])	389
7.9.46	(dhc-check-symbol source)	390
7.9.47	(dhc-internal-error str)	390
7.9.48	dhc-type	390
7.9.49	(==> dhc-type print)	390
7.9.50	(==> dhc-type hashcode)	390
7.9.51	(==> dhc-type access [v])	390
7.9.52	(new dhc-type class [a1] [a2])	390
7.9.53	(dhc-desc-to-type desc)	391
7.9.54	(new dhc-symbol name lex [fmt])	391
7.9.55	(dhc-search-symtable symbolname table)	391
7.9.56	(dhc-add-to-symtable table symbolobject)	391
7.9.57	(dhc-add-symbol-table symb lex)	391
7.9.58	(dhc-add-global-table symb)	391
7.9.59	(new t-node t-node-list type [source] [symbol])	391
7.9.60	(dhc-make-t-node expr)	392
7.9.61	(dhc-copy-source-tree source)	392
7.9.62	(dhc-get-treotype source)	392
7.9.63	(dhc-get-type source)	393
7.9.64	(dhm-t func (source) . body)	393
7.9.65	(dhm-t-declare model f1 ... fn)	393
7.9.66	(dhm-c symb (source treotype replace) . body)	393
7.9.67	(dhm-c-declare model f1 ... fn)	393
7.9.68	(get-dhm-t symb)	393
7.9.69	(get-dhm-c symb)	393
7.9.70	(get-dhm-target source)	394
7.9.71	(dhm-p func (source) . body)	394
7.9.72	(get-dhm-p symb)	394
7.9.73	(dhc-pp body)	394
7.9.74	(dhc-parse-replacement-source-t source newsource)	394
7.9.75	(process-numerical-args-t arglist retype)	394
7.9.76	u-nodes	395
8	Standard Libraries	397
8.1	Shell Commands	397
8.1.1	String List Utilities	397
8.1.2	File Information	398
8.1.3	Variables	399
8.1.4	Directories	399
8.1.5	Shell Commands	400
8.2	Tensor/Matrix/Vector/Scalar Libraries	401
8.2.1	Generic IDX Macros and Functions	401
8.2.2	IDX of Single Precision Floating Point Numbers	405
8.2.3	IDX of Double Precision Floating Point Numbers	411
8.2.4	IDX of Integers	418

8.2.5	IDX of unsigned bytes	420
8.2.6	Arrays of vectors and matrices	422
8.2.7	IDX reading and writing	424
8.2.8	IDX mapping from files	444
8.2.9	Convolution, Subsampling, Oversampling.	447
8.2.10	IDX sorting and binary search functions	453
8.2.11	IDX operations on squares of elements	455
8.3	C Interface Utilities	456
8.3.1	Libc interface	456
8.3.2	Control Structures	457
8.3.3	Pointer Manipulation/Dereferencing	457
8.4	Plotting	459
8.4.1	Plotter	460
8.4.2	low level utility functions	467
8.5	C-Style Input / Output	468
8.5.1	(fprintf file-pointer args...)	468
8.5.2	(fwrite-str file-pointer s)	469
8.5.3	(stdout)	469
8.5.4	(fopen filename type)	469
8.5.5	(fclose file-pointer)	470
8.5.6	(popen filename type)	470
8.5.7	(pclose file-pointer)	470
8.5.8	(ftell file-pointer)	470
8.5.9	(fseek file-pointer pos)	471
8.5.10	(fseek-from-end file-pointer pos)	471
8.5.11	(fseek-from-current file-pointer pos)	471
8.5.12	(fgetc file-pointer)	472
8.5.13	(putc file-pointer val)	472
8.5.14	(fread-ubyte file-pointer)	472
8.5.15	(fwrite-ubyte file-pointer val)	472
8.5.16	(fread-byte file-pointer)	473
8.5.17	(fwrite-byte file-pointer val)	473
8.5.18	(fread-short file-pointer)	473
8.5.19	(fwrite-short file-pointer val)	473
8.5.20	(fread-int file-pointer)	474
8.5.21	(fwrite-int file-pointer val)	474
8.5.22	(fread-flt file-pointer)	474
8.5.23	(fwrite-flt file-pointer val)	475
8.5.24	(fread-real file-pointer)	475
8.5.25	(fwrite-real file-pointer val)	475
8.5.26	(reverse_n ptr sz n)	475
8.5.27	(fscan-int file-pointer)	476
8.5.28	(fscan-flt file-pointer)	476
8.5.29	(fscan-str file-pointer)	476
8.5.30	(fgets file-pointer max-size)	476
8.5.31	(file-size file-name)	477

8.5.32	(rewind f)	477
8.5.33	(skip-comments start f)	477
8.6	Graph Library	477
8.6.1	Building and Manipulating Grammars	477
8.6.2	Manipulating Graphs.	483
8.6.3	Searching and Pruning Graphs	486
8.6.4	Transformation Operators on Graphs	487
8.6.5	Other Graph Search Functions	491
8.7	Image Processing Libraries	491
8.7.1	Demos and examples	491
8.7.2	Reading/Writing Image Files to/from IDX.	491
8.7.3	Reading/Writing PPM/PGM/PBM Image Files	493
8.7.4	Reading/Writing PBM Image files	496
8.7.5	Reading/Writing Run-Length-Encoded Image Files	497
8.7.6	RGBA Images of ubytes	498
8.7.7	RGBA Images of Floats	504
8.7.8	Greyscale Images of ubytes.	508
8.7.9	Greyscale Images of shorts	512
8.7.10	Greyscale Images of Floats	515
8.7.11	Computing Image Warping Maps	515
8.7.12	Connected Component Analysis	516
8.7.13	Morphological Operations	520
8.7.14	Morphological Operation on Images of Shorts.	521
8.7.15	Miscellaneous Image Macros and Conversions	523
8.7.16	Constants used in run-length encoded images.	524
8.7.17	Converting RLE Image to greyscale pixelmap.	524
8.7.18	Color Clustering and Quantization	526
9	Packages	533
9.1	GBLearn2: Machine Learning Library	533
9.1.1	GBLearn2: Basic Concepts	534
9.1.2	Examples and Demos	536
9.1.3	Learning Algorithms	536
9.1.4	Data Sources	540
9.1.5	Meters	544
9.1.6	gb-states	545
9.1.7	gb-param	548
9.1.8	Module Libraries	550
9.2	GBLearning: Gradient-Based Learning Library	567
9.2.1	Examples and Demos	567
9.2.2	modules	567
9.2.3	workbench	576
9.2.4	trainer	577
9.2.5	Databases	580
9.2.6	param	584
9.2.7	sn-states	586

9.2.8	libnets	588
9.2.9	net-ccc	590
9.2.10	net-cfe	591
9.2.11	net-cff	591
9.2.12	net-cf	591
9.2.13	net-c	592
9.2.14	net-cscf	592
9.2.15	net-cscfe	593
9.2.16	net-cscf	593
9.2.17	net-cscscf	593
9.2.18	net-cscscfe	594
9.2.19	net-cscscscf	594
9.2.20	net-fff	594
9.2.21	net-ff	595
9.2.22	lclass-meter	595
9.2.23	ccc-tdnn	595
9.2.24	conx	597
9.2.25	ftdnn	598
9.2.26	seq3-modules	599
9.3	SDL: Simple DirectMedia Layer	600
9.3.1	Installing SDL	600
9.3.2	SDL Demos	600
9.3.3	LibSDL: High-level Interface to SDL	601
9.3.4	Low-Level Interface to SDL	610
9.4	OpenGL	656
9.5	OpenRM	656
9.6	Video4Linux: video grabbing	656
9.6.1	Requirements and Installation	656
9.6.2	Video4Linux API	656
9.6.3	Demos	661
9.7	BLAS: Basic Linear Algebra Subroutine	663
9.8	LAPACK	664
9.9	GSL: GNU Scientific Library	664
9.10	OpenCV: Intel Computer Vision Library	664
9.11	ALSA: Advanced Linux Sound Architecture	664
10	Applications	665
10.1	Graphic Tools	665
11	Tutorials	667
11.1	Tutorial: Writing Games with Lush and SDL	667
11.1.1	A Quick Reminder on Basic Lush Programming	667
11.1.2	A Simple Lunar Lander	668
11.1.3	SpaceWar: missiles and collision detection	680

12 FAQ (Frequently Asked Questions)	681
12.1 What to do when I get this error?	681
12.1.1 Module has undefined references	681
12.1.2 compiler : Unknown Type in: ()	681
12.2 How Do I ... ?	682
12.2.1 Read lines from a file into a list	682
12.2.2 Apply a function to all elements of a vector	682
12.2.3 Get a pointer to the raw data of an idx	682
12.2.4 Get a pointer to a function written in Lisp	682
12.2.5 Know if a function can be used in compiled code	683
13 Local Packages	685
<< document file not found >>	

Chapter 1

Introduction to Lush

1.1 What is Lush

Lush is an object-oriented programming language with features designed to please researchers, experimenters, and engineers interested in large-scale numerical and graphic applications. Lush is designed to be used in situations where one would want to combine the flexibility of a high-level, loosely-typed interpreted language, with the efficiency of a strongly-typed, natively-compiled language, and with the easy integration of code written in C, C++, or other languages.

Lush can be used advantageously for projects where one would otherwise a combination of an interpreted language like Python, Perl, Matlab, S+, or even (gasp!) BASIC, and a compiled language like C. Lush brings the best of both worlds by wrapping three languages into one: (1) a loosely-typed, garbage-collected, dynamically scoped, interpreted language with a simple Lisp-like syntax, (2) a strongly-typed, lexically-scoped compiled language that uses the same Lisp-like syntax, and (3) the C language, which can be freely mixed with Lush code within a single program, even within a single function. It sounds complicated, but it is not. In fact, Lush is designed to be very simple to learn and easy to use.

Lush's main features are as follows:

- A very clean, simple, and easy to learn Lisp-like syntax.
- A very efficient native compiler (through C).
- An easy way to interface C functions and libraries, and a powerful dynamic loader for object files of libraries (.o, .a and .so files) written in other compiled languages.
- The ability to freely mix Lisp and C in a single function.
- A powerful set of vector/matrix/tensor operations.

- A huge library of numerical routines , including GSL , LAPACK , and BLAS.
- An extensive set of graphic routines (including an object-oriented GUI toolkit , an interface to OpenGL , and an interface to SDL).
- Sound and video grabbing (using ALSA and Video4Linux).
- A library of image and signal processing routines.
- Several libraries for machine learning , neural net and statistical estimation.

This combination of flexibility , efficiency , and extensive libraries makes Lush an ideal platform for research and development in signal processing , image processing , machine learning , computer vision , bio-informatics , data mining , statistics , and artificial intelligence. Its speed and extensive libraries allow such things as real-time sound , image , and video processing. Many use Lush as a general purpose script language , or as a general language for application development. Some have been known to use Lush to develop 2D and 3D games. A few have even used Lush to develop commercial software for embedded processors.

Compared to Matlab , Lush is a "real" object-oriented programming language with data structures , typed matrices , a simple syntax , and a native compiler. Compared to Python , Lush has an efficient native compiler (compiled numerical code is as fast as C) , and the ability to practice functional programming. The ability to easily call C libraries , and the ability to freely mix C and Lisp is somewhat unique to Lush.

If at any time you have:

- wished you had simple interpreted language with which you could quickly try out ideas , implement efficient numerical algorithms , or prototype GUI-based applications.
- written a piece of software in C and wished you could control it from a simple , interpreted script language.
- written such a script language interpreter yourself and wished it were a full-blown programming language.
- used an interpreted script-like language , such as Perl , MatLab , Mathematica , Tcl , Python , or BASIC , and wished the compiler generated efficient code , and wished it were easier to call C routines from it.
- written a program with one of the above languages and wished you could generate a portable standalone application from it.
- wished you could combine two languages for your programming projects: an efficient , compiled , no-frills language such as C or C++ for the low-level , and an interpreted , weakly typed language with smart memory management such as Lisp for the high level and the user interface.

then , Lush is for you.

Many software projects , particularly research projects , require two languages: an efficient compiled language such as C or C++ for implementing the low-level or computationally expensive function , and a flexible , possibly interpreted language for high-level control , scripting , experimentation , and tinkering. Popular research-oriented interpreters such as Matlab are somewhat inefficient , lack support for complex data structures , lack the power of a full-fledged object-oriented programming language , and lack simple interfacing to C and other compiled languages.

Lush uses a very simple Lisp syntax. If the word "Lisp" sends shivers down your spine , be advised that the dialect of Lisp that Lush implements is extremely simple to learn , with possibly one the simplest syntax around. Most scientists , engineers , and software developers who have learned Lush have become proficient at it in a few days , even if they had no prior exposure to Lisp. It's very simple. Really. In fact , Lush has been used to teach programming to kids.

The Lush compiler has several interesting properties , and a few limitations. Its main advantage is that it generate very efficient C code which is then compiled with the best available C compiler for the machine considered.

Lush Currently runs on various Unix platforms including Linux/x86 , Solaris , and SGI/Irix.

1.2 A tour of the features

Lush is built around a compact , portable , and intentionally simple Lisp interpreter written in C. It features all the usual functionalities and constructs found in every decent object-oriented programming language such as conditional statements , loops , local variables , functions , macros , objects , classes , methods , and inheritance , but it also provides a large number of functions for manipulating lists , strings , vectors , matrices , and tensors. String functions include such things as regular expression matching and substitutions.

The Lush interpreter is quite similar to some lisps of the mid 80's like "Le_Lisp" and "UCI-Lisp". It is significantly simpler than "Common Lisp" and quite different from Scheme.

An unusual feature of Lush is its compiler to C , its dynamic loader , and the ease with which one can interface existing C functions and libraries. En even more unusual feature is the ability to intermix Lisp and C source code within a single function.

The vector and matrix manipulation engine is quite powerful and efficient , making Lush ideal for heavy numerical applications and applications such as signal and image processing. It includes functions to create , resize , and convert vectors , matrices , and tensors with up to eight dimensions. It includes basic matrix operations such as scalar operations on all elements of a matrix , dot products , outer products , transposition , highly optimized 1D and 2D convolutions. A set of iterators is also provided to access all matrix elements

without requiring costly bound checking.

Lush is an object-oriented language. Classes with slots and methods can be defined and compiled. Derived classes inherit the slots and methods of their parent class (similar to the C++ semantics).

Lush includes all the functions familiar to Lisp aficionados , such as list functions , list iterators , physical list manipulation , macros , splicing macros , symbol manipulation , etc , but most casual users will probably prefer to stay away from those.

Lush provides a set of simple-to-use graphic functions to draw lines , polygons , rectangles , pictures , and text in color , with automatic refresh and double-buffering capability (for simple animations). Graphics can be drawn in an X-windows window (on Unix) , or sent to a PostScript file. Those low-level portable graphic functions are used to build high-level functionalities , such as function plotting.

Lush also contains a very compact and easy-to-use object-oriented graphical user interface generator called “Ogre”. Ogre is entirely written in Lisp on top of the low-level graphic functions mentioned above. It contains predefined classes for buttons , sliders , radio buttons , menus , string editors , etc... Ogre includes an automatic mechanism for placing objects in the window , thereby greatly simplifying the design and implementation of GUIs. Simple GUIs can be written in extremely short times and are very compact.

Lush provides two models for Input/Output. One is a set of “Lispy” I/O functions that allow easy input and output of ASCII data , lisp expressions , lisp objects , and matrices. It includes such goodies as pipes and sockets. The other model essentially provides access to the standard C I/O library , including `fopen` , `popen` , `fprintf` , `fscanf` , `fgetc` , `fputc` , `fgets` , and various functions for reading and writing matrices. Large matrices can be mapped in the virtual addressing space , instead of explicitly loaded in memory. This allows to access very large datasets efficiently.

1.3 A tour of the Lush Libraries

Lush has a huge collection of libraries and utilities available to the user. Some are written in Lisp , some are written in C and interfaced to Lush , others are pre-existing libraries that have been interfaced to Lush.

For our numerically enclined friends , Lush has a full interface to the GNU Scientific Library (GSL) , and to the LAPACK and BLAS Linear Algebra libraries. This gives access to an extensive set of numerical and statistical functions (several thousand in fact).

Lush includes a full interface to the industry-standard OpenGL library that allows the creation of 3D graphics and animations. This interface include an interface to GLUT , OpenGLU , and OpenRM (scene graph rendering engine). This makes Lush an excellent platform to write interactive VR applications and computer games.

Another popular library interfaced to Lush is the video game API Simple

Directmedia Layer (SDL). It is enhanced by a high-level library that allows the easy manipulation of sprites and movable screen objects with pixel-accurate collision detection. This library, combined with Lush's simple syntax, makes it ideal for developing simple video games and for teaching programming to children.

Lush contains a library of image processing functions to load, save, resize and resample, warp, filter, and analyze images. Mathematical morphology operations such as connected component analysis, distance transform, erosion, and dilation, are also available for bitonal images. It also includes classes and functions for easily grabbing video using the Video4Linux API.

Lush includes a graph library that allows grammar and finite state machine construction, graph transduction and composition, and viterbi search algorithm.

Lush includes an extensive library for gradient-based machine learning, including neural networks, radial basis functions, support vector machines, and many others. This library is based on an innovative object-oriented design that allows to build large learning machines out of multiple learning modules and to train them cooperatively. Commercially used optical recognition systems have been built with this library.

In addition, Lush provides various interfaces to multimedia libraries, including Video4Linux (video grabbing) and ALSA (audio recording/playing).

1.4 Why Lisp?

In some segments of the software industry, Lisp is sometimes perceived as an oddity of essentially academic interest. Its theoretical computer science heritage also causes common misconceptions about Lisp, such as being inefficient and difficult to learn. So why choose Lisp? We chose Lisp because contrary to the common prejudice, it is extremely easy to learn, in addition to being flexible, efficient, and compact. Lisp has such a simple, clean and flexible syntax that it is probably the easiest language to learn (at the other extreme, you will find Perl). Our experience with teaching C developers to use Lush is that they become proficient at it in a few days. Script language designers often make the mistake of not only designing their own language functionalities, but also designing their own syntax. Examples of this include MatLab, Mathematica, S+, and many others. Why invent a new syntax when a good one such as Lisp already exists? Lush is an object-oriented dialect of Lisp that puts the emphasis on ease of use, efficiency for numerical operations (unlike many traditional Lisp implementations), and easy interfacing with existing code written in C.

Another unique advantage of Lisp is that it is a "programmable programming language". In Lisp, a program is just a data structure that can be created and manipulated just like any other data structure. In other words, Lisp programs can create other Lisp programs. This allows users (and not just language designers) to extend the language themselves ad infinitum. This feature also makes it easy to write such things as syntax transformers, programs that op-

imize themselves , compilers , automatic differentiators , etc. Most users will probably stay away from writing such things , but they will profit from the work of others in these areas (the CLush compiler is an example).

1.5 What does LUSH stand for?

LUSH is supposed to stand for "Lisp Universal SHell" or something like that. But English dictionaries tell another tale. According to the 1914 edition of the Century Dictionary (<http://century-dictionary.com>), Lush has 4 main meanings in English:

LUSH (1)

- Adjective
 1. Slack; limp; flexible.
 2. Mellow; easily turned, as ground.
 3. Fresh, luxuriant, and juicy; succulent, as grass or other vegetation.
- Noun
 1. A twig for thatching. [Prov. Eng.]

LUSH (2)

- Verb Intransitive
 1. To rush violently.
 2. To splash in water.

LUSH (3)

- Noun
 1. Beer; intoxicating drink. [Slang.]
- Verb
 1. To drink; tipple on. [Slang.]
 2. To drink intoxicating liquor. [Slang.]

LUSH (4)

- Noun
 1. The Burbot (fresh water cod fish).

To some of us, Lush certainly appears fresh, luxuriant, juicy, succulent, flexible, mellow, and intoxicating.

To a few others, it might appear slack, limp and fishy.

You be the judge.

1.6 History of Lush

Lush is the direct descendent of the SN system which was originally developed by Leon Bottou and Yann LeCun as the front-end of a neural network simulator. Various incarnations of SN have been developed continuously since 1987 , some of which were sold commercially by Neuristique S.A. in France , and eventually grew into a full-fledged prototyping and development environment.

The versions developed at AT&T Bell Labs , and then at AT&T Labs and at the NEC Research Institute were used to build many succesful technologies and products. The most notable ones are:

- a handwriting recognition system used by many banks across the world to read checks automatically. In fact , some ATM machines made by NCR (that can read checks) run compiled SN code on embedded DSP boards.
- the prototype of the DjVu image compression system.
- numerous machine learning algorithms developed at AT&T since 1988 , including the "LeNet" family of convolutional neural networks and some early implementations of the Support Vector Machine algorithm.

SN was primarily used internally at AT&T Bell Labs for many research projects in machine learning , pattern recognition , and image processing. But its various incarnations were used at AT&T Labs , Lucent , the Salk Institute , the University of Toronto , Universite of Montreal , UC Berkeley , and many other research institutions. The commercial versions of SN were used in several large companies as a prototyping tool: Thomson-CSF , ONERA ,

Contributors include: Leon Bottou , Yann LeCun , Jie Huang Fu , Patrice Simard , Yoshua Bengio , Jean Borelly , Patrick Haffner , Pascal Vincent , Sergey Ioffe , and many others.

In 2001 , AT&T and Neuristique released their respective versions under the GPL , allowing the development and distribution of Lush. Turning SN into Lush was done by Yann LeCun , Leon Bottou and Jie Huang-Fu at the NEC Research Institute.

Here is a family tree of the various incarnations of SN and Lush:

```

SN(1987) neural network simulator for AmigaOS (Leon Bottou, Yann LeCun)
|
SN1(1988) ported to SunOS. added shared-weight neural nets and graphics (LeCun)
|  \ \
|   SN1.3(1989) commercial version for Unix (Neuristique)
|   /
SN2(1990) new lisp interpreter and graphic functions (Bottou)
|  \ \
|   SN2.2(1991) commercial version (Neuristique)
|   |
|   SN2.5(1991) ogre GUI toolkit (Neuristique)
|   / \ \
\ \ /  SN2.8(1993+) enhanced version (Neuristique)
|     \ \
|     TL3(1993+) lisp interpreter for Unix and Win32 (Neuristique)
|       [GPL]
|       \ \-----

```

```

|
SN27ATT(1991) custom AT&T version
|           (LeCun, Bottou, Simard, AT&T Labs)
|
SN3(1992) IDX matrix engine, Lisp->C compiler/loader and
|           gradient-based learning library
|           (Bottou, LeCun, AT&T)
|
SN3.1(1995) redesigned compiler, added OpenGL and SGI VL
|           support (Bottou, LeCun, Simard, AT&T Labs)
|
SN3.2(2000) hardened/cleanup SN3.x code,
|           added SDL support (LeCun)
|-----|
|/
|
ATTLUSH(2001) merging of TL3 interpreter + SN3.2 compiler
[GPL]           and libraries (Bottou, LeCun, AT&T Labs).
|
LUSH(2002) rewrote the compiler/loader (Bottou, NEC Research Institute)
[GPL]
|
LUSH(2002) rewrote library, documentation, and interfaced packages
[GPL]           (LeCun, Huang-Fu, NEC)
|

```

1.7 What is Lush good for?

Lush is a good tool for a variety of applications. It was originally developed as an environment for experimentation and development of machine learning, neural nets, and pattern recognition applications, but over the years, it has grown into a full-fledged language and rapid development environment.

Here are examples of situations in which Lush (or its predecessor SN) have been used:

- as a simulation environment for neural network and machine learning experiments.
- as a Matlab-like prototyping tool for numerical computation, signal processing, image processing, statistical estimation,...
- as a "control/script language" for software projects.
- as a super debugger and diagnostic tool for large applications
- as a quick prototyping tool for GUI-based applications.

- as a script language for quick hacks.

Here is a small subset of the research projects that have been done with Lush:

- numerous projects in handwriting recognition
- many projects in neural network, machine learning, and statistical estimations
- data mining, fault detection, and database marketing projects.
- image processing research

A few full-fledged commercial applications have been built with Lush, including:

- a complete check amount reader now integrated in NCR's automatic teller machines and large back-office check reading machines (Lush-generated code runs on DSP boards). This is a huge piece of complicated code (60,000 lines of Lush Lisp code automatically converted to C).
- Early version of the foreground/background segmentation module for the DjVu image compression system.
- a neural network simulator and neural network training tool.
- a pen-based data entry systems with handwriting recognition

Chapter 2

Getting Started with Lush: a quick tutorial

Author(s): Yann LeCun

2.1 Installing Lush on Unix

get the tarball from [<http://sourceforge.net/projects/lush>], cd to a convenient place (say `/home/your-home-directory` or `/usr/local`) and unpack the tarball. Then cd to `lush` and type `./configure` then `make` . If you want to install the Lush executable so that all users on your machine have access to Lush , su to root , and type `make install` . Otherwise , simply add `lush/bin` to your shell path.

Compiling and installing lush requires that you have certain packages and libraries available. In particular , you need `libiberty.a` , which is not installed by default on certain Linux distributions (e.g. Mandrake 8.2 has it in `libbinutils2-devel-2.11.92.0.12-6mdk`).

To enjoy the full capabilities of Lush , you may also want to install SDL and its development libraries (e.g. on Mandrake 8.2 `libSDL1.2-1.2.3-6mdk` and `libSDL1.2-devel-1.2.3-6mdk`). It is also preferable to set up OpenGL is correctly on your machine (including the "development" packages that contain header files etc).

In addition , you may optionally install:

- GSL: the GNU Scientific Library available at [<http://sources.redhat.com/gsl>]
- The LAPACK and BLAS numerical libraries (which many Linux distros have pre-packaged)
- The Intel Computer Vision Library available at [<http://www.intel.com/software/products/opensource/libra>] [<http://opencv.sourceforge.net>]

2.2 starting Lush

If the lush executable was installed in your path , you can start lush by simply typing `lush` at the Unix prompt. Otherwise you can make a symbolic link from anywhere in your path to where the lush executable resides (most likely `/usr/local/lush/bin/lush` or `/wherever-you-installed-lush/bin/lush`)

```

amiga\\% lush
LUSH Lisp Universal Shell (compiled on Oct 17 2002)
  Copyright (C) 2002 Leon Bottou, Yann Le Cun, AT&T Corp, NECI.
  Includes parts of TL3
  Copyright (C) 1987-1999 Leon Bottou and Neuristique.
  Includes selected parts of SN3.2:
  Copyright (C) 1991-2001 AT&T Corp.
This program is free software; it is distributed under the terms
of the GNU Public Licence (GPL) with ABSOLUTELY NO WARRANTY.
Type '(helptool)' for details.

```

```

... loading startup file "/home/yann/lush/sys/sysenv.lsh".
... loading library file "/home/yann/lush/lsh/lushenv.lsh"
[compat.lsh]
[graphenv.lsh]
[brace.lsh]
[lushrc.lsh]
?

```

To exit Lush , simply hit Control-D at the prompt.

If you have `emacs` installed on your system , it is highly recommended to run Lush from within Emacs so as to enjoy such niceties as command-line editing , history , and IDE-like features. See the next section "setting things up" for details on how to do this.

2.3 Setting things up

Lush can be run from the shell prompt. It provides command line editing , history , and symbol and file name completion.

Many will prefer to run Lush under Emacs. To make this convenient , the following line should be added to the `.emacs` file in your home directory:

```
(load "your-lush-dir/etc/lush.el")
```

This file provides two functionalities in Emacs:

- The ability to start Lush within Emacs by typing `Meta-X lush` (META-X is emacs-speak for ESC-X or ALT-X). The lush executable must be

in your shell path for this to work. Lush will run in Emacs's so-called "inferior-lisp" mode. The current directory in which Lush is started is the directory of the file in the buffer you were editing right before starting Lush. It is convenient to split the Emacs window into two. The bottom window will be the Lush execution buffer (called "*inferior-lisp*") while the top window will contain the Lush source file being worked on.

- A Lush mode for Emacs that is automatically activated when a `.lsh` file is being edited, which provides keyword highlighting, automatic indentation, matching paren flashing, etc.

Emacs's Lisp/Lush mode provides a set of nice commands to facilitate Lisp programming. Among them are C-M-e (Ctrl-Meta-e) to jump to the end of a function, C-M-a to jump to the beginning, C-M-q to indent/format a function, and C-M-x to send the function or expression in which the cursor is to the interpreter (assuming it is running in the *inferior-lisp* buffer).

2.4 Reporting Bugs and Problems

If you encounter a bug or want to signal a problem, go to: [http://sourceforge.net/tracker/?atid=410163&group_id=410163], click "Submit New", and enter a bug report

If you want to request a new feature, go to [http://sourceforge.net/tracker/?atid=410166&group_id=34223&func=detail], and click "Submit New".

2.5 Getting Help and Browsing the Documentation

Lush has an on-line documentation system. It can be invoked by typing (`helptool`) at the Lush prompt, which will popup a graphic browser window with the manual. A non-graphic version of the on-line help can be invoked by typing `^Atopic` for help on functions or concepts whose name include `topic` (the `^A` can be typed either as Caret and A or as Control-A, but the latter form is inconvenient when running Lush within Emacs). Here is an example:

```
? ^Aregex
Search Results for: regex
1. Regular Expressions (regex).
2. (regex-match <r> <s>)
3. (regex-extract <r> <s>)
4. (regex-peek <r> <s> [<start>])
5. (regex-subst <r> <s> <str>)
6. (regex-rseek <r> <s> [<n> [<gr>]])
7. (regex-split <r> <s> [<n> [<gr> [<neg>]]])
8. (regex-skip <r> <s> [<n> [<gr> [<neg>]]])
```

```

9. (regex-count <r> <s>)
10. (regex-tail <r> <s> [<n> [<gr> [<neg>]])
11. (regex-member <rl> <s>)
12. (glob <regex> <l>)
choice? 2

```

```
-----
(regex-match <r> <s>)                                [DX]

```

Returns <t> if regular expression <r> exactly matches the entire string <s>. Returns the empty list otherwise.

Example:

```
? (regex_match "(+|-)?[0-9]+(\\|\\. [0-9]*)?" "-56")
= t

```

```
= ()

```

```
?

```

The list of symbols whose name contains a particular string can be listed using `"^Ssymbolstring"`.

```
? ^Sregex
cleanup-regex      regex          regex-count      regex-member
regex-rseek       regex-skip     regex-split      regex-tail
regex_extract     regex_match   regex_seek       regex_subst

= ()

```

2.6 Basic Syntax

The Lush syntax is probably one of the simplest you will find in any programming languages. It is based on Lisp. Don't be scared away by the word "Lisp"! Lush is probably one of the easiest-to-learn language you will ever encounter. The big idea is that every expression is a list enclosed in parentheses which, when "evaluated", returns a result. The first element of the list is the function being called, and the remaining elements are its arguments. So instead of typing `sqrt(4)` to call the square root function like in C, you type `(sqrt 4)`.

Here is an example of how an expression typed at the Lush prompt is evaluated and its result printed:

```
? (sqrt 4)
= 2

```

The list notation is also used for what we generally think of as infix operators, which in Lush (and other Lisps) are just functions like any other. Instead of writing `3+4`, we write `(+ 3 4)`:

```
? (+ 3 4)
= 7
```

Arguments of most functions are evaluated before the function is called, which allows nested expression:

```
? (sqrt (+ 3 4))
= 2.6458
```

Words like `sqrt` or `+` in the previous example are called symbols. Symbols can be assigned a value. The value of the symbol `sqrt` in the above example is a function that returns the square root of its argument. The value of a symbol can be assigned to any lisp object, and therefore is used as the basic mechanism for variables. Assigning a value to a symbol is generally performed with the function `setq` :

```
? (setq x 5)
= 5
? (* x x)
= 25
? (setq x "cou")
= "cou"
? (setq x (concat x x))
= "coucou"
```

When not assigned to a value, symbols contain `nil` , which is printed as (and equivalent to) the empty list " () ". The empty list `nil` is used to represent the value "FALSE" returned by a boolean expression like

```
(> (- 2 4) 1)
= ()
```

Variables can be set to any type of Lush object. In interpreted mode, the types of the variables need not be declared, but they must be declared if the code is to be compiled (more on that later). Variables can be assigned numbers, matrices and vectors, strings, lists, functions, macros, classes, hash table, and several other predefined objects types (such as windows and files), as well as user-defined objects (instances of user-defined classes).

2.6.1 Symbol Names and Special Characters

Symbol names in Lush are case-insensitive (i.e. `SQRT` and `sqrt` are the same variable). Lush makes no difference between underscores and dashes (hyphens) in symbol names. So `idx-sum` is the same as `idx_sum` . Symbol names cannot contain certain characters: parentheses , spaces , forward and backward quotes

, commas , curly braces , square brackets , hash , caret , tilde , and control characters. Those characters play special roles in the Lush reader.

By convention , sub-words of a symbol name are often delimited by hyphens and occasionally by dots:

```
gsl-my-new-pointer  brace.read
```

2.7 Basic Types

Lush has a small number of basic types, a number of not-so-basic built-in types, and a mechanism for creating new types (classes). The basic types include numbers (stored as 8-byte double-precision floating point numbers), generic pointers, and list cells. The not-so-basic types include strings, symbols, seven major types of functions (called DE, DF, DM, DX, DY, DH, and DZ), nine types of vector storage structures (for storing bytes, unsigned bytes, shorts, ints, floats, doubles, generic pointers, and lisp objects), one type of vector/matrix/tensor access structure, two types of file descriptors, hash tables, graphic windows, classes, and dates. New types with slots and methods (a la C++) can also be defined by deriving from the base class `object` .

Some of the types can be entered as literals:

- numbers: `(setq x -3.4e6)`
- strings: `(setq x "coucou")`
- symbols: `(setq x 'coucou)`
- lists: `(setq x '(1 2 3))`

Numbers and strings have the interesting property that they evaluate to themselves, i.e. they return themselves when typed on the line or in an expression. On the other hand, symbols evaluate to their assigned value, and lists evaluate to the result of applying their first element (interpreted as a function) to the rest of the elements (the arguments). That's why the symbol and the list in the above example are preceded by a quote. The quote prevents what follows from being evaluated. Without the quote `coucou` would return its value (or the empty list `()` if it has no value), while `(1 2 3)` would produce an error because `1` is not a function.

2.8 defining functions

Functions can be defined in many ways, but the most common one is to use the `de` construct (similar to the `defun` construct in most other Lisps). The general form is:

```
(de <name> (<arg1>...<argn>) <body1>...<bodym>)
```

Here is a specific example:

```
? (de square (x) (* x x))
= square
? (square 4)
= 16
```

This definition should be interpreted as follows. `de` means "define function", it is followed by the function name " `square` ", a list of formal arguments (`x`), and the body of the function (`* x x`), which means multiply `x` by itself and return the result. The return value of the function is the value returned by the last evaluation in the body `bodym`.

The definition of a particular function can be displayed by typing `^Pfunctionname`:

```
? ^Psquare
(de square (x)
  (* x x) )
= t
```

The directive "`^P`", like "`^A`" introduced earlier is a so-called macro-character. Macro-characters are special combinations of characters that are translated into function calls by the interpreter. Lush has several predefined macro-characters to reduce typing for several commonly used interactive functionalities. A particularly useful one is "`^Lfilename`" which loads a Lisp source file (more on this later).

2.9 Loops, conditionals, local variables, and such

Lush provides traditional loop constructs such as `for`, `while`, `repeat`, and so on. Here is an example of how to compute the sum of the inverses of the first 10 integers (i.e. the 10th term of the harmonic series).

```
? (setq z 0)
= 0
? (setq i 0)
= 0
? (while (< i 10) (incr i) (incr z (/ i)))
= 2.929
```

the function `incr` increments its first argument by the value of its second argument, or by 1 if the second argument is not present, and returns the result. The function `/` called with one argument returns its inverse.

The above example is quite awful because it leaves two global variables `i` and `z` laying around for no reason. A cleaner way of doing things consists in

using temporary variables that are automatically destroyed when they are no longer needed. This is done with the `let*` construct:

```
(let* ((z 0) (i 0))
  (while (< i 10) (incr i) (incr z (/ i))))
```

In its most general form, the `let*` construct is used as follows:

```
(let* ((<var1> <value1>) ... (<varN> <valueN>))
  <body1> ... <bodyP>)
```

It creates temporary local variable `var1 ... varN`, sets them to the values `value1 ... valueN` respectively, evaluates `body1 ... bodyP` and returns the result of the evaluation of `bodyP`. The values of the variables are destroyed upon exiting the `let*`. In fact, if a temporary variable created in a `let*` has the same name as a preexisting variable, the preexisting value will be temporarily "masked" during the life of the temporary variable. The preexisting value will be restored upon exit:

```
(setq z 345)
= 345
? (let* ((z 0) (i 0)) (while (< i 10) (incr i) (incr z (/ i))))
= 2.929
? z
= 345
```

For the interested reader, Lush uses "dynamic variable binding", in contrast with many other languages (particularly compiled ones) that use lexical scoping.

The code above can be encapsulated in a function using the `de` construct:

```
? ;; compute the n-th term of the harmonic series
? (de harmonic (n)
  (let* ((z 0) (i 0)) (while (< i n) (incr i) (incr z (/ i)))))
= harmonic
? (harmonic 10)
= 2.929
? (printf "\\%18.14f\\n" (harmonic 10))
2.92896825396825
= ()
```

In Lush, the text that follows one or several semi-colons on a line is taken as comments and ignored by the interpreter.

Conditional constructs include `if`, `when`, `cond`, and `selectq`. The basic form of the `if` construct is as follows:

```
(if <conditional-expression>
  <evaluate-if-non-nil>
  <evaluates-if-nil-1>
  <evaluates-if-nil-2>
  .....
  <evaluates-if-nil-n>)
```

If the conditional expression returns a non-nil result , the expression `evaluate-if-non-nil` is evaluated and its result returned. If it returns nil , the remaining expressions are evaluated , and the result of the last one is returned. Multiple expressions can be grouped in the `evaluate-if-non-nil` part by grouping them in a `progn` construct: `(progn expr1 ... exprN)` . To illustrate the use of `if` , here is a particularly stupid but didactic implementation of a function that compute the n-th Fibonacci number:

```
(de fibo (n)
  (if (= n 0)                ; test if n equals 0
      0                      ; if yes, return 0
      (if (= n 1)           ; else test if n equals 1
          1                  ; if yes return 1
          (+ (fibo (- n 1)) (fibo (- n 2))))))
```

This is an example of recursive function (a function that calls itself) , a concept dear to the heart of many Lisp aficionados , but whose usage is somewhat discouraged in Lush (the Lush interpreter likes recursive functions , but the Lush compiler doesn't like them at all).

2.10 Compiling Lush functions

To improve the execution speed , Lush functions can be compiled. The Lush compiler has certain features and limitations that conspire to make Compiled Lush (CLush) a distinct language , different from interpreted Lush in several ways. For starters , to avoid inefficient run-time type-checking in compiled code , one must define the type of all the variables before a function can be compiled. The compilable version of `harmonic` is:

```
? ;; define the function
? (de harmonic (n)
  ((-double-) n)
  (let* ((z 0) (i 0))
    ((-double-) z i)
    (while (< i n) (incr i) (incr z (/ i)))))
= harmonic
```

The lines `((-double-) n)` and `((-double-) z i)` declare the type of the argument `n` and the local variables `z`, and `i` as double precision floating point numbers. Before we compile the function, let's measure the number of CPU seconds required to compute the sum of the million-th harmonic sum in interpreted mode (this is on a 800MHz Pentium running Linux):

```
? (time (harmonic 1000000))
= 2.05
```

We can compile the function by simply typing `(dhc-make "junk" harmonic)`. Lush will translate the Lisp code into C and write the C source into the file `"C/junk.c"`. The C compiler will then compile the code and generate the object file `"C/ architecture /file.o"`, where `architecture` is the name of the architecture on which Lush is currently running (e.g. `i686-pc-linux-gnu`). This allows to maintain object files for different architectures in a single source directory tree. Then, Lush's dynamic loader will automatically load the object code back into the interpreter.

```
(dhc-make "junk" harmonic)
Preprocessing and parsing harmonic ...
Generating C for harmonic ...
gcc -DHAVE_CONFIG_H -I/home/yann/lush/include -DNO_DEBUG -Wall -O3 -funroll-loops
= "/tmp/C/i686-pc-linux-gnu/junk.o"
```

Now let's see how many CPU seconds it takes to execute the compiled version of the code. We need to repeat the code a few times to get a meaningful figure:

```
? (time (repeat 50 (harmonic 1000000)))
= 2.37
```

The speedup over the interpreted version is around 45. An important point to remember is that the interpreter always treats numbers as double precision floating point, while numbers in compiled code are treated as specified by their type declaration. Lush understands the following scalar types: `-ubyte-`, `-byte-`, `-short-`, `-int-`, `-float-`, `-double-`, `-bool-`, `-gptr-`. they correspond respectively to the C types `unsigned char`, `signed char`, `short`, `int`, `float`, `double`, `unsigned char`, `void*`.

Just for curiosity, let's look at the C code generated by Lush's Lisp to C compiler in `current-dir/C/junk.c`:

```
/*
 * FUNCTION harmonic
 */
extern_c real
```

```

C_harmonic (real L1_n)
{
  TRACE_PUSH ("C_harmonic");
  {
    real L_Tmp0;
    {
      real L2_1_z;
      real L2_2_i;
      L2_1_z = 0;
      L2_2_i = 0;
      L_Tmp0 = 0;
    /* While loop */
      goto L_1;
    L_0:
      {
        L2_2_i = (L2_2_i + 1);
        L2_1_z = (L2_1_z + (1 / (real) L2_2_i));
        L_Tmp0 = L2_1_z;
      }
    /* While loop test*/
    L_1:
      {
        if ((L2_2_i < L1_n))
          goto L_0;
      }
    }
    TRACE_POP ("C_harmonic");
    return L_Tmp0;
  }
}

```

It looks a bit funny with all these `gotos` , but it's reasonably simple and straightforward C code.

2.11 Writing Lush program files

The normal way to write Lush code is to open a file , say "toto.lsh" , in your favorite editor , and to write all the necessary function definitions and compilation instructions in it. The file can then be loaded into the interpreter with `(load "toto")` . For example , a Lush file for the above function would contain:

```

;; a function that computes the sum of the inverses of the first n integers.
(de harmonic (n)
  ((-double-) n)
  (let* ((z 0) (i 0))

```

```

((-double-) z i)
(while (< i n) (incr i) (incr z (/ i))))

;; compile the function to toto.c and toto.o
;; if the first arg is nil, the .c file has the same
;; base name as the file being loaded.
(dhc-make () harmonic)

```

When loading this file, Lush processes its content as if it were typed at the prompt. Therefore, loading this file will define `harmonic`, and then compile it. If you subsequently reload the same file, the compilation will only occur if the source file `toto.lsh` was modified since the last compilation.

The use of GNU Emacs to edit Lush files is very strongly recommended as it has a special Lisp mode that flashes matching parentheses, highlights keywords and comments, and automatically indents Lisp expressions (using ALT-CTRL-Q). More details on how to setup Emacs are given in the "setting things up" section of this tutorial.

2.12 On-line Documentation

As we said earlier, Lush possesses an on-line documentation system. This system can be used by Lush users to document their own programs. Its use is strongly recommended. Documenting functions or other objects can be done with the "#?" macro-character. Here is what `toto.lsh` should have looked like in the first place:

```

#? (harmonic <n>)
;; computes the sum of the inverses of the first <n> integers
? (de harmonic (n)
  ((-double-) n)
  (let* ((z 0) (i 0))
    ((-double-) z i)
    (while (< i n) (incr i) (incr z (/ i))))

;; compile the function to toto.c and toto.o
(dhc-make () harmonic)

```

The string after the "#?" is what will be searched when search feature of `(helptool)` is invoked on this file. Contiguous comment lines immediately following the "#?" line constitute the text of the help. After this file is loaded, the on-line help for `harmonic` and all the functions in the same file can be displayed with:

```
(helptool "toto.lsh")
```

2.13 Multiple file programs and search path

Sometimes, a Lisp program requires other Lisp files to be loaded in order to run. For example imagine that we have written a set of signal processing functions in the file `"/home/yann/lsh/dsp.lsh"`, and now we are writing a program `"tutu.lsh"` that uses some of those functions. The beginning of `"tutu.lsh"` would look like the following

```
#? *** My Cool Signal Processing Application
;; this is my really cool dsp application.
(libload "/home/yann/lsh/dsp")
.....
```

(the `.lsh` extension is automatically added if required).

The function `libload` has an interesting features: it remembers which files were already libloaded and loads them only once into the interpreter (subsequent calls to `libload` with a previously libloaded file will have no effect). There is an exception to this behavior: `libload` keeps track of the tree of dependencies between Lush program files, and also keeps track of the times of last modification of each file. When a file is loaded with `libload`, all the files on which this file depends are also loaded, but only if they have been modified since last time they were libloaded. Here is a possible scenario: file A libloads files B and C. File A is libloaded, as a result file B and C are libloaded. File B is modified. File A is libloaded again, as a result B is re-libloaded, and A is re-libloaded. C is not re-libloaded since it was previously libloaded and was not modified. This behavior is recursive (it applies to all files directly or indirectly libloaded by A).

In certain cases, the lower-level function `load` may be preferable. `load` simply loads a file:

```
? (load "/home/yann/lsh/dsp")
```

Putting absolute path names in a program is generally a bad idea since programs have a pernicious tendency to move from time to time. The solution is to rely on Lush's search path mechanism. When any loading directive is executed, Lush looks for the file in the current directory, and then in the search path. The default search path includes `lush/local`, `lush/packages`, `lush/lsh`, `lush/sys` and several subdirectories of `lush/lsh`. The search path can be user-extended with the `addpath` function:

```
(addpath "/home/yann/lsh")
```

now, loading `"dsp.lsh"` only requires:

```
(libload "dsp")
```

2.14 Calling C functions from Lush

The ability to freely mix Lisp and C code is one of the most interesting features of Lush.

Let's say you have written a C file called `titi.c` with the following content:

```
float sq(float x)
{
    return x*x;
}
```

You have compiled the file and produced the object file `titi.o`. Calling `sq` from the Lush interpreter is as simple as the following. First, dynamically load the object file into Lush

```
? (mod-load "titi.o")
```

then, write a lisp function whose only purpose is to call `sq`.

```
? (de square (x) ((-float-) x)
    (cphheader "extern float sq(float);")
    (to-float #{ sq( $x ) #} ))
```

Here is what the above means: `(de square (x) ...)` means define a new Lisp function called `square` with a single argument `x`. `((-float-) x)` simply declares `x` as float. `(to-float ...)` converts its argument to float (like a cast in C). The sequence `# ... #` allows to insert C code within Lisp code. Therefore `# sq($x) #` simply calls the C function `sq` and returns the result. Lisp variables can be inserted in in-line C code by prepending a Dollar sign, thus `$x` refers to the Lisp float variable `x`. The `cphheader` directive allows one to include a string in the "header" section of the C files generated by the Lisp to C compiler. We use it here to specify the type of the argument and the return value of `sq`.

We can now compile the above function using:

```
? (dhc-make () square)
```

Two file `C/square.c` and `C/i686-pc-linux-gnu/square.o` will be generated with the C source and the object code for `square`. The object code will be automatically loaded into Lush. Now `square` can be called from the Lisp interpreter:

```
? (square 5)
= 25
```

In the above example , the `to-float` casting function was used to tell the compiler what type the returned value has. The following "cast" functions are available: `to-float`, `to-int`, `to-double`, `to-bool`, `to-gptr` . These functions cast their argument to the corresponding type.

2.15 Mixing Lisp and inline C code

If we do not have a ready-made C file for our function , we can use the inline C capability to write code in C right inside the Lisp function. Here is how we could write the familiar `harmonic` function this way:

```
(de harmonic (n)
  ((-double-) n)
  (let ((r 0))
    ((-double-) r)
    #{ { double i=0;
        while (i<$n){ i+=1.0; $r+=1/i; };
      } #} r))
(dhc-make () harmonic)
Preprocessing and parsing harmonic ...
Generating C for harmonic ...
gcc -DHAVE_CONFIG_H -I/home/yann/lush/include -DNO_DEBUG -Wall -O3 -funroll-loops -mcpu=i686
= "/tmp/C/i686-pc-linux-gnu/junk.o"
? (time (repeat 50 (harmonic 1000000)))
= 2.36
```

That's about the same time as the version written in compiled Lisp. In this case , writing directly in inline C was not particularly advantageous , but there are cases where writing in Lisp is inefficient , awkward , or even outright impractical. In these cases one can resort to inline C implementations.

2.16 Lists

Lush belongs to the Lisp family , and therefore has lots of functions to manipulate lists. In fact , Lush programs themselves are lists. Lists are sequences of Lush objects enclosed in parentheses. Things like `(+ 1 2)` is a list. However , when the list `(+ 1 2)` is entered at the Lush prompt , it is immediately evaluated and discarded right away:

```
? (+ 1 2)
= 3
```

To create a list , we must prevent the evaluation from happening. This is done with the `quote` function or the `'''` (quote) macro-character:

```
(quote (+ 1 2))
= (+ 1 2)
? '(+ 1 2)
= (+ 1 2)
```

Lists can be put into variables like any other Lisp object:

```
(setq x '(+ 1 2))
```

List elements can be any lisp object, including other lists. Here is a list with a number, a string, a list of numbers, a symbol, and a list, and a number:

```
(setq x '(1 "coucou" (4 5 6) asd (7 (8 9) 10) 20))
```

Here is another example where the same list as above is assembled from its members using the `list` function:

```
(setq x (list 1 "coucou" '(4 5 6) 'asd '(7 (8 9) 10) 20))
```

The first element (the head) of a list can be obtained with the `car` function, and the rest of the list (tail) with the `cdr` function (pronounced kudder). The `cons` function can be used to put a head and a tail together:

```
? (setq x '(+ 1 2))
= (+ 1 2)
? (car x)
= +
? (cdr x)
= (1 2)
? (setq z (cons '+ '(1 2)))
= (+ 1 2)
```

A list can be evaluated:

```
(setq x '(+ 1 2))
= (+ 1 2)
? (eval x)
= 3
```

Operations can be performed on all the elements of a list using the `all` construct:

```
? (setq z '(1 3 5 7 9))
= (1 3 5 7 9)
? (all ((x z)) (- x 1))
= (0 2 4 6 8)
```

This says , make a temporary variable `x` , and evaluate the body of the `all` with `x` set to each element of `z` in sequence. Then return a list of all the results.

This section only scratched the surface of what can be done with lists. The reader is referred to the List section in the "Core Interpreter" chapter of the Lush manual.

2.17 Strings and Regular Expressions

[under construction]

2.18 Scalars, Vectors, Matrices, Tensors, and Storages

Lush's piece de resistance is its matrix engine. Lush can operate on scalars , vectors , matrices , or high-dimensional tensors from 0 to 8 dimensions. creating a tensor of `float` is done simply with:

```
? (setq m (matrix 10 8 4)) ; create 3D matrix
= ::INDEX3:<10x8x4>
? (m 3 4 2 45.6) ; set value of element (3,4,2) to 45.6
= ::INDEX3:<10x8x4>
? (m 3 4 2) ; get value of element (3,4,2).
= 45.6
```

Tensors of various basic types can be created with the functions listed below. Each function has two versions , the regular version initializes all the elements to zero , while the version with `-nc` at the end do not (no clear) and are a bit faster. All of these functions take 0 to 8 integer arguments that are the sizes in each dimension:

- `double-matrix`, `double-matrix-nc`, `real-matrix`, `real-matrix-nc` : doubles
- `float-matrix`, `float-matrix-nc`, `flt-matrix`, `flt-matrix-nc` : floats
- `int-matrix`, `int-matrix-nc` : 32-bit ints
- `short-matrix`, `short-matrix-nc` : 16-bit shorts
- `byte-matrix`, `byte-matrix-nc` : 8-bit bytes
- `ubyte-matrix`, `ubyte-matrix-nc` : 8-bit unsigned bytes
- `gptr-matrix`, `gptr-matrix-nc` : generic pointers (`void*`)

A matrix with 0 dimension is called a scalar. It is different from a regular number in that it behaves like a pointer to a number. (the value of a scalar passed as argument to a function can be modified by that function).

Matrices of doubles can be entered literally using square brackets:

```
(setq m [[0 2 3 4][5 6 7 8]])
```

Matrices of other types can be specified by adjoining an appropriate character after the open bracket. Here is how to create a vector of floats:

```
(setq m [f 0 2 3 4])
```

Characters `f` , `i` , `s` , `b` , `u` , `a` specify float , int , short , byte , ubyte , and atom respectively. Here is an example with atoms:

```
(setq m [a "choucroute" "garnie"])
```

Matrices are really composed of two separate entities:

- the `storage` or `srg` which contains the following fields: a pointer to the actual data , an element type identifier (and the size thereof) , and flags that indicate if the data is writable or not , if it is in RAM or memory-mapped from a file.
- the `index` or `idx` points to a `srg` and contains the following fields: the number of dimensions of the tensor , the size in each dimension , the address increment from one element to the next in each dimension (called `modulo`) , and the offset of the first tensor element in the `storage` . This structure allows multiple `idx` to point to the same `srg` thereby allowing the data to be accessed in multiple ways without copying.

2.18.1 Working with Storages

`srg` and `idx` can be created independently of each other. So for example `(setq s (double-storage 8))` creates a new storage of doubles with 8 elements. The call `(double-storage)` returns an unsized storage. Creating an `idx` on a particular storage `s` can be done with `(new-index s '(3 4))` . This resizes the storage `s` to 12 elements (if it has less than that) , and creates an `idx` of size (3 , 4) on this storage. Creating other `idxs` on the same storage allows to access the same piece of data in multiple ways.

When directly accessing storage elements , storages behave like 1D `idx` (vectors). Storages can be created , allocated in memory , or memory-mapped to a file. Storage creation function are provided for the following types: doubles , floats , ints , shorts , bytes , unsigned bytes , generic pointers , and lisp object pointers.

2.18.2 Manipulating Tensors and IDX

Several functions are provided to create and manipulate the "idx" structure and its various fields.

- `(new-index s dimlist)` creates a new idx on storage `s`. The dimensions of the idx are given in `dimlist`. The storage is resized if needed.
- `(idx-clone m)` make a clone of idx `m` (a copy of `m` that points to the same storage).
- `(copy-matrix m [r])` returns a fresh copy of tensor `m`. if `r` is present, `m` is copied into `r` with appropriate type conversion if necessary.
- `(idx-ndim m)` return number of dimensions.
- `(idx-dim m n)` return size of n-th dimension.
- `(idx-modulo m n)` return number of elements that separate two successive elements in the n-th dimension.
- `(idx-offset m)` return offset of first element of `m` in storage.
- `(idx-size m)` return size of memory are used by idx data in bytes.
- `(idx-storage m)` return storage pointed to by `m`.
- `(contiguousp m)` returns true if the elements of `m` are contiguous in memory.

2.18.3 tensor iterators

Lush provides efficient iterators to cycle over dimensions of a tensor. It is almost always better to use those iterators, rather than a regular loop with regular access to array elements because the latter method causes an array bound checking at each access, while the former doesn't. The main idx iterator is `idx-bloop` whose synopsis is as follows:

```
(idx-bloop ((s1 idx1) [(s2 idx2) [...(sn idxn) ]]) body)
```

Each `si` will be an idx with one less dimension than the corresponding `idxi`, and will simultaneously loop over the successive "slices" of `idxi` for each possible value of the first index. In other words, applying function `myfunc` to each element of a vector `v1` and putting the result in the corresponding element in `v2` can be done with:

```
(idx-bloop ((x1 v1) (x2 v2)) (x2 (myfunc (x1))))
```

`x1` and `x2` are scalars (i.e. zero-dimensional tensors). The above function work just as well if `v1` and `v2` are `n`-dimensional tensors and `myfunc` accepts `n-1`-th dimensional tensors as arguments. It should not be assumed that the multiple evaluations of the body are executed in sequence. They may be executed in parallel on some implementations of Lush. A similar function `idx-eloop` iterates on the last dimension of a tensor instead of the first dimension. For example, the matrix product operation $C = A \cdot B$ can be written as follows:

```
(de idx-m2timesm2 (A B C)
  (idx-eloop ((Bj B)(Cj C)) (idx-m2dotm1 A Bj Cj)))
```

where `idx-m2dotm1` is the usual matrix-vector product. The `idx-eloop` construct simultaneously iterates over all columns of `B` and `C` (naturally, the function `idx-m2timesm2` is predefined).

Another `idx` iterator allows to easily write the inner loop of a tensor operation in `C`, while leaving all the bookkeeping to the Lisp:

```
(cidx-bloop (i_1 [i_2...i_n] (c_1 l_1) [(c_1 l_1)...(c_m l_m)] p_1 [p_2...])
```

This is somewhat equivalent to `n` nested `idx-bloop`s, which will simultaneously loop over the first `n` dimensions of `idxs l_1` to `l_m`. The arguments `i_1` to `i_n` are strings containing names of `C` local variables that will be set to the loop index in each of the `n` dimensions. At each iteration, the `C` variables provided in strings `c_1` to `c_m` will point to the appropriate values in the `idxs l_1` to `l_m`. For example, the following function will fill matrix `a` with $\cos(i+j)$.

```
(de foo (a)
  ((-idx2- (-flt-)) a)
  (cidx-bloop ("i" "j" ("a" a)) #{ *a = cos(i+j); #} a)
```

The return value is (like in `idx-bloop`) the last `idx` specified in the declaration (in the example above, the return value is superfluous).

2.18.4 `idx` manipulations

Several functions are provided to conveniently manipulate the `idx` structure itself. These functions do not manipulate the actual data (and therefore are efficient), they merely derive new `idx` from existing `idx` to access the data in new ways.

`(narrow m n s [o])` returns a clone of `idx m` where the size in the `n`-th dimension is reduced to `s` elements, offset by `o`. This can be used to select bands or rectangular areas in matrices.

2.18. SCALARS, VECTORS, MATRICES, TENSORS, AND STORAGE_S47

```
? (setq m [[0 2 3 4][5 6 7 8]])
= [[ 0.00  2.00  3.00  4.00 ]
   [ 5.00  6.00  7.00  8.00 ]]
? (narrow m 1 2 1)
= [[ 2.00  3.00 ]
   [ 6.00  7.00 ]]
```

(`select m n s`) return a clone of `m`, with the `n`-th dimension removed, and which is the `s`-th "slice" of `m`, in the `n`-th dimension. For example, the 3rd column of matrix `m` in the above example (which is a vector) can be obtained with

```
? (select m 1 2)
= [ 3.00  7.00 ]
```

(`unfold m n ksize step`) returns an `idx` with an added dimension at the end of `m` obtained by "unfolding" `n`-th dimension. More explicitly, successive rows in the last dimension will contain successive, possibly overlapping, subsets of size `ksize` of consecutive elements in the `n`-th dimension. The separation between successive subsets is `step`. Here is an example:

```
? (setq m [0 1 2 3 4 5 6 7])
= [ 0.00  1.00  2.00  3.00  4.00  5.00  6.00  7.00 ]
? (unfold m 0 3 1)
= [[ 0.00  1.00  2.00 ]
   [ 1.00  2.00  3.00 ]
   [ 2.00  3.00  4.00 ]
   [ 3.00  4.00  5.00 ]
   [ 4.00  5.00  6.00 ]
   [ 5.00  6.00  7.00 ]]
```

No data is copied or replicated in the process, i.e., the three occurrences of the number 2 in the above example actually come from a single memory location. The size of the new dimension is `ksize`. If `dimn` is the size of the `n`-th dimension of `m`, the size of the returned `idx` in the `n`-th dimension is $(\text{dimn} - \text{ksize}) / \text{step} + 1$. The values of `dimn`, `ksize`, and `step` must be such that the new size is an integer.

The `unfold` function essentially allows to reduce a discrete convolution to a matrix-vector product. Here is an example:

```
? (setq m [1 1 0 2 3 4 2 0])
= [ 1.00  1.00  0.00  2.00  3.00  4.00  2.00  0.00 ]
? (setq kernel [-1 2 -1])
= [-1.00  2.00 -1.00 ]
? (setq output (matrix 6))
```

```

= [ 0.00 0.00 0.00 0.00 0.00 0.00 ]
? (idx-m2dotm1 (unfold m 0 3 1) kernel output)
= [ 1.00 -3.00 1.00 0.00 3.00 0.00 ]

```

A subsampled convolution can be implemented by unfolding with a `step` larger than 1:

```

? (setq m [0 1 2 3 4 5 6 7 8])
= [ 0.00 1.00 2.00 3.00 4.00 5.00 6.00 7.00 8.00 ]
? (setq z (unfold m 0 3 2))
= [[ 0.00 1.00 2.00 ]
   [ 2.00 3.00 4.00 ]
   [ 4.00 5.00 6.00 ]
   [ 6.00 7.00 8.00 ]]
? (setq kernel [1 2 1])
= [ 1.00 2.00 1.00 ]
? (idx-m2dotm1 z kernel output)
= [ 4.00 12.00 20.00 28.00 ]

```

Naturally, there is no real need for most programmers to use the `unfold` construct because the standard lush library contains efficient 1D and 2D convolutions.

`(idx-transclone m dimlist)` returns a clone of `idx m` where the dimensions have been permuted according to the list of dimension indices `dimlist`. For example `(idx-transclone m '(0 2 1))` permutes the second and third dimensions (again, no data is moved or copied by this operations).

A few functions are provided to physically hack an `idx` to change a dimension, a modulo, or the offset (repectiveley `(idx-changedim m n v)`, `(idx-changemod m n v)`, `(idx-changeoffset m n)`).

2.18.5 simple tensor operations

Lush has lots of functions that operate on tensors of all types. These include component-wise unary operations (applying a function to all the elements of a tensor), such as `abs`, `atan`, `clear`, `copy`, `cos`, `sin`, `exp`, `log`, `minus`, `sqrt`, They all take the form

```
(idx-XXX m [r])
```

where `XXX` is one of the above function, `m` is the input matrix (from 0 to 8 dimensions) and `r` is an optional result matrix which must have the same dimensions as `m`. If `r` is not present, the result is created and returned.

Also available are component-wise dyadic operations such as `add`, `sub`, `mul`, `div`. They take the form

```
(idx-XXX m1 m2 [r])
```

Next come dyadic operations where the second argument is a scalar `s`, such as `(idx-addm0 m s [r])` which adds scalar `s` to each element of tensor `m`, or `(idx-addm0acc m s r)` which does the same thing but accumulates the result in `r`, or `(idx-dotm0 m1 s [r])` which multiplies each element of `m` by scalar `s`.

2.18.6 operations between idx and scalars

Multiplying all the elements of an `idx` by a scalar can be done with:

```
(idx-dotm0 m s [r])>
```

where `m` is an `idx` of any type and any number of dimensions, `s` is a scalar (i.e. an `idx` with zero dimension), and `r` is an optional `idx` in which the result will be written. Turning a number into a scalar can be performed with `((double-matrix) number)`. So for example, multiplying all the elements of a matrix by 10 is done as follows:

```
? (idx-dotm0 [1 2 3] ((double-matrix) 10))
= [10.00 20.00 30.00 ]
```

Similarly, function `idx-addm0` adds a scalar to all the elements of an `idx`.

2.18.7 contracting operations with scalar result

A number of functions are provided that take one or two input tensors and write the result in a scalar. Such functions include dot products, distances, sums of terms. They "contract" all the dimensions, i.e. the generalized dot product of two tensors is the sum of all the products of corresponding terms in the two tensors:

```
(idx-dot m1 m2 [r])>
```

Other such functions include `idx-inf` (smallest element), `idx-sqrdist` (sum of square difference between two tensor), `idx-sum` (sum of all terms), `idx-sumsqr` (sum of squares), `idx-sup` (largest element)

2.18.8 special inner and outer products

There are several functions for performing matrix-vector products, outer products, and such. They include: `(idx-m1extm1 m1 m2 [r])` (outer product of vectors), `(idx-m2dotm1 m v [r])` (matrix-vector multiply), `(idx-m2extm2 m1 m2 [r])` (outer product of matrices which gives a 4D tensor), `(idx-m4dotm2 m v [r])` (product of a 4D-tensor by a matrix where the last 2 dimension of the tensor are contract with the first two dimensions of the matrix). `idx-m4dotm2`

can be conveniently combined with appropriate calls to `unfold` to perform 2D convolutions, as demonstrated in the library file `lush/lsh/libidx/idx-convol.lsh`.

2.18.9 Simple Linear Algebra

Lush provides a full interface to the Gnu Scientific Library as well as to BLAS and LAPACK, which contain all the linear algebra anyone could wish for. Using those functions directly may be a bit complicated, so lush provides a library with simple forms of the most common functions. This library can be loaded with `(libload "libnum/linalg")`. It contains the following functions:

- `(eigen-symm m)` : eigenvalues of real symmetric matrix
- `(eigen-symmv m)` : eigenvalues and eigenvectors of real symmetric matrix
- `(eigen-herm m)` : eigenvalues of complex hermitian matrix
- `(eigen-hermv m)` : eigenvalues and eigenvectors of complex hermitian matrix
- `(svd a v s)` : singular value decomposition of real matrix
- `(solve-hh a b)` : linear system solver (householder method)
- `(solve-lu a b)` : linear system solver (LU decomposition method)
- `(solve-lu-complex a b)` : complex linear system solver (LU decomposition method)
- `(solve-sv a b)` : linear system solver/least square solution (SVD method)
- `(inverse-lu a i)` : matrix inverse (LU decomposition method)
- `(determinant a)` : determinant
- `(determinant-log a)` : log determinant
- `(determinant-sign a)` : sign of determinant

2.19 Complex Numbers

Complex numbers in Lush are represented as (and totally equivalent to) vectors of doubles with 2 elements. Lush's complex numbers library is not automatically loaded at startup and should be loaded with:

```
(libload "libnum/libcomplex")
```

Creating a complex number is done just like creating a vector of 2 doubles, or using the macro `new-complex`, which is equivalent to `(double-matrix 2)`:

```
(setq z [3 4])           ; complex literal [real-part imaginary-part]
(setq z (new-complex))   ; allocation of complex [0 0]
```

Complex vectors and matrices can be created with `complex-matrix` :

```
(setq z (complex-matrix 3 3))
```

which is exactly equivalent to:

```
(setq z (double-matrix 3 3 2))
```

Several complex arithmetics and functions are provided. The function names all begin with the letter `c` . Complex dyadic functions include `c+`, `c-`, `c*`, `c/`, `c**`, `clogb` (addition , subtraction , multiplication , division , power , complex log with complex base). All commonunary operations on complex are also provided , including conjugate , inverse , square root , exponential , log , and trigonometric functions. More details can be obtained in the packages/libnum section of the Lush manual (["libnum: Numerical Library"]).

The meory layout of complex vectors and matrices so defined are compatible with the LAPACK and GSL functions that take double complex vectors and matrices as arguments.

2.20 Object-Oriented Lush Programming

Lush is an object-oriented language that allows users to define hierarchies of classes with inheritance , slots , and methods by deriving from the root class `object` .

Naturally , Lush has a number of built-in immutable classes such as numbers , list cells , strings , symbols , etc , to which users can add methods (but not slots). This section will concentrate on user-defined classes.

2.20.1 defining a new class

User defined classes are derived from the root class `object` with constructs like this:

```
(defclass rectangle object width leng)
```

This can be interpreted as: define the class `rectangle` as a subclass of `object` with two slots `width` and `leng` . Subclasses inherit the slots of their superclass. (which is a moot point here since `object` has no slot).

A new instance of the class `rectangle` can be created with the function `new` :

```
(setq r (new rectangle))
```

By default , the slots are left unfilled (set to the empty list). Accessing slots of an object can be done with the macro-character " : "

```
? (setq :r:width 10)
= 10
? (setq :r:leng 3)
= 3
? ^Pr
;;*** ::rectangle:819fa08, INSTANCE OF ::class:rectangle
;; FROM CLASS: rectangle
;;      width=10
;;      leng=3
;; FROM CLASS: object
```

Classes can be given methods (or virtual member functions in the C++ semantics) defined as in the following example:

```
(defmethod rectangle area () (* leng width))
```

This defines the method `area` for the class `rectangle` which takes no argument , and returns the product of the slots `leng` and `width` . This is essentially similar to defining a function with `de` , except that within the body of a method , the slots of the object are directly accessible by name as if they were local variables. Calling the above method can be done as follows:

```
? (==> r area)
= 30
```

which can be read as "send to `r` the message `area` ".

Methods can be passed arguments , and can refer to the current object through the special variable `this` . Here is an example of how to compute the volume of a box whose base is the current object and height is passed as argument:

```
? (defmethod rectangle volume (height) (* height (==> this area)))
= volume
? (==> r volume 10)
= 300
```

2.20.2 inheritance

Lush allows to create subclasses of existing classes. Subclasses inherit the slots and methods of their parent class. For example , we can define a `box` class ,

derived from the `rectangle` class , and containing a `height` slot in addition the inherited `width` and `leng` :

```
? (defclass box rectangle height)
= box
```

Setting the slots of an object directly as we did in the `rectangle` example is considered bad practice. A better alternative is to fill the slots when the object is created. This can be done by defining a constructor method. The constructor is a method that has the same name as the class. It is called whenever the `new` function is called to create an object of that class. Here is a constructor for `box` :

```
? (defmethod box box (l w h) (setq leng l) (setq width w) (setq height h))
= box
? (setq monolith (new box 9 4 1))
= ::box:8138958
? ^Pmonolith
;;*** ::box:8138958, INSTANCE OF ::class:box
;; FROM CLASS: box
;;   height=1
;; FROM CLASS: rectangle
;;   width=4
;;   leng=9
;; FROM CLASS: object
= ()
? (defmethod box volume () (* width leng height))
= volume
? (==> monolith volume)
= 36
```

An alternative way to define the `volume` method would have been to multiply the `height` by whatever result the `area` method inherited from `rectangle` returns:

```
(defmethod box volume () (* height (==> this area)))
```

Class hierarchies are nice , but beginner should be careful not to create deep hierarchy trees with more than 3 or 4 levels of inheritance. Experience proves that deeper trees result in hard-to-maintain code.

Unlike C++ , Lush neither support multiple inheritance , nor private slots and methods.

2.21 Input and Output

[under construction]

2.22 A quick tutorial on Lush Graphics

Lush offers several ways to output graphics. The most common one is the Lush graphic driver , which provides a small set of graphic primitive to display data in a window or write it to a PostScript file. Lush also provides interfaces to the OpenGL API and to the SDL API , but these will not be described in this section.

Opening a window is simply done with:

```
(new-window [<x> <y> [<w> <h> [<name>]]])
```

where the arguments are respectively the position , size and name of the new window.

Drawing in a window can be done with a collection of simple primitives:

- drawing a line: `(draw-line 10 10 100 200)`
- setting the drawing color: `(color-rgb 0.4 0.1 0.8)`
- drawing a rectangle: `(draw-rect 10 10 50 80)`
- drawing a filled rectangle: `(fill-rect 10 10 50 80)`
- clearing the screen: `(cls)`
- displaying an image: `(rgb-draw-matrix 10 10 im)` where `im` must be an `NxM` , `NxMx1` , `NxMx3` , or `NxMx4` matrix that will be interpreted as an `NxM` image with greyscale or RGB pixels (whose individual values must be between 0 and 255).
- changing the font and drawing a text:

```
(font "Helvetica-Bold-20") (gprintf 100 100 "salut les potes")
```

Double-buffered and flicker-free redrawing can be done with the `graphics-batch` function. Used in the form `(graphics-batch exp1 exp2 ... exp3)` , it will evaluate all its argument , but any graphic operation performed in that time will be performed into an off-screen graphic buffer. This graphic buffer will be slapped on the window as the call to `graphics-batch` terminates. This allows smooth double-buffered animation:

```
(while not-finished
  (compute-stuff)
  (graphics-batch
    (draw-stuff-1)
    (draw-stuff-2)
    (draw-stuff-3)))
```

Here is an example:

```
(libload "libimage/image-io")
(setq m
  (image-read-rgb
    (concat-fname lushdir "packages/video4linux/demos/jabba-320x240.jpg")))
(new-window)
(for (i 0 99) (graphics-batch (cls) (rgb-draw-matrix i i m)))
```

2.23 Writing Executable "Shell Scripts" with Lush

Lush can be used to write executable scripts that can be invoked from the Unix prompt just like shell scripts , Perl scripts , or executable programs.

Rather than a long explanation , here is an example of a Lush script file:

```
#!/usr/bin/lush @@
(printf "Capitalizing the arguments:\n")
(each ((arg (cdr argv))) (printf "\\%s \\%s\n" arg (upcase arg)))
```

The @@ at the end of the first line tells Lush to run in script mode. In that mode , Lush does not display the usual startup banner. It load `sys/sysenv.lsh` . It sets the variable `argv` to a list of strings , where the strings are the command line arguments (including the command itself). Finally it reads and evaluates the content of the script file. It is important to note that `lsh/lushenv.lsh` is not loaded automaticall in script mode. To have the same environment as in regular Lush , the first line of a script should probably be `(libload (concat-fname lushdir "lsh/lushenv"))` .

The library `libc/shell.lsh` contains several functions to facilitate the writing of shell scripts. Among other things , it contains Lisp versions of common shell functions such as `ls` , `cp` , `mv` , etc.

Here is another script example: a photo slideshow program in 12 lines of Lush. Put this in a file (say `lush-slideshow`) , make it executable with `chmod a+x lush-slideshow` . Then , at the shell prompt , type: `lush-slideshow photo1.jpg photo2.jpg photo3.jpg` :

```
#!/usr/bin/lush @@
(libload (concat-fname lushdir "lsh/lushenv"))
(libload "libimage/image-io")
(libload "libimage/rgbaimage")
(setq window (x11-window 1 1 640 480 "Slide Show"))
(while t
  (each ((f (cdr argv)))
```

```
(rgb-draw-matrix
 0 0 (rgbaim-resize (image-read-rgba f) (xsize) (ysize) 0))
(sleep 4))
```

This assumes that your lush executable is in `/usr/bin/lush`, which should be the case if Lush was installed with "make install".

2.24 Writing Standalone Binary Executables with Lush

[under construction]

Using the function `dump`, it is possible to create a binary image of the lisp structures contained in a Lush process. Such binary dump files are useful for distributing Lush programs to users who may not have installed Lush. Details are given in a section of this manual entitled Lush Runtime. [Lush Runtime]

2.25 The Ogre GUI Toolkit

Lush contains a fairly complete Graphical User Interface toolkit built on top of the standard graphic library. Ogre allows very rapid development of GUI applications. An introduction and a tutorial are available in the Standard Library/Ogre section of this manual.

See: Introduction to Ogre.

See: Ogre Tutorial.

See: The Ogre Class Library.

Here we will just show a couple of very short examples of how simple GUIs can be created in just a few lines. First, here is the famous one-liner button "call a function when clicked":

```
(ogre)
(de coucou (c) (printf "coucou\n"))
(new autowindowobject 10 10 () () "ouch" (new stdbutton "click me" coucou))
```

The `autowindowobject` is an ogre object that opens a new window and adjustst itself to fit the size of its content. In this case, the content is a single standard button (a.k.a. `stdbutton`). The `stdbutton` constructor takes a string and a function with one argument which is called each time the button is pressed. Now, here is a slightly more sophisticated example:

```
(ogre)
(libload "libimage/image-io")
(libload "libimage/rgbaimage")
;; load image
(defvar image
```

2.26. USING THE SIMPLE DIRECTMEDIA LAYER INTERFACE (SDL)57

```
(image-read-rgba
  (concat-fname lushdir "packages/video4linux/demos/jabba-320x240.jpg")))
;; make an imageviewer object
(setq image-viewer (new imageviewer 240 220 image t))
;; define a few functions to tell the imageviewer to display rotated images
(de rot90-right (c) (==> image-viewer setttext (rgbaim-rot90-right image)))
(de rot00 (c) (==> image-viewer setttext image))
(de rot90-left (c) (==> image-viewer setttext (rgbaim-rot90-left image)))
;; create the GUI window
(new autowindowobject 10 10 () () "ouaf"
  ;; insert a column
  (new column
    ;; make a row of buttons
    (setq button-row
      (new row
        (new stdbutton "Left" rot90-left)          ; rotate left when clicked
        (new stdbutton "Straight" rot00)          ; straighten up when clicked
        (new stdbutton "Right" rot90-right)))      ; rotate right when clicked
    ;; insert imageviewer
    image-viewer))
```

2.26 Using the Simple Directmedia Layer interface (SDL)

Lush includes an interface to the well-known SDL library (Simple Directmedia Layer) , which has platform-independent functions for 2D graphics , timers , low-level keyboard , mouse , and joystick input , sound playing , direct read from CDROM , and many others features.

Lush provides a high-level interface to some of those functions in `sdl/libSDL` . It includes higher-level classes , such as screens and sprites , with pixel-accurate collision detection. This allows rapid development of graphic-intensive interactive applications such as real-time 2D games.

A tutorial on writing simple games with SDL is available in the Tutorial section of the Lush manual. A full documentation of the SDL interface is available at [”SDL: Simple DirectMedia Layer”]

The present section is a quick introduction to the `libSDL` library. Before you start , make sure SDL and all its development libraries and header files are installed on your system. If you are on Linux , those things are available as packages with your distro.

Loading and initializing the `libSDL` library is performed with:

```
(libload "sdl/libSDL")
(sdl-initialize)
```

Next , we can open an SDL screen. `libSDL` provides three kinds of SDL screen classes.

- The first kind , `sdl-screen` is preferred if you intend to draw on the screen using `libSDL` 's sprite class , or using SDL's low-level drawing functions.
- The second kind , `sdlgl-screen` is preferred if you intend draw 3D graphics using OpenGL.
- The third kind , `sdlidx-screen` is preferred if you want to draw your images into a regular Lush matrix , and quickly update an SDL screen with the content of that matrix.

2.26.1 `sdlidx-screen` class: Mapping the content of an IDX to an SDL screen

Let's start by describing the `sdlidx-screen` class. Opening an `sdlidx-screen` is done by first allocating a `ubyte-matrix` of size `HxWx4` , and then creating an `sdlidx-screen` with this matrix as argument. This will open an SDL screen of width `W` and height `H`:

```
(setq depth 3)
(setq m (ubyte-matrix 480 640 depth))
(setq screen (new sdlidx-screen m "my title"))
```

If the `depth` parameter is 3 , each element of matrix `m` is a component of an RGB pixel. So for example `(m y x 0)` is the R component of pixel `(x , y)` (from the top left of the screen). If the `depth` parameter is 4 , each element of `m` is a component of an RGBA pixel. The A component is the Alpha channel (transparency). An alpha value of 255 is opaque , while 128 is half transparent. The matrix `m` can be seen as the "back buffer" of a double-buffered screen. Changing values in `m` will not immediately change the corresponding pixel on the screen. Once the desired RGBA image has been written in `m` , slapping it on the screen can be done with `(=> screen flip)` . If the screen was opened with `depth =4` , the image displayed will be an alpha-blended mixture of the previously displayed image and the content of `m` . Here is a complete example:

```
(libload "sdl/libSDL")
(libload "libimage/image-io")
(libload "libimage/rgbaimage")
(sdl-initialize)
(setq m (ubyte-matrix 480 640 3))
(setq screen (new sdlidx-screen m "my picture"))
(setq photo (rgbaim-resize (image-read-rgba "mypicture.jpg") 640 480 0))
(idx-copy (narrow photo 2 3 0) m)
(=> screen flip)
```

Subsequent calls to `(=> screen flip)` will refresh the screen with the content of `m`. Note that SDL screens (unlike regular Lush graphic windows) are not automatically refreshed when uncovered from behind other windows. That is why SDL screens should be used only for animated graphics and games. Refresh rates on SDL screens are likely to be faster than with regular Lush graphic windows.

2.26.2 sdl-screen class: doubled buffered animations

The `sdl-screen` class is intended for double-buffered animations using sprites (movable screen objects). This situation is primarily encountered when writing video games and other interactive multimedia applications.

Opening an `sdl-screen` is simply done with `(setq screen (new sdl-screen 800 600 "the title"))`. Drawing into an `sdl-screen` is made very simple by the `sdl-sprite` class (as explained below), but low-level SDL drawing can be performed by creating `SDL_surface` objects and calling `SDL_BlitSurface` on the `p` slot of the `sdl-screen`.

2.26.3 sdl-sprite: movable animated objects

Interactive multimedia applications often involve drawing movable objects on the screen. This can be performed using Lush's `sdl-sprite` class, defined also in `sdl/libSDL`. Creating a sprite for drawing on the `sdl-screen` `screen`, and loading an image into it is done as follows:

```
(setq sprite (new sdl-sprite screen 1))
(=> sprite load-frame "mysprite.png" 0 0 0)
```

This creates a new `sdl-sprite` with id number 1 on screen `screen`, and loads frame 0 of this sprite with the image in file "mysprite.png" (the image can be in any format accepted by the SDL-image library distributed with SDL). The three numerical arguments (all zeros in the example) are the frame number to be loaded, and the coordinates of the "handle" or "hot-point" of the sprite relative to the top-left corner.

Sprites can be moved around and drawn on the screen with:

```
(=> sprite move x y)
(=> sprite draw)
(=> screen flip)
```

All drawing operations occur in the back buffer (not on the visible screen) until the call to `(=> screen flip)`.

The main loop of a typical application will look something like this:

```
(while (not stop)
```

```
[get user input, update coordinates, and such]
(=> screen clear)
(=> sprite1 move x1 y1)
(=> sprite2 move x2 y2)
(=> sprite1 draw)
(=> sprite2 draw)
(=> screen flip))
```

Sprites can have multiple frames. Switching between frames provides a mechanism for animating sprites. More details on how to write such applications (including reading keyboard events , animation and such) are given in the tutorial [”Tutorial: Writing Games with Lush and SDL.”] . The code for that tutorial is available at .

2.26.4 sdlgl-screen class: OpenGL 3D graphics in SDL screens

The `sdlgl-screen` class is designed to allow OpenGL 3D graphics in an SDL screen. OpenGL graphics can be done in such an SDL screen , or can be done directly in GLUT windows. The `sdlgl-screen` should be used when SDL functionalities are used. Opening an `sdlgl-screen` can be done just like opening an `sdl-screen` :

```
(setq screen (new sdlgl-screen 640 480 "my title"))
```

An example of use of `sdlgl-screen` is available at .

2.27 Using OpenGL

Lush is also interfaced to OpenGL , and OpenRM which allows the development of 3D graphic applications.

See: OpenGL

Chapter 3

Lush Basics

3.1 Lush Interpreter Basics.

The Lush language works on data structures referred to as Lush objects. The main data structure is the “list”. A list is an ordered sequence of Lush objects of any type including other lists.

All other Lush objects are collectively referred to as “atoms”. There are in fact many kind of atoms used for implementing various Lush features. For instance:

- A number represents a numerical value.
- A symbol is just a name used for naming other lisp objects.
- A string is a sequence of characters useful for storing text.

A Lush program is stored as a collection of “functions” which are just another kind of Lush object. When a function is called, a predefined sequence of operations is applied to a couple of lisp objects (the function arguments) and a new lisp object is returned (the function result).

Once Lush is started, a loop is entered which “reads” a Lush object, “evaluates” it and displays the resulting Lush object. The evaluation of a list consists in calling the function named after the head of the list with the rest of the list as arguments.

Examples:

- Typing `(+ 3 4)` prints 7 which is the result of calling a function named `+` on the two numbers 3 and 4.
- Typing `(de sqr(x) (* x x))` defines a function named `sqr` with one argument referred to as symbol `x`. This functions calls the multiplication function `*` with two arguments both equal to `x`. The result of this call is then returned.

- Typing (`sqr 4`) then calls function `sqr` on the number `4`. The number `16` is then returned.

3.2 Lush Startup.

See: (`dump fname`)

See: (`startup ...argv...`)

See: (`toplevel`)

At startup time, Lush locates the executable file and goes up in the directory hierarchy until it finds a suitable initialization file.

- If there is a file named "`sys/sysenv.dump`" Lush loads this binary dump file (created with function `dump`) and proceed with the startup procedure.
- If there is a file named "`sysenv.lshc`" or "`sysenv.lsh`" in directory "`sys`", Lush loads this text or binary lisp file and proceed with the startup procedure. Lush refuses to start without one of the above files.

Lush then calls the lisp function `startup` with the Lush command line argument as arguments. This function is defined in the system file loaded above. In the standard `sysenv` file, this procedure loads the files given as argument to the Lush command. If no files are specified, file "`lushenv.lsh`" is searched and loaded. This file loads default environment (including a file "`.lush/lushrc.lsh`" in the user's home directory of Unix systems.). "`lushenv.lsh`" is searched in the following directories: "`local`", "`packages`", and "`lsh`".

Lush then calls function `toplevel`, which is currently defined in the Lush language by the "`sysenv.lsh`" file. This function sets up some debugging code, and waits for user commands.

The `toplevel` function continuously prints a question mark prompt, reads a lisp expression, evaluates this object and prints the result on the standard output. This process is repeated until you enter an end of file character. The global variable `result` always contains the last result printed by `toplevel`.

If the outermost list has not been completed when you enter a carriage return, `toplevel` prompts you for rest of the list, by displaying a few spaces instead of displaying a question mark prompt.

Example: Starting Lush

```
\\% lush/bin/lush
LUSH Lisp Universal Shell (compiled on Aug 14 2002)
  Copyright (C) 2002 Leon Bottou, Yann Le Cun, AT&T Corp, NECI.
  Includes parts of TL3
  Copyright (C) 1987-1999 Leon Bottou and Neuristique.
  Includes selected parts of SN3.2:
  Copyright (C) 1991-2001 AT&T Corp.
This program is free software; it is distributed under the terms
of the GNU Public Licence (GPL) with ABSOLUTELY NO WARRANTY.
```

Type '(helptool)' for details.

```
... loading startup file "/home/yann/lush/sys/sysenv.lsh".
... loading library file "/home/yann/lush/lsh/lushenv.lsh"
[compat.lsh]
[graphenv.lsh]
[.lushrc]
?
```

The startup process can be customized by specifying a magic first argument on the command line. The magic first argument starts with character "@", optionally followed by a second character "@", or followed by a filename,

- The optional filename indicates an alternate initialization file that will be loaded instead of the default "sys/sysenv.lsh". This file might be a binary dump file (".dump") or a lisp file (".lshc" or ".lsh"). This file must define suitable `startup` and `toplevel` functions. How these functions are defined determines the rest of the startup sequence. This allows to take complete control of Lush's startup sequence.
- The second character "@" indicates that Lush must start in script mode. Lush will not display the usual banners. It will silently load the "sysenv" initialization file, call the `startup` function, and exit. The default startup function in "sysenv" detects when Lush is called in this mode. When that happens, it stores the command line arguments in variable `argv` and loads the file indicated as second command line argument. This is suitable for writing Lush scripts just like we write shell scripts:

```
#!/usr/bin/lush @@
(printf "Capitalizing the arguments:\n")
(each ((arg argv)) (printf "\\%s \\%s\n" arg (upcase arg)))
```

3.3 The Lush Reader.

At this point you can type a textual representation of certain lisp object. The reader is a program which converts this textual representation into an actual lisp object.

Lush input is composed of words separated by blank delimiters (i.e. `space`, `tab` or `end-of-line`) by parentheses, or macro-characters. Any character between a semicolon (;) and an end-of-line is considered as a comment and is ignored by the interpreter.

- A numerical word is interpreted as a number.
- A word surrounded by double quotes is considered as a string.

- Any other words is interpreted as a symbol name.
- A list is read as a sequence of words enclosed within parenthesis. The empty list is written () .

The exact syntax for lisp objects is defined later in this manual. Moreover , the behavior of the reader can be modified by defining “macro-characters”. Whenever the Lush reader reaches a macro-character , it executes immediately an attached function which often calls the Lisp reader recursively. For instance , the character quote (') is a macro-character.

3.4 The Lush Evaluator.

The evaluator is a program which computes a new lisp object from another lisp object using a very simple rule. Every kind of object is in fact defined by a class (yet another Lush object). The class defines what happens when this object is evaluated and what happens when a list starting with this object name is evaluated.

The Lisp evaluator thus does not need to know anything about the evaluation behavior of all the different kinds of objects. In addition , it does not even need to test the nature of the objects. An indirect call is considerably more efficient than series of tests. Here is the evaluation algorithm for evaluating a Lisp object *p* :

```
eval(p) :
  if (p is the empty list)
    return the empty list
  else if (p is a list)
    q = eval(car p)
    if (q is an atom)
      call the class function 'listeval' for q with args (q,p)
    else
      call the default 'listeval' function with args (q,p)
  else if (p is an atom)
    call the class function 'selfeval' for p with arg (p)
```

3.5 Errors in Lush.

See: Interruptions in Lush.

See: (on-error *p* *l1* ... *ln*)

See: (errname)

See: (debug-hook)

The evaluation of the lists sometimes leads to an error. When an error occurs , Lush then executes the `debug-hook` function , which has been set on startup by function `toplevel` .

This function prints an error message , followed by the top elements of the evaluation stack (i.e. the list whose evaluation has caused the error , the list whose evaluation has caused the evaluation of the list that causes the error , etc...).

You are then prompted for a “Debug toplevel”.

- If you answer `n` , or simply type a carriage return , Lush aborts any current evaluation , resets the stack , and restarts the toplevel function.
- If you answer `y` , a prompt "`[Debug] ?`" is displayed , and you are able to type Lush commands for examining the current variables , display the current evaluation stack or perform any useful debugging tasks. You can leave this debugging toplevel by typing `Ctrl-D` or using the function `exit` .

3.6 Interruptions in Lush.

Lush can be interrupted at any time by typing `Ctrl-C` or `Ctrl-Break` . Interruption handling is essentially similar to error handling. First , Lush runs the function `break-hook` function which typically has been set up on startup by function `toplevel` .

This function prints an error message. You are then prompted for a “Break toplevel” which is basically a special kind of debug toplevel. When you leave the break toplevel by typing `Ctrl-D` or by using function `exit` , you get a chance to resume the execution at the interruption point.

Example:

```
? (for (i 1 100)
    (for (j 1 100)
        (if (= i j) (print (* i j))))))
1
4
^C9

*** Break
** in: (if (= i j) (print (* i j)))
** from: (for (j 1 100) (if (= i j) (print (* i j))))
** from: (for (i 1 100) (for (j 1 100) (if (= i j) (print (* i j))))) ...
** from: (load "$stdin" "$stdout")
** from: (let ((break-hook (lambda () (beep) (if (not (ask "Break top ...
Break toplevel [y/n] ?y
[Break] ? i
= 3
[Break] ? j
= 4
```

```

[Break] ? (* i j)
= 12
[Break] ? (where)
Current eval stack:
** in: (where)
** from: (load "$stdin" "$stdout" "[Break] ? ")
** from: (if (not (ask "Break toplevel")) (error "Stopped") (load "$s ...
** from: (if (= i j) (print (* i j)))
** from: (for (j 1 100) (if (= i j) (print (* i j))))
** from: (for (i 1 100) (for (j 1 100) (if (= i j) (print (* i j))))) ...
** from: (load "$stdin" "$stdout")
** from: (let ((break-hook (lambda () (beep) (if (not (ask "Break top ...
= ()
[Break] ? ^D
Resume execution [y/n] ?y
Resuming execution
16
25
36
49
64....

```

Break toplevels and debug toplevels cannot be nested. If an error or an interruption occurs when you are working in a break or debug toplevel, Lush will clean up the evaluation stack, and restart the toplevel function.

On Unix systems, Lush is inconditionnally interrupted by typing `Ctrl-\`. This method however may leave Lush with inconsistent internal states. Use it only when `Ctrl-C` does not work for some reason.

See: Errors in Lush.

See: (on-break p 11 ... ln)

See: (errname)

See: (break-hook)

3.7 Leaving Lush.

For leaving Lush, just type `Ctrl-D` or use the `exit` function. Type `y` and a carriage return when the confirmation request is displayed.

Example:

```

? (exit)
= ()

Really quit [y/n] ?y
Bye
\\%

```

3.8 Lush Memory Management.

Under normal use , Lush uses a fast reference counting garbage collector. Cells are allocated and deallocated on the fly. Each cell has a counter which contains the number of pointers that can access it , when the counter is equal to 0 , the cell is free; when it is greater than 0 , the cell is left for later use. The free cells are linked in a free list.

Careful use of these counters and of the free list allow the current Lush to be two to three times faster than the previous versions. You will probably notice the lack of bothersome pauses while the garbage collector takes effect.

If an error occurs , these counters are no longer right , and a classical garbage collector rebuilds the free-list and recomputes the counter values. Cells are allocated in the core memory by chunks. Once all allocated cells are used , Lush gets a new chunk from the system , and appends them in its free-list. This allocation scheme is used by all objects in Lush.

3.9 Lush Libraries.

Lush program files and libraries customarily end with the ".lsh" extension.

Lush reads library files from three directories:

- **sys** : contains the system libraries without which Lush cannot run
- **lsh** : contains libraries that are part of the standard Lush package
- **packages** : contains libraries that are contributed by users , which may not (necessarily) be supported to the same extent as the ones in lsh , and may not (necessarily) be available to all installation/versions of Lush.
- **local** : is meant to contain libraries that are specific to you local installation of Lush.

Loading a program or a library is done with the `load` or the `libload` functions , as in (`libload "cmacro"`) . This will search the above three directories for a file named "`cmacro.lsh`" . In other words , the above three directories constitute the default search path for finding programs and libraries. The search order is `local`, `packages`, `lsh`, `sys` .

The core libraries (in the `sys` directory) contains "`sysenv.lsh`" , which includes functions , macros , and macrocharacters that define the core language.

Some files in the library directories have the extension ".lshc" instead of ".lsh". These are "tokenized" version of the corresponding ".lsh" file. Tokenized files are binary files that load faster than the equivalent ".lsh".

The standard library directory tree (`lsh` directory and its subdirectories) contains many libraries and utilities without which Lush would not be nearly as much fun. These libraries can be relied on to be present with all distributions of Lush on all platforms.

When Lush is started , it looks for a file named "lushenv.dump, lshc, lsh" in the following directories: current directory , "local" , "packages" , and "lsh" . If it is found , it is loaded. There is a default "lushenv.lsh" in the "lsh" directory. This default "lushenv.lsh" sets the search path to , defines a few autoload functions , and reads the ".lushrc" file in the home directory of the user.

The standard library contains many niceties including numerous numerical and graphic functions and objects.

A notable example is "ogre" , an object-oriented GUI toolkit which makes it very easy to build mouse driven graphical user interfaces. A couple lines of code are enough to setup a window with a button that calls a lush function when clicked upon:

```
;; start ogre
(ogre)
;; define function that will be called when button is clicked
(de print-coucou (c) (printf "coucou\n"))
;; open ogre window and insert a button in it
(new windowobject 0 0 100 40 "hello"
  (new stdbutton "coucou" print-coucou))
```

3.10 Lush Runtime.

Using function `dump` , it is possible to create a binary image of the lisp structures contained in a Lush process. Such binary dump files are useful for distributing Lush based programs without providing an extensive list of Lush library files.

Lush indeeds undumps file "sysenv.dump" before even searching for the regular startup library "sysenv.lsh" . You can also request Lush to load a specific dump file by typing "@<filename >" as the first command line argument to the Lush program.

It is then useful to redefine functions `startup` and `toplevel` in order to redefine the startup sequence and the interactive loop. The following function dumps a binary file for a typical Ogre application:

```
(de dump-cold(s)
  ;; save functions <toplevel> and <startup>.
  (unlock-symbol toplevel)
  (unlock-symbol startup)
  (setq old-toplevel toplevel)
  (setq old-startup startup)

  ;; Redefining <startup> gives you a chance
  ;; to initialize your application and parse the
  ;; command line arguments
```

```

(de startup argv
  ;; Disable <CTRL-C>
  (unlock-symbol break-hook)
  (setq break-hook (lambda() t))
  ;; Initialize the OGRE library
  (ogre)
  ;; Parse arguments and create the main window.
  (init-my-application argv) )

;; Redefining <toplevel> gives you a chance to
;; alter the user interaction loop. This function
;; is reentered whenever an error occurs.
;; The default <toplevel> waits for lisp commands.
;; The following function just accepts OGRE events.
(de toplevel()
  (process-pending-events)
  (while my-application-mainwindow
    (waitevent)
    (process-pending-events) )
  (exit 0) )

;; We can then dump a binary file.
(dump s)

;; And restore the previous status.
(setq toplevel old-toplevel)
(setq startup old-startup)
(lock-symbol toplevel)
(lock-symbol startup) )

```

Redefining function `toplevel` prevents the user to enter Lush commands and gives total control to your program. The above example , for instance , restricts user interaction to graphics interfaces. The user cannot enter lisp expressions.

See: Lush Startup.

See: `(dump fname [exec])`

See: `(startup ...argv...)`

See: `(toplevel)`

Chapter 4

Core Interpreter and Startup Library

This section includes all the function available at startup with the default environment `lsh/lushenv.lsh` . Included are the core interpreter functions and the library functions defined in `sys/sysenv.lsh` , all the files in `lsh/libstd` (that are loaded at startup).

4.1 Lists.

List are the elementary data structure in Lush. Several functions are designed for handling lists. Lists are stored as pairs (`car` . `cdr`). The first element of the pair (`car`) is a pointer to the first element of the list. The second element of the pair (`cdr`) is a pointer to the list of the remaining elements.

The textual representation of a list is composed of an opening parenthesis, a sequence of Lush objects separated by blanks and closing parenthesis. The empty list is thus written as `()` .

The second element (`cdr`) of a pair is not necessary a list. A special dotted pair notation is used for representing such a pair. This notation is composed of an opening parenthesis, the first Lush objects, a blank separated dot, the second Lush object and a closing parenthesis.

Examples:

```
(1 2 3)           ; car=2 and cdr=(2 3)
((q w) 2)        ; car=(q w) and cdr=2
(1 2 (e r))      ; car=1 and cdr=(2 (e r))
(a . b)          ; car=a and cdr=b
```

4.1.1 List Manipulation Functions

The following functions are useful for accessing or building lists.

(car l)

Type: DX

Returns the first element of a list or dotted pair *l* .

Example:

```
? (car '(a b c d))
= a
```

(cdr l)

Type: DX

Returns the next elements of a list or dotted pair *l* .

Example:

```
? (cdr '(a b c d))
= (b c d)
```

(c...r l)

Letters “...” stands for any combination of two or three letters *a* or *d* . All the combinations of two *car* or *cdr* functions are written in *C*. The combination of three are written in Lisp in “*sysenv.lsh*” .

Example:

```
? (cadr '(a b c d))
= b
```

(cons a1 a2)

Type: DX

Builds a list whose *car* is *a1* and *cdr* *a2* .

Example:

```
? (cons '+ '(2 3))
= (+ 2 3)
```

```
? (cons 'a 'b)
= (a . b)
```

(list a1 ... an)

Type: DY

Builds a list with its arguments **a1** to **an** .

Example:

```
? (list 'a '(b c) 'd)
= (a (b c) d)
```

(makelist n v)

Type: DX

Returns a list containing **n** times element **v** .

Example:

```
? (makelist 4 'e)
= (e e e e)
```

(range [n1] n2 [delta])

Type: DX

Returns a list of all the numbers between **n1** and **n2** , stepping by **delta** .
The default value for both arguments **n1** and **delta** is 1 .

Example:

```
? (range 2 5)
= (2 3 4 5)
```

(length l)

Type: DX

Returns the length of the list **l** .

Example:

```
? (length '(a b c d e))
= 5
```

This function is able to detect that list **l** is circular. The value **-1** is returned when such a condition occurs.

(append l1 ... ln)

Type: DX

Concatenates lists *l1* to *ln* . The idiom `(append l ())` may be used to get a fresh copy of the list *l* .

Example:

```
? (append '(1 2 3) '(4 5 6))
= (1 2 3 4 5 6)
```

(reverse l)

Type: DX

Returns a list *l* in reverse order.

Example:

```
? (reverse '(1 2 3))
= (3 2 1)
```

(nth n l)

Type: DX

Returns the *n* th elements of the list *l* . The first element is numbered 0 . For compatibility with the previous versions of Lush, the syntax `(nth 1 n)` is allowed. In this case, the first element is numbered 1 .

Example:

```
? (nth 3 '(a b c d e))
= d
```

(nthcdr n l)

Type: DX

Returns the *n* th pair of the list *l* . The first pair is numbered 0. This function is equivalent to *n* calls of the `cdr` functions.

Example:

```
? (nthcdr 3 '(a b c 3 e f))
= (3 e f)
```

(nth% i l)

Type: DE(miscenv.lsh)

This function is similar to function `nth`, except it considers `l` as a virtually circular list. In other words, it reduces `i` to the rest of its division by the length of list `l`. Since it calls function `length`, no effective circular list should be given as argument.

(last l)

Type: DX

Returns the last element of the list `l`.

Example:

```
? (last '(a b c 3 e f))
= f
```

(lastcdr l)

Type: DX

Returns the last pair of the list `l`.

Example:

```
? (lastcdr '(a b c 3 e f))
= (f)
```

(nfirst n l)

Type: DX

When `n` is 0, the empty list is returned.

When `n` is greater than 0, this function returns the `n` first elements of `l`. When the length of the list is lower than `n`, a copy of the list is returned.

When `n` is lower than 0, this function returns all elements but the `n` last. When the length of the list is lower than the absolute value of `n`, the empty list is returned. If `l` is a circular list, an error occurs.

(member e l)

Type: DX

Searches list `l` for element `e`. If element `e` is not found, function `member` returns the empty list `()`. Otherwise function `member` returns the first sublist of `l` whose first element is equal to `e`.

There is a single equality test in Lush that is able to recursively compare lists and strings.

Example:

```
? (member 'e (range 1 4))
= ()
```

```
? (member 'e '(a b c d e f g h i j))
= (e f g h i j)
```

(flatten x)

Type: DX

Returns all the atoms in *x* , linked in a single list.

Example:

```
? (flatten '2)
= (2)
```

```
? (flatten '(1 2 (3 (6 7)) 4))
= (1 2 3 6 7 4)
```

(filter f x)

Type: DE

Returns a list of those elements of *x* for which the application of function *f* to the element returns a non-nil result.**(assoc key alist)**

Type: DX

See: (alist-add *key value alist*)See: (alist-get *key alist*)

See: Hash Tables.

Function **assoc** is useful for searching an “alist”. An alist is a list of pairs (*key . value*) representing associations between a search key *key* and a value *value* .

Function **assoc** returns the first pair of *alist* whose first element is equal to *key* . The empty list is returned when no matching pair is found. Once the pair is located, you can access the value associated with the key using function **cdr** . You can also change the value using function **rplacd** .

Remark: Alists should be only used for small numbers of key-value pairs. We suggest using hash tables as soon as the association involves more than twelve key-value pairs. Hash tables indeed require more memory but are much faster and much more convenient.

(alist-add key value alist)

Type: DE(miscenv.lsh)

See: (assoc key alist)

This function (based on function `assoc`) adds or changes a pair (`key value`), in a alist `alist` and returns the new alist.

This function is defined as follows:

```
(de alist-add(key value alist)
  (let ((pair (assoc key alist)))
    (if pair
      (rplacd pair value)
      (setq alist (cons (cons key value) alist)) )
    alist ) )
```

(alist-get key alist)

Type: DE(miscenv.lsh)

See: (assoc key alist)

This function (based on function `assoc`) returns the value associated to a key `key` in alist `alist` .

This function is defined as follows:

```
(de alist-get(key alist)
  (let ((pair (assoc key alist)))
    (when pair
      (cdr pair) ) ) )
```

(sort-list l comp)

Type: DE(sysenv.lsh)

Returns a sorted copy of list `l` according the order relation achieved by the diadic function `comp` .

```
? (sort-list '(12 3 1 2 3 4 5 6) >)
= (1 2 3 3 4 5 6 12)
```

(sort-index l)

Type: DE(miscenv.lsh)

The function `sort` ranks the component of its argument. Argument `l` may be a list or a vector.

This function returns a list wich elements of are indices of `l` such that its components are ranked from the biggest to the lowest.

(lbound1 l)

Type: DE(miscenv.lsh)

The function `lbound1` returns the dimensions of a list `l`.

The list `l` is assumed to be homogeneous, which means that the length of its member sublists have all the same length, and this in a recursive way. In other word, list `l` is the list counterpart of a multidimensional array.

```
? (lbound1 '(1 2 (3 4) 5))
= (4)
```

(lboundn l)

Type: DE(miscenv.lsh)

The function `lboundn` returns the dimensions of the list `l`.

All lists, even unhomogeneous, are supported.

```
? (lboundn '(1 2 (3 4) 5))
= (4 2)
```

(sublist l n s1...sn)

Type: DE(miscenv.lsh)

This function extracts sublists from a flat list `l`. The procedure follows three steps.

- 1 List `l` may be splitted into several sublists of equal size. When `n` is `()`, list `l` is used as it is (without being splitted). Otherwise argument `n` defines a format for splitting list `l` which is then considered as a flat list. When `n` is an integer, `l` is splitted into sublists of length `n`. When `n` is a list, it describes the dimensions of the splitted list following the semantic of `msplit-dim`.
- 2 From the shaped list is extracted a rectangular sublist. Arguments `s1` ... `sn` define the selected subscripts. They follow the same subscripts semantic as `submatrix`.
- 3 The sublist is returned in a flat format.

The following example takes a list of numbers from 0 to 99, considered as an array `5 * 20`, and extracts the items in the intersection of lines 1, 2 and columns 4, 5, 6, 7.

```
? (sublist (range 100) '(5 ()) '(1 2) '(4 7)))
= (25 26 27 28 45 46 47 48)
```

```
? (sublist (range 0 29) 5 () '(2 3)))
= (2 3 7 8 12 13 17 18 22 23 27 28)
```

See: (submatrix mat s1 ...[sn])

See: (msplit-dim m targetdim newdims)

4.1.2 Physical List Manipulation Functions

Instead of building new pairs, a few fast functions directly change the pointers stored in the pairs given as argument. This is safe if the programmer knows that the function's arguments are not used elsewhere.

Side effects should be expected when another lisp object points to the argument of such a physical list modification function. The modification then appears in both places.

Physical list manipulation functions can be used to construct lists that reference themselves. Some Lush functions are able to test this condition (e.g. `length`) and avoid an infinite loop. Some Lush functions however (e.g. `print` , `flatten`) will loop forever when processing such a list. You can interrupt these loops by typing `Ctrl-C` or `Break` .

See: (length l)

See: (== n1 n2)

(rplaca l e)

Type: DX

Physically replaces the "car" of list l by e .

Example:

```
? (setq a '(1 2 3))
= (1 2 3)
? (setq b (cdr a))
= (2 3)
? (rplaca b 'e) ; this is fast
= (e 3)
? a ; but causes side effects
= (1 e 3)
```

(rplacd l e)

Type: DX

Physically replaces the cdr of list l by e .

(displace l1 l2)

Type: DX

Replaces both the car and cdr of `l1` by the car and cdr of `l2` . The main purpose of this function is the implementation of DMD functions.

(nconc l1 ... ln)

Type: DE(sysenv.lsh)

Function `nconc` returns the concatenation of the lists `l1` to `ln` . It uses physical replacement and therefore causes side effects.

(nconc1 l e)

Type: DE(sysenv.lsh)

Function `nconc` returns a list formed by appending element `e` to the end of list `l` . It uses physical replacement and therefore causes side effects.

Example:

```
? (setq a '(a b c d e))
= (a b c d e)
? (nconc1 a 'f)
= (a b c d e f)
```

(list-insert l pos x)

Type: DE(miscenv.lsh)

Function `list-insert` returns a list formed by inserting element `x` at position `pos` of list `l` . It uses physical replacement and therefore causes side effects.

(list-delete l pos)

Type: DE(miscenv.lsh)

Function `list-insert` returns a list formed by deleting the element located at position `pos` from list `l` . It uses physical replacement and therefore causes side effects.

(list-merge l1 l2)

Type: DE(miscenv.lsh)

Function `list-merge` returns a list formed by appending to list `l1` the elements of `l2` that are not already in `l1` . This function uses physical replacement and therefore causes side effects.

4.2 Numbers.

See: Special Numerical Values (IEEE754).

Numbers are the simplest kind of atoms. They are stored as double precision floating point numbers.

Numerical computations in Lush are usually performed in double precision. Single precision however is used by some number crunching functions, like the matrix calculus operations.

You can type in decimal number using the traditional exponential notation or using hexadimal numbers preceded by characters "0x" . Here are a few examples of legal numbers:

```

23
3.14
12e6           ;; equivalent to 12000000
1.2E6         ;; equivalent to 1200000
0x1A          ;; equivalent to 26

```

4.2.1 Elementary Numerical Functions.

The following elementary functions operate on numbers:

(+ n1 ... nm)

Type: DX

Computes the sum of numbers n1 to nm .

Example:

```
? (+ 1 2 3 4)
= 10
```

(1+ n)

Type: DX

Adds 1 to n .

Example:

```
? (1+ 4)
= 5
```

(- [n1] n2)

Type: DX

Subtracts n2 from n1 . If one argument only is provided, returns the opposite of n2 .

Example:

```
? (- 2 5)
= -3
```

(1- n)

Type: DX
Subtracts 1 from **n** .
Example:

```
? (1- (+ 4 5))
= 8
```

(* n1 ... nm)

Type: DX
Computes the product of **n1** to **nm** .
Here is a memory intensive way of computing a factorial:

```
? (apply '* (range 1 5))
= 120
```

(2* n)

Type: DX
Multiplies **n** by 2.
Example:

```
? (2* (+ 4 5))
= 18
```

(/ [n1] n2)

Type: DX
Divides **n1** by **n2** . If one argument only is provided, Function / returns the inverse of **n2** .
Example:

```
? (/ 3)
= 0.3333
```

(2/ n)

Type: DX

Divides *n* by 2.

Example:

```
? (2/ 5)
= 2.5
```

(n m)**

Type: DX

Returns *n* raised to the *m* th power.

Example:

```
? (** 2 5)
= 32
```

(next2 n)**

(miscenv.lsh)

This function returns the first power of 2 equal or greater than *n* .**(max l1 ... ln)**

Type: DX

Return the maximum element within *l1* ... *ln* . Arguments *li* may be numbers or strings.

Example:

```
? (max 2 3 (* 3 3) (+ 4 3))
= 9
```

(min l1 ... ln)

Type: DX

Return the minimum element within *l1* ... *ln* . Arguments *li* may be numbers or strings.

Example:

```
? (min 2 3 (* 3 3) (+ 4 3))
= 2
```

(round x [n]))

(miscenv.lsh)

This functions rounds the number *x* respecting the desired number of significant digits. When the number of significant digit is 0 or lower, it returns the first power of 10 smaller than or equal to *x* .

The number *x* is expected to be greater than or equal to 0.

```
? pi
= 3.1416
? (round pi 3)
= 3.14
? (round (* 1000 pi) 3))
= 3140
? ;; for most numbers, (round x 0) is not equal to (round x 1)
? (round pi 1)
= 3
? (round pi 0)
= 1
? ;; for powers of ten, (round x 0) is equal to (round x 1)
? (round 1000 1)
= 1000
? (round 1000 0)
= 1000
```

4.2.2 Integer Arithmetic Functions.

Besides the ordinary number operation functions , a few functions deal with integer numbers only. For that purpose we consider only the first 32 bits of the integer part of the number.

(div n1 n2)

Type: DX

Returns the quotient of the Euclidian division of *n1* by *n2* .

Example:

```
? (div 5 2)
= 2
```

(mod n1 n2)

Type: DX

Returns the remainder of the Euclidian division of *n1* by *n2* .

Example:

```
? (mod 5 2)
= 1
```

(bitand n1 n2)

Type: DX

Returns the bitwise AND of integer numbers **n1** and **n2** .

(bitor n1 n2)

Type: DX

Returns the bitwise OR of integer numbers **n1** and **n2** .

(bitxor n1 n2)

Type: DX

Returns the bitwise exclusive or (XOR) of integer numbers **n1** and **n2** .

(bitshl n1 n2)

Type: DX

Shifts the bits of integer **n1** by **n2** positions to the left and returns the result. This is similar to the C operator `>>` .

(bitshr n1 n2)

Type: DX

Shifts the bits of integer **n1** by **n2** positions to the right and returns the result. This is similar to the C operator `>>` .

4.2.3 Mathematical Numerical Functions.

Most mathematical functions are implemented as DZ functions. DZs are a remnant of a now obsolete compiler for numerical expressions. For more details about DZ functions the corresponding section.

(sgn n)

Type: DZ

Returns +1 if **n** is equal or greater than 0 , -1 otherwise.

```
? (sgn -2)
= -1
```

```
? (sgn 1.2)
= 1
```

(abs n)

Type: DZ

Returns the absolute value of n .

? (abs 123.3)
= 123.3

? (abs -23.3)
= 23.3

(int n)

Type: DZ

Returns the integral part of n .

? (int -4.5)
= -5

? (int 4.5)
= 4

(sqrt n)

Type: DZ

Returns the square root of n

? (sqrt 15)
= 3.873

(0-x-1 n)

Type: DZ

This function implements a piecewise saturated linear function. It returns 0 if n is smaller than -1. It returns 1 if n is larger than +1. It returns n if n is in the -1 to +1 range.

? (0-x-1 -2)
= -1

? (0-x-1 0.7)
= 0.7

```
? (0-x-1 1.3)
= 1
```

(0-1-0 n)

Type: DZ

This function implements the indicator function of the -1 to +1 range. It returns 1 if *n* is in the -1 to +1 range. It returns 0 otherwise.

```
? (0-1-0 -2)
= 0
```

```
? (0-1-0 0.7)
= 1
```

(sin n)

Type: DZ

Returns the sine of *n* radians.

```
? (sin (/ 3.1415 3))
= 0.866
```

(cos n)

Type: DZ

Returns the cosine of *n* radians.

```
? (cos (/ 3.1415 3))
= 0.5
```

(tan n)

Type: DZ

Returns the tangent of *n* radians.

```
? (tan (/ 3.1415 3))
= 1.7319
```

(atan n)

Type: DZ

Returns the arc tangent of n , in radians.

```
? (* 4 (atan 1))  
= 3.1416
```

(exp n)

Type: DZ

Returns the exponential of n .

```
? (exp 1)  
= 2.7183
```

(exp-1 n)

Type: DZ

Returns the exponential of n minus 1. This function gives an accurate value of $\exp(n)-1$ even for small values of n .

```
? (exp-1 0.5)  
= 0.6487
```

(log n)

Type: DZ

Returns the natural logarithm of n .

```
? (log 2)  
= 0.6931
```

(log1+ n)

Type: DZ

Returns the natural logarithm of n plus 1. This function gives an accurate value of $\log(1+n)$ even for small values of n .

```
? (log1+ 1)  
= 0.6931
```

(sinh n)

Type: DZ

Returns the hyperbolic sine of n .

? (sinh 1)

= 1.1752

(cosh n)

Type: DZ

Returns the hyperbolic cosine of n .

? (cosh 1)

= 1.5431

(tanh n)

Type: DZ

See: (qtanh n)

Returns the hyperbolic tangent of n .

? (tanh 1)

= 0.7616

(qtanh n)

Type: DZ

See: (tanh n)

Returns a rational approximation of the hyperbolic tangent of n . This function is orders of magnitude faster than `tanh` . Its accuracy however is quite low (about 0.00001.)

? (qtanh 1)

= 0.7616

The formula of the approximation is: $\tanh(x) \approx \frac{P(x)^{16-1}}{P(x)^{16+1}}$ where $P(x) = 1 + 0.125 x + 0.0078125 x^2 + 0.00032552083 x^3$

(qdtanh n)

Type: DZ

Returns a rational approximation of the derivative of the hyperbolic tangent of n . The accuracy of this function is quite low (about 0.00001), but its computation time is very small.

? (qdtanh 1)
= 0.42

The formula of the approximation is: $\left[\frac{d}{dx} (1 - \tanh(x)^2) \approx 1 - \left(\frac{P(x)^{16}-1}{P(x)^{16}+1} \right)^2 \right]$ where $\left[P(x) = 1 + 0.125x + 0.0078125x^2 + 0.00032552083x^3 \right]$

(qstdsigmoid n)

Type: DZ

See: (qtanh n)

Returns the value of standard sigmoid at point n , computed with a rational approximation to the hyperbolic tangent. The standard sigmoid is defined as $\left[1.7159 \operatorname{tanh} \left(\frac{2}{3}x \right) \right]$

(qdstdsigmoid n)

Type: DZ

See: (qdtanh n)

Returns the value of the derivative of standard sigmoid at point n , computed with a rational approximation to the hyperbolic tangent. The standard sigmoid is defined as $\left[1.7159 \operatorname{tanh} \left(\frac{2}{3}x \right) \right]$

(qexpmx n)

Type: DZ

Returns a rational approximation of the exponential of minus the absolute value of n . The accuracy of this function is quite low (about 0.00001), but its computation time may be smaller on certain computers.

? (qexpmx 1)
= 0.3679

The formula of the approximation is: $\left[e^{-x} \approx \frac{1}{P(-x)^8} \right]$ where $\left[P(x) = 1 + 0.125x + 0.0078125x^2 + 0.00032552083x^3 + 0.000010172526x^4 \right]$

(qdexpmx n)

Type: DZ

Returns a rational approximation of the derivative of the exponential of minus the absolute value of n . The accuracy of this function is quite low (about 0.00001), but its computation time may be smaller on certain computers.

? (qdexpmx 1)
= 0.3679

The formula of the approximation is: $\left[\frac{d}{dx} e^{-x} \approx -\operatorname{sgn}(x) \frac{1}{P(-x)^8} \right]$ where $\left[P(x) = 1 + 0.125x + 0.0078125x^2 + 0.00032552083x^3 + 0.000010172526x^4 \right]$

(qexpmx2 n)

Type: DZ

Returns a rational approximation of the exponential of minus the square of n . The accuracy of this function is quite low (about 0.00001), but its computation time may be smaller on certain computers.

? (qexpmx 1)
= 0.3679

The formula of the approximation is: $\left[e^{-x^2} \approx \frac{1}{P(x^2)^8} \right]$ where $\left[P(x) = 1 + 0.125x + 0.0078125x^2 + 0.00032552083x^3 + 0.000010172526x^4 \right]$

(qdexpmx2 n)

Type: DZ

Returns a rational approximation of the exponential of minus the square of n . The accuracy of this function is quite low (about 0.00001), but its computation time may be smaller on certain computers.

? (qdexpmx 1)
= 0.3679

The formula of the approximation is: $\left[\frac{d}{dx} e^{-n^2} \approx -2x \frac{1}{P(n^2)^8} \right]$ where $\left[P(x) = 1 + 0.125x + 0.0078125x^2 + 0.00032552083x^3 + 0.000010172526x^4 \right]$

(solve min max f)

Type: DX

Simple dichotomical root finder. `f` is a function with one argument only. `solve` will return a numerical approximation of the solution of $f(x)=0$ between `min` and `max`. To ensure the existence of this solution, `f(min)` and `f(max)` must have different signs.

Example:

```
? (solve 1 2 (lambda (x)
              (- 2 (* x x)) ))
= 1.4142
```

4.2.4 Random Numbers

Statistical functions are provided for computing random numbers, computing various statistics on list of numbers or performing simple linear regressions.

(seed n)

Type: DX

Sets the random number generator seed. The random generator is replicable. The same seed will produce the same sequence of random numbers. Argument `n` must be a legal integer number.

(rand [[a] b])

Type: DX

Returns a uniform random number:

- If no arguments are provided, function `rand` returns a random number between 0 and 1.
- If argument `b` only is given, function `rand` returns a random number between `-b` and `b`.
- If both arguments `a` and `b`, function `rand` returns a random number between `a` and `b`.

Example:

```
? (rand)
= 0.1861
```

(**gauss** [[*m*] *s*])

Type: DX

Returns a gaussian number with mean *m* and standard deviation *s* . The default value for argument *m* is 0 . The default value for argument *s* is 1 .

Example

```
? (gauss 2)
= 3.3327
```

4.2.5 Constant Numbers.

PI

Type: CONST(miscenv.lsh)

The well-known transcendental trigonometric number.

```
? pi
= ()
```

4.2.6 Hashing

(**hashcode** *expr*)

Type: DX

Returns a string containing a hash code for the lisp expression *expr* . Hash codes are guaranteed to be equal when two expressions are logically equal , that is to say , equal as defined by function = .

4.3 Booleans Operators.

Expression in Lush are considered “false” when their evaluation returns the empty list. All other expression are considered “true”. The symbol **t** however is preferred for this meaning.

4.3.1 Boolean Arithmetic

A few functions are useful for programming tests and returning boolean values

(**=** *n1* *n2*)

Type: DX

Test if *n1* is equal to *n2* . Two objects are equal if they have the same type and if they convey the same information. The criterion for deciding equality therefore depends on the type of the object.

This function supports any Lush object.

- It is able to compare the value of numbers, strings, dates, lists, matrices, arrays, hash tables and high level objects (those created by `new` and belonging to classes created by `defclass`).
- The comparison of other low-level objects (i.e. from user defined classes written in C language following the Open protocole) usually rely on the physical equality. However it is possible to put other comparison functions into the C structures defining the classes.

Example:

```
? (= '(1 2 (3)) (cons 1 '(2 (3))))
= t
```

```
? (= [ 1.00 2.00 ] [ 1.00 2.00 ])
= t
```

Testing the equality between special numerical values may be tricky. For example, the IEEE754 specification (supported by most of the industry) defines special bit patterns named NaN (Not a Number). Comparing two NaNs should always return false. Major operating systems and compilers however do not respect this.

Lush expressly specifies that the result returned by the equality test is **always** true when objects `n1` and `n2` share the same memory location. In other words, if the function `==` returns `t` then the function `=` returns `t` as well.

See: Special Numerical Values (IEEE754).

See: `(== n1 n2)`

See: Comparison of User Defined Objects.

```
(<> n1 n2)
```

Type: DX

Tests if `n1` is different from `n2` . This function is equivalent to

```
(de <> (n1 n2) (not (= n1 n2)))
```

Example:

```
? (<> 2 "abcd")
= t
```

```
? (<> 2 2)
= ()
```

See: `(= n1 n2)`

```
(== n1 n2)
```

Type: DX

See: (= n1 n2)

Test if **n1** and **n2** are physically equal.

This function does not even look at the information conveyed by the lisp objects **n1** or **n2** . It just tests that the pointers returned by the expression **n1** and **n2** are equals.

Pointer equality not only means that objects referred to by expression **n1** and **n2** are equal, but also means that they are located at the same memory addresses.

This information is meaningful as soon as you use functions that modify objects (as opposed to functions returning a modified copy of the object and leave the initial object unchanged). These functions include the following:

- Physically changing list component using `rplaca` or `rplacd` .
- Setting the contents of an array or a matrix.
- Setting the value of a slot of an object

Function `==` may be used in conjunction with these functions. Modifying indeed the object referred to as **n1** will therefore also modify the object referred to as **n2** . Function `==` may also be used as a fast way to compare objects and matrix when you know that the only possibility of equality is physical equality.

Example:

```
? (setq a 3)
= 3
? (== a 3)
= ()
? (= a 3)
= t
? (== a a)
= t
? (setq a [1 2])
= [1 2]
? (setq b [1 2])
= [1 2]
? (= a b)
= t
? (== a b)
= ()
? (a 0)
= [0 2]
? (= a b)
= ()
```

? b
= [1 2]

(0= n)

Type: DX

Test if **n** is equal to 0.

See: Special Numerical Values (IEEE754).

(0<> n)

Type: DX

Test if **n** is different from 0.

See: Special Numerical Values (IEEE754).

(<= n1 n2)

Type: DX

Test if **n1** is less or equal than **n2** . Arguments **n1** or **n2** may be dates, strings or numbers. User defined objects may be supported as well.

See: Special Numerical Values (IEEE754).

See: Comparison of User Defined Objects.

(< n1 n2)

Type: DX

Test if **n1** is less than **n2** . Arguments **n1** or **n2** may be dates, strings or numbers. User defined objects may be supported as well.

See: Special Numerical Values (IEEE754).

See: Comparison of User Defined Objects.

(>= n1 n2)

Type: DX

Test if **n1** is greater or equal than **n2** . Arguments **n1** or **n2** may be dates, strings or numbers. User defined objects may be supported as well.

See: Special Numerical Values (IEEE754).

See: Comparison of User Defined Objects.

(> n1 n2)

Type: DX

Test if **n1** is greater equal than **n2** . Arguments **n1** or **n2** may be dates, strings or numbers. User defined objects may be supported as well.

See: Special Numerical Values (IEEE754).

See: Comparison of User Defined Objects.

(and l1 ... ln)

Type: DX

Evaluates sequentially `l1 .. ln` . If a result is the empty list, function `and` immediately returns `()` . The result of the evaluation of `ln` is returned otherwise.

Example:

```
? (and (= 2 2) (= 2 3) (print (= 2 2)))
= ()
```

(or l1 ... ln)

Type: DX

Evaluates sequentially `l1 .. ln` . If a result is not the empty list, this result is immediately returned. The empty list is returned otherwise.

Example:

```
? (or (= 2 2) (= 2 3))
= t
```

(null l)

Type: DX

Returns `t` if the result of evaluating `l` is the empty list. Returns the empty list otherwise. The macro character `| |` may be used for that purpose.

Example:

```
? (null t)
= ()
```

(not l)

Type: DX

The `null` and `not` functions are identical.

`l`

Type: MCHAR(sysenv.lsh)

This macro-character expand to `(null l)` .

4.3.2 Predicates.

Predicates test the type of any lisp object.

(listp 1)

Type: DX

Returns a non-nil result when 1 is a list, including the case of the empty list.

Example:

```
? (listp ())  
= t
```

```
? (listp "abc")  
= ()
```

```
? (listp '(2 3))  
= (2 3)
```

(consp 1)

Type: DX

Returns a non-nil result when 1 is a non-empty list.

Example:

```
? (consp ())  
= ()
```

```
? (consp "abc")  
= ()
```

```
? (consp '(2 3))  
= (2 3)
```

(atomp 1)

Type: DX

Returns a non-nil result when 1 is an atom.

Example:

```
? (atomp ())  
= t
```

(numberp 1)

Type: DX

Returns a non-nil result when 1 is a number.

Example:

```
? (numberp 3.14)
= 3.14
```

(symbolp 1)

Type: DX

Returns a non-nil when 1 is a symbol.

Example:

```
? (symbolp 'a)
= a
```

(externp 1)

Type: DX

Returns a non-nil result item when 1 is a non numerical atom.

Example:

```
? (externp car)
= ::DX:car
```

(classp 1)

Type: DE

Returns a non-nil item when 1 is a class.

Example:

```
? (classp object)
= t
```

(functionp n)

Type: DX

Returns t if n is a function.

Example:

```
? (functionp "abc")
= ()
```

```
? (functionp functionp)
= ::DX:functionp
```

```
(stringp s)
```

Type: DX

Returns a non nil result if **s** is a string.

Example:

```
? (stringp "abc")
= "abc"
```

4.4 Symbols.

Symbols are the only named objects. Symbol names may be up to 40 characters long.

The evaluation of a symbol returns the “value” of the symbol. Function `setq` changes the value of a symbol. The value of a new symbol is always the empty list.

During a call to a Lush function, or during the execution of certain special functions (e.g. function `let`) predefined symbols take a temporary value. The previous values are then restored when the function exits.

For instance, assume that we define a function `(def sqr(x) (* x x))` . Evaluating list `(sqr 4)` will perform the following actions:

- save the old values of symbol `x` which is the name of a function argument,
- set the value of symbol `x` to `4` ,
- execute the function body `(* x x)` and save the result,
- restore the previous value of symbol `x` ,
- and return the saved result.

This dynamical binding makes the interpreter getting faster, but somewhat precludes the development of efficient compilers.

The textual representation of a symbol is composed by the characters of its name. The reader usually converts symbol names to lowercase and replaces underscores by dashes. This can be prevented by quoting the symbol name with vertical bars. Such vertical bars are useful for defining symbols whose name contains any character usually forbidden: parenthesis, spaces, semi-colon, macro-characters, upper-case characters etc....

Examples :

```

gasp GasP          ; are the same symbol named "gasp"
|GasP|            ; is different symbol named "GasP"
|);( '|          ; is a symbol named ");(' '"
"gasP"           ; is a string
12               ; is a number
|12|            ; is a symbol named "12"

```

Unlike several dialect of lisp, Lush does not provide other fields in a symbol for storing a specific functional value or a property list. The value of a symbol is the only storage provided by symbols.

4.4.1 (defvar name val)

Type: DM(sysenv.lsh)

Defines a new global variable named **name** (a symbol) and initialized with value **val** . Traditionally function **setq** was used for that purpose in SN/TLisp. There are cases however when **setq** refuses to create a new global variable. Eventually we will completely disallow creating global variables with **setq** and require the use of **defvar** for that purpose.

4.4.2 (setq v1 a1 ... [vn an])

Type: DX

Sets the value of symbols **v1** ... **vn** to **a1** ... **an** . It is good practice to create the symbol being setq'ed beforehand using **defvar** , **let** , **let*** or other constructs that create global or local variables.

The **setq** function has a special behavior when the **vi** are not symbols. This behavior is documented later, with the scope function.

Example:

```

? (setq a 3)
= 3
? (setq b 6)
= 6
? a b
= 3
= 6
? (setq a b b a)
= 3
? a b
= 6
= 3

```

4.4.3 (set v a)

Type: DX

Sets the value of the symbol *v* to *a* . *set* is different from *setq* because *v* is evaluated first.

Example:

```
? (setq s 'a)
= a
? (set s 3)
= 3
? s
= a
? a
= 3
```

4.4.4 (incr v [n])

Type: DX

Increments *v* by *n* (default 1). *v* is not evaluated, and must be a symbol whose value is a number. Number *n* may be a positive or negative.

Example:

```
? (setq s 0)
= 0
? (incr s)
= 1
? s
= 1
```

4.4.5 (named s)

Type: DX

Returns a new symbol whose name is the string *s* .

Example:

```
? (named "a")
= a
```

```
? (named "A")
= |A|
```

4.4.6 (nameof s)

Type: DX

Returns a string containing the name of symbol `s` . This is the converse of `named` .

Example:

```
? (nameof 'a)
= "a"
```

```
? (nameof '|A|)
= "A"
```

4.4.7 (namedclean s)

Type: DX

Returns a new symbol whose name is computed by normalizing the string `s` using the same algorithm as the Lisp reader. This is the converse of function `pname` .

Example:

```
? (namedclean "a")
= a
```

```
? (namedclean "A")
= a
```

```
? (namedclean "|A|")
= |A|
```

4.4.8 (lock-symbol s1 ... sn)

Type: DX

Locks symbols `s1` to `sn` . Symbols may be locked. You can no longer change the value of locked symbols , but you may still modify them temporarily by using the `let` function. Some Lush functions and all C functions are stored in locked symbols. This avoids the accidental loss of a C function.

4.4.9 (unlock-symbol s1 ... sn)

Type: DX

Unlocks symbols `s1` to `sn` .

4.4.10 (symlist)

Type: DX

Returns the list of all the symbol names: i. e. a list of strings.

Example:

```
? (length (symlist))
= 1799
```

4.4.11 (oblist)

Type: DX

Returns the list of all the symbols: i. e. a list of symbols.

Example:

```
? (length (oblist))
= 1799
```

4.4.12 (macrochp s)

Type: DX

Returns *t* if *s* symbol defines a macro-character.

Example:

```
? (macrochp ' '|)
= t
```

4.4.13 (putp anything name value)

Type: DE(sysenv.lsh)

Every Lush object (atoms, cons, etc.) may be enriched by defining properties identified by a symbolic name. Function `putp` is used to define such properties. Its first argument is an arbitrary Lush object `anything`. Function `putp` sets the property named by symbol `name` to value `value`. Properties can be later retrieved using `getp`.

4.4.14 (getp anything name)

Type: DE(sysenv.lsh)

Retrieves the property `name` for the Lush object `anything`.

4.5 Namespaces and Packages

4.5.1 (`private ..symbols..`)

Type: DM(sysenv.lsh)

This function must be used inside a package definition. It indicates that the symbols `..symbols..` must be considered private to the package.

See: (`package name ..expressions..`)

4.5.2 (`package name ..expressions..`)

Type: DM(sysenv.lsh)

Author(s): Leon Bottou

This construct provides a simple way to make symbols "semi-private" to a group of functions so as to avoid name clashes. It basically has the effect of prepending the package `name` to all the symbols defined within its scope and marked as private (using the `private` construct).

Example:

```
(package foobar
  (private foo)
  (de foo(x) (print x))
  (de bar(x) (foo x)) )
```

defines a global symbol `bar` . Symbol `foo` however is renamed `foobar.foo` in all expressions occurring after (`private foo`) . Hence the above expression would be equivalent to:

```
(de bar (x)
  (toto.foo x) )
(de toto.foo (x)
  (print x) )
```

Internals: Symbol `name` is used to identify the group of functions. Function `package` maintains a list of hidden symbols declared with function `private` . Each expression is first transformed to replace all occurrences of a hidden symbol by a new symbol whose name is composed by prepending the package name `name` , then evaluated as usual.

4.6 Control structures.

Control structure are lisp functions specially designed for controlling how a program is executed. There are control structures for starting the execution (e.g. `eval`) and control structures for programming loops (e.g. `while`).

4.6.1 (eval l1...ln)

Type: DX

See: The Tlisp Evaluator.

The function `eval` calls the evaluator on objects `l1` to `ln`, and returns the result of the last evaluation, i.e. `(eval ln)`.

Example:

```
? (eval (cons '+ '(3 4)))
= 7
```

4.6.2 (apply f l)

Type: DX

Function `apply` applies function `f` to the argument list `l`.

First of all, `f` is evaluated. The result can be a function, a symbol whose value is a function or even a list whose result is a function. This function is then applied to arguments `l`.

Evaluating `(apply f l)` is quite similar to `(eval (cons f l))`. It blocks however the evaluation of the arguments stored in list `l`.

Example:

```
? (apply 'append '((1 2 3) (1 3 5)))
= (1 2 3 1 3 5)
```

```
? (apply append '(list 2 3) (list 4 (+ 1 2)))
= (list 2 3 list 4 (+ 1 2))
```

```
? (eval (cons append '(list 2 3) (list 4 (+ 1 2))))
= (2 3 4 3)
```

4.6.3 (quote a)

Type: DX

See: ' a

See: ' expr , a ,@ a

Function `quote` returns the Tlisp object `a` without evaluating it.

In most cases indeed, functions evaluate their arguments and operate on the results of these evaluations. In `(car (range 2 5))` for instance, function `car` evaluates the arguments `(range 2 5)` giving `(2 3 4 5)` and returns the first element of this list.

```
? (car (range 2 5))
= 2
```

Using the function `quote` prevents a function from evaluating a Lisp object passed as argument. It evaluates the list `(quote a)` instead, returning the object `a` itself.

```
? (car '(range 2 5))
= range
```

4.6.4 'a

Type: DMC(sysenv.lsh)

See: `(quote a)`

See: `'expr , a ,@ a`

The function `quote` usually is called by using the quote `'` macro-character. Writing `'a` is a shorthand for `(quote a)`.

Example:

```
? (car '(range 2 5))
= range
```

```
? '(range 2 5)
= (quote (range 2 5))
```

4.6.5 'expr ,a ,@a

The backquote `'` macro provides a handy syntax for producing complex lists depending of external conditions. The backquote macro works like the quote macro with the following additions:

- Within a backquoted expression, a comma `,` indicates that the next Lisp object must be evaluated.
- If the comma is followed by an AT sign `(,@)`, the list returned by the evaluation of the next Lisp object will be merged within the current list.

Example:

```
? '(1 2 ,( + 4 5) (+ 4 5 ,(range 1 4)))
= (1 2 9 (+ 4 5 (1 2 3 4)))
? '(1 2 ,( + 4 5) (+ 4 5 ,@(range 1 4)))
= (1 2 9 (+ 4 5 1 2 3 4))
```

Implementation: Combinations of backquote and comma macros are expanded into lisp expressions composed of calls to functions `list`, `append` and `cons`, which builds the result.

Example:

```
;; note the quote <'> followed by a backquote <'>
? ' '(1 2 ,(+ 4 5) (+ 4 5 ,@(range 1 4)))
= (list '1 '2 (+ 4 5) (append '(+ 4 5) (range 1 4)))
```

4.6.6 (`progn l1...ln`)

Type: DY

Evaluates `l1` to `ln` and returns the result of the evaluation of `ln`.

Example:

```
? (progn
  (print (+ 2 3))
  (* 2 3) )
5
= 6
```

4.6.7 (`prog1 l1...ln`)

Type: DY

Evaluates `l1` to `ln` and returns the result of the evaluation of `l1`.

Example:

```
? (prog1
  (print (+ 2 3))
  (* 2 3) )
5
= 5
```

4.6.8 (`let ((s1 v1) ... (sn vn)) l1 ... ln`)

Type: DY

The function `let` is used to define local variables.

Function `let` performs the following operations:

- Expressions `v1 ... vn` are evaluated.
- The values of symbols `s1 ... sn` are saved.
- The results of the evaluations of `v1 ... vn` are stored into symbols `s1 ... sn`.

- Expressions `l1 ... ln` are evaluated. If these expressions refer to symbols `s1 ... sn`, they will see the values returned by expressions `v1 ... vn` (unless you change these values using `setq` !)
- The saved values of `s1 ... sn` are restored.
- The result of the evaluation of `ln` is returned.

Example:

```
? (let ((s 0))
    (for (i 1 10)
      (setq s (+ s i)) ) )
= 55
```

4.6.9 (let* ((s1 v1) ... (sn vn)) l1 ... ln)

Type: DY

Function `let*` is also used for defining local variables. Unlike function `let` however, the evaluation of the variable initial values are performed after setting the previous variables. You can therefore refer to the previously defined variables.

Function `let*` performs the following operations:

- For each pair `(si vi)`, the value of symbol `si` is saved, the expression `vi` is evaluated and the result is stored into symbol `si`.
- Expressions `l1 ... ln` are evaluated. If these expressions refer to symbols `s1 ... sn`, they will see the values returned by expressions `v1 ... vn` (unless you change these values using `setq` !)
- The saved values of `s1 ... sn` are restored.
- The result of the evaluation of `ln` is returned.

Example:

```
? (let* ((i 1)
        (j (1+ i)) )
    (print i j)
    (* 2 j) )
1 2
= 4
```

4.6.10 (let-filter ((filter data)) l1...ln)

Type: DM(sysenv.lsh)

The function `let-filter` attempts to match the data expression `data` with the filter expression `filter`. A filter expression is a lisp expression composed uniquely with lists and symbols. For instance, argument `filter` can be a single symbol, a list of symbols, or something more complex such as `(a ((b c)) d . e)`.

- A symbol in the filter expression matches anything in the data expression. When such a match occurs, the previous value of the symbol is saved and the symbol is assigned the matching part of the data expression.
- A list in the filter expression matches a list in the data expression if all its elements match the corresponding elements in the data expression.

If there is a match, function `let-filter` evaluates the lists `l1 ... ln` with the new values of the symbols listed in the match expression, restores the initial value of the symbols, and return the result of the last evaluation. Otherwise, function `let-filter` simply restores the initial values of the symbol and returns the empty list.

4.6.11 (if cond yes [no1...non])

Type: DY

First, expression `cond` is evaluated. If the result is not the empty list, expression `yes` is evaluated, otherwise the remaining expressions `no1 ... non` are evaluated. Function `if` returns the result of the last evaluation.

Example:

```
? (if (> 3 4)
      3
      4 )
= 4
```

4.6.12 (when cond yes1...yesn)

Type: DY

First, expression `cond` is evaluated. If the result is not the empty list, expressions `yes1` to `yesn` are evaluated and the result of `yesn` is returned.

Example:

```
? (when (> 3 4)
      (print "error") )
= ()
```

4.6.13 (while cond l1...ln)

Type: DY

While the evaluation of `cond` gives a result different from the empty list, expressions `l1` to `ln` are evaluated in a loop.

If the result of the first evaluation of `cond` is the empty list, expressions `l1` to `ln` are never evaluated and function `while` returns the empty list. Otherwise function `while` returns the result of the last evaluation of `ln`.

Example:

```
? (let ((l (range 2 5)))
  (while l
    (print (car l) (sqrt (car l)))
    (setq l (cdr l)) ) )
2 1.4142
3 1.7321
4 2
5 2.2361
= ()
```

4.6.14 (do-while cond l1...ln)

Type: DY

Function `do-while` evaluates lists `l1` to `ln` and loops until the evaluation of condition `cond` returns the empty list. The result of the last evaluation of `ln` is then returned.

Expressions `l1` to `ln` are evaluated before testing for the loop condition. If the result of the first evaluation of `cond` is the empty list, expression `l1 ... ln` are evaluated exactly once.

Example:

```
(de input(prompt regex)
  (let ((answer ()))
    (do-while (not (regex-match regex answer))
      (printf "\\%s" prompt)
      (flush)
      (setq answer (read-string)) )
    answer) )
```

4.6.15 (repeat n l1...ln)

Type: DY

Repeats `n` times the evaluation of `l1` to `ln`. Function `repeat` returns the result of the last evaluation of `ln`.

Example:

```
? (repeat 4
   (prin 1) )
1111= 1
```

4.6.16 (mapwhile cond l1...ln)

Type: DM(miscenv.lsh)

See: (while cond l1 ... ln)

Function `mapwhile` first evaluates expression `expr` . If this evaluation returns a non nil result, it evaluates lists `l1` to `ln` . This process is repeated until the evaluation of `cond` returns the empty list.

Unlike function `while` however, function `mapwhile` returns a list containing all the results of the successive evaluation of `ln` .

```
? (let ((i 0))
    (mapwhile (< i 5)
              (incr i)
              (* i i) ) )
= (1 4 9 16 25)
```

4.6.17 (for (symb start end [step]) l1...ln)

Type: DY

Function `for` implements a classical “for” loop. The value of symbol `symb` is first saved. Then `symb` takes numeric values from `start` to `end` stepping by `step` (default 1). For each value, expression `l1` to `ln` are evaluated. The value of `symb` is then restored.

Function `for` returns the result of last evaluation of `ln` .

Example:

```
? (for (i 2 5)
      (print i (sqrt i)) )
2 1.4142
3 1.7321
4 2
5 2.2361
= 2.2361
```

4.6.18 (mapfor (symb start end [step]) l1...ln)

Type: DMD(sysenv.lsh)

See: (for (symb start end [step]) l1 ... ln)

Function `mapfor` implements a loop like function `for`. It returns however a list containing the results of the evaluation of `ln` for all values of the loop index.

Example:

```
? (mapfor (i 2 5)
      (sqrt i) )
= (1.4142 1.7321 2 2.2361)
```

4.6.19 (cond l1...ln)

Type: DY

Function `cond` implements the standard lisp “cond” form.

Arguments `l1 ... ln` are lists of the form `(cond . body)`. The conditions `cond` are evaluated sequentially until one returns a result different from the empty list. The function `progn` is applied to the corresponding `body` and the result is returned.

Example:

```
;; a function for returning the sign of a number
(de sign(x)
  (cond
    ((< x 0) -1)
    (> x 0) +1)
    (t 0) ) )
```

4.6.20 (selectq s l1...ln)

Type: DY

The arguments `l1` to `ln` are lists of the form `(case . body)`.

The argument `s` is first evaluated. The lists `l1` to `ln` are then checked until `s` is equal to `case`, or `s` is a member of the list `case`, or `case` is equal to `t`. The function `progn` is applied to the corresponding `body` and the result is returned.

Example:

```
(selectq x
  (0 (print "zero"))
  ((2 4 6 8) (print "even"))
  ((1 2 3 5 7) (print "prime"))
  (t (print "nothing interesting")) )
```

4.6.21 (mapc f l1...ln)

Type: DX

For $i=1$ till (length $l1$) this function constructs a list containing the i -th element of each list $l1 \dots ln$ and applies the function f to to it. Th lists $l1$ till ln should have the same length. This function returns the result of the last evaluation.

Example:

```
? (mapc print '(1 2 3))
1
2
3
= 3
```

4.6.22 (mapcar f l1...ln)

Type: DX

Like `mapc` but returning the list of all evaluation results rather than just the last one.

Example:

```
? (mapcar '+ '(1 2 3) '(4 5 6))
= (5 7 9)
```

4.6.23 (rmapc f l1...ln)

Type: DX

See: (mapc f l1 ... ln)

`rmapc` works like `mapc` , but is able to recursively scan sublists in $l1 \dots ln$

.

Example:

```
? (rmapc 'prin '(1 2 (3 (4)) 5))
12345= 5
```

4.6.24 (rmapcar f l1...ln)

Type: DX

See: (mapcar f l1 ... ln)

`rmapcar` works like `mapcar` , but is able to recursively scan sublists in $l1 \dots ln$.

Example:

```
? (rmapcar 'prin '(1 2 (3 (4)) 5))
12345= (1 2 (3 (4)) 5)
```

4.6.25 (each ((s1 v1) ... (sn vn)) l1 ... ln)

Type: DY

For $i=1$ till (length $v1$) this function performs the let function

- (let ((s1 i-th element < $v1$ >) ... (sn i-th element < vn >)) l1 ... ln)

The lists $v1$ till vn should have the same length. This function returns the result of the last evaluation. Example:

```
? (each ((i '(1 2 3 5 7)))
      (print i (sqrt i)) )
1 1
2 1.4142
3 1.7321
5 2.2361
7 2.6458
= 2.6458
```

4.6.26 (all ((s1 v1) ... (sn vn)) l1 ... ln)

Type: DY

Like `each` but returning the list of all evaluation results rather than just the last one.

Example:

```
? (all ((i '(1 2 3 5 7)))
      (sqrt i) )
= (1 1.4142 1.7321 2.2361 2.6458)
```

4.7 Functions.

A Lush function is an unnamed object which defines how to evaluate a list whose first element evaluates to this function. Functions are usually stored in the value field of a symbol usually referred to as the function name.

There are thus two equivalent ways to define a function named `square` which computes the square of its argument:

```
(setq square (lambda(x) (* x x)))
```

or

```
(de square(x) (* x x))
```

The evaluation of the list `(lambda...)` returns a function that `setq` stores into symbol `square`. The function `de` is a shorthand for this manipulation.

A function behaves as any atom does. Its class however defines how to evaluate a list whose first element evaluation returns that function. The evaluation of a list whose first element is a symbol, the value of which is a function, thus simply calls that function. The evaluation of the following list then returns 25

```
(square 5)
```

But the evaluation of the following list also returns 25

```
((lambda(x) (* x x)) 5)
```

There are several kinds of functions in Lush; briefly described hereafter:

- DX functions are written in C, using the standard Lisp to C interface package. Most of the C functions are written as DX.
- DY functions are directly written in C. These functions are usually control structures, like `cond`, `progn`, etc...
- DE functions are written in Lisp. These functions evaluate their arguments before calling their body. Most Lisp functions are DE functions.
- DF functions are also written in Lisp. They do not however evaluate their arguments before giving control to the function body.
- DM functions are also known as “macro functions”. They are written in Lisp. Their argument list includes the function name itself, and its elements are not evaluated. The body of the function is supposed to return an evaluable list, known as the macro expansion. This list is then evaluated in the calling context, and gives the result of the function call.
- DMD functions are a refinement of DM functions. Unlike the DM functions, their argument list does not include the function name. The result of the body evaluation replaces the calling form itself in the calling context. The evaluation process is then restarted. This replacement process happens just once, because the code is physically replaced by the result of the macro expansion during the first call.
- DMC are known as “macro characters”. These are not real function. They are bound to a single character name and take no arguments. Once the Lush reader reaches such a character, it calls the function body and substitutes the result to the macro-character.

4.7.1 Argument List.

A formal argument list is associated to each Lush function (i.e. DE, DF, DM or DMD) This argument list defines which symbols will be bound to the actual argument of the function, while evaluating its body.

Valid formal argument lists may be a single symbol, a list of symbols, or something more complex like (a b (c d) . f) , or (a b &optional c d) .

If the formal argument list is a single symbol, it will be bound to the list of arguments when the function is called.

If the formal argument list is a list, its “car” will be matched against the “car” of the actual arguments, and its “cdr” will be matched against the “cdr” of the actual arguments.

Example:

```
? (de surface((x1 y1) (x2 y2) )
      (* (- x1 x1) (- y2 y2)) )
= surface
? (surface '(4 5) '(8 10))
= 40
```

Moreover, two symbols &optional and &rest in argument lists have a special meaning, similar to the Common Lisp conventions.

Only symbols and lists composed of a symbol and a default value may follow the &optional symbol in the formal argument list. These symbols are optional arguments. If they are omitted, they are bound to their default value, or the empty list.

Example:

```
? (de printline(s &optional (indent 0) (terminate "."))
      (tab indent)
      (printf "\\%s\\%s\\n" s terminate) )
= printline
? (printline "hello")
hello.
= ()
? (printline "hello" 6)
      hello.
= ()
? (printline "hello" 6 "@")
      hello@
= ()
```

Finally, the formal argument list may be terminated by &rest followed by another symbol. The list of all remaining actual arguments will be bound to that symbol.

4.7.2 Defining a Function.**(de symb args . body)**

Type: DY

See: Argument List.

Creates a new function which evaluates its arguments (DE) and stores this function into the value field of symbol **symb** .

A valid argument list **args** may be a single symbol, a list of symbols or something more complex like **(a b (c d) . f)** . A valid function body **body** is a list of Lush expressions which are evaluated whenever the function is called.

Whenever a DE function is called,

- The previous values of the argument names are saved,
- The actual arguments are evaluated and stored in the corresponding argument names,
- Each expression in the function body is evaluated and the last result is saved,
- The previous values of the argument names are restored,
- The saved result is returned.

Example:

```
? (de square (x) (* x x))
= square
? (square 9)
= 81
? (de first-arg 1
   (car 1) )
= first-arg
? (first-arg 'a 'b)
= a
? (first-arg 'prems 'deuse 'troise 'quatrz)
= prems
```

(df symb args . body)

Type: DY

See: Argument List.

See: **(dm symb args . body)**)

Creates a new function which does not evaluates its arguments (DF) and stores this function into the value field of symbol **symb** .

A valid argument list **args** may be a single symbol, a list of symbols or something more complex like **(a b (c d) . f)** . A valid function body **body** is a list of Lush expressions which are evaluated whenever the function is called.

Whenever a DF function is called,

- The previous values of the argument names are saved,
- The actual arguments are not evaluated but directly stored in the corresponding argument names,
- Each expression in the function body is evaluated and the last result is saved,
- The previous values of the argument names are restored,
- The saved result is returned.

Example: Defining a control structure

```
? (df ifn(cond no . yes)
    (if (not (eval cond))
        (eval no)
        (apply progn yes))) )
= ifn
? (ifn (= 2 3) 1 2)
= 2
```

Such a definition causes problems if some subexpression depend on a the value of a symbol named like a function argument.

```
? (ifn (= 2 3) cond 2)
= (= 2 3)
```

The previous expression should indeed return the value of symbol `cond` in the current context rather than the value of symbol `cond` in the context of function `ifn` .

This problem is alleviated by using “macro functions”.

```
(dm symb args . body)
```

Type: DY

See: Argument List.

Creates a macro-function (DM) and stores it into the value field of symbol `symb` . The `body` of a DM function actually computes the expression which will be evaluated in the calling context. This expression can be found with the function `macro-expand` .

A valid argument list `args` may be a single symbol, a list of symbols or something more complex like `(a b (c d) . f)` . A valid function body `body` is a list of Lush expressions which are evaluated whenever the function is called.

Whenever a DM function is called,

- The previous values of the argument names are saved,
- The elements of the entire calling expression, including the function name, are stored in the corresponding argument names,
- Each expression in the function body is evaluated and the last result (known as the macro expansion) is saved,
- The previous values of the argument names are restored,
- The saved macro-expansion is evaluated,
- The result is returned.

Example:

```
? (dm ifn(fname cond . rest)
   (list 'if (list 'not cond) . rest) )
= ifn
? (ifn (= 2 3) "yes" "no")
= "yes"
? (setq cond "yes")
= "yes"
? (ifn (= 2 3) cond)
= "yes"
```

(macro-expand macrocall)

Type: DY

See: (dm symb args . body)

Returns the “expansion” of a call to a DM function.

Example:

```
? (dm ifn(fname cond . rest)
   (list 'if (list 'not cond) . rest) )
= ifn
? (macro-expand (ifn (= 2 3) "yes" "no"))
= (if (not (= 2 3)) "yes" "no")
```

(dmd symb args . body)

Type: DM(sysenv.lsh)

See: Argument List.

See: (dm symb args . body)

Creates a DMD function and stores it into symbol `symb` . DMD functions display two differences with DM functions:

- Unlike the DM functions, their argument list does not include the function name.
- The result of the body evaluation replaces the calling form itself in the calling context.

A valid argument list `args` may be a single symbol, a list of symbols or something more complex like `(a b (c d) . f)`. A valid function body `body` is a list of Lush expressions which are evaluated whenever the function is called.

Whenever a DMD function is called,

- The previous values of the argument names are saved,
- The arguments are stored in the corresponding argument names,
- Each expression in the function body is evaluated and the last result (known as the macro expansion) is saved,
- The previous values of the argument names are restored,
- The macro expansion is physically installed in the calling form using function `displace`, replacing forever the call to the DMD function in the calling code,
- The saved macro-expansion is evaluated,
- The result is returned.

Example:

```
? (dmd ifn(cond . rest)
    (list 'if (list 'not cond) . rest) )
= ifn
? (de if-test(n)
    (ifn (= n 2) "yes" "no") )
= if-test
? (if-test 2)
= "no"
? (pretty if-test)
(de if-test(n)
  (if (not (= n 2)) "yes" "no") )
= (if-test)
```

(dmc symb . body)

Type: DY

Defines a macro-character `symb`. A macro-character is not really functions, since it interacts only with the Lisp reader.

There are three kinds of macro-characters:

- The first ones are single character macro-characters. When the Lush reader encounters one of them, it immediately calls the associated DMC function and returns its result.
- Caret macro-characters have a two character name, whose first character is a caret `^`. A caret DMC behaves exactly like a single characters DMC: When the reader encounters a caret followed by another character, it immediately calls the associated DMC function and returns its result.
Caret macro-character may also be abbreviated by typing the corresponding control character (if this control character exists and is not intercepted by the operating system). If you are running WinLush, the fastest method consists in typing `Ctrl+Shift+Letter`. It is always possible however to type the caret followed by the character.
- Hash macro-characters are also two character macro-characters, whose first character is a hash sign `#`. If their body returns `()`, nothing is read. If it returns a one element list, this element is read. If it returns anything else, an error is signaled.

Hash macro characters are useful for conditionally reading certain pieces of code.

Since the reader performs special actions whenever it encounters a macro-character, it is advisable to surround the name of symbol `symp` with vertical bars `|`.

Examples:

The macro character `'` which expands into a call to function `quote` is defined as :

```
(dmc |'| (list 'quote (read)))
```

The macro character `^P` which expands into a call to function `pretty` is defined as :

```
(dmc |^P| (list 'pretty (read)))
```

The following macro-character `#!` could be used to signal debug instructions that will be not read if variable `ndebug` is set.

```
(dmc |#!| (let ((inst (read))) (when (not ndebug) inst)))
```

```
(lambda args . body)
```

Type: DY

See: Argument List.

See: `(de symp args . body)`

Returns a function which evaluates its arguments (DE). This function operates like function `de` but does not store the function in a particular symbol.

Example:

```
? ((lambda (x)
    (* x x) ) (+ 4 5))
= 81
```

(flambda args . body)

Type: DY

See: Argument List.

See: (df symb args . body)

Returns a function which does not evaluate its arguments (DF). This function operates like function `df` but does not store the function in a particular symbol.

Example:

```
? ((flambda (x)
    (print x) ) (+ 4 5))
(+ 4 5)
= (+ 4 5)
```

(mlambda args . body)

Type: DY

See: Argument List.

See: (dm symb args . body)

Returns a macro-function (DM). This function operates like function `dm` but does not store the function in a particular symbol.

(funcdef n)

Type: DX

Returns a list which defines the function `n` . This works for DEs, DFs, DMs and other interpreted functions, but returns nil for DH, DX, DY, and other compiled or intrinsic functions.

Example:

```
? (funcdef caddr)
= (lambda (l)
    (car (caddr l)) )
```

(pretty f)

Type: DM(sysenv.lsh)

Display a nicely indented definition of function `f` .

Actually, `pretty` sends a pretty message to `f` . Each class defines how an instance of that class will be displayed. In particular, functions are displayed by printing an indented version of the definition of the functions `f1` to `fn`.

Example

```
? (pretty caddr)
(de caddr (l)
  (car (caddr 1)) )
= t
```

^P function

Type: DMC(sysenv.lsh)

See: (pretty f)

This macro character expands into a call of function `pretty` .

Example:

```
? ^Paddpath
(de addpath (dir)
  (setq dir (concat_fname dir))
  (let ((oldpath (path))
        (newpath (list dir)) )
    (while oldpath
      (when (<> dir (car oldpath))
        (setq newpath (nconc1 newpath (car oldpath))) )
      (setq oldpath (cdr oldpath)) )
    (apply path newpath) ) )
= ()
```

4.8 Strings.

Strings are sequence of characters.

The textual representation of a string is composed of the characters enclosed between double-quotes. A string may contain macro-characters, parentheses, semi-colons, as well as any other character. A line terminating backslash indicates a multi-line string.

The following “C style” escape sequences are recognized inside a string:

- `\\` for a single backslash,
- `\"` for a double quote,

- `\n` , `\r` , `\t` , `\b` , `\f` respectively for a linefeed character (Ascii LF), a carriage return (Ascii CR), a tab character (Ascii TAB), a backspace character (Ascii BS), and a formfeed character (Ascii FF),
- `\e` for a `end-of-file` character (Stdio's EOF),
- `\~?` for a control character `control-?` ,
- `\+?` for a meta character `meta-?` (high bit is set),
- `\000` where 000 are octal digits for any ascii character,
- `\xHH` where HH are hexadecimal digits for any ascii character.

4.8.1 Basic String Functions

Like most Lush functions, the basic functions operating on strings do not modify their arguments. They create instead a new string on the basis of their arguments.

See: (`> n1 n2`)

See: (`<n1 n2`)

(concat s1 ... sn)

Type: DX

Concatenates strings `s1` to `sn` .

Example:

```
? (concat "hello" " my friends")
= "hello my friends"
```

(len s)

Type: DX

Returns the length of the string `s` .

Example:

```
? (len "abcd")
= 4
```

(mid s n [l])

Type: DX

Returns `l` characters of string `s` starting at the `n`th position. The position `n` is a number between 1 and the length of the string minus 1. When argument `l` is omitted, function `mid` returns characters until the end of the string `s` .

Example:

```
? (mid "alphabet" 3 2)
= "ph"
```

```
? (mid "alphabet" 3)
= "phabet"
```

(right s n)

Type: DX

Returns the *n* rightmost characters of *s* .

Example:

```
? (right "alphabet" 3)
= "bet"
```

(left s n)

Type: DX

Returns the *n* leftmost characters of *s* .

Example:

```
? (left "alphabet" 3)
= "alp"
```

(strins s1 n s2)

Type: DX

Insert string *s2* at position *n* into the string *s1* , and returns the result.

When *n* is equal to 0, the `strins` function actually concatenates *s2* and *s1* .

Example:

```
? (strins "alphabet" 3 "***")
= "alp***habet"
```

(strdel s1 n l)

Type: DX

Removes *l* character at position *n* from the string *s1* .

Example:

```
? (strdel "alphabet" 3 2)
= "alabet"
```

(index s r [n])

Type: DX

Searches the first occurrence of the string **s** in the string **r**, starting at the **n** th position. **index** returns the position of the first match. If such an occurrence cannot be found, it returns the empty list.

Example:

```
? (index "pha" "alpha alphabet alphabetical" 4)
= 9
```

(upcase s)

Type: DX

Returns string **s** with all letters converted to uppercase.

Example:

```
? (upcase "alphabet")
= "ALPHABET"
```

(upcase1 s)

Type: DX

Returns string **s** with first letter converted to uppercase.

Example:

```
? (upcase1 "alphabet")
= "Alphabet"
```

(downcase s)

Type: DX

Returns string **s** with all letters converted to lowercase.

Example:

```
? (downcase "ALPHABet")
= "alphabet"
```

(val s)

Type: DX

Returns the numerical value of **s** considered as a number. Returns the empty list if **s** does not represent a decimal or hexadecimal number.

Example:

```
? (val "3.14")  
= 3.14
```

```
? (val "abcd")  
= ()
```

```
? (val "0xABCD")  
= 43981
```

(str n)

Type: DX

Returns the decimal string representation of the number *n* .

Example:

```
? (str (2* 3.14))  
= "6.28"
```

(strhex n)

Type: DX

Returns the hexadecimal string representation of integer number *n* .

Example:

```
? (strhex 18)  
= "0x12"
```

(strgptr p)

Type: DX

Returns the hexadecimal string representation of pointer *p* preceded by an ampersand.

(asc s)

Type: DX

Returns the numeric ASCII code of the first character of string *s* . This function causes an error if *s* is an empty string.

Example

```
? (asc "abcd")  
= 97
```

(chr n)

Type: DX

Returns a one-character string containing the ASCII character of code `n` . Integer `n` must be strictly positive, and smaller than 255.

Example

```
? (chr 48)
= "0"
```

(isprint s)

Type: DX

Returns `t` if string `s` contains only printable characters.

Example:

```
? (isprint "alpha bet")
= t
```

```
? (isprint "alpha\\`Cbet")
= ()
```

(pname l)

Type: DX

Returns a string representation for the lisp object `l` . `pname` is able to give a string representation for numbers, strings, symbols, lists, etc...

Example:

```
? (pname (cons 'a '(b c)))
= "(a b c)"
```

(sprintf format ... args ...)

Type: DX

Like the C language function `sprintf` , this function returns a string similar to a format string `format` . The following escape sequences, however are replaced by a representation of the corresponding arguments of `sprintf` :

- `"%"` is replaced by a single `\%`.
- `"%1"` is replaced by a representation of a lisp object.
- `"%[-] [<n]s">` is replaced by a string, right justified in a field of length `n` if `n` is specified. When the optional minus sign is present, the string is left justified.

- "%[-] [<n]d"> is replaced by an integer, right justified in a field of *n* characters, if *n* is specified. When the optional minus sign is present, the string is left justified.
- "%[-] [<n [. m]] c "> where *c* is one of the characters *e* , *f* or *g* , is replaced by a floating point number in a *n* character field, with *m* digits after the decimal point. *e* specifies a format with an exponent, *f* specifies a format without an exponent, and *g* uses whichever format is more compact. When the optional minus sign is present, the string is left justified.

Example:

```
? (sprintf "\\%5s(\\%3d) is equal to \\%6.3f\\n" "sqrt" 2 (sqrt 2))
= " sqrt( 2) is equal to 1.414\\n"
```

(strip *s*)

Type: DE(miscenv.lsh)

This function deletes the leftmost and rightmost spaces in string *s* .

```
(strip " This sentences is full of spaces. ")
```

(stripl *s*)

Type: DE(miscenv.lsh)

This function deletes the leftmost spaces in string *s* .

```
(stripl " This sentences is full of spaces. ")
```

(stripr *s*)

Type: DE(miscenv.lsh)

This function deletes the rightmost spaces in string *s* .

```
(stripr " This sentences is full of spaces. ")
```

4.8.2 Regular Expressions (regex).

A regular expression describes a family of strings built according to the same pattern. A regular expression is represented by a string which “matches” (using certain conventions) any string in the family. TL provides four regular expression primitives (`regex-match` , `regex-extract` , `regex-peek` , and `regex-subst`) and several library functions.

The conventions for describing regular expressions in Lush are quite similar to those used by the `egrep` unix utility:

- An ordinary character matches itself. Some characters, () \ [] | . ? * and \ have a special meaning, and should be quoted by prepending a backslash \ . The string "\\\\" actually is composed of two backslashes (because backslashes in strings should be escaped!), and thus matches a single backslash.
- A dot . matches any character.
- A caret ^ matches the beginning of the string.
- A dollar sign \$ matches the end of the string.
- A range specification matches any specified character. For example, regular expression [YyNn] matches Y y N or n , regular expression [0-9] matches any digit, regular expression [^0-9] matches any character except a digit, regular expression [A-Za-z] matches a closing bracket, or any uppercase or lowercase letter.
- The concatenation of two regular expressions matches the concatenation of two strings matches regular expression. Regular expressions can be grouped with parenthesis, and modified by the ? + and * characters.
- A regular expression followed by a question mark ? matches 0 or 1 instance of the single regular expression.
- A regular expression followed by a plus sign + matches 1 or more instances of the single regular expression.
- A regular expression followed by a star * matches 0 or more instances of the single regular expression.
- Finally, two regular expressions separated by a bar — match any string matching the first or the second regular expression.

Parenthesis can be used to group regular expressions. For instance, the regular expression "(+|-)?[0-9]+(\\. [0-9]*)?" matches a signed number with an optional fractional part. Furthermore, there is a “register” associated with each parenthesized part of a regular expression. The matching routines use these registers to keep track of the characters matched by the corresponding part of the regular expression. This is useful with functions `regex-extract` and `regex-subst` .

(`regex-match r s`)

Type: DX

Returns `t` if regular expression `r` exactly matches the entire string `s` . Returns the empty list otherwise.

Example:

```
? (regex-match "(+|-)?[0-9]+(\\. [0-9]*)?" "-56")
= t
```

(`regex-extract` *r s*)

Type: DX

If regular expression *r* matches the entire string *s*, this function returns a list of strings representing the contents of each register, that is to say the characters matched by each section of the regular expression *r* delimited by parenthesis. This is useful for extracting specific segments of a string.

If the regular expression *r* does not match the string *s*, function `regex-extract` returns the empty list. If the regular expression *r* matches the string but does not contain parenthesis, this function returns a list containing the initial string *s*.

Example:

```
? (regex-extract "(+|-)?([0-9]+)(\\|\\. [0-9]*)?" "-56.23")
= ("-" "56" ".23")
```

(`regex-seek` *r s* [*start*])

Type: DX

Searchs the first substring in *s* that matches the regular expression *r*, starting at position *start* in *s*. If the argument *start* is not provided, string *s* is searched from the beginning.

If such a substring is found, `regex-seek` returns a list (*begin length*), where *begin* is the index of the first character of the substring, and *length* is the length of the subscript. The instruction (`mid s begin length`) may be used to extract this substring.

If no such substring exists, `regex-seek` returns the empty list.

Example:

```
? (regex-seek "(+|-)?[0-9]+(\\|\\. [0-9]*)?," "a=56.2, b=57,")
= (3 5)
```

(`regex-subst` *r s str*)

Type: DX

Replaces all substring matching regular expression *r* in string *str* by string *s*.

A “register” is associated to each piece of the regular expression *r* enclosed within parenthesis. Registers are numbered from %0 to %9. During each match, the substring of *str* matching each piece of the regular expression is stored into the corresponding register.

During the replacement process, characters %0 to %9 in the replacement string *s* are substituted the content of the corresponding register. (A single % is denoted as %%).

Example:

```
? (regex-subst "[a-h]([1-8])" "\\%1\\%0" "e2-e4, d7-d5, d2-d4, d5xd4?")
= "2e-4e, 7d-5d, 2d-4d, 5dx4d?"
```

(regex-rseek r s [n [gr]])

Type: DE(miscenv.lsh)

This function seeks recursively the first occurrence of **r** in **s** . and returns the list made of the locations.

When argument **n** is provided, it seeks and returns the locations of the **n** first occurrences and it returns **()** on failure.

Optional regex **gr** defines the allowed garbage stuff before and between occurrences. When **n** is not provided, this function checks the garbage stuff after the occurrences too. If unallowed garbage stuff is found, the function returns **()** . By default, any garbage stuff is allowed.

Since even void garbage is checked, a caret **"^"** is often added to **gr** .

(regex-split r s [n [gr [neg]]])

Type: DE(miscenv.lsh)

This function splits a string **s** into occurrences of **r** .

When integer **n** is provided, this function provides only the **n** first occurrences.

When regex **gr** is provided, garbage is checked (see function **regex-rseek**).

When **neg** is provided and non nil, this function returns the garbage stuff instead. When both **n** and **neg** are provided and non nil, the **n** garbages before and between the **n** first occurrences are returned.

(regex-skip r s [n [gr [neg]]])

Type: DE(miscenv.lsh)

This function skips the **n** first occurrences of regex **r** in a string **s** .

When **n** is equal to 0, it returns **s** . When **n** is lower than 0, it generates an error. When **n** is either nil or undefined, it is set to 1.

When **neg** is either nil or undefined, it returns the right residual of **s** just following the **n** th occurrence.

When **neg** is not nil, it returns the right residual of **s** beginning with the **n** th occurrence.

When regex **gr** is provided, garbage is checked (see function **regex-rseek**).

(regex-count r s)

Type: DE(miscenv.lsh)

This function recursively seeks the occurrences of regex **r** in string **s** and returns the number of occurrences found.

(regex-tail r s [n [gr [neg]]])

Type: DE(miscenv.lsh)

This function seeks recursively the occurrences of regex **r** in string **s** .

When **neg** is either nil or undefined, it returns the right residual of **s** beginning before the **n** th last occurrence.

When **neg** is non nil, it returns the right residual of **s** beginning after the **n** th last occurrence (and thus beginning before the **n** th garbage.

When **n** is either nil or undefined, it is set to 1.

When regex **gr** is provided, garbage is checked (see function **regex-rseek**).

(regex-member r1 s)

Type: DE(miscenv.lsh)

This function returns the first member of list **r1** which is a matching regex for string **s** .

4.9 Storages

Storages are vectors of identical objects that are used as data storage areas for scalars, vectors, matrices, tensors, and similar structures.

- Data in a storage is rarely accessed directly, but often accessed through an **IDX** object. Several **IDX**s can point to the same data in a storage, allowing access of the data in multiple ways simultaneously.
- The data in a storage can reside in memory (a memory storage) or on disk (a memory-mapped storage). Disk files can be memory mapped into a storage and accessed element by element or byte by byte through storage or **IDX** access functions.
- Storages can be resized/reallocated without adverse effects on the **IDX** that point to them.

4.9.1 Storage Creation and Allocation

Each of the following functions has several names. For example , **int-storage** can also be invoked with **new_i32_storage** as well as **new-int-storage** . However , the names described below are preferred.

(float-storage [n])

Type: DE(storage.c)

Creates a storage of class **|FSTORAGE|** whose elements are 4-byte floats. If **n** is present, **n** elements are allocated in memory and initialized to 0 (or **NIL**). If **n** is not present, an unsized storage is returned that can be subsequently allocated

in memory with `storage_malloc`, `storage_malloc_nc`, or memory-mapped from a file with `storage_mmap`.

(double-storage [n])

Type: DE(storage.c)

Creates a storage of class `|DSTORAGE|` whose elements are double precision floats (8 bytes). If `n` is present, `n` elements are allocated in memory and initialized to 0 (or NIL). If `n` is not present, an unsized storage is returned that can be subsequently allocated in memory with `storage_malloc`, `storage_malloc_nc`, or memory-mapped from a file with `storage_mmap`.

(int-storage [n])

Type: DE(storage.c)

Creates a storage of class `|I32STORAGE|` whose elements are 4 byte signed integers. If `n` is present, `n` elements are allocated in memory and initialized to 0 (or NIL). If `n` is not present, an unsized storage is returned that can be subsequently allocated in memory with `storage_malloc`, `storage_malloc_nc`, or memory-mapped from a file with `storage_mmap`.

(short-storage [n])

Type: DE(storage.c)

Creates a storage of class `|I16STORAGE|` whose elements are 2 byte signed integers. If `n` is present, `n` elements are allocated in memory and initialized to 0 (or NIL). If `n` is not present, an unsized storage is returned that can be subsequently allocated in memory with `storage_malloc`, `storage_malloc_nc`, or memory-mapped from a file with `storage_mmap`.

(byte-storage [n])

Type: DE(storage.c)

Creates a storage of class `|I8STORAGE|` whose elements are single-byte signed integers. If `n` is present, `n` elements are allocated in memory and initialized to 0 (or NIL). If `n` is not present, an unsized storage is returned that can be subsequently allocated in memory with `storage_malloc`, `storage_malloc_nc`, or memory-mapped from a file with `storage_mmap`.

(ubyte-storage [n])

Type: DE(storage.c)

Creates a storage of class `|U8STORAGE|` whose elements are single-byte unsigned integers. If `n` is present, `n` elements are allocated in memory and initialized to 0 (or NIL). If `n` is not present, an unsized storage is returned that can be subsequently allocated in memory with `storage_malloc`, `storage_malloc_nc`, or memory-mapped from a file with `storage_mmap`.

(gptr-storage [n])

Type: DE(storage.c)

Creates a storage of class |GPTRSTORAGE| whose elements are generic memory pointers. If *n* is present, *n* elements are allocated in memory and initialized to 0 (or NIL). If *n* is not present, an unsized storage is returned that can be subsequently allocated in memory with `storage_malloc`, `storage_malloc_nc`, or memory-mapped from a file with `storage_mmap`.

(atom-storage [n])

Type: DE(storage.c)

Creates a storage of class |ATSTORAGE| whose elements are Lisp objects. If *n* is present, *n* elements are allocated in memory and initialized to 0 (or NIL). If *n* is not present, an unsized storage is returned that can be subsequently allocated in memory with `storage_malloc`, `storage_malloc_nc`, or memory-mapped from a file with `storage_mmap`.

(packed-storage [n])

Type: DE(storage.c)

Creates a storage of class |PSTORAGE| whose elements are packed fixed point numbers (one-byte signed numbers with 4 bit integer part and 4 bit fractional part). If *n* is present, *n* elements are allocated in memory and initialized to 0 (or NIL). If *n* is not present, an unsized storage is returned that can be subsequently allocated in memory with `storage_malloc`, `storage_malloc_nc`, or memory-mapped from a file with `storage_mmap`. This type of storage is somewhat obsolete, but is there for backward compatibility with SN2.x.

4.9.2 Storage allocation**(storage-malloc srg n)**

Type: DE(storage.c)

Set the size of storage *srg* to *n*, allocate *n* elements in memory, and clear the content (to 0 or NIL depending on the type of storage). Storage *srg* must be unsized (as returned by `new-xxx-storage`).

(storage-malloc-nc srg n)

Type: DE(storage.c)

Set the size of storage *srg* to *n*, allocate *n* elements in memory. The elements are not cleared to save time. Storage *srg* must be unsized (as returned by `new-xxx-storage`).

(storage-realloc srg n)

Type: DE(storage.c)

Change the size of storage **srg** to **n** , and reallocate **n** elements in memory. **IDX**s that points to **srg** are unaffected (i.e. they point to the right place and their content data is unchanged). The newly allocated data segment is initialize to 0 or NIL.

(storage-realloc-nc srg n)

Type: DE(storage.c)

Change the size of storage **srg** to **n** , and reallocate **n** elements in memory. **IDX**s that points to **srg** are unaffected (i.e. they point to the right place and their content data is unchanged). The newly allocated data segment is NOT initialized.

(storage-mmap srg file offset)

Type: DE(storage.c)

maps the content of file **file** into memory. **file** can be either a file descriptor or a string containing a filename. Each element of the file starting at byte **offset** is made accessible through the corresponding element of the storage **srg** . Storage of Lisp objects (`|ATSTORAGE|`) cannot be memory mapped. Mapped storages are read-only. Whether a storage is read-only can be determined with the function `writablep` .

4.9.3 Storage access

Storage elements can be read and set as if the storage were an **IDX** with one dimension (a vector). In other words `(s 3)` returns the value of element 3 of storage **s** . `(s 3 5)` sets it to 5.

4.9.4 Miscellaneous storage functions**(storage-save srg file)**

Type: DX

save the content of storage **srg** into file **file** . **file** may be a string containing a filename or a file descriptor. Data is written raw , without a header , and in the native format of the machine (multibyte elements will not be portable between big and small-endian machines).

(storage-load srg file)

Type: DX

load storage **srg** with the content of file **file** . **file** may be a string containing a filename or a file descriptor.

(storagep srg)

Type: DX

return true if `srg` is a storage , and nil otherwise.**(writablep srg)**

Type: DX

return true if `srg` is writable , and nil otherwise. Memory mapped storages are generally not writable.**(storage-read-only srg)**

Type: DX

marks storage `srg` as being read-only. No Lisp function can write into it after this.**(storage-size srg)**

Type: DX

return the number of element in storage `srg` . Return 0 if the storage is unsized.**(storage-clear srg)**

Type: DX

Fills up all the element of storage `srg` with 0 or NIL.

4.10 Matrices, Vectors, Tensors, Indexes and IDX

Lush has a powerful mechanism for manipulating tabular data such as scalars, vectors, matrices, and tensors from 0 to 8 dimensions.

The basic Lush type for tabular data is called an INDEX or an IDX. The IDX is merely an access structure, the actual data is stored in a STORAGE associated with the IDX.

The data contained in a storage, and accessed with an IDX, can be of type double, float, int, short, byte, ubyte, gptr (generic pointer), atom (lisp object), and packed (one-byte fixed point number).

There are many IDX-related functions in Lush. Some are described in this section, but the vast majority is described in the Standard Libraries section.

creating a tensor of float is done simply with:

```
? (setq m (float-matrix 10 8 4)) ; create 3D matrix
= ::INDEX3:<10x8x4>
? (m 3 4 2 45.6) ; set value of element (3,4,2) to 45.6
```

```
= ::INDEX3:<10x8x4>
? (m 3 4 2)           ; get value of element (3,4,2).
= 45.6
```

Tensors of various basic types can be created with the functions listed below. Each function has two versions, the regular version initializes all the elements to zero, while the version with `-nc` at the end do not (no clear) and are a bit faster. All of these functions take 0 to 8 integer arguments that are the sizes in each dimension:

- `double-matrix`, `double-matrix-nc`, `real-matrix`, `real-matrix-nc` : doubles
- `float-matrix`, `float-matrix-nc`, `flt-matrix`, `flt-matrix-nc` : floats
- `int-matrix`, `int-matrix-nc` : 32-bit ints
- `short-matrix`, `short-matrix-nc` : 16-bit shorts
- `byte-matrix`, `byte-matrix-nc` : 8-bit bytes
- `ubyte-matrix`, `ubyte-matrix-nc` : 8-bit unsigned bytes
- `gptr-matrix`, `gptr-matrix-nc` : generic pointers (void*)

Tables of lisp objects (atoms) are created with the `array` function.

4.10.1 IDX Structure

An IDX is a structure that contains a pointer to the STORAGE where the data resides, and offset indicating where the IDX data starts in the storage, a number of dimensions (between 0 and 8), a size for each dimension, and a modulo for each dimension. The modulo determines by how much a pointer would have to be incremented to go from one element of the IDX to the next element in the given dimension.

Multiple IDXs can point to the same STORAGE and provide access to the same data in multiple ways.

an IDX with 0 dimension is called a scalar, and can be seen as a pointer to a number.

4.10.2 IDX Creation and Allocation

Tensors of various basic types can be created with the functions listed below. Each function has two versions, the regular version initializes all the elements to zero, while the version with `-nc` at the end do not (no clear) and are a bit faster. All of these functions take 0 to 8 integer arguments that are the sizes in each dimension:

(double-matrix s1...sn)

Type: DX

create an n -dimensional tensor of `double`. The arguments are the sizes in each dimension. The elements are initialized to 0.

(double-matrix-nc s1...sn)

Type: DX

create an n -dimensional tensor of `double`. The arguments are the sizes in each dimension. The elements are not initialized and may contain non-sensical values.

(float-matrix s1...sn)

Type: DX

create an n -dimensional tensor of `float`. The arguments are the sizes in each dimension. The elements are initialized to 0.

(float-matrix-nc s1...sn)

Type: DX

create an n -dimensional tensor of `float`. The arguments are the sizes in each dimension. The elements are not initialized and may contain non-sensical values.

(int-matrix s1...sn)

Type: DX

create an n -dimensional tensor of 32-bit `int`. The arguments are the sizes in each dimension. The elements are initialized to 0.

(int-matrix-nc> s1...sn)

create an n -dimensional tensor of 32-bit `int`. The arguments are the sizes in each dimension. The elements are not initialized

(short-matrix s1...sn)

Type: DX

create an n -dimensional tensor of 16-bit `short`. The arguments are the sizes in each dimension. The elements are initialized to 0.

(short-matrix-nc> s1...sn)

create an n -dimensional tensor of 16-bit `short`. The arguments are the sizes in each dimension. The elements are not initialized

(byte-matrix s1...sn)

Type: DX

create an n -dimensional tensor of signed **byte** . The arguments are the sizes in each dimension. The elements are initialized to 0.

(byte-matrix-nc> s1...sn)

create an n -dimensional tensor of signed **byte** . The arguments are the sizes in each dimension. The elements are not initialized

(ubyte-matrix s1...sn)

Type: DX

create an n -dimensional tensor of unsigned **byte** . The arguments are the sizes in each dimension. The elements are initialized to 0.

(ubyte-matrix-nc> s1...sn)

create an n -dimensional tensor of unsigned **byte** . The arguments are the sizes in each dimension. The elements are not initialized

(gptr-matrix s1...sn)

Type: DX

create an n -dimensional tensor of generic pointers (void *). The arguments are the sizes in each dimension. The elements are initialized to NULL.

(gptr-matrix-nc> s1...sn)

create an n -dimensional tensor of generic pointers (void *). The arguments are the sizes in each dimension. The elements are not initialized and may cause memory faults if dereferenced carelessly.

(array m s1...sn)

Type: DX

create an n -dimensional tensor lisp objects. The arguments are the sizes in each dimension. The elements are initially nil.

(new-index srg slist)

Type: DX

Creates a new **IDX** pointing to the storage **srg** . **slist** is a list of integers that specifies the size in each of the dimensions. The **srg** is allocated and resized if necessary.

4.10.3 **IDX Literals**

The IDX literal syntax can be used to create an IDX and fill it up with values in one fell swoop. Here is how to create a vector of doubles:

```
(setq m [3 4 5])
```

A two-dimensional matrix of doubles can be created as follows:

```
(setq m [[1 2 3] [3 4 5]])
```

Vectors and matrices of other types than doubles can be created by specifying the type with a single character right after the first open bracket. Here is how a vector of floats can be created:

```
(setq m [f 1 2 3])
```

Similarly, IDX of int, short, bytes, ubytes, and atoms can be created with the characters *i*, *s*, *b*, *u*, *a* respectively. Here is another example:

```
(setq m [a "choucroute" "garnie"])
```

scalar literals

Scalars literals can be entered with the scalar macros described below. Here is an example:

```
(idx-dotm0 vector (new-scalarf 34) output)
```

```

(new-scalar value)
creates a new scalar of double and fills it with value )
(new-scalarf value)
creates a new scalar of float and fills it with value )
(new-scalari value)
creates a new scalar of int and fills it with value )
(new-scalars value)
creates a new scalar of short and fills it with value )
(new-scalarb value)
creates a new scalar of byte and fills it with value )
(new-scalaru value)
creates a new scalar of ubyte and fills it with value )
(new-scalara value)
creates a new scalar of atom and fills it with value )
(new-scalarp value)
creates a new scalar of gptr and fills it with value )

```

4.10.4 IDX Copying and Cloning

(idx-copy m1 [m2])

Type: DX

Copy `m1` into `m2` . If `m1` and `m2` are of different numerical types , appropriate conversion will be performed. `m1` and `m2` must have the same dimensions. If `m2` is not present , a fresh copy of `m1` is returned (the data is copied).

(copy-matrix m1 [m2])

Type: DX

Copy `m1` into `m2` . Unlike `idx-copy` , this function does not require that `m1` and `m2` have the same structure as long as they have the same number of elements. In other words , `m1` can be a 4 by 5 matrix , and `m2` a 20-dimensional vector.

(idx-clone m)

Type: DX

Return a clone of `IDX m` . The returned `IDX` is identical to `m` and points to the same storage at the same place. This does not actually copy any data.

(copy-any-matrix m1 [m2])

Type: DX

obsolete function. Works like `copy-matrix` (sort of). If `m2` is omitted return a new copy of `m1` (a new storage is created and `m1` is copied in it). If `m2` is present , the storage of `m1` is copied into the storage of `m2` . Return `m2` .

4.10.5 IDX Iterators

(idx-bloop ((symb1 idx1) [(symb2 idx2) [...(symbn idxn)]]) body)

Type: DY

make each `symbi` be an `idx` that loops over the first dimension of its corresponding `idxi` . Execute `body` for each value.

More precisely, each `si` will be an `idx` with one less dimension than the corresponding `ixi` , and will simultaneously loop over the successive "slices" of `ixi` for each possible value of the first index. In other words, applying function `myfunc` to each element of a vector `v1` and putting the result in the corresponding element in `v2` can be done with:

```
(idx-bloop ((x1 v1) (x2 v2)) (x2 (myfunc (x1))))
```

`x1` and `x2` are scalars (i.e. zero-dimensional tensors). The above function work just as well is `v1` and `v2` are `n` -dimensional tensors and `myfunc` accepts `n-1` -th dimensional tensors as arguments.

```
(idx-pbloop ((symb1 idx1) [ (symb2> idx2) [...(symbn idxn) ]]) body)
```

Works just like `idx-bloop` , but the multiple evaluations of the body may be executed in parallel on some implementations of Lush.

```
(idx-eloop ((symb1 idx1) [ (symb2> idx2) [...(symbn idxn) ]]) body)
```

Type: DY

Make each `symbi` be an `idx` that loops over the last dimension of its corresponding `idxi` . Execute `body` for each value.

This is like `idx-bloop`, but it loops on the last dimension, instead of the first. For example, the matrix product operation $C = A*B$ can be written as follows:

```
(de idx-m2timesm2 (A B C)
  (idx-eloop ((Bj B)(Cj C)) (idx-m2dotm1 A Bj Cj)))
```

where `idx-m2dotm1` is the usual matrix-vector product. The `idx-eloop` construct simultaneously iterates over all columns of `B` and `C` .

```
(idx-peloop ((symb1 idx1) [ (symb2> idx2) [...(symbn idxn) ]]) body)
```

Works just like `idx-pbloop` , but the multiple evaluations of the body may be executed in parallel on some implementations of Lush.

```
(cidx-bloop (i_1 [i_2...i_n] (c_1 l_1) [(c_1 l_1)...(c_m l_m)] p_1 [p_2...])
```

Type: DF

This iterator is designed to facilitate the implementation of inner loops of tensor functions in C, while leaving all the bookkeeping to the Lisp.

A call to `cidx-bloop` as shown in the synopsis is somewhat equivalent to `n` nested `idx-bloop` s, which will loop over the first `n` dimensions of `idxs l_1` to `l_m` simultaneously. The arguments `i_1` to `i_n` are strings containing names of C local variables that will be created and set to the loop index in each of the `n` dimensions. At each iteration, the C variables provided in strings `c_1` to `c_m` will point to the appropriate values in the `idxs l_1` to `l_m` . For example, the following function will fill matrix `a` with `cos(i+j)`.

```
(de foo (a)
  ((-idx2- (-flt-)) a)
  (cidx-bloop ("i" "j" ("a" a)) #{ *a = cos(i+j); #} a)
```

The return value is (like in `idx-bloop`) the last `IDX` specified in the declaration (in the example above, the return value is superfluous).

(idx-gloop (p1...[pn]) body)

Type: DM

an enhanced version of `bloop` which allows to "bloop" through a pointer table , and to have access to the current index value of the loop. Each `pi` is a list with 1, 2, or 3 elements. if it has 2 elements the meaning is like a in regular `bloop`. if it has 1 element , which must be a symbol , it will be used as a local variable which contains the current index of the loop. if it has 3 elements. It must be of the form (`symbol p m`) , where `p` is a 1D index and `m` an index of at least one dimension. `symbol` will take the values (`idx-select m 0 (p i)`) for all possible values of `i`.

4.10.6 IDX Information

(indexp m)

Type: DX

Return true if `m` is an IDX.

(matrixp m)

Type: DX

Return true if `m` is an IDX whose storage is of a numerical or generic pointer type , as opposed to lisp objects.

(arrayp m)

Type: DX

return true if `m` is an IDX over a storage of lisp objects (as opposed to numerical types). For historical reasons , this kind of IDX is also known as an array.

(idx-storage m)

Type: DX

return the storage on which IDX `m` points.

(idx-size m)

Type: DX

return the memory size occupied by the elements of `m` in bytes.

(idx-nelements m)

Type: DX

return the number of elements of IDX `m` . This function is not compilable (no reason , we just forgot to write the macros).

(idx-ndim m)

Type: DX

return the number of dimensions of `IDX m` .**(idx-offset m)**

Type: DX

return the offset of the first element of `IDX m` in its storage. The offset is in number of elements , not bytes.**(idx-dim m [n])**

Type: DX

With two arguments, return the size of `IDX m` in its `n` -th dimension (starting at zero). With one argument, return the size of `IDX m` in each dimension in the form of a list of integers.

Example:

```
? (idx-dim [[2 3 4] [5 6 7]])
= (2 3)
```

(idx-bound m)

Type: DX

Return the largest index in each dimension in the form of a list. This function is mostly here for historical reasons.

(idx-modulo m [n])

Type: DX

With two arguments: return the number of elements that separate two elements that have successive indices in the `n` -th dimension (starting at zero). With one argument , return a list of all the modulus in all the dimensions.**(contiguousp m)**

Type: DE

returns true if the elements of `m` are contiguous in memory. `IDX` resulting from `idx-transclone` or `narrow` are generally non-contiguous.**(idx-ptr m)**

Type: DX

return a `gptr` that points to the first element of `m` . This must be used with caution because the data pointed to by the pointer may be garbage collected (when `m` is destroyed) , or even displaced (when `m` is resized). `idx-ptr` is primarily used when calling C function that require a pointer to numerical data.

4.10.7 IDX Manipulations

Cloning manipulations

These functions take an IDX as argument and return a modified clone of this IDX pointing to the same data (or a subset of it).

(`idx-transclone`)

(`select m n s`) Type: DM

Return a clone of IDX `m` (on the same storage), with the `n`-th dimension removed, and which is the `s`-th "slice" of `m`, in the `n`-th dimension.

```
? (setq m [[0 1 2 3 4][10 11 12 13 14]])
= [[ 0.00  1.00  2.00  3.00  4.00 ]
   [10.00 11.00 12.00 13.00 14.00 ]]
? (select m 1 2)
= [ 2.00 12.00 ]
```

(`narrow m n s [o]`) Type: DM

Make a clone of IDX `m`, and reduce its size in the `n`-th dimension to `s` elements, offset by `o`.

```
? (setq m [[0 1 2 3 4][10 11 12 13 14]])
= [[ 0.00  1.00  2.00  3.00  4.00 ]
   [10.00 11.00 12.00 13.00 14.00 ]]
? (narrow m 1 3 2)
= [[ 2.00  3.00  4.00 ]
   [12.00 13.00 14.00 ]]
```

(`unfold m n ksize step`) Type: DM

Return an idx on the same storage as `m` (pointing to the same data) with an added dimension at the end obtained by "unfolding" the `n`-th dimension. The size of the new dimension is `ksize`. This essentially manipulates the modulus to make convolutions look like matrix-vector multiplies. For example, a one-dimensional convolution between vector `v` and kernel `k` can be done as follows:

```
? (setq v [0 1 2 3 4 5 6])
= [ 0.00  1.00  2.00  3.00  4.00  5.00  6.00 ]
? (setq k [1 2 1])
= [ 1.00  2.00  1.00 ]
? (setq uv (unfold v 0 3 1))
= [[ 0.00  1.00  2.00 ]
   [ 1.00  2.00  3.00 ]
   [ 2.00  3.00  4.00 ]
   [ 3.00  4.00  5.00 ]
   [ 4.00  5.00  6.00 ]]
```

```

? (setq r (matrix 5))
= [ 0.00 0.00 0.00 0.00 0.00 ]
? (idx-m2dotm1 uv k r)
= [ 4.00 8.00 12.00 16.00 20.00 ]

```

A subsampled convolution (where the kernel is stepped by more than one element) can be performed by setting the `step` parameter to a value other than 1:

```

? (setq uv (unfold v 0 3 2))
= [[ 0.00 1.00 2.00 ]
   [ 2.00 3.00 4.00 ]
   [ 4.00 5.00 6.00 ]]
? (setq r (matrix 3))
= [ 0.00 0.00 0.00 ]
? (idx-m2dotm1 uv k r)
= [ 4.00 12.00 20.00 ]

```

Here are other amusing examples:

```

? (unfold [3 4 5 6 7] 0 5 1)
= [[ 3.00 4.00 5.00 6.00 7.00 ]]
? (unfold [3 4] 0 1 1)
= [[ 3.00 ]
   [ 4.00 ]]

```

Naturally, there is no real need for most programmers to use the `unfold` construct directly because the standard library contains efficient predefined 1D and 2D convolutions.

(idx-transclone m dimlist) Type: DX

return a clone of `idx m` where the dimensions have been permuted according to the list of dimension indices `dimlist` . For example:

```
(idx-transclone m '(0 2 1))
```

permutes the second and third dimensions.

(transpose m) Type: DE

Simple transpose function for 1D vectors and 2D matrices. This does not copy any data , but simply returns a new `IDX` on the same data with the appropriate structure.

(diagonal m d) Type: DM

return a clone of `m` where the modulus and dimensions have been hacked so as to grab the diagonal of the last `d` dimensions. Those dimensions must all have the same size. The result has `d-1` less dimensions than the original `IDX`.

```
? (setq m [[[111 112 113][121 122 123][131 132 133]]
           [[211 212 213][221 222 223][231 232 233]])
= [[[[111.00 112.00 113.00 ]
     [121.00 122.00 123.00 ]
     [131.00 132.00 133.00 ]]
    [[211.00 212.00 213.00 ]
     [221.00 222.00 223.00 ]
     [231.00 232.00 233.00 ]]]
? (diagonal m 2)
= [[111.00 122.00 133.00 ]
   [211.00 222.00 233.00 ]]
```

Compilable Direct IDX manipulations

(idx-changedim m n v) Type: DX

change the size of the n -th dimension of m to v . This generates an error if the resulting IDX overflows its storage.

```
? (setq m (matrix 3 4))
= [[ 0.00  0.00  0.00  0.00 ]
   [ 0.00  0.00  0.00  0.00 ]
   [ 0.00  0.00  0.00  0.00 ]]
? (m () () (range 0 11))
= [[ 0.00  1.00  2.00  3.00 ]
   [ 4.00  5.00  6.00  7.00 ]
   [ 8.00  9.00 10.00 11.00 ]]
? (idx-changedim m 1 3)
= ()
? m
= [[ 0.00  1.00  2.00 ]
   [ 4.00  5.00  6.00 ]
   [ 8.00  9.00 10.00 ]]
```

(idx-changemod m n v) Type: DX

change modulo of n -th dimension of m to v .

(idx-changeoffset m n) Type: DX

change offset of m in its storage to n

Non-compilable Direct IDX manipulations

These functions directly hack into the IDX passed as argument. Because of those bad side effects, they cannot be called in compiled code.

(idx-redim m dimlist) Type: DX

Redimensions IDX m to the list of dimensions passed as argument.

(idx-undim m) Type: DX

Turn `m` into an un-dimensioned `IDX`.

(idx-unfold m n ksize step) Type: DX

same as `unfold` , but hacks `m` directly instead of cloning it.

(idx-diagonal m d) Type: DX

same as `diagonal` , but hacks `m` directly instead of cloning it.

(idx-narrow m n s [o])>

Works like `narrow` , but hack `m` "in place" instead of returning a hacked clone.

(idx-select m n s) Type: DX

same as `select` , but hacks `m` directly instead of cloning it.

(idx-transpose m perm-list) Type: DX

Permute the dimensions of `idx` using the list of int's `perm-list` as a permutation matrix.

(idx-transpose2 idx d1 d2) Type: DX

permute dimesnions `d1` and `d2` if `IDX idx` .

4.10.8 Loading/Saving

There are two ways to load and save matrices and tensors in Lush which deal with two incompatible formats. The functions listed below use the so-called "classic" format , but cannot be used in compiled code. The other set of functions , described in the standard library section uses a different format (so-called `IDX` format) and are based on the C `stdio` library. This discrepancy exists largely for historical reasons and will be fixed eventually.

Loading, Saving, and Mapping Matrices.

There are two types of "classic" matrix files: ASCII files and binary files. Binary files are a more accurate , more compact , and more efficient way of saving matrix data. ASCII matrix files should be used only reserved to

(save-matrix mat file) Type: DX

Stores the matrix `mat` into the binary file `file` . Argument `file` may be a filename string or a file descriptor created with `open-write` or `open-append` . When `file` is a file name , a suffix ".mat" is added when needed. [Binary Matrix File Format.]

(save-ascii-matrix mat file) Type: DX

Stores the matrix `mat` into the text file `file` . Argument `file` may be a filename string or a file descriptor created with `open-write` or `open-append` . When `file` is a file name , a suffix ".mat" is added when needed. [Ascii Matrix File Format.]

(load-matrix [mat] file) Type: DX

Loads a matrix from file `file` and returns it. This function senses the file format and performs the adequate actions. If a symbol `mat` is specified , it stores the new matrix in symbol `mat` .

Argument `file` may be a filename string or a file descriptor created with `open-read`. When `file` is a file name, a suffix `".mat"` is added when needed. [Binary Matrix File Format.][Ascii Matrix File Format.]

(`map-matrix` [`mat`] `filename`) Type: DX

This function is available on certain systems only. It loads a binary matrix file by mapping the file to certain memory addresses using the virtual memory facilities provided by the operating system. This function is especially handy for handling very large matrices. Mapped matrices however are always read-only.

Native Matrix File Formats

This section describes the "classic" or "native" file formats used for Lush matrices. The IDX format (used by compilable matrix I/O functions) is not described here.

Ascii Matrix File Format.

ASCII matrix files are generated with function `save-ascii-matrix`. These machine independent files are handy for transferring data between different computers. They tend however to eat a lot of disk space.

The first line of the file is a matrix header. It is composed of the letters ".MAT", the number of dimensions and the size of each dimension. The array elements are written then, separated by spaces and newlines characters.

The following lines display the content of a valid ascii matrix file.

```
.MAT 2 3 4
 1 -1 1 2 -1 1 -3
 1 -1
 4 0 0
```

The same matrix could have been written in the following, more natural way.

```
.MAT 2 3 4
 1 -1 1
 2 -1 1
-3 1 -1
 4 0 0
```

Binary Matrix File Format.

Binary matrix files are generated with function `save-matrix`. Binary matrix files begin with a header which describes the type and the size of the matrix. Then comes a binary image of the matrix.

The core header is a C structure defined as follows :

```
struct header {
    int magic;
```

```

        int ndim;
        int dim[3];
};

```

It can be followed by further `int` when the matrix has more than 3 dimensions.

The first member, `magic`, of this structure is a “magic” number, which encodes the type of the matrix. This number must be:

- 0x1E3D4C51 for a single precision matrix
- 0x1E3D4C53 for a double precision matrix
- 0x1E3D4C54 for an integer matrix
- 0x1E3D4C56 for a short matrix
- 0x1E3D4C55 for a byte matrix
- 0x1E3D4C52 for a packed matrix

The second member of this structure, `ndim`, is the number of dimensions of the matrix. Then come the dimensions themselves in the array `dim`. If there are more than three dimensions, the array `dim` is extended to accommodate the extra dimensions.

When the number of dimensions (specified in `ndim`) is greater than 3, the header is completed by `ndim - 3` further integers.

This header is followed by a binary image of the matrix. Elements are stored with the last index changing faster, i.e.

(0,0,0) (0,0,1) (0,0,2) ... (0,1,0) (0,1,2) etc...

- as float numbers for single precision matrix
- as double numbers for double precision matrix
- as int numbers for integer matrix
- as short numbers for short matrix
- as unsigned char numbers for byte matrix
- and as char for packed matrix.

In this latter case, each byte represents a fixed point number between -8 and +8 (+8 not included), the first 4 bits (most significant nybble) contain the integral part and the remaining 4 bits contain the fractional part. It uses a two’s complement format. Here are two C functions which convert a “compact fixed point” number into a floating point number and back.

```

/* Converts a packed number to a float */
float unpack(b)
int b;
{
    if (b & 0x80)
        b |= ~0x7f;
    else
        b &= 0x7f;
    return (float)b / 16.0;
}
/* Converts a float into a single byte packed number */
unsigned char pack(x)
float x;
{
    if (x > 8.0-1.0/16.0)
        return 0x7f;
    else if (x < -8.0)
        return 0x80;
    else
        return x*16;
}

```

Foreign Matrix Files.

Most programs write their results in specific files under a specific format. These files always have specific headers which describe the file type and size. The matrix file formats of TLisp are just a typical example. The following functions are for reading and writing information stored in files created by other programs. These functions can read binary or text files into a Lush matrix.

Foreign Binary Matrices.

(import-raw-matrix mat file [offset]) Type: DX

This function just reads the contents of file `file` into matrix `mat`. Accessing the matrix then access this data as single precision numbers (float), double precision numbers (double), integers (int), short integers (short), bytes (unsigned char) or packed numbers according to the type of matrix `mat`.

Argument `mat` must be a matrix freshly created with functions `matrix`, `dmatrix`, `imatrix`, `smatrix`, `bmatrix` or `pmatrix`. Submatrices are usually rejected. Argument `file` might be a file name or a file descriptor created with `open-read`. When argument `offset` is specified, a header of offset bytes is skipped.

After executing this function, the file descriptor `file` points to the first byte following the matrix data.

(export-raw-matrix mat file) Type: DX

Stores the data of the matrix `mat` into the binary file `file`. Argument `file` may be a filename string or a file descriptor created with `open-write` or

`open-append` . No header is written.

Foreign Ascii Matrices.

(import-text-matrix mat file) Type: DX

This function reads numbers in the text file `file` and stores them in matrix `mat` . Argument `mat` must be a freshly created matrix (submatrices are not accepted). Argument `file` might be or a file descriptor created with `open-read` .

After executing this function, the file descriptor `file` points to the first non blank character following the matrix data.

(export-text-matrix mat file) Type: DX

Stores the data of the matrix `mat` into the text file `file` . Argument `file` may be a filename string or a file descriptor created with `open-write` or `open-append` . No header is stored.

4.10.9 Component-wise Unary Operations

All of these functions apply a unary function to each element and either write the result in the elements of the second argument if it is present , or return a new tensor with the result if the second argument is not present (except for `idx-clear`). The two argument must have the same dimensions. If their numerical types are different , appropriate conversions are performed.

(idx-clear src)

Type: DX

Set elements of `src` to 0.

(idx-minus src [dst])

Type: DX

negate all elements of `src` .

(idx-abs src [dst])

Type: DX

absolute value of elements of `src` .

(idx-sqrt src [dst])

Type: DX

square root of elements of `src` .

(idx-inv src [dst])

Type: DX

inverse of elements of `src` .

(idx-sin src [dst])

Type: DX

apply sine to elements of **src** .

(idx-cos src [dst])

Type: DX

apply cosine to elements of **src** .

(idx-atan src [dst])

Type: DX

apply arctangent to elements of **src** .

(idx-log src [dst])

Type: DX

apply log to elements of **src** .

(idx-exp src [dst])

Type: DX

apply exp to elements of **src** .

(idx-qtanh src [dst])

Type: DX

apply rational approximation to hyperbolic tangent to elements of **src** .

(idx-qdtanh src [dst])

Type: DX

apply derivative of the rational approximation to hyperbolic tangent to elements of **src** .

(idx-stdsigmoid src [dst])

Type: DX

apply the "standard" neural-net sigmoid function to elements of **src** .

(idx-dstdsigmoid src [dst])

Type: DX

apply derivative of the "standard" neural-net sigmoid function to elements of **src** .

(idx-expmx src [dst])

Type: DX

apply a rational approximation of $(\exp -x)$ to elements of **src** .**(idx-dexpmx src [dst])**

Type: DX

apply derivative of **expmx** to elements of **src** .

4.10.10 Component-wise dyadic operations

All of these functions apply a dyadic function to each pair of corresponding elements in the first two arguments. They either write the result in the elements of the third argument if it is present , or return a new tensor with the result if the third argument is not present. All the arguments must have the same dimensions. If their numerical types are different , appropriate conversions are performed.

(idx-add m1 m2 [r])

Type: DX

component-wise addition of **m1** and **m2** . Result in **r** if present , or returned if not present.**(idx-sub m1 m2 [r])**

Type: DX

component-wise subtraction of **m1** and **m2** . Result in **r** if present , or returned if not present.**(idx-mul m1 m2 [r])**

Type: DX

component-wise multiplication of **m1** and **m2** . Result in **r** if present , or returned if not present.**(idx-div m1 m2 [r])**

Type: DX

component-wise division of **m1** and **m2** . Result in **r** if present , or returned if not present.

4.10.11 Contracting Operations with Scalar Result.

The following functions include dot products , distances , sums of terms , min , max , etc and return scalars. These operations "contract" all the dimensions. The dyadic ones can be seen as generalized dot product of two tensors (e.g. the sum of all the products of corresponding terms in the two tensors). These function have a base form and an accumulating form (which accumulates the result in `idx0` past as last argument). If the last (and optional) `idx0` argument is present , the result is written in it. If it is not present , a number (not an `idx0`) is returned.

(idx-sum m [r])

Type: DX

sum of all terms of `m` .

(idx-sup m [r])

Type: DX

max of all terms of `m` .

(idx-inf m [r])

Type: DX

min of all terms of `m` .

(idx-sumsq m [r])

Type: DX

sum of squares of all terms of `m` .

(idx-dot m1 m2 [r])

Type: DX

generalized dot product of `m1` and `m2` , i.e. the sum of all products of corresponding terms in `m1` and `m2` .

(idx-sqrdist m1 m2 [r])

Type: DX

generalized Uclidean distance between `m1` and `m2` , i.e. the sum of squares of all the differences between corresponding terms of `m1` and `m2` .

(idx-sumacc m r)

Type: DX

sum of terms of `m` . Result accumulated in `idx0 r` .

(idx-supacc m r)

Type: DX

max of terms of m . Result accumulated in $\text{idx0 } r$.**(idx-infacc m r)**

Type: DX

min of terms of m . Result accumulated in $\text{idx0 } r$.**(idx-sumsq racc m r)**

Type: DX

sum square of terms of m . Result accumulated in $\text{idx0 } r$.**(idx-dotacc m1 m2 r)**

Type: DX

generalized dot product of $m1$ and $m2$, i.e. the sum of all products of corresponding terms in $m1$ and $m2$. Result accumulated in $\text{idx0 } r$.**(idx-sqrdistacc m1 m2 r)**

Type: DX

generalized Uclidean distance between $m1$ and $m2$, i.e. the sum of squares of all the differences between corresponding terms of $m1$ and $m2$. Result accumulated in $\text{idx0 } r$.

4.10.12 Operations between tensors and scalars

(idx-dotm0 m s [r])

Type: DX

multiply all terms of m by scalar s (an idx0).**(idx-addm0 m s [r])**

Type: DX

add scalar s (an idx0) to all terms of m .**(idx-dotm0acc m s r)**

Type: DX

multiply all terms of m by scalar s (an idx0). Result accumulated in r .**(idx-addm0acc m s r)**

Type: DX

add scalar s (an idx0) to all terms of m . Result accumulated in r .

4.10.13 Matrix/Vector and 4-Tensor/Matrix products**(idx-m2dotm1 m1 m2 [r])**

Type: DX

matrix-vector multiply.

(idx-m4dotm2 m1 m2 [r])

Type: DX

4-tensor by 2-matrix multiplication: $R_{ij} = \text{sum}_{kl} M1_{ijkl} * M2_{kl}$ **(idx-m2dotm1acc m1 m2 r)**

Type: DX

matrix-vector multiply. Result accumulated in r .**(idx-m4dotm2acc m1 m2 r)**

Type: DX

4-tensor by 2-matrix multiplication with accumulation: $R_{ij} += \text{sum}_{kl} M1_{ijkl} M2_{kl}$ **4.10.14 Outer Products****(idx-m1extm1 m1 m2 [r])**

Type: DX

outer product between vectors: $R_{ij} = M1_i * M_j$ **(idx-m2extm2 m1 m2 [r])**

Type: DX

outer product between matrices. Gives a 4-tensor: $R_{ijkl} = M1_{ij} * M2_{kl}$ **(idx-m1extm1acc m1 m2 r)**outer product between vectors with accumulation: $R_{ij} += M1_i * M_j$ **(idx-m2extm2acc m1 m2 [r])**

Type: DX

outer product between matrices with accumulation. Gives a 4-tensor: $R_{ijkl} += M1_{ij} * M2_{kl}$ **4.10.15 Backward Compatibility Functions**

These functions are here for historical reasons.

(matrix [n0...np])

Type: DX

same as (float-matrix [n0 ... np]): allocate a matrix of floats.

(submatrix m '(low0 hi0) '(low1 hi1)...)

Type: DX

this is superseded by narrow. Return a new index that points to a subset of the elements of m . The lists passed as arguments contain lower and upper index of the selected elements in each dimension.**(bound m)**

Type: DX

Return a list of upper bound for the index in each dimension.

(nrvectorp m)

Type: DX

Obsolete function.

(nrmatrixp m)

Type: DX

Obsolete function.

(dot-product a b)

Type: DM

return the dot product of a and b **(m*m a b [c])**

Type: DE

obsolete and non-compilable matrix product function.

(m+m a b [c])

Type: DE

obsolete and non-compilable matrix addition function.

(mxm a b [c])

Type: DE

obsolete and non-compilable component-wise matrix multiply function.

(m-m a b [c])

Type: DE

obsolete and non-compilable matrix subtraction function.

(m*c a b [c])

Type: DE

obsolete and non-compilable multiplication of a matrix by a scalar.

(m+c a b [c])

Type: DE

obsolete and non-compilable addition of a scalar to the elements of a matrix.

4.11 Objects.

All atoms are stored as pairs (class , instance). The “class” component provides informations about the general properties of that kind of atom. The “instance” component contains the properties of this specific atom in particular. Classes themselves are unnamed atoms , stored as value of certain symbols. For example , classes are instances of the class `class` , i.e. of the class stored as value of the symbol `class` .

4.11.1 Object Terminology.

A class defines two major kind of informations:

- An atom can contain storage for one or several lisp objects. Such storage areas are called “slots”, and have symbolic names. A class defines the number of slots allocated in a new atom as well as their names. This slot information is statically defined by the `defclass` function during the class definition.
- Atoms can receive messages sent with the function `==>`. Messages are identified by a symbol called the message selector. A class defines the possible message selectors and the functions to execute when it receives such messages. These “methods” are dynamically defined with the `putmethod` function.

Predefined classes associated to standard Lush objects however often have hidden properties. A symbol, for example, has some storage for its value; this storage is not a slot. A symbol also returns its value when it is evaluated; this behavior is not a method.

4.11.2 Inheritance.

New classes defined with the `defclass` function always are subclasses of another class. This latter superclass, may again have its own superclasses.

- “A class inherits the slots of its superclasses”. When a new object is created, space is allocated for the slots defined by its class, and also for the slots defined by all its superclasses.
- “A class inherits the methods of its superclasses”. When an object receives a message with a given selector, it searches the methods defined by its class for the selector. If no method is found that matches the selector, it searches the methods of its superclass, the methods of the superclass of its superclass, and so on. When it reaches a matching method, it executes the associated function like an instance of the class that defines the matching method.

4.11.3 Predefined Classes

Here is a compact list of the main predefined classes. Most of them are named by uppercase symbols, which ought to be surrounded with bars for avoiding the usual lowercase conversion.

- `object` : The root class of most user defined classes.
- `class` : Classes actually are lisp objects, instance of the `class` class. Class `class` thus is an instance of itself !
- `|NUMBER|` : The class of the numbers.
- `|LIST|` : The class of the pairs used for building lists.
- `|SYMB|` : The class of the symbols.
- `|STRING|` : The class of the strings.
- `|DE| |DF| |DM| |DX| |DY|` : The classes of the functions.
- etc...

Hidden properties of complex predefined classes are also inherited. For example, `setq` shall work on any instance of a subclass of class `|SYMB|` .

All the user defined classes must be direct or indirect subclasses of the class `object` . This `object` class defines a vanilla lisp object with no hidden properties. Instances of subclasses of the class `object` are easily created with the function `new` .

4.11.4 Defining a Class.

```
(defclass name superclass s1 ... sn)
```

Type: DM

Defines a subclass of class `superclass` with slots `s1 ... sn` . The class is stored into symbol `name` , which is returned. Slot specifications `s1 ... sn` can be symbols or lists (`symbol default`) which indicate initial values assigned to the slots of a new object.

Example:

```
;;; Creates a class <rect> with two slots <width> and <height>
? (defclass rect object
    width height )
= rect
;;; Creates a subclass of <rect>,
;;; with a an additional slot <name> slot with default value <"noname">
? (defclass namedrect rect
    (name "noname") )
= namedrect
```

```
(makeclass classname superclass slotnamelist slotdefaultlist)
```

Type: DX

This a lower-level function to create a new class named `classname` . Most user will prefer to use `defclass` . The new class inherits class `superclass` . List `slotnamelist` contains the names of the additional slots defined by the new class. List `slotdefaultlist` contains the default values for these slots. Both lists are in reverse order (i.e. the first slot in the list appears last when pretty-printing the class). This function does not set the value of symbol `classname` . Use the macro function `defclass` for that.

```
(classname c)
```

Type: DX

Returns the name of class `c` , that is to say the symbol whose value was set by function `defclass` .

Example:

```
? (classname rect)
= rect
```

```
(slots c)
```

Type: DX

Returns the list of the slots defined for class `c` .

Example:

```
? (slots rect)
= (width height)
? (slots namedrect)
= ((name "noname"))
```

(methods c)

Type: DX

Returns the list of all method names defined in class *c* .

(super c)

Type: DX

Returns the superclass of class *c* , or the empty list if *c* has no superclass.

Example:

```
? (super rect)
= ::class:object
? (classname (super namedrect))
= rect
```

(subclasses c)

Type: DX

Returns the list of all subclasses of class *c* . When argument *c* is the empty list, function `subclasses` returns the list of root classes, i.e. classes without superclass.

Example:

```
? (subclasses object)
= (::class:timer ::class:eventhandler ::class:eventlock ::class:brace-renderer
  ::class:ldoc ::class:ltree-node ::class:help-book )
```

(rsubclasses c)

(miscenv.lsh)

Returns the recursive list of all subclasses of class *c* . When argument *c* is the empty list, function `subclasses` returns the recursive list of all classes.

Example:

```
? (rsubclasses button)
((::class:stdbutton (::class:tinybutton))
 (::class:checkbox (::class:radiobutton)) (::class:filereqbutton)
 (::class:c-iconbutton) (::class:c-logo) )
= ()
```

4.11.5 Creating Objects.

(new class ... args ...)

Type: DY

Creates a new instance of a subclass `class` of the class `object` .

If no constructor method is defined, no arguments are allowed. If a constructor method is defined for class `class` , a construction message is sent to the new object, with arguments `args` .

Example:

```
;; Classes <rect>, <namedrect>, are defined as examples for <defclass>
;; Creates instances of classes <rect> and <namedrect>
? (setq r (new rect))
= ::rect:06fa0
? (setq nr (new namedrect))
= ::namedrect:06fe8
```

This function works only for creating instances of subclass of the class `object` . It is not possible to directly create instances of subclasses of other predefined classes.

(new-empty class)

Type: DX

Creates a new instance of class `class` without calling the constructor.

(delete x)

Type: DX

Deletes the Lush object `x` .

The operation of this function is somewhat tricky because all other references to object `x` must be converted to references to the empty list.

First of all , if destructor methods are defined , destruction messages are sent. Object `x` is then converted to an instance of the class `|ZOMBIE|` . During its normal operation , the interpreter recognizes these zombies and always replaces a reference to a zombie by a reference to the empty list.

(classof x)

Type: DX

Returns the class of an object `x` .

Example:

```
? (classof (new object))
= ::class:object
```

(is-of-class x class)

Type: DX

Tests if the class of object `x` or any of its superclass is equal to class `class`.

Example:

```
? (is-of-class nr rect)
= t
```

4.11.6 Accessing Slots.

There are various ways to read or change the values stored in the slots of an object. Changing the value stored in a slot changes the object in a way comparable to `rplaca` and `rplacd` in the case of lists. This change will be reflected through all pointers referring to this same object.

See: `(== n1 n2)`

:obj:slot1...:slotn

Type: DMC(sysenv.lsh)

See: `(scope obj slot1 ... slotn)`

See: `(scope symb)`

This macro character is expanded as a call to the `scope` function.

(scope obj slot1 ... slotn)

Type: DY

See: `: obj : slot1 ...: slotn`

See: `(scope symb)`

The simplest method for setting or getting the slots of an object is called the “scope macro”. The syntax `:obj:slot` refers to the slot `slot` of object `obj`. Actually, the scope macro character converts expression `:obj:slot` into a list `(scope obj slot)`.

When this list is evaluated, the `scope` function returns the value of slot `slot` of object `obj`. Moreover, most functions that affect the value of a symbol recognize such a list as a reference to slot `slot` of object `obj`.

Nested scope macros are allowed: `:obj:slot1:slot2` refers to slot `slot2` of the object contained in slot `slot1` of object `obj`.

Example:

```
; ; Classes <rect>, <namedrect>, are defined as examples for <defclass>
; ; Objects <r> and <nr> are defined as examples for <new>
? (setq :nr:name "myrectangle")
= "myrectangle"
? (setq :nr:width 8
```

```

      :nr:height 6 )
= 6
? (print :nr:width :nr:height :nr:name)
8 6 "myrectangle"
= "myrectangle"

```

(scope symb)

Type: DY

See: (scope obj slot1 ... slotn)

See: : obj : slot1 ...: slotn

The scope macro has another important use, the “unary scope macro”. Expression `:symb` is converted into list `(scope symb)` which refers to the global value of symbol `symb` .

Example:

```

? (setq x 3)
= 3
? (let ((x 8))
  (print :x)
  (setq :x 6)
  (print x) )
3
8
= 8
? x
= 6

```

(letslot obj l1 ... ln)

Type: DY

See: (scope obj slot1 ... slotn)

See: this

Calls function `progn` on `l1 ... ln` within the scope of object `obj` , and returns the last result. Within the object scope, each slot of the object, either defined by its class or inherited from its superclasses, can be directly accessed as value of its symbolic slot name.

Example:

```

;; Classes <rect>, <namedrect>, are defined as examples for <defclass>
;; Objects <r> and <nr> are defined as examples for <new>
? (setq width ())
= ()
? (letslot r
  (setq width 4)

```

```

      (setq height 5) )
= 5
? width
= ()
? (letslot r (* width height))
= 20

```

Within an object scope, the symbol `this` always refers to object itself. The behavior of the others symbols remains unaffected, with the following restriction:

The function `setq` causes an error if you create a new global variables from an object scope. This feature prevents many obscure bugs in Lush programs, that may occur if you mistype the name of a slot. However, the warning is not issued if you explicitly address a global variable by using the unary scope macro.

Good practice actually consists in explicitly create all global variables with the unary scope macro before using them, as the "`graphenv.lsh`" does.

4.11.7 Defining Methods.

```
(defmethod class symb args . body)
```

Type: DM(sysenv.lsh)

See: Argument List.

See: (`==>` obj symb ... args ...)

Defines a method named `symb` for class `class` . Argument `symb` must be a symbol, and is not evaluated. Argument `args` must be a valid argument list.

When an instance of class `class` receives a message whose selector is `symb` , the body `body` of the method is executed with a call to function `progn` . During this execution, the slots defined by `class` and its superclasses are bound to their symbolic names, as in `letslot` , and the arguments of the message are evaluated and bound to the symbols in argument list `args` .

Example:

```

;; Classes <rect>, <namedrect>, are defined as examples for <defclass>
;; Objects <r> and <nr> are defined as examples for <new>
;;
;; This is a method for computing the surface of a rect.
? (defmethod rect surface()
  (* width height) )
= surface
;; a method for printing n times the name of a namedrect
? (defmethod namedrect showname(n)
  (repeat n (print name)) )
= showname

```

(demethod class symb args . body)

Type: DM(sysenv.lsh)

Identical to `defmethod` .

(dfmethod class symb args . body)

Type: DM(sysenv.lsh)

Defines an flambda-method named `symb` for class `class` . flambda-methods do not evaluate their arguments, unlike regular methods defined with `demethod` or `defmethod`.

(dmmethod class symb args . body)

Type: DM(sysenv.lsh)

Defines a mlambda-method named `symb` for class `class` . mlambda-methods are to regular methods what macros (defined with `dm` or `mlambda`) are to regular functions (defined with `de` or `lambda`).

(putmethod class symb func)

Type: DX

Add method `symb` to class `class` . `func` must be a lambda, an flambda, or an mlambda.

(pretty-method class symb)

Type: DE(sysenv.lsh)

Prints a nicely indented definition of method `symb` of class `class` .

4.11.8 Sending Messages.

(==> obj symb ... args ...)

Type: MSG

Sends a message whose selector is `symb` to object `obj` , with arguments `args`

.

Example

```
;; Classes <rect>, <namedrect>, are defined as examples for <defclass>
;; Objects <r> and <nr> are defined as examples for <new>
;; Method <surface> is defined as an example for <defmethod>
;;
;; Send a surface message to rect <r>
? (==> r surface)
= 20
;; Send a showname message to namedrect <nr>
```

```

? (==> nr showname 2)
"myrectangle"
"myrectangle"
= "myrectangle"
;; Send a surface message to namedrect <nr>
;; The method is inherited from superclass <rect>
? (==> nr surface)
= 48

```

When an object receives a message, the method is executed in the object scope defined by the class that owns the method: If the method is defined by a superclass, only the slots of that superclass and of its superclasses are directly referred by their names.

When no method `symbol` is defined in the class or superclasses, the kernel searches a method named `-unknown` and executes this method with two arguments: the initial method name `sybm` and the list of the evaluated arguments `args`. There is an error if method `-unknown` is not defined.

```
(==> obj (class . symb) ... args ... )
```

Type: MSG

This construction is known as a “cast-and-send”. A message with selector `symb` and arguments `args` is sent to the object `obj`, considered as an instance of class `class`. Methods indeed are searched in class `class` and its superclasses, instead of the class of `obj` and its superclasses.

Argument `class` is evaluated and must return a superclass of the class of `obj`. Argument `symb` is not evaluated and must be a valid method selector for class `class`.

Example:

```

;; Classes <rect>, <namedrect>, are defined as examples for <defclass>
;; Objects <r> and <nr> are defined as examples for <new>
;; Method <surface> is defined as an example for <defmethod>
;;
;; Override surface method of class rect
? (defmethod namedrect surface()
  ;; prints the name by sending a showname message
  (==> this showname 1)
  ;; returns 1+ the result of the surface method of class rect
  (1+ (==> this (rect . surface)))) )
= surface
? (==> nr surface)
"myrectangle"
= 49
;; Still call the surface method of class rect
? (==> nr (rect . surface))

```

= 48

this

See: (letslot obj l1 ... ln)

See: (==> obj symb ... args ...)

While Lush is executing a method, you can use symbol **this** to refer to the object receiving the message.

(**check==>** class symb)

Type: DX

Searches class **class** and its superclasses for a method whose selector is **symb** . If a matching method exists, this function returns the method as a DF function.

Example:

```
;; Classes <rect>, <namedrect>, are defined as examples for <defclass>
;; Objects <r> and <nr> are defined as examples for <new>
;; Method <surface> is defined as an example for <defmethod>
;;
;; There is no method yoyo in rect
? (check==> rect 'yoyo)
= ()
;; there is a method surface in rect
? (check==> rect 'surface)
= DF:0604c
```

(**apply==>** obj symbol args)

Type: DX

Sends message **symbol** to object **obj** with the list **args** as arguments.

4.11.9 Constructors and Destructors.

See: (new class ... args ...)

See: (delete x)

Constructors are specific methods whose selector is the name of their class, as returned by the function **classname** . Constructors are useful for easily initializing an object during its creation.

Whenever a new object is created with the **new** function, a construction message is sent to the new object. The optional arguments of the **new** function are sent as arguments to the construction message.

Example:

```

;; Classes <rect>, <namedrect>, are defined as examples for <defclass>
;; Method <surface> is defined as an example for <defmethod>
;;
;; Define a constructor for class rect
? (defmethod rect rect(w h)
   (setq width w)
   (setq height h) )
= rect
;; creates a new rectangle
? (setq r (new rect 8 9))
= ::rect:06fb8
;; sends a surface message
? (==> r surface)
= 72

```

The function `new` only calls a constructor when it is defined by the class of the newly created object. Constructors however can call the constructor method of their superclass.

Example:

```

;; defines a constructor for class namedrect
? (defmethod namedrect namedrect(w h n)
   ;; explicitly calls the constructor of class rect.
   (==> this rect w h)
   (setq name n) )

```

Destructors are specific methods whose selector is `-destructor`. When an object is deleted, all the destructors of its class and superclasses are successively called. Deletion occurs when the function `delete` is called on an object, or when the garbage collection mechanism proves that an object will no longer be accessed.

Note: Only objects created with function `new` receive destruction messages. Numbers, strings, symbols, usually do not receive such messages.

Example:

```

;; Classes <rect>, <namedrect>, are defined as examples for <defclass>
;; Objects <r> and <nr> are defined as examples for <new>
;; Method <surface> is defined as an example for <defmethod>
;;
;; Defines a destructor for the rect class
? (defmethod rect -destructor()
   (printf "rect destructor on \\%l\\n" this) )
= -destructor
;; Defines a destructor for the namedrect class

```

```

? (defmethod namedrect -destructor()
    (printf "namedrect destructor on \\%l (\\%s)\\n" this name) )
= -destructor
;; deletes objects r and nr
? (delete r)
rect destructor on ::rect:06fb8
= ()
? r
= ()
? (delete nr)
namedrect destructor on ::namedrect:06fe8 (myrectangle)
rect destructor on ::namedrect:06fe8
= ()
? nr
= ()

```

Destructors allow objects to perform some housekeeping before they disappear.

4.11.10 Predefined Methods.

(==> obj pname)

Type: MSG

See: (pname 1)

The function `pname` returns a string that describes a lisp object. When executing this function, Lush searches the object for a method `pname` before applying the hardcoded defaults. Overriding method `pname` let you redefine how Lush prints certain objects.

Example:

```

;; Classes <rect>, <namedrect>, are defined as examples for <defclass>
;; Objects <r> and <nr> are defined as examples for <new>
;; Method <surface> is defined as an example for <defmethod>
;;
;; defines a pname function for the rect class
? (defmethod rect pname()
    (sprintf "<rectangle \\%d by \\%d>" width height) )
= pname
? (new rect 8 6)
= <rectangle 8 by 6>

```

(==> obj print)

Type: MSG

See: (prin l1 ... l1n)

When Lush prints an object with function `prin` or `print` , it searches a method `print` before applying hardcoded defaults based on function `pname` .

It is a good idea to override method `print` instead of method `pname` when large amounts of text are necessary (in the case of matrices, for example).

(==> obj `pprint`)

Type: MSG

See: (`pprin l1 ... ln`)

When Lush prints an atom with function `pprin` or `pprint` , it searches a method `pprint` before applying the hardcoded defaults.

(==> obj `pretty`)

Type: MSG

See: (`pretty f`)

Function `pretty` displays an object by sending them a message `pretty` . The default method `pretty` defined by class `object` displays the slots of the object. Specific `pretty` method also are defined for functions, symbols, classes, etc ...

4.12 Hash Tables.

Hash tables provide a powerful way to maintain associations between lisp objects. Hash tables behave like arrays except that any lisp object is an acceptable subscript. This simplicity makes hash tables extremely convenient.

Hash tables are so convenient that you will be quickly tempted to use them everywhere. You must remember however that a simple array requires less memory (about five times) and less processor cycles.

4.12.1 Representing Associations with Hash Tables.

See: (`= n1 n2`)

See: (`== n1 n2`)

A hash table is an instance of class `|HTABLE|` . Like most lisp objects, hash tables can be compared using function `=` and saved using function `bwrite` .

A hash table maintains a set of associations between two lisp objects, namely a key and a value. This set is organized in a way that allows fast retrieval of the value associated to a given key.

Hash table queries search an association whose key is equal to the queried key. Since there are two ways to test equality (i.e. functions `=` and `==`), there are two kinds of hash tables.

•

- The most useful hash tables use the logical equality, as implemented by function `=`. Two objects are logically equal if they convey the same useful information. These hash table work properly as long as you do not modify the objects used as keys by the hash table. This can happen if you used arrays or custom objects as keys in your hash table. See the description of function `htable-rehash` for more information about this point.
- Hash tables can also rely on pointer equality, as implemented by function `==`. These hash tables are useful to associate a piece of information with a particular lisp object. Functions `getp` and `putp`, for instance, make use of pointer-equality hash tables. The associations of a pointer-equality hash table are automatically removed when the key object is deallocated.

`(htable [nelem [flag]])`

Type: DX

Function `htable` returns a new empty hash table. Optional arguments `nelem` and `flag` are seldom useful:

- Argument `nelem` specifies an estimate of the number of elements in the hash table. This estimate is useful for saving a few processor cycles when filling the table.
- Argument `flag` is a boolean value specifying whether the hash table will use logical equality (the default) or pointer equality (when argument `flag` is not nil.)

`(htable key)`

This expression returns the value associated with key `key` in hash table `htable`. Expression `key` can be any lisp object. If no association is defined for key `key`, the empty list is returned.

`(htable key value)`

This expression associates value `value` to key `key` in the hash table `htable`. The value can be subsequently retrieved by expression `(htable key)`.

This operation is performed in two logical steps: (a) any existing association matching key `key` is first deleted, and (b) a new association is created if argument `value` is not the empty list.

Example:

```
? (setq a (htable))
= ::HTABLE:6f3cac
? (a 'hello "hi")
= ::HTABLE:6f3cac
? (a '(3 4) "ho")
```

```

= ::HTABLE:6f3cac
? (a 'hello)
= "hi"
? (a '(3 4 5))
= ()
? (a '(3 4))
= "ho"

```

(htable-size htable)

Type: DX

Function `htable-size` returns the number of associations managed by the hash table.

(htable-keys htable)

Type: DX

Function `htable-keys` returns a list of the keys of all the associations managed by the hash table. It is guaranteed that querying these keys will return a non nil result.

```

? (setq a (htable))
= ::HTABLE:6f43c
? (for (i 0 10) (a i (sqrt i)))
= ::HTABLE:6f43c
? (htable-keys a)
= (7 6 3 10 5 9 8 4 2 0 1)

```

(htable-alist htable)

Type: DX

See: `(assoc key alist)`

Function `htable-alist` returns an alist representing all associations managed by the hash table. This function also provide a way to iterate over all associations:

```

? (setq a (htable))
= ::HTABLE:6f43c
? (for (i 0 10) (a i (sqrt i)))
= ::HTABLE:6f43c
? (each (((key . value) (htable-alist a)))
        (printf "\\%4d \\%6.4f\\n" key value))))
7 2.6458
6 2.4495
3 1.7321

```

```

10 3.1623
 5 2.2361
 9 3.0000
 8 2.8284
 4 2.0000
 2 1.4142
 0 0.0000
 1 1.0000
= ()

```

(htable-rehash htable)

Type: DX

Each association managed by a hash table contains a pointer to the key object and a pointer to the value object. Functions that modify the values of these objects also modify the state of the hash table.

Hash table queries search an association whose key is logically equal to the queried key (unless pointer equality was specified when creating the hash table). The hash table works by storing each association in locations depending on the information conveyed by the key objects.

Modifying the object used as a key creates a lot of problems. The corresponding association is no longer stored in a location corresponding to the new information conveyed by the key objects. The hash table state is no longer consistent.

Function `htable-rehash` triggers the relocation of all associations of the hash table `htable` to the location corresponding to the new values of the key objects. You can then reliably use the hash table again.

```

? (setq a (htable))
= ::HTABLE:6f32c
? (setq m [1 2])
= [1 2]
? (a m 3)
= ::HTABLE:6f32c
? (a [1 2])
= 3
? (m 0 0)
= [ 0.00 2.00 ]
? (a [1 2])
= () ;; we lost it
? (a [0 2])
= () ;; result is undefined
? (htable-rehash a)
= 1
? (a [0 2])

```

```

= 3                                    ;; entry is no longer lost
? (a [1 2])
= ()                                    ;; this key no longer exists

```

The relocation of the hash table entries actually happens when you first access the hash table after calling `htable-rehash` . You can call function `htable-rehash` several times in a row. The cost of relocating all the hash table entries will be only incurred once when you will access the hash table again.

(htable-info htable)

Type: DX

Function `htable-info` returns a list containing statistical information about the hash table. This list is an alist with the following form:

```
((size . 1294) (buckets . 3841) (hits 1132 145 14 3))
```

In this example, the hash table contains 1294 associations. Associations are stored into 3841 buckets. Out of the 1294 associations, 1132 can be retrieved directly, 145 require one extra search iteration, 14 require two extra search iteration, etc...

4.12.2 Representing Sets with Hash Tables.

Sets are easily represented by creating a hash table and defining associations whose keys are the set elements and whose value is non nil.

- Expression `(htable element)` returns a non nil result if element `element` belongs to the set represented by `htable` .
- Expression `(htable element t)` adds element `element` to the set represented by `htable` .
- Expression `(htable element ())` removes element `element` from the set represented by `htable` .
- Expression `(htable-keys htable)` returns a list of all the elements of the set represented by `htable` .

The following functions perform additional operations on sets represented by hash tables.

(hset eltlist)

Type: DE(sysenv.lsh)

Function `hset` returns a hash table representing the set of all elements of list `eltlist` .

(hset-and htable1 htable2)

Type: DE(sysenv.lsh)

Function **hset-and** returns a new hash table representing the intersection of the sets represented by htables **htable1** and **htable2** .

```
? (htable-keys (hset-and (hset '(1 2 3 4)) (hset '(2 4 5 6))))
= (4 2)
```

(hset-or htable1 htable2)

Type: DE(sysenv.lsh)

Function **hset-or** returns a new hash table representing the intersection of the sets represented by htables **htable1** and **htable2** .

```
? (htable-keys (hset-or (hset '(1 2 3 4)) (hset '(2 4 5 6))))
= (4 2 1 6 3 5)
```

4.13 Dates.

Dates are instances of a special class `[DATE]` . This facility is built on top of various system functions which are themselves based on a count of seconds since January 1st, 1970. Most systems are able to handle a date between 1910 and 2030.

Date are stored using several formats depending of their accuracy. Each format is defined as a range over various components of a date: **year** , **month** , **day** , **hour** , **minute** , **second** .

- For instance, a date “**year to day** ” identifies any single day between 1910 and 2030.
- Similarly, a date “**year to second** ” identifies any single second in the same interval.
- A date “**hour to minute** ” identifies any single minute within a day.

Finally we have defined a standard way to numerically encode dates: The possibly fractional number of days (daynumber) since January 1st, 1970 at zero hour local.

4.13.1 (date-to-day date)

Type: DX

Returns the possibly fractional number of days between January 1st, 1970 and the specified **date** .

```
? (date-to-day (string-to-date "1993-12-04" 'year 'day))
= 8738
```

```
? (date-to-day (now))
= 12059.8523
```

4.13.2 (day-to-date daynumber)

Type: DX

Builds a date string from a possibly fractional day number.

```
? (day-to-date 8738)
= ::DATE:year-second:(1993-12-04 00:00:00)
```

```
? (day-to-date 8738.23)
= ::DATE:year-second:(1993-12-04 05:31:11)
```

4.13.3 (today)

Type: DX

Returns today's date at zero hour as a date "year to day "

```
? (today)
= ::DATE:year-day:(2003-01-07)
```

4.13.4 (now)

Type: DX

Returns today's date and now's time as a date "year to second "

```
? (now)
= ::DATE:year-second:(2003-01-07 20:27:15)
```

4.13.5 (split-date date)

Type: DX

Returns an alist containing all the components of a date `date` . Each component can be accessed using function `assoc` . The following components are provided:

```

<year>:      The year (minus 1900).
<month>:     The month number (0..11).
<day>:       The day of the month (1..31).
<hour>:      The hour (0..23).
<minute>:    The number of minutes (0..59).
<second>:    The number of seconds.
<wday>:      The day of the week (0..6).
<yday>:      The day of the year (1..366).

```

These items are present if the resolution of the date object allows it. The last two items require at least a date “year to day ”

```

? (split-date (now))
= ((year . 103) (month . 1) (day . 7) (hour . 20) (minute . 27) (second . 15)
   (wday . 2) (yday . 6) )

```

```

? (split-date (today))
= ((year . 103) (month . 1) (day . 7) (wday . 2) (yday . 6))

```

4.13.6 (date-type date)

Type: DX

This function returns a string describing the current resolution of a date `date` as a list of two symbols chosen in `year` , `month` , `day` , `hour` , `minute` , `second` .

```

? (date-type (today))
= (year day)

```

```

? (date-type (now))
= (year second)

```

4.13.7 (date-extend date from to)

Type: DX

This function returns a new date equivalent to `date` with the resolution specified by the two symbols `from` and `to` which may be any of `year` , `month` , `day` , `hour` , `minute` , `second` .

Reducing the resolution of a date is achieved by truncating unneeded components. Extending the resolution of a date towards more accuracy is achieved by inserting zeroes. Extending the resolution of a date towards fields of larger magnitude is achieved by inserting the current year, month, day, etc...

```
? (date-extend (now) 'year 'minute)
= ::DATE:year-minute:(2003-01-07 20:27)
```

```
? (date-extend (date-extend (now) 'year 'minute) 'hour 'second)
= ::DATE:hour-second:(20:27:00)
```

4.13.8 (date-to-string date [format])

Type: DX

Format a date `date` according to `format` .

Argument `format` is a format string reminiscent of `printf` . When argument `format` is omitted, a regular ANSI string is produced.

YYYY-MM-DD HH:MM:SS	for <year> to <second>
YYYY-MM-DD	for <year> to <day>
MM-DD HH:MM	for <month> to <minute>

The format for `date-to-string` is a character string that consists of field descriptors and text characters, reminiscent of `printf` . Each field descriptor consists of a % character followed by another character that specifies the replacement for the field descriptor. All other characters are copied from `fmt` into the result. The following field descriptors are supported:

- %%: Same as %.
- %a: Day of week, using locale's abbreviated weekday names.
- %A: Day of week, using locale's full weekday names.
- %b,%h: Month, using locale's abbreviated month names.
- %B: Month, using locale's full month names.
- %c: Date and time as %x %X.
- %C: Date and time, in locale's long-format date and time representation.
- %d: Day of month (01-31).
- %D: Date as %m/%d/%y.
- %e: Day of month (1-31; single digits are preceded by a blank).
- %H: Hour (00-23).
- %I: Hour (00-12).
- %j: Day number of year (001-366).

- %k: Hour (0-23; single digits are preceded by a blank).
- %l: Hour (1-12; single digits are preceded by a blank).
- %m: Month number (01-12).
- %M: Minute (00-59).
- %n: Same as \n.
- %p: Locale's equivalent of AM or PM.
- %r: Time as %I:%M:%S %p.
- %R: Time as %H:%M.
- %S: Seconds (00-59).
- %t: Same as \t.
- %T: Time as %H:%M:%S.
- %U: Week number of year (01-52), Sunday is the first day of the week.
- %w: Day of week; Sunday is day 0.
- %W: Week number of year (01-52), Monday is the first day of the week.
- %x: Date, using locale's date format.
- %X: Time, using locale's time format.
- %y: Year within century (00-99).
- %Y: Year, including century (for example, 1988).

The difference between %U and %W lies in which day is counted as the first day of the week. Week number 01 is the first week with four or more January days in it.

```
? (date-to-string (today))
= "2003-01-07"
```

```
? (date-to-string (now))
= "2003-01-07 20:27:15"
```

```
? (date-to-string (now) "\\%C \\%X")
= "20 20:27:15"
```

4.13.9 (string-to-date string [from to])

Type: DX

Function `string-to-date` converts the ansi date string `string` to a date using the resolution specified by `from` and `to` .

Argument `from` and `to` are symbols specifying the resolution of the target date. They may be any of the following symbols: `year` , `month` , `day` , `hour` , `minute` , `second` . When these arguments are omitted, a date “ `year to second` ” is assumed.

Abbreviated dates are handled by filling the leftmost fields with the current date and the rightmost fields with zeroes.

```
? (string-to-date "1993-12-12 8:30" 'year 'minute)
= ::DATE:year-minute:(1993-12-12 08:30)
```

4.13.10 (date-add-second date count)

Type: DX

This function adds `count` seconds to date `date` .

```
? (date-add-second (now) 67)
= ::DATE:year-second:(2003-01-07 20:28:22)
```

4.13.11 (date-add-minute date count)

Type: DX

This function adds `count` minutes to date `date` .

```
? (date-add-minute (now) 67)
= ::DATE:year-second:(2003-01-07 21:34:15)
```

4.13.12 (date-add-hour date count)

Type: DX

This function adds `count` hours to date `date` .

```
? (date-add-hour (now) 67)
= ::DATE:year-second:(2003-01-10 15:27:15)
```

4.13.13 (date-add-day date count)

Type: DX

This function adds `count` days to date `date` .

```
? (date-add-day (now) 67)
= ::DATE:year-second:(2003-03-15 20:27:15)
```

4.13.14 (date-add-month date count [opt])

Type: DX

This function adds `count` months to date `date` .

This is tricky because months have different lengths. Adding an integral number of month to certain dates may produce an illegal day number for the target month. (eg. `(date-add-month "1996-05-31" 1)`). The behavior of this function is therefore further specified by the value of optional argument `opt` .

- When `opt` is nil or omitted, this function generates a Lush error whenever the answer would be an illegal date.
- When `opt` is `'no-error` , this function returns the empty list `()` whenever the answer would be an illegal date.
- When `opt` is `'end-of-month` , this function always returns the last day of the target month regardless of the specified day number.
- When `opt` takes any other value (eg. `'loosely` or `t`) this function always returns a legal date. Answers with oversized day numbers are corrected by returning the last day of the target month.

```
? (date-add-month (now) -2 t)
= ::DATE:year-second:(2002-11-07 20:27:15)
```

4.13.15 (date-add-year date count [opt])

Type: DX

This function adds `count` months to date `date` .

This is tricky because adding an integral number of years to february 29th dates may produce an illegal date. The behavior of this function is therefore further specified by the value of optional argument `opt` .

- When `opt` is nil or omitted, this function generates a Lush error whenever the answer would be an illegal date.
- When `opt` is `'no-error` , this function returns the empty list `()` whenever the answer would be an illegal date.

- When `opt` takes any other value (eg. `'loosely` or `t`) this function always returns a legal date. The day number is changed to 28 whenever the answer would be the february 29th during a non leap year.

```
? (date-add-year (now) 2 t)
= ::DATE:year-second:(2005-01-07 20:27:15)
```

4.13.16 (date-code date flags)

Type: DX

This function returns a list of numbers encoding a date `date` suitably for a statistical analysis. Argument `flag` is an integer composed by oring (using `bitor` for instance) the following constants controlling which values to generate:

- 1: One value representing the number of days since January 1st, 1970.
- 16: Two values encoding a point moving on the unit circle once a day.
- 8: Two values encoding a point moving on the unit circle once a week.
- 4: Two values encoding a point moving on the unit circle once a month.
- 2: Two values encoding a point moving on the unit circle once a year.

These components are generated only if they make sense with the resolution of the date `date` .

```
? (date-code (now) 24)
= (0.5992 -0.8006 -0.8357 0.5492)
```

```
? (date-code (today) 31)
= (12059 Nan Nan -0.2225 0.9749 0.3473 0.9378 0.9947 0.1031)
```

4.14 Loading and Executing Lush Program Files.

4.14.1 (load in [out [prompt]])

Type: DX

See: (open-write `s` [`suffix`])

See: (open-read `s` [`suffix`])

See: `^L filename`

This function is used for loading a Lush source or binary file.

This function starts a new toplevel with the file named `in` as input file, and the file named `out` as the output file. Strings `in` and `out` represent filenames according to the general filename conventions explained with functions `open-read` and `open-write` .

Function `load` searches its input file `in` along the current search path. It also tries various suffixes usually associated with Lush program files (ie. `.lsh` `.lshc` as well as `.sn` `.tl` `.snc` `.tlc` for backward compatibility with TL3 and SN).

Function `load` then enters a loop which reads a Lush expression on the input file, evaluates this expression, and prints the result on the output file. This loop terminates when an end-of-file condition is detected or when function `exit` is called. Function `load` returns the full filename of the input file.

When the output file argument is omitted `out`, function `load` does not display the evaluation results. This is useful for silently loading a Lush program file. This file however can produce an output because it may contain calls to some printing function (eg. `print`.)

When string `prompt` is provided, function `load` uses this string as a prompt when it is waiting for a new expression on the input file.

4.14.2 `^L filename`

See: (load in [out [prompt]])

This macro-character expands into a call to function `load`. If a filename `filename` is not provided, the last file name used with macro-character `^L` is used. This file name is saved in variable `last-loaded-file`.

4.14.3 (libload libname [opt])

Type: DE(sysenv.lsh)

See: file-being-loaded

See: (load in [out [prompt]])

This is the primary function used to load a library or a Lush file into another Lush file or library. A typical Lush program file will begin with a bunch of `libload`. The difference between `load` and `libload` is that `libload` remembers when a particular file was previously libloaded, and only loads it once. `libload` also keeps track of the dependencies between Lush program files (which file libloads which other file). This is the basis of the automatic and transparent on-demand compilation mechanism as described in detail below.

`libload` first searches the file `libname`. When `libload` is called from an existing Lush file, it first searches file `libname` relative to the location of the existing Lush file. If no such file exists, it uses function `filepath` to locate the file along the lush search path.

`libload` then updates its internal database of dependencies between Lush files. When function `libload` is called from an existing Lush file, it records the fact that this Lush file depends on the Lush file `filename`.

`libload` then tests whether it is necessary to load the file. It is necessary to load the file if

- file `libname` has never been loaded,
- file `libname` has been modified since the last time it was loaded,

- file `libname` has not been modified since the last time it was loaded, but depends on a Lush file that has been modified since the last time it was loaded.
- argument `opt` is non nil.

When this is the case, `libload` calls `load` to load the file, and stores a timestamp for the file in its internal database. Function `libload` always returns the full name of `libname` .

4.14.4 file-being-loaded

Type: VAR

See: (load in [out [prompt]])

This variable is set when you load a Lush program file using function `load` . It contains the full name of the file being loaded. It provides a way to know the full filename of the Lush program file in which this variable is checked.

4.14.5 (find-file dirlist filename extlist)

Type: DE(sysenv.lsh)

Finds the first file that exists fo the form `dir / file . ext` where `dir` is one of the directories in list `dirlist` and `ext` is one of the suffixes in `extlist` .

Example:

```
(find-file '("local" "lsh") "lushenv" '(".lshc" ".lsh"))
```

Note: This is somewhat similar to `filepath` but allows searching through any directory tree.

4.14.6 (save f s1 ... sn)

Type: DM(sysenv.lsh)

This function just saves indented definitions of the symbols `s1` to `sn` into the file named `f` . A suffix `".lsh"` is automatically added to `f` if it has no suffix.

4.14.7 (autoload f s1 ... sn)

Type: DM(sysenv.lsh)

Defines functions `s1` to `sn` with a code that first loads file `f` , and then restart the evaluation.

The first time one of these functions is called, its definition is loaded from file `f` , and the execution continues. Such autoloading functions often are defined by `"stdenv.lsh"` . They avoid loading many files while allowing the user to call those functions at any time.

4.15 Lisp-Style Input / Output.

4.15.1 Input.

(read)

Type: DX

See: The Lush Reader.

Calls the Lisp reader, and returns the next Lush object available on the current input. Function `read` is able to read both text and binary Lush files.

(read-string [spec])

Type: DX

Reads a string on the current input.

- When argument `spec` is a number, it indicates how many character should be read from the current input and returned as a string.
- When argument `spec` is a string, it specifies which characters are to be read. For example, "0-9" matches any number, " 0-9" matches anything but a number, "\t\n\r\f\e" anything but a blank character or end of file, "yYnN" reads only one of these characters.

The default `spec` for function `read-string` is "\n\r\f\e" and means "read anything until the end of line".

(skip-char [spec])

Type: DX

Skips the characters matching string `spec` and returns the next available character. The default value of argument `spec` is "\t\n\r\f" and means "skip any blank characters".

Example:

This function reads the input stream and builds a list with each line, until it reaches an end-of-file character

```
(de read-lines()
  (let ((ans ()))
    (while (<> (skip-char "\\n\\r\\f") "\\e")
      (setq ans (nconc1 ans (read-string)))) ) )
```

This function reads the input stream and builds a list with each word, until it reaches an end-of-file character

```
(de read-words()
  (let ((ans ()))
    (while (<> (skip-char "\\e")
      (setq ans (nconc1 ans (read-string "~ \\t\\n\\r\\f\\e")))) ) )
```

(ask s)

Type: DX

Prints the string `s` on the standard output and waits for a "yes" or "no" answer on the standard input. Function `ask` returns `t` or `()` according to the user's answer. Function `ask` is not affected by the current input or output file.

Example:

```
? (ask "Your answer")
Your answer [y/n] ?
```

4.15.2 Output.

(prin l1 ... ln)

Type: DX

Prints the Lush objects `l1` to `ln` separated by spaces. Unlike `print`, the function `prin` does not terminate the line with a linefeed. Function `prin` returns the result of the last evaluation.

(print l1 ... ln)

Type: DX

Prints the Lush objects `l1` to `ln` separated by spaces and output a newline character. Function `print` returns the result of the last evaluation.

Example:

```
? (print "hello" (+ 3 4) '(a b c d))
"hello" 7 (a b c d)
= (a b c d)
```

(pprin l1 ... ln)

Type: DE(sysenv.lsh)

Prints the Lush objects `l1` to `ln` with a nice indentation. The indentation is defined according many heuristics suitable for printing nicely function definitions. Unlike function `pprint`, `pprin` does not terminate the line with a linefeed.

(pprint l1 ... ln)

Type: DE(sysenv.lsh)

Prints the evaluations of `l1` to `ln` with a nice indentation. The indentation is defined according many heuristics suitable for printing nicely function definitions. Function `pprint` terminates the line with a linefeed.

Although the current implementation of `pretty` is a more complicated, a simple implementation of `pretty` would be :

```
(df pretty(f) (pprint (funcdef f)))
```

```
(printf format ... args ... )
```

Type: DX

This function is similar to the `printf` function in C language. String `format` will be printed verbatim, except for the following escape sequences, which refer to `args`.

- `"%"` is replaced by a single `\%`.
- `"%l"` is replaced by a representation of a lisp object.
- `"%[-] [<n]s">` is replaced by a string, right justified in a field of length `n` if `n` is specified. When the optional minus sign is present, the string is left justified.
- `"%[-] [<n]d">` is replaced by an integer, right justified in a field of `n` characters, if `n` is specified. When the optional minus sign is present, the string is left justified.
- `"%[-] [<n [. m]] c ">` where `c` is one of the characters `e`, `f` or `g`, is replaced by a floating point number in a `n` character field, with `m` digits after the decimal point. `e` specifies a format with an exponent, `f` specifies a format without an exponent, and `g` uses whichever format is more compact. When the optional minus sign is present, the string is left justified.

Function `printf` is the only way to print the real contents of a character string, without adding surrounding double quotes or escaping the special characters.

Example:

```
? (printf "\\%5s\\%3d is equal to \\%6.3f\\n" "sqrt" 2 (sqrt 2))
sqrt 2 is equal to 1.414
= ()
```

```
(tab [n])
```

Type: DX

When a numeric argument `n` is given, `tab` outputs blanks until the cursor reaches the `n`-th column on the current line. The cursor position is unchanged if the cursor is located after the `n`-th column. Function `tab` always returns the current cursor position.

```
? (tab)
= 0
```

4.15.3 File Names.

(basename path [suffix])

Type: DX

Returns the terminal component of the pathname `path` by removing any prefix ending with `/`. If string `suffix` is provided, a possible suffix `suffix` is removed from the pathname.

This function makes no reference to the actual contents of the filesystem. It just chops the filename according to the file naming conventions.

(dirname path)

Type: DX

Returns the name of the directory containing the file named `path`.

This function makes no reference to the actual contents of the filesystem. It just chops the filename according to the file naming conventions.

(concat-fname [filename1] filename2)

Type: DX

Returns an absolute filename equivalent to `filename2` parsed from current directory `filename1`. Argument `filename1` defaults to the current directory.

This function makes no reference to the actual contents of the filesystem. It just concatenates the filenames according to the file naming conventions. Nevertheless, if the resulting filename refers to an actual directory, a directory separator `/` is appended to the returned filename.

Example:

```
? (concat-fname "..")
= "/home/yann"
```

```
? (concat-fname (getenv "HOME") "myfile.lsh")
= "/home/yann/myfile.lsh"
```

(tmpname [[dir] suffix])

Type: DX

This function creates a temporary file and returns its name. When Lush exits, an automatic cleanup code removes these files.

The optional string `dir` specifies the directory where the temporary file will be created. A value of `()` is interpreted as the default system directory for temporary files.

This function generates file names using the optional suffix `suffix`. If you do not specify this argument, the generated filename will have no suffix.

Examples:

```
(tmpname)           ;; creates a temporary file.
(tmpname () "sn")   ;; creates a temporary file whose name ends with ".lsh".
(tmpname ".")       ;; creates a temporary file in the current directory.
```

4.15.4 Searched Path.

See: Lush Startup.

When Lush is looking for a file, it searches for it among a list of directories. Several functions are provided to handle this feature.

(path s1...sn)

Type: DX

The function `path` allows you to get and set this list. Without arguments, `path` returns the current path. Otherwise, the strings `s1` to `sn` become the current path.

The initial path is automatically set while Lush is starting up by retrieving the Lush executable file and scanning up the directories until it finds a suitable library directory.

```
? (progn
  (printf "--> The current path is:\n")
  (path) )
--> The current path is:
= ("/home/yann/lush/local" "/home/yann/lush/packages"
  "/home/yann/lush/lsh/libogre" "/home/yann/lush/lsh/compiler"
  "/home/yann/lush/lsh/libidx" "/home/yann/lush/lsh/libstd"
  "/home/yann/lush/lsh" "/home/yann/lush/sys" )
```

(addpath s)

Type: DE(sysenv.lsh)

See: (path s1 ... sn)

The function `addpath` adds the directory `s` to the file search path. Directory `s` is added at the head of the search path. Other occurrences of this directory in the path are removed, in order to keep the search path small and non redundant.

Example:

```
? (addpath (concat (getenv "HOME") "/mysnlib"))
```

(filepath filename [suffixes])

Type: DX

See: (path s1 ... sn)

See: (concat-fname [filename1] filename2)

See: (load in [out [prompt]])

The function `filepath` searches file `filename` along the current search path, using the suffixes specified by argument `suffixes` . Function `filepath` returns the absolute path name of the file. The empty list is returned if no file matches the argument along the current search path.

Argument `suffixes` is a string containing a sequence of possible suffixes (including the initial period) separated by vertical bars “—”. The sequence represents the priority order. A suffix may have zero characters.

```
".dump"      Try suffix <.dump>.
".lshc|.lsh|" Try suffixes <.lshc> and <.lsh>, then try without suffix.
"|.lsh"      Try without suffix, then try suffix <.lsh>.
".lsh|"      Try suffix <sn>, then try without suffix.
```

Argument `suffixes` may specify a single suffix without the initial period. This is a shorthand for testing first the specified suffix and then testing without a suffix.

```
"sn"          Equivalent to ".lsh|"
```

When argument `suffixes` is omitted, `filepath` uses the same suffixes than function `load` . This default provides a way to determine which file would be loaded by a given call to function `load` .

```
? (filepath "sysenv")
= "/home/yann/lush/sys/sysenv.lsh"
```

4.15.5 Files and Directories Management.

(files [directory])

Type: DX

Returns the list of the files located in directory `directory` . If argument `directory` is omitted, the current directory is assumed.

```
? (files)
= ( "." "." "CVS" "config" "0-CVS-INFO" "COPYING" "COPYRIGHT" "Makefile.in"
  "Makefile" "README" "aclocal.m4" "configure" "configure.ac" "debian" "etc"
  "include" "lsh" "src" "sys" "config.log" "bin" "config.status" "packages"
  "NOTES.~1.27.~" "local" "TODO" "TODO~" "doc" "RELEASES" "README.~1.2.~" )
```

(fileinfo filename)

Type: DX

Returns an a-list containing lots of information about the file named **filename** in the form of an association list. Information includes type, dates, permissions, size,... This function returns nil if the named file does not exist.

```
? (fileinfo ".")
= ((type . dir) (size . 4096) (mode . 16877) (dev . 775) (ino . 229389)
   (nlink . 15) (uid . 501) (gid . 501)
   (atime . ::DATE:year-second:(2002-08-14 05:12:36))
   (mtime . ::DATE:year-second:(2003-01-06 21:26:59))
   (ctime . ::DATE:year-second:(2003-01-06 21:26:59)) )
```

(dirp path)

Type: DX

Returns **t** if string **path** is the filename of a directory. This function also returns **t** if string **path** is the filename of a symbolic link referring to a directory.

```
? (dirp "..")
= t
```

(filep path)

Type: DX

Returns **t** if string **path** is the filename of a regular file (not a directory). This function also returns **t** if string **path** is the filename of a symbolic link referring to a regular file.

```
? (filep "..")
= ()
```

(lockfile path)

Type: DX

Attempts to create file **path** .

- If this file already exists, function **lockfile** returns the empty list **()** .
- If this file does not exist already, function **lockfile** writes the current date into the file and returns **t** .

This function is useful for implementing a simple locking scheme.

(mkdir dirname)

Type: DX

This function creates a new directory `dirname` .See: (mkdir-on-need [`s`])See: (chdir [`s`])**(mkdir-on-need path)**

Type: DE(miscenv.lsh)

This function is much like the ‘`mkdir -p`’ command found on many Unix systems. It creates a directory named `path` . This function also creates all required parent directories specified by string `path` .

This function does nothing if directory `path` already exists already. You can safely call this function multiple times. It is assumed that the caller has the appropriate rights to create the `path` .

The following expression checks makes sure that all directories `"/tmp"` , `"/tmp/just"` , `"/tmp/just/an"` and `"/tmp/just/an/example"` exist:

```
(mkdir-on-need "/tmp/just/an/example")
```

(unlink path)

Type: DX

This function removes the file or directory `path` . It does not complain if the file or directory `path` does not exist.

(rename path1 path2)

Type: DX

This function renames file `path1` as `path2` .**(copyfile file1 file2)**

Type: DX

Copy file `file1` onto file `file2` . Arguments `file1` and `file2` may be file descriptors or file names (as described in `open-read` and `open-write` .

4.15.6 Files and Descriptors.**(script [s])**

Type: DX

Creates a file named `s` and records all the input/output operations on that file. The suffix `".script"` is added to the name `s` . This function is handy for keeping a trace of the Lush session. If function `script` is called without arguments , the current script file is closed.

(open-read s [suffixes])

Type: DX

See: (sys `shellcmd`)See: (filepath `filename` [`suffixes`])

This function returns a read file descriptor of the file named `s` . The empty list is returned when an error occurs while opening the file. The file descriptor is closed when the garbage collector detects that the file is no longer necessary. File descriptors however can be manually closed with the function `delete` .

Filenames have the following form:

- A legal operating system filename will be searched along the current path.
- "\$stdin" stands for the Lush standard input.
- On some systems, "| `shell command`" may be used to open a pipe to another process. See the description of function `sys` to have comments about the implementation of this feature.

When string `suffixes` is specified, the suffixes specified by this string are tried while searching the file along the search path. The possible values of the string `suffix` are documented with function `filepath` .

(open-write s [suffix])

Type: DX

See: (sys `shellcmd`)

This function returns a write file descriptor of the file named `s` . The empty list is returned when an error occurs while opening the file. The file descriptor is closed when the garbage collector detects that the file is no longer necessary. File descriptors however can be manually closed with the `delete` function.

Filenames have the following form:

- A legal operating system filename.
- "\$stdout" stands for the Lush standard output stream.
- "\$stderr" stands for the Lush standard error stream.
- On some systems, "| `shell command`" may be used to open a pipe to another process. See the description of function `sys` to have comments about the implementation of this feature.

An optional suffix can be specified as argument `suffix` . This suffix is appended to the filename unless the filename already contains a suffix.

(open-append s [suffix])

Type: DX

This function returns a “append” file descriptor of the file named `s` . All output on this file descriptor shall be appended to the file. The empty list is returned when an error occurs while opening the file. The file descriptor is closed when the garbage collector detects that the file is no longer necessary. File descriptors however can be manually closed with the `delete` function.

An optional suffix can be specified as argument `suffix` . This suffix is appended to the filename unless the filename already contains a suffix.

(filteropen cmd fin fout)

Type: DX

This command is available under Unix. Check function `winsys` under Windows.

This function runs the shell command `cmd` in the background , stores into symbol `fin` a file descriptor for reading its standard output , and stores into symbol `fout` a file descriptor for writing into its standard input.

This commands allows for launching another program , sending commands and receiving the results. Arguments `fin` and `fout` are evaluated and must probably be quoted symbols.

(socketopen host port fin fout)

Type: DX

This command is available under Unix.

This function connects a server process listening on port number `port` on the remote computer named `host` . It stores into symbol `fin` a file descriptor for reading from the server , and stores into symbol `fout` a file descriptor for writing to the server.

Arguments `fin` and `fout` are evaluated and must probably be quoted symbols.

(writing f l1 ... ln)

Type: DY

See: (open-write s [suffix])

See: (open-append s [suffix])

Calls `progn` on `l1` to `ln` , while redirecting the current output to `f` . Argument `f` may be a filename or a file descriptor obtained by `open-write` or `open-append` .

(reading f l1... ln)

Type: DY

See: (open-read s [suffix])

Calls `progn` on `l1` to `ln` , while redirecting the current input to `f` . Argument `f` may be a filename or a read file descriptor obtained by `open-read` .

Example: This function returns the list of the files in directory `s`

```
(de directory(s)
  (let ((ans ()))
    (reading (concat "|/bin/ls " s)
      (while (<> (skip-char) "\\e")
        (setq ans (nconc1 ans (read-string)))) ) )
  ans ) )
```

(read8 f)

Type: DX

Reads a byte on file descriptor `f` . This function returns a number in the range 0 to 255.

(write8 f b)

Type: DX

Writes a byte `b` on file descriptor `f` . Argument `b` must be a number in the range 0 to 255. This function does not flush the file `f` . The function `flush` must be used in order to write out all remaining data.

(fsize f)

Type: DX

Returns the number of bytes remaining in the file. When the file is not exhaustible or not rewindable (such as pipes and terminals) it causes an error.

(flush [f])

Type: DX

Flushes the file descriptor `f` .

When no arguments are provided, both the current input and current output are flushed: All characters pending on the current input are removed until a newline is encountered, all characters buffered on the current output are printed out.

Example:

```
;;; This function ask a question and returns a string.
(de input(question)
  (printf "\\%s" question)
  (flush)
  (read-string) )
```

4.15.7 Binary I/O.

(**bwrite** l1 ... ln)

Type: DX

See: (bread)

Function **bwrite** writes the lisp objects **l1** to **ln** on the current output stream as a sequence of binary tokens. This function is able to write correctly lisp objects with circular references. It cannot however write lisp objects useful for their side effects like open file descriptors or windows.

(**bwrite-exact** l1 ... ln)

Type: DX

See: (bwrite l1 ... ln)

See: (bread)

Function **bwrite-exact** is similar to **bwrite** with a few twists that makes it more accurate but less frequently useful.

- Function **bwrite** writes objects by first specifying the class name. This implies the class must be defined and have the same name when reading the object. Function **bwrite-exact** instead writes a complete definition of the class. This implies that **bread** will create a new unnamed class regardless of the existing classes.
- Function **bwrite** writes indexes by saving their content regardless of how the index is laid out in the corresponding storages. This implies that **bread** will create a new storage for each of them. Function **bwrite-exact** instead writes the storage and the index layout. But this is wasteful if the storage is much larger than the index.

(**bread**)

Type: DX

See: (read)

Function **bread** reads on the current input stream a binary sequence of tokens and builds the corresponding lisp object. Functions **read** and **load** are also able to recognize binary tokens in the middle of a text file. This allows for reading a file containing a mixture of lisp objects representing as binary sequence of tokens or as text.

(**bread-exact**)

Type: DX

See: (bread)

Function **bread-exact** is similar to **bwrite** with one twists that makes it more accurate but less frequently useful.

- Saving an instance of a user defined class with `bwrite` saves only the slots of the instance and a reference to the class name. Function `bread` will be able to read this binary file if the class has already been defined in the interpreter. This class must define the same slots with the same order as the class of the actual object stored in the file. Additional slots will just be skipped and left with their default value. Function `bread-exact` instead causes an error when there are additional slots.

Note that both `bread` and `bread-exact` can read files produced by either `bwrite` or `bwrite-exact` . There is no difference between `bread` and `bread-exact` when reading a file produced by `bwrite-exact` .

(`tokenize fin fout`)

Type: DE(`sysenv.lsh`)

Function `tokenize` reads all lisp expressions from file `fin` and writes their binary equivalent into file `fout` .

This function is useful for converting Lush program files (SN files) into prepared binary file (SNC files). Loading a prepared file is much faster than loading a text file. Prepared copies of certain system libraries are located in directory "`<lushdir /lib>`".

(`dump fname [exec]`)

Type: DX

See: Startup.

See: Runtime.

Writes the status of the lisp interpreter into file `fname` . If `fname` has no suffix, suffix "`.dump`" is assumed. This function works by creating a binary file containing a binary description of the state of the current Lush session.

Binary dump files are more useful for launching Lush with preloaded libraries and application. During startup indeed, Tlisp can load a binary dump file specified by the command line arguments. Moreover, if Tlisp does not find the system library "`sysenv.lsh`" , it attempts to undump a file names "`sysenv.dump`" .

When specified, the optional argument `exec` specifies a lisp expression which will be evaluated immediately before loading the dump file. Variable `file-being-loaded` is set to the name of the `dump` file. This feature is useful for loading, using function `mod-load` , all TL/Open extensions required by the dump file.

4.16 Miscellaneous.

4.16.1 Debug.

A few functions provides debugging support and statistical informations about the resources used by Lush.

(symbols str)

Type: DE(sysenv.lsh)

See: `^S str`

Prints all the symbols whose name contains the string `str` . This is especially useful if you don't remember the exact spelling of a function or variable.

^S str

Type: DMC(sysenv.lsh)

See: `(symbols str)`Macro-character for function `symbols`**(sizeof c_type)**

Type: DX

Returns the number of bytes needed to implement a C item of C type named `c_type` . Supported types are signed numerical types and pointer type `"gptr"` .

```
? (sizeof "double")
= 8
```

```
? (sizeof "gptr")
= 4
```

```
? (sizeof "short")
= 2
```

This function is mainly used for formatted I/O.

(used)

Type: DX

Returns the number of currently used Lush cells. This function is especially useful to track underlocks and overlocks when debugging a new Lush primitive written in C.

(error [symb] string [l])

Type: DX

Causes an error whose text is given by `symb` , `string` and `l` . When `symb` is provided, `error` rewinds the evaluation stack until it reaches a call to a function named `symb` . The error message then displays the stack relative to that function call. When `symb` is omitted, the stack is not displayed.

Example:

```
? (error 'myfunc "bad number" 4)
*** myfunc : Bad number : 4
Debug toplevel [y/n] ?
```

(pause string)

Type: DE

print **string** and start a toplevel prompt. Inserting a **(pause <string >**) expression in a function allows to stop the execution and get a prompt in the context of the **pause** . This allows to examine and/or modify variable values etc.

(debug l1 ... ln)

Type: DY

See: **(trace-hook level line expr info)**

Evaluates **l1** to **ln** , in debug mode. When in debug mode, Lush calls function **trace-hook** before and after each evaluation. The default **trace-hook** function just prints an indented trace of each evaluator call, as well as each value.

Example:

```
? (debug (* (+ 3 4) a)))
-> (* (+ 3 4) a)
  -> (+ 3 4)
    -> 3
      <- 3
        -> 4
          <- 4
            <- 7
              -> a
                <- 8
                  <- 56
                    = 56
```

(nodebug l1 ... ln)

Type: DY

Calls function **progn** on expression **l1** to **ln** without displaying a trace , even if this function is called inside a call to function **debug** . Functions like **pretty** and **save** use this function to produce clean output in debug mode.

(where)

Type: DX

See: Interruptions in Lush.

See: Errors in Lush.

This function prints the entire evaluation stack. Called from a debug toplevel, function `where` gives useful hints for undertsanding what happened and when.

(trace-hook level line expr info)

Type: DE

See: (debug 11 ... 1n)

This function is called before and after each evaluation when the Lush kernel runs in debug mode (see function `debug`). Redefining this function allows interactive Lush debugger to execute programs step by step.

- When `trace-hook` is called before an evaluation, argument `level` is a positive number indicating the number of recursive calls to the evaluator since Lush entered the debug mode. Argument `line` is a string displaying the first line of the expression being evaluated. Argument `expr` is the expression being evaluated. Argument `info` is the list of the expressions whose recursive evaluation led to the evaluation of `expr`.

If function `trace-hook` returns `()`, the evaluation will continue without calling `trace-hook` for recursive evaluations. The next call to `trace-hook` will occur after the current evaluation returns. If function `trace-hook` returns `t`, the recursive calls to the evaluator (calls performed while evaluating the current expression) will be traced.

- When `trace-hook` is called after an evaluation, argument `level` is a negative number representing the opposite of the number of recursive calls to the evaluator since Lush entered the debug mode. Argument `line` is a string representing the first line of the returned value. Argument `expr` is the expression being evaluated. Argument `info` is the value returned by the evaluation.

If function `trace-hook` returns `()`, the remaining evaluations will continue without calling `trace-hook` until an evaluation with a smaller recursion level returns. If function `trace-hook` returns `t`, the next evaluations will be traced normally.

result

Type: VAR

The last result produced by the toplevel is always stored in the variable `result`.

Example:

```
? (sqrt 81)
= 9
? (* result result)
= 81
```

version

Type: VAR

This symbol contains a string, whose contents is the full name of the Lush based program you are running.

lushdir

Type: VAR

This variable contains the filename of the root of the Lush tree.

(exit [n])

Type: DX

See: Lush Startup.

This functions exits the current toplevel, as if a `Ctrl-D` had been typed. When the optional argument `n` is given, function `exit` quits the Lush process with return code `n` .

(discard l1 ... ln)

Type: DY

Evaluates `l1` to `ln` like `progn` , but prevents the entire result to be printed by the toplevel. It will print its first line instead. This function is useful for working quietly on very long lists.

(startup ...argv...)

Type: DE(sysenv.lsh)

See: Lush Startup.

Lush calls this function once during the startup procedure as soon as the system library "`sysenv.lsh`" or "`sysenv.dump`" has been loaded. The behavior of the default `startup` procedure is explained in section "Lush Startup."

The arguments `argv` are the command line arguments. Neither the executable name nor the magic first command line argument are passed to `startup` . However, a fake argument "`@@`" is prepended when Lush runs in script mode.

(toplevel)

Type: DE(sysenv.lsh)

See: Lush Startup.

This function usually is defined by "`sysenv.lsh`" and is called after function `startup` during the startup process. This function is restarted whenever an error occurs.

(on-error p l1 ... ln)

Type: DM

See: Interruptions in Lush.

This function evaluates lists `l1` to `ln` and returns the last result like function `progn`. If however an errors occur during these evaluations, the expression `p` is evaluated before the usual error processing.

The function `on-error` does not stop the error processing, but allows for performing cleanup tasks, or for printing a custom error message.

(on-break p l1 ... ln)

Type: DM

See: Errors in Lush.

This function evaluates lists `l1` to `ln` and returns the last result like function `progn`. If however the user interrupts the evaluation by depressing `Ctrl-C`, the expression `p` is evaluated before the usual interruption processing.

The function `on-break` does not stop the interruption processing, but allows for performing cleanup tasks, or for printing a custom error message.

(on-error-macro pm l1 ... ln)

Type: DM

See: Errors in Lush.

See: `(on-error p l1 ... ln)`

This function first evaluates list `pm` and stores the result which is usually an expression to be evaluated if an error occurs. It evaluates then lists `l1` to `ln` and returns the last result like function `progn`. If however an errors occur during these evaluations, the result of the initial evaluation of `pm` is evaluated before the usual error processing.

Actually, this function works like `on-error`, except the expression which will be executed in case of error is the result of the evaluation of `pm` in the call context of `on-error-macro` i.e. before `l1 ... ln` are evaluated.

This allows a safer behaviour than `on-error`, since the evaluation of `pm` is made in the call context instead of being made in the error context. Indeed, the error context may have values pushed on the stacks by functions called in `l1 ... ln`, therefore hiding the expected values.

```
(let* ((my-window (new windowobject 100 100 200 100 "Handler"
                                (new string "processing...")))
      (errq (new errorrequester my-window)))
  (on-error-macro
   '(=> ,errq popup (errname))
   (let* ((errq ()))
     (error "this is is a programmed error\\nused as test."))))
```

(on-break-macro pm l1 ... ln)

Type: DM

See: Errors in Lush.

See: (on-break p l1 ... ln)

This function first evaluates list `pm` and stores the result which is usually an expression to be evaluated in case the user interrupts the execution with the `Ctrl-C` or `Ctrl-Break` key. It evaluates then the lists `l1` to `ln` and returns the last result like function `progn` . If however the user interrupts the evaluation by depressing `Ctrl-C` or `Ctrl-Break` , the result of the evaluation of `pm` stored previously is evaluated before the usual interruption processing.

Actually, this function works like `on-break` , except the expression which will be executed in case of interruption is the result of the evaluation of `pm` in the call context of `on-break-macro` i.e. before `l1 ... ln` are evaluated.

This allows a safer behaviour than `on-break` , since the evaluation of `pm` is made in the call context instead of being made in the interruption context. Indeed, the interruption context may have values pushed on the stacks by functions called in `l1 ... ln` , therefore hiding the expected values.

(errname)

Type: DX

This function returns the last error message as a string. This function is usually called by the `break-hook` or `debug-hook` functions , whenever an error or an interruption occurs.

(debug-hook)

Type: DE(sysenv.lsh)

See: Errors in Lush.

This function is defined by "`sysenv.lsh`" and is called whenever an error occurs. Use care when modifying this function, since an incorrect `debug-hook` function may result into a deadly infinite loop!.

(break-hook)

Type: DE(sysenv.lsh)

See: Interruptions in Lush.

These functions is defined by "`sysenv.lsh`" and is called whenever the user causes an interruption by depressing `Ctrl-C` .

4.16.2 System.

A few functions are strongly related to the operating system. The quality of the implementation however depends on the quality of the underlying operating system.

Some functions are not available under all operating systems:

- The function `bground` is only available on Unix system.
- The function `xmnc` is only available on Apollo Domain/OS computers.
- The functions `winlushp` , `winsys` and `winedit` are only available under Windows/95 and Windows/NT.

(beep)

Type: DX

Function `beep` is self explanatory.**(sleep n)**

Type: DX

Function `sleep` waits `n` seconds and returns.**(sys shellcmd)**

Type: DX

See: `#$ shellcmd`See: (open-write `s` [`suffix`])See: (open-read `s` [`suffix`])Executes the shell command defined in string `shellcmd` .

Example (Unix):

`? (sys "pwd")``/home/toto/bip``= 0`

The implementation of function `sys` and of the pipe convention in filenames (see `open-read` and `open-write`) is highly system dependent.

- This implementation is extremely reliable on Unix systems. Using pipes and forking background processes has been the tradition of Unix since the origins.

Function `sys` will always wait until the spawned process exits and return the process exit code.

- This implementation is much less reliable under Windows 95 or NT. You can indeed run very different kinds of programs, including legacy MS-DOS and WIN16 programs. A number of programs running under Windows have been written with the assumption that they would run using a MSDOS style console. Running them as background processes or with standard handles redirected to pipes will cause problems.

If you are running the Console based version of Lush, function `sys` replicates as well as possible the behavior of the Unix version. The spawned

process inherits the Lush console and may be controlled by the user. If you are running a MSDOS or WIN32 Console application, function `sys` returns when the application exits. If you are running a Windows 3.1 or WIN32 GUI application, function `sys` returns immediately.

If you are running the GUI based version of Lush, the spawned process no longer inherits a MSDOS console. We have chosen to handle MSDOS and WIN32 Console application like a GUI application. Function `sys` will create a new console for the process and return immediately regardless of the application type.

Using pipes is reasonably reliable under the Console based version of Lush. It becomes very adventurous under the GUI version of Lush because we do not want to create a console and let the user interact with the program. Bugs in Windows 95 make pipes so dangerous that we just forbid them.

Please investigate function `winsys` to start a new process with a collection of refined options (redirecting standard handles, using the command interpreter, detaching the process, waiting until the process exits). You should nevertheless remember that these options do not work well with all application types.

```
#$ shellcmd
```

Type: DMC(sysenv.lsh)

See: (sys `shellcmd`)

Macro-character for function `sys` .

```
? # $pwd
```

```
/home/toto/bip
```

```
= 0
```

```
(chdir [s])
```

Type: DX

See: (path `s1 ... sn`)

See: (sys `shellcmd`)

If the string `s` is provided, this function sets the current working directory to directory `s` . This function returns the current working directory name.

The current directory specified with `chdir` is used by all shell commands launched with function `sys` . When accessing a file, Lush also searches the current directory before the directories specified by function `path` .

Example:

```
? (chdir)
```

```
= "/home/yann/lush"
```

(getpid)

Type: DX

Returns the process ID of this process.

Example:

```
? (getpid)
= 15623
```

(getuid)

Type: DX

Returns the user ID of this process.

Example:

```
? (getuid)
= 501
```

(getusername)

Type: DX

Returns the user name for this process.

Example:

```
? (getusername)
= "yann"
```

(edit file)

Type: DE(sysenv.lsh)

See: `^E file`

Calls the standard editor on file `file` . The editor name is stored in the symbol `edit-call` , and is initialized by looking at the environment variable `EDITOR` or `VISUAL` .

`^E file`

Type: DMC(sysenv.lsh)

See: `(edit file)`Macro-character for function `edit` .**edit-call**

Type: VAR(sysenv.lsh)

This variable contains the name of the editor command.

(fedit func)

Type: DE(sysenv.lsh)

See: `^F func`

Copies an indented definition version of `func` in a temporary file and calls the function `edit` on this file. When the editor returns, the edited definition of `func` is loaded.

^F func

Type: DMC(sysenv.lsh)

See: `(fedit func)`Macro-character for function `fedit` .**(getconf varname)**

Type: DX

Obtains the value of variables determined by the shell script `configure` . String argument `varname` can take the following values:

- "SHELL" The shell interpreter, e.g. `"/bin/sh"` .
- "OPTS" The optimization options for the compilers, e.g. `"-DNO_DEBUG -O2"` .
- "DEFS" The symbol definition options for the compilers, e.g. `"-DHAVE_CONFIG_H"`
- "LIBS" Libraries used for linking the main executable, e.g. `"-lbfd -liberty -ldl -lm"` .
- "host" String describing the host computer, e.g. `"i686-pc-linux-gnu"` .
- "CPP" Command for running the C preprocessor, e.g. `"gcc -E"` .
- "CPPFLAGS" The preprocessing options for the compilers.
- "CC" Command for running the C compiler, e.g. `"gcc"` .
- "CFLAGS" The C specific compiler options besides those specified in `CPPFLAGS` , `OPTS` and `DEFS` .
- "CXX" Command for running the C++ compiler, e.g. `"g++"` .
- "CXXFLAGS" The C++ specific compiler options besides those specified in `CPPFLAGS` , `OPTS` and `DEFS` .
- "F77" Command for running the Fortran compiler, e.g. `"g++"` .
- "FFLAGS" The Fortran specific compiler options besides those specified in `CPPFLAGS` , `OPTS` and `DEFS` .

- "LDCC" Command for linking executables, e.g. "g++" .
- "LDFLAGS" Options for linking executables.
- "PTHREAD_FLAGS" Flags for compiling or linking programs that use the Posix thread library.
- "PTHREAD_LIBS" Libraries for linking programs that use the Posix thread library.
- "AR" Command for creating static libraries, e.g. "/usr/bin/ar" .
- "MV" Command for renaming and moving files, e.g. "/bin/mv" .
- "CP" Command for copying files, e.g. "/bin/cp" .
- "LN_S" Command for creating symbolic links, e.g. "ln -s" .
- "TOUCH" Command for updating the modification time of a file, e.g. "/bin/touch" .
- "INDENT" Command for running the GNU indent program.
- "RANLIB" Command for creating the table of contents of a static library created with AR , e.g. "ranlib" .
- "X_LIBS" Options for linking executables using the X11 library, e.g. "-L/usr/X11R6/lib -lSM -lICE -lX11" .
- "X_CFLAGS" Options for compiling programs using the X11 library, e.g. "-I/usr/X11R6/include" .
- "CC_PIC_FLAG" Compiler flag for producing position independent code appropriate for creating shared libraries, e.g. "-fPIC" .
- "CC_EXP_FLAG" Linker flag for linking executables that export their symbol tables to resolve undefined symbols in shared libraries, e.g. "-Wl,-export-dynamic" .
- "MAKESO" Command for creating a shared library, e.g. "gcc -shared -o" .
- "EXEEXT" Filename extension for executable programs, e.g. "" or "exe" .
- "OBJEXT" Filename extension for object files, e.g. "o" or "obj" .
- "SOEXT" Filename extension for shared libraries, e.g. "so" or "sl" or "dll" .
- "LUSH_MAJOR" The major version of Lush specified in loadable modules.
- "LUSH_MINOR" The minor version of Lush specified in loadable modules.

(getenv s)

Type: DX

Returns the value of the environment variable whose name is `s` .

Example:

```
? (getenv "HOME")
= "/home/yann"
```

(getconf s)

Type: DX

Returns the value of the autoconf variable determined at compilation time.

These variables indicate how the lisp interpreter was compiled. See file "`include/autoconf.h.in`" for a list of variables.

Example:

```
? (getenv "CC")
= ""
```

```
? (getenv "CFLAGS")
= ""
```

(time l1 ... ln)

Type: DY

Returns the CPU time in seconds spent to evaluate `l1` to `ln` . Without arguments, function `time` returns the real time spent since you started Lush.

Example:

```
? (time (repeat 40000
          (sqrt 2) ))
= 0.04
```

(ctime)

Type: DX

This function is identical to the Unix `ctime` function. It returns date and time as a 26 character string.

Example:

```
? (ctime)
= "Tue Jan  7 20:27:18 2003"
```

(localtime)

Type: DX

This function is identical to the Unix `localtime` function. It returns a list of the following form:

```
( tm_isdst <day light saving time>
  tm_yday <day of the year (0..365)>
  tm_wdat <day of the week (sunday=0)>
  tm_mon <month number (0..11)>
  tm_mday <day of the month (1..31)>
  tm_hour <hour (0..23)>
  tm_min <minutes (0..59)>
  tm_sec <seconds (0..59)> )
```

Example:

```
? (localtime)
= (tm-isdst 0 tm-yday 6 tm-wday 2 tm-year 103 tm-mon 0 tm-mday 7 tm-hour 20
   tm-min 27 tm-sec 18 )
```

(isatty filename)

Type: DX

Returns `t` if file `filename` is an interactive terminal.

(background s l1 ... ln)

Type: DY

Creates a background process which evaluates expressions `l1` to `ln` and exits. The output of this process is redirected to a file named `s`. A suffix `".job"` is concatenated to filename `s` when necessary.

(winlushp)

This function is only available under Windows. It returns `()` if you are running the console version of Tlisp. It returns `t` if you are running the GUI version of Tlisp.

Examples:

```
;; You can test if you are running under windows as follows.
(cond
  ((not winlushp)
   (printf "Not running under Windows\\n" )
  ((winlushp)
```

```
(printf "Running WinLush (GUI version of Lush under WIN32)\n" )
(t
  (printf "Running Lush (Console version of Lush under WIN32)\n" ) )
```

(winedit filename [as-untitled])

This function is only available when you are running the GUI version of Lush for Windows (WinLush).

Function **winedit** opens a text editor window for the file named **filename** . If the name **filename** ends with suffix **".lsh"** , WinLush will allow syntax coloring and automatic indentation.

The optional flag **as-untitled** tells WinLush to open the file as an untitled window.

- If flag **as-untitled** is the empty list, WinLush first searches an already open text editor window for this file. If such a window already exists, WinLush just pops up the window. Otherwise, WinLush creates a new window and loads the file. This mode is convenient for editing well defined files.
- If flag **as-untitled** is not nil, WinLush creates a new untitled window and transfers the contents of the file. The new editor window will not remember the file name. This mode is convenient when you are using a temporary text file for transferring some text into an editor.
- If you do not specify argument **as-untitled** , WinLush makes a decision on the basis of the file location. If the file is located under the system temporary directory, WinLush creates an untitled editor window. Otherwise, WinLush creates a regular text editor window.

(winsys stdin stdout stderr commandp detachp waitp cmdline)

See: (sys **shellcmd**)

This function is available when you are running Lush under Windows 95 or NT.

Function **winsys** creates a process for the program specified by command line **cmdline** . The other arguments allow you to tune finely the process parameters.

The legal and reliable combinations depend on:

- whether you are running the console based (Lush) or GUI based (WinLush) version of Lush,
- whether you are running Windows 95 or Windows NT,
- whether you are executing a MSDOS application, WIN16 application, WIN32 Console based application, WIN32 GUI application, OS2 2.1 application (under NT) or Posix application (under NT).

Arguments `stdin` , `stdout` and `stderr` allow you to define the standard input, output and error handles of the program. These arguments may be `()` , `t` or a variable name.

- Value `()` represents an invalid handle. Therefore the process will not be able to use the corresponding standard handle. This may cause Windows to create a MSDOS console as a replacement.
- Value `t` represents the handle inherited from the Lush process. Under WinLush, this value redirects the standard output or standard error to the WinLush console.
- A variable name tells Lush to pass one end of the pipe to the process. The other end of the pipe will be stored into the specified variable as a Lush file descriptor. If you specify the same variable name for the standard output and error, Lush will use a single pipe for both standard handles.

Flag `commandp` is used to determine how the command line is processed. If this flag is set, the command line is passed to the command interpreter. You can thus run the internal commands implemented by `CMD.EXE` (Windows NT) or `COMMAND.COM` (Windows 95 - an MSDOS program!!). If this flag is not set, the command line is directly interpreted by WIN32's `CreateProcess` function.

Flag `detachp` is used to run a process in the background. This flag should prevent Windows to create a console or inherit the console handles. It is rather buggy under Windows 95.

Flags `waitp` tells if function `winsys` must wait until termination of the process and return the process exit code. If this flag is unset, the function returns immediately.

Examples:

```
;; Here is an alias for function SYS under WinLush
(winsys () () () t () () cmdline)
;; Here is an alias for function SYS under Console Lush
(winsys t t t t () t cmdline)
;; Here is a way to open a read pipe
(winsys () 'readpipe t () t () cmdline)
```

(win-show-workbench flag)

See: `(wbm-window [[[x y] w h] name])`

See: `(wbm-toplevel-window [[[x y] w h] name])`

This function is available when you are running Lush under Windows 95 or NT. This function controls the visibility of the graphical interface of the GUI version Lush for Windows (WinLush). Calling this function does nothing if you are running the console version of Lush for Windows.

When argument `flag` is the empty list, WinLush switches to “hidden” mode. The main WinLush window (containing the Lush console, the text editor

windows, and possibly a few graphical windows) disappears from the computer screen.

WinLush reverts to normal mode when function `win-show-workbench` is called again with a non nil argument. There is also a safety feature that automatically reverts WinLush normal mode when some user input is required on the Lush console.

When WinLush runs in “hidden” mode, function `wbm-window` works like function `wbm-toplevel-window` . Instead of creating a child window of the invisible main WinLush window, `wbm-window` creates a toplevel window which directly appears on the Windows desktop.

It is possible to directly start WinLush in “hidden” mode by specifying option `"-hide"` as the first command line argument. The remaining argument are interpreted as usual. This command line argument allows for creating self contained WinLush applications.

You can for instance create a file `"c:/calctoolapp.lsh"` containing the following three lines:

```
(setq mainwnd (calctool))
(while mainwnd (waitevent) (process-pending-events))
(exit 0)
```

You can then start WinLush with the following command line:

```
winlush -hide stdenv c:/calctoolapp.lsh
```

This command starts WinLush in hidden mode (`"-hide"`). WinLush will first load the standard environment (`"stdenv"`) and then load our example file (`"c:/calctoolapp.lsh"`). This file will create a calculator window on the Windows desktop, processes graphical events, and terminates execution when the user closes the calculator window.

The user will only see the calculator window. The WinLush main window never appears on the screen because WinLush remains in “hidden” mode as long as there is no need for user input on the Lush console.

4.16.3 Special Numerical Values (IEEE754).

This section is valid on computers supporting the IEEE754 floating point representation scheme. Most computers and compilers nowadays support this.

Numerical computations may encounter three kinds of problems:

- Problems arise when we attempt to perform a clearly forbidden operation, such as taking the logarithm of a negative operand. This attempt causes a low level floating exception which generates a Lush error.
- A computation may return a number too large or too small to be represented by the usual floating point data types. Too small number are

approximated by zero. Too large numbers are represented using special bit patterns for representing infinities.

- Special bit patterns are used to represent infinities or missing data. The special bit pattern representing missing data is called “NaN” (NaN stands for “Not a Number”). Computations involving NaNs usually return NaNs.

The special bit patterns representing infinities or NaNs are still Lush numbers and `numberp` will always return `t` when applied on a NaN. However NaNs and Infinities do not behave like regular numbers.

Special bit patterns have no literal representation for the Lush reader. Most machines however print them as "Nan" or "Inf" or "-Inf" . You should not save these bit patterns into ascii format files because the reader will not be able to read them back. You must save them into binary files.

According to IEEE specifications, testing the equality of two special bit patterns should always return false. Comparing two bit patterns should cause a floating exception. Major operating systems and compilers do not abide with this specification. For instance the Microsoft Visual C++ 4.2 compiler considers that NaNs are equal to anything. It is unclear if this is a bug or a speed compromise (comparisons of actual number are much more frequent).

These problems make us unable to comply fully with the standard. Portable programs can only rely on the following guarantees:

- NaNs and Infinities are different from regular numbers. The equality test will return false (we had to hand code this with Visual C++ in the Win-doze version of TL3!).
- The Negative Infinity is smaller than any regular number. The Positive Infinity is greater than any regular number.
- Overflows may produce a floating point exception or return an Infinity. This is operating system and compiler dependent.
- Illegal operations (eg. comparing NaNs) are likely to cause a floating exception. They may however return NaNs without causing a floating exception on certain systems.

NaNs and Infinities however can be safely manipulated and tested by the following functions:

(nan)

Type: DX

When NaNs are supported , this function returns a NaN number. This is the only way provided by Lush for generating a NaN.

When NaNs are not supported , this function returns (). This is the only way provided by Lush for checking if NaNs are supported by this machine.

(nanp x)

Type: DX

This function returns `t` when `x` is a NaN. It returns `()` otherwise.**(not-nan x)**

Type: DX

This function returns `()` when `x` is a NaN. It returns `x` otherwise.**(infinity)**

Type: DX

When Infinities is supported, this function returns the Positive Infinity. The Negative Infinity can be safely computed as `(- (infinity))`.**(infinityp x)**

Type: DX

This function returns `t` when `x` is an infinite number (wether it is positive or negative).**(progn-without-fpe l1...ln)**

Type: DY

Temporarily disables the generation of floating point exceptions, evaluates `l1` to `ln` and returns the result of the evaluation of `ln`.

Example

```
? (progn-without-fpe (log 0))
= -Inf
```

4.17 Graphics.

Lush features a collection of graphics functions for creating windows, drawing figures, plotting curves, and handling graphical events.

A window descriptor is a lisp object of class `|WINDOW|`. Most graphical commands implicitly operate on the current window defined by the window descriptor stored in the symbol `window`.Each window descriptor forwards all drawing requests to a “Graphic Driver” which performs the system dependent calls required to display the drawings on the corresponding device. The most common type of window (and driver) is the `x11-window`, which displays on your screen under Xwindows. Another one is the `ps-window` which sends all graphic command to an Encapsulated PostScript file. A particularly useful one is the `comdraw-window` which sends graphics to a `comdraw` graphic editor. This type of window can be used to

produce "editable" graphics that can be modified using `comdraw` of one of the other applications of the `ivtools` package.

4.17.1 Creating a Window

The standard function for creating a window on your screen is `new-window`. This function determines which system you are running and calls the appropriate low level function.

Function `print-window` performs the same task for creating a window descriptor that accesses your printer.

(new-window [[x y] w h] [name])

Type: DX

See: (`x11-window` [[[x y] w h] name])
 See: (`comdraw-window` [[[x y] w h] name])
 See: (`ps-window` [[w h] name])
 See: (`wbm-window` [[[x y] w h] name])
 See: `window`

This is the device independent function for creating a window. It figures out which drivers are available, and call the appropriate function for creating a window.

Function `new-window` function creates a window named `name`. Arguments `w` and `h` specify the size of the window. Arguments `x` and `y` specify the location of the window.

The resulting window descriptor is then stored into variable `window`. This operation makes the window current.

Windows are closed when they are no longer referenced by the lisp interpreter, when they are deleted with the `delete` function, or when the windowing system sends a deletion request.

(print-window w h [destination])

Type: DX

See: (`ps-window` [[w h] name])
 See: (`wpr-window` [[w h] name])
 See: (`wpr-printers`)
 See: Print Requester.
 See: Producing Encapsulated PostScript Files.
 See: `window`

This is the device independent function for printing graphics on the system printer. This function figures out the best way to send graphics on the printer or the file specified by string `destination`. Arguments `w` and `h` specify the size of the printing area. The drawings will be scaled to fit a standard size page.

The resulting window descriptor is then stored into variable `window`. This operation makes the window current.

Function `print-window` actually calls a system dependent function (e.g. `wpr-window` under Windows 95 and NT, `ps-window` under Unix operating systems). Argument `destination` is passed verbatim to these functions. There are therefore two ways to write a portable program. The simplest way consists in omitting argument `destination` . The best way consists in using the `PrintRequester` class provided with the Ogre library.

The window descriptors returned by `print-window` are very similar to the window descriptors returned by `new-window` . They print pages instead of rendering graphics on your screen. A page is output when you call function `cls` or when the window descriptor is closed (because it is no longer referenced by the interpreter or because you called function `delete`).

Note: Calling function `cls` while a clip rectangle is active (see function `clip`) will not output a page. It will simply clear the clipping zone. Before using `cls` , you can use function `clip` to make sure that no clipping rectangle is active.

4.17.2 Drawing.

The coordinate system of a window has its origin in the upper left hand corner. Positive x and y coordinates go to the right and downwards.

(cls)

Type: DX

Clears the current graphics window.

(draw-line x1 y1 x2 y2)

Type: DX

Draws a line from point (`x1` , `y1`) to the point (`x2` , `y2`) in the current window , with the current color.

(draw-rect x y w h)

Type: DX

Draws a hollow rectangle in the current window , with the current color. Its top-left corner is located at coordinates (`x` , `y`); its width is `w` , and its height is `h` .

(draw-round-rect x y w h [r])

Type: DX

Draws a hollow rectangle with rounded corners of radius `r` .

(draw-circle x y r)

Type: DX

Draws a hollow circle in the current window , with the current color. The center of the circle is located at position (x , y). The radius of the circle is radius is r .

(draw-arc x y r fromangle toangle)

Type: DX

Draws a segment of a circle in the current window , with the current color. The center of the circle is located at position (x , y). The radius of the circle is radius is r . The segment is delimited by angle **fromangle** and **toangle** which are values between -360 and 360 inclusive. Argument **toangle** must be larger than **fromangle** .

(fill-rect x y w h)

Type: DX

Fills a rectangle in the current window , with the current color. Its top-left corner is located at coordinates (x , y); its width is w , and its height is h .

(fill-round-rect x y w h [r])

Type: DX

Fills a rectangle whith rounded corners of radius r .

(fill-circle x y r)

Type: DX

Fills a circle in the current window , with the current color. The center of the circle is located at position (x , y). The radius of the circle is radius is r .

(fill-arc x y r fromangle toangle)

Type: DX

Fills a segment of a circle in the current window , with the current color. The center of the circle is located at position (x , y). The radius of the circle is radius is r . The segment is delimited by angle **fromangle** and **toangle** which are values between -360 and 360 inclusive. Argument **toangle** must be larger than **fromangle** .

(fill-polygon x1 y1 ... xn yn)

Type: DX

Fills the polygon defined by the points (x_1 , y_1) to (x_n , y_n) , in the current window , with the current color.

(rgb-draw-matrix x y mat [sx sy])

Type: DX

This is the main function to draw an RGB or greyscale image.

This function displays the pixels contained in `idx mat` at location `(x , y)` in the current window. The optional arguments `sx` and `sy` specify the horizontal and vertical zoom factors.

The matrix `mat` can be an `idx2` or `idx3`. If it is an `idx2`, each value is interpreted as a greyscale value where 0 is black and 255 is white. If it is an `idx3`, the last dimension must be 1, 3, or more. If the last dimension is 1, the values are interpreted as greyscale values. If the last dimension is 3 or more, the first 3 values in the last dimension are interpreted as RGB intensities (between 0 and 255).

(draw-value x y val maxv maxs)

Type: DX

This function displays a real value as a white or black square.

Arguments `x` and `y` are the coordinates of the center of the square. Argument `val` is the value to be displayed. Argument `maxv` is the maximum absolute value for `val`. Argument `maxs` is the maximum size of the square. When `val` is larger than `maxv`, a square of size `maxs` is drawn.

Note: The square will not be visible if it is drawn in the same color as the background.

(draw-list x y l ncol maxv apart maxs)

Type: DX

High level drawing function to graphically represent a list of real numbers. A list of real numbers `l` will be represented as a series of black or white squares on a grey background (actual colors depend on the implementation).

Arguments `x` and `y` are the coordinates of the top left edge of the background area. The squares are arranged in an array with `ncol` columns, the number of lines is then defined by the length of the list `l`. The grey background is always rectangular, even if `ncol` is not a divisor of the length of `l`.

As with `draw-value`, argument `maxv` is the maximum absolute value for the list elements. The size of the square is however bounded by `maxs`.

Argument `apart` defines the size of the space occupied by a single square, i.e., the centers of two neighboring squares will be `apart` pixels apart. It is generally suitable to choose a value for `apart` slightly greater than `maxs` so that no squares will overlap. For best results, the difference between `apart` and `maxs` should be an even number.

Example:

```
; draw 6 values on 3 columns and 2 lines, in 50 pixel squares
```

```
(draw-list 50 50 '( 4 5 -10
                  0 4 16 ) 3 16 52 50)
; draw the values in l on a single line
(draw-list 100 100 l (length l) 10 20 18)
```

(gray-draw-list x y l ncol minv maxv apart)

Type: DX

High level drawing function to graphically represent a list of real numbers. A list of real numbers `l` will be represented as a series of gray squares whose gray level is related to the represented value. This function uses a clever dithering algorithm on black&white displays.

Arguments `x` and `y` are the coordinates of the topleft edge of the first square (representing the first element in the list). The squares are arranged in an array with `ncol` columns, the number of lines is defined by the length of the list `l`.

Values between `minv` and `maxv` will be displayed as gray levels. A value of `minv` will be rendered as a black square. A value of `maxv` will be rendered as a white square. Argument `minv` can be defined to be greater than `maxv` in order to produce a reverse video effect. The sizes of the squares are defined by the `apart` parameter.

(gray-draw-matrix x y mat minv maxv apartx aparty)

Type: DX

High level drawing function to graphically represent a 2D matrix `mat` as a series of gray rectangles whose gray level is related to the represented value. This function uses a clever dithering algorithm on black and white displays.

Arguments `x` and `y` are the coordinates of the topleft edge of the first square (representing the first element in the list). The values between `minv` and `maxv` will be displayed as gray levels. A value `minv` will be rendered as a black rectangle. A value `maxv` will be rendered as a white rectangle. Value `minv` can be defined to be greater than `maxv` in order to produce a reverse video effect. The sizes of the rectangles are defined by the `apartx` and `aparty` parameter.

(draw-mat m [x y])

Type: DE

Function `draw-mat` is easy-to-use wrapper for function `gray-draw-matrix`. Argument `m` is the matrix to draw. Optional arguments `x` and `y` define the position of the upper left corner of the matrix display.

The function `draw-mat` computes statistical information about the matrix values in order to adjust the gray scales. This function automatically creates a window if no window is current when the function is called.

(color-draw-list x y l ncol minv maxv apart cmap)

Type: DX

See: (alloccolor *r g b*)

This function is essentially similar to `gray-draw-list` . Argument `cmap` however must be a one-dimensional matrix of 64 color numbers returned by `alloccolor` . Values between `minv` and `maxv` are rendered using the ramp of colors specified by this matrix.

(**color-draw-matrix** *x y mat minv maxv apartx aparty cmap*)

Type: DX

See: (alloccolor *r g b*)

This function is essentially similar to `gray-draw-matrix` . Argument `cmap` however must be a one-dimensional integer matrix of 64 color numbers returned by `alloccolor` . Values between `minv` and `maxv` are rendered using the ramp of colors specified by this matrix.

4.17.3 Drawing with Colors.

Color selection is highly dependent on the capabilities of your graphics hardware and drivers.

- There are still computers equipped with black and white screens. You can check that your hardware can display enough colors using function `colorp` .
- Many graphic devices use color palettes. The color palette defines a small set of colors that can be displayed simultaneously on the screen. Lush will allocate palette entries as soon as you start using a new color. Since palette entries are rare resources, you should use colors cautiously.

Lush provides two elementary functions for dealing with colors. Function `alloccolor` takes a set of three numbers (ranging from 0 to 1) representing the red, green and blue component of the color. This function allocates a palette entry if necessary and returns a hardware dependent color number. These color numbers can be used with function `color` to set the current drawing color.

It happens sometimes that it is not possible to allocate a color because the color palette is full. Lush will first try to create its own color palette and install it whenever you activate a Lush window. Windows created by other programs can display funny colors until you activate one of these windows. If the private color palette is full, function `alloccolor` simply returns a default color.

The use of a limited number of colors is therefore highly encouraged. The function `color-stdmap` provides several convenient sets of standard colors for this purpose.

See: Ogre Color Palette.

(**colorp**)

Type: DE

The function `colorp` is a predicate which returns `t` when the display is known to support color. This function returns `()` when the display does not support color.

(color [c])

Type: DX

See: (allocolor `r g b`)

If `c` is specified, sets the current color to color number `c`. Function `color` always returns the current color. Color numbers are obtained with function `allocolor`. A few predefined color numbers however are defined by the file `"graphenv.lsh"`:

- `color-fg` or -1 for the system foreground color.
- `color-bg` or -2 for the system background color.
- `color-gray` or -3 for a 50 percent dithered gray level.

(allocolor `r g b`)

Type: DX

See: (color [`c`])

Ask the driver for a color number defined by its three primitive values. Arguments `r`, `g`, `b` are reals between 0 and 1 that define the red, green and blue components of the desired color.

Returns a color number suitable for using with the `color` function. This function returns the constant `color-gray` if the system is not able to display the requested color.

color-fg

Type: VAR(graphenv.lsh)

See: (color [`c`])

System foreground color.

color-bg

Type: VAR(graphenv.lsh)

See: (color [`c`])

System background color.

color-gray

Type: VAR(graphenv.lsh)

See: (color [`c`])

A 50 percent dithered gray level.

(color-shade x)

Type: DE(graphenv.lsh)

This function calls `alloccolor` then `color` for setting a grayscale color. Argument `x` is a real between 0 and 1. Function `color-shade` provides support for grayscale monitors. In addition, drivers often use dithering for filling circles or rectangles on black and white monitors .

(color-rgb r g b)

Type: DE(graphenv.lsh)

This function calls `alloccolor` then `color` for setting the current color according to its three primitive values. Arguments `r` , `g` , `b` are reals between 0 and 1. Of course, this function has a poor effect without a color screen.

(color-stdmap [symb])

Type: DE

The function `color-std` provides a convenient way to create a set of predefined matching colors. This function returns an integer matrix containing color identifiers that you can use with function `color` .

Using this function maximizes the chances that several part of your Lush programs will use the same palette entries and therefore avoid palette saturation.

The argument `symb` is a symbol which specifies out of the predefined color sets:

- Symbol `rainbow` defines 64 colors located on a circle linking the pure red, green and blue hues. The first 48 colors actually represent a rainbow. The last 16 colors directly interpolate purple to red.
- Symbol `spread` defines 7 colors that are easy to discriminate.
- Symbol `shade` defines a set of 64 gray levels going from black to white.

(color-std (x [symb]))

Type: DE

The function `color-std` provides access to the convenient sets of colors allocated by `color-stdmap` .

- Argument `x` is number whose decimal part selects a color in the set returned by function `color-stdmap` . Value 0 selects the first color of the set. Value 0.9999 selects the last color of the set.
- Optional argument `symb` is a symbol which specifies a color set for function `color-stdmap` .

See: (color-stdmap [symb])

4.17.4 Drawing Text.

Text may be printed in graphic window by using the functions `draw-text` and `gprintf`. The current font may be changed by the function `font`. The size and/or rectangle of a graphic text may be scanned by the functions `text-width`, `text-height` and `rect-text`.

(draw-text x y s)

Type: DX

Draws the text of the string `s` in the current window at position `(x , y)`, using the current font and the current color.

(gprintf x y fmt ... args ...)

Type: DE(graphenv.lsh)

See: `(draw-text x y s)`

Function `gprintf` behaves like function `printf`, but prints out the text on the current graphic window, with the current font and color, starting at location `(x , y)`.

(font [fontname])

Type: DX

See: `(draw-text x y s)`

Function `font` sets the font used for rendering characters in subsequent calls of function `draw-text`. It always returns the current font name.

Font `"default"` is recognized by all the drivers, and sets a reasonably small default font. Always using font `"default"` is the only guaranteed way to write fully portable programs.

All drivers however recognize PostScript(tm) style font names. These font names have the following form:

```
"<family>[-<style>] [-<size>]"
```

- The recognized font family `family` usually depends heavily on the operating system configuration. It is reasonable to assume that the basic families `"Times"`, `"Helvetica"` and `"Courier"` are available (or translated) everywhere.
- The optional style is preceded by a dash. Style `"Roman"` lets you print plain text. More complex style are composed by concatenating an optional weight specification (e.g. `Light`, `DemiBold`, `Bold`, `Black`, etc...) and an optional slant specification (i.e. `Italic` or `Oblique`).
- The optional size is a number between 1 and 128.

Not all combination of families, styles and size are supported on all systems. A good compromise between fancy display and portable graphics consists in using PostScript font names, but to limit yourself to the following fonts:

Courier-X, Courier-Bold-X, Courier-Italic-X, Courier-BoldItalic-X
 Times-Roman-X, Times-Bold-X, Times-Italic-X, Times-BoldItalic-X,
 Helvetica-X, Helvetica-Bold-X, Helvetica-Oblique-X,
 Helvetica-BoldOblique-X, Symbol-X

where size X is one of 8 , 10 , 11 , 12 , 14 , 18 , 24 .

Most drivers also understand window system dependent font names (when such font names are defined by the window system). For instance, the X11 driver recognizes all font names listed by the utility program `xlsfonts` . These font names are very long, but a star `*` can be used as a wildcard character. A typical example is `"*-times-medium-r*--18-"` . Function `x11-fontname` can be used to translate a Postscript(tm) style font name into a system specific font name.

(rect-text x y s)

Type: DX

Returns a list `(x y w h)` describing the rectangle that would be affected by a call to `(draw-text x y s)` .

A few driver , including the PS driver , do not implement this function. In that case , `rect-text` returns the empty list.

(text-width s)

Type: DE(graphenv.lsh)

See: `(rect-text x y s)`

Returns the width of the text in string `s` , if it is printed with the current font. This function uses `rect-text` if it is available, or uses some crude heuristics.

(text-height s)

Type: DE(graphenv.lsh)

Returns the height of the text in string `s` , if it is printed with the current font. This function uses `rect-text` if it is available, or uses some crude heuristics.

4.17.5 The Drawing Context.

Several features define the drawing context. The most usefull ones are the current driver , the current font , the current color and the clipping rectangle.

(gsave l1 ...ln)

Type: DY

Function **gsave** saves the graphics state of the current window, evaluates lists **l1** to **ln** with a call to **progn**, and restores the saved graphics state. Function **gsave** returns the result of the last evaluation.

(gdriver)

Type: DX

Returns the name of the graphic driver of the current window.

(xsize)

Type: DX

Returns the width of the current window.

(ysize)

Type: DX

Returns the height of the current window.

window

Type: VAR

This variable **window** defines the current window. Most graphics functions refers implicitly to the window descriptor stored in this variable.

(clip [x y w h])

Type: DX

When a clip rectangle has been set, graphics output is restricted to the inner part of the clip rectangle. the **clip** function allows for manipulating the clip rectangle.

- **(clip x y w h)** sets a new clip rectangle whose top left corner is located at position (**x** , **y**), whose width is **w** , and whose height is **h** .
- **(clip ())** cancels clipping and unset the clipping rectangle.
- **(clip)** just returns the current clipping rectangle.

Function **(clip)** always returns the previous clip rectangle, as a list (**x y w h**) , or the empty list if no previous clipping rectangle was set.

(addclip rect)

Type: DX

Argument `rect` must be a list of the form `(x y w h)`, where `x`, `y`, `w`, and `h` are numbers.

Function `addclip` sets the current clip rectangle to the intersection of the current clipping rectangle with a rectangle whose top left corner is located at position `(x, y)`, whose width is `w`, and whose height is `h`.

If this intersection is empty, the empty list is returned. Otherwise, `addclip` returns the new clip rectangle.

4.17.6 Double Buffering and Synchronization.**(graphics-batch l1 ... ln)**

Type: DY

Most graphics drivers can work with an off-screen bitmap to update the screen after several graphics commands. The `graphics-batch` function provides a support for this double buffering abilities.

This function evaluates the lists `l1` to `ln`, but if a graphic instruction is executed, the screen update is delayed until the end of these evaluations, and is performed only once.

(graphics-sync)

Type: DX

Inside a `graphics-batch` construct, this function immediatly updates the screen without waiting for the completion of the `graphics-batch` instruction.

4.17.7 Plotting (obsolete).

CAUTION: The functionalities described in this section are obsolete and superseded by the `libplot/plotter` library (see the plotting section in the standard library documentation).

Functions in the present section are maintained for backward compatibility, but their use should be avoided.

The following functions provide a way to plot curves. The basic abstraction is the plot port data structure, a Lush list which contains:

- The position of the current point,
- A pointer to the graphics window,
- Information for mapping real coordinates to pixel positions.

The plotting functions perform their drawing action in the current plot port. Function, `setup-axes` is provided for quickly initializing a plot port, making it current and drawing the axes.

Currently, plot ports are implemented as lists rather than a user defined class. In fact, the "`graphenv.lsh`" package has existed prior to the object extensions of Lush. This silly implementation has been kept for preserving backwards compatibility.

(new-plot-port brect rect object)

Type: DE(`graphenv.lsh`)

(OBSOLETE) Initializes a new plot port in the current window, and returns it.

Argument `brect` is a list (`bxmin bymin bxmax bymax`) whose elements are the pixel coordinates of the target rectangle in the current window.

Argument `rect` is a list (`xmin ymin xmax ymax`) whose elements are the extremal real coordinates. No clipping is performed, except window clipping.

Argument `object` is a function generally called by `plt-plot` for drawing a small object at the current pen location. A few predefined functions are documented with `plt-plot` .

(copy-plot-port p [object])

Type: DE(`graphenv.lsh`)

(OBSOLETE) Returns a new plot port, with the same real to pixel coordinates and window data as the plot-port `p` , but with a different current point and a different object. This function allows the user to plot a second curve simultaneously, using the same axis system.

(draw-axes brect x1 y1 name [x21] [y21])

Type: DE(`graphenv.lsh`)

(OBSOLETE) Draws axes on the leftmost and bottommost side of rectangle `brect` .

Argument `brect` is a list (`bxmin bymin bxmax bymax`) whose elements are the pixel coordinates of the target rectangle in the current window.

Arguments `x1` and `y1` are lists of label values for the `x` and `y` axis. These labels are numbers expressed with the same scale than the axes. The optional functions `x21` and `y21` are used to convert these numbers into strings to display at the proper positions along the axes. Both functions default to `str` which just prints the textual representation of the number.

Finally, argument `name` is a string that will be printed at the top of the `brect` using a 18 points font.

plot-port

Type: VAR(`graphenv.lsh`)

(OBSOLETE) The current plot port is stored in this global variable. Most of the plotting functions refers implicitly to the plot port stored in `plot-port` .

(setup-axes x1 y1 x2 y2 [xst yst [s [object]])

Type: DE(graphenv.lsh)

See: (new-plot-port **brect** **rect** **object**)

See: (draw-axes **brect** **x1** **y1** **name** [**x21**] [**y21**])

(OBSOLETE) Calls both **new-plot-port** and **draw-axes** , for drawing a curve in the current window. Makes the resulting **plot-port** current. The curve will occupy the entire window. Function **setup-axes** is the simplest way to initialize the plotting package.

Arguments **x1** and **y1** are the lower bounds of the actual coordinates; arguments **x2** and **y2** are the upper bounds of the actual coordinates. The axes will be labeled every **xst** units horizontally and every **yst** units vertically. Argument **s** is a string that will be printed at the top of the window. **object** is the default plot object used by **plt-plot** .

When arguments **xst** or **yst** are not provided (or are the empty list), function **setup-axes** computes a reasonable value using the extrema.

(plot-lists lx ly)

Type: DE(graphenv.lsh)

(OBSOLETE) Plots a curve taking using in list **lx** as abscissas and values in **ly** as ordinates in the current plot port. This function uses **plt-draw** and **plt-plot** to draw a symbol at each data point.

(graph-xy lx ly name)

Type: DE(graphenv.lsh)

(OBSOLETE) Creates a convenient plot port in the current window, draws the axes, and calls the function **plot-lists** .

(plot-lists-sd lx ly lsd)

Type: DE(graphenv.lsh)

(OBSOLETE) Plots a curve taking values in list **lx** as abscissas and values in **ly** as ordinates. Argument **lsd** is a list of values that will be displayed as uncertainty bars around each data point.

This function uses **plt-draw** , **plt-sd** and **plt-plot** to draw a symbol and an uncertainty bar at each data point.

(graph-xyv lx ly lsd name)

Type: DE(graphenv.lsh)

(OBSOLETE) Creates a convenient plot port in the current window, draws the axes, and calls the function **plot-lists-sd** .

(plt-clear)

Type: DE(graphenv.lsh)

(OBSOLETE) Clears the current point. The the next `plt-draw` will thus not draw a line.

(plt-move x y)

Type: DE(graphenv.lsh)

(OBSOLETE) Sets the current point to real coordinates `x` and `y` .

(plt-draw <x y)

Type: DE(graphenv.lsh)

(OBSOLETE) Draws a line from current point to actual coordinates `x` and `y` . Sets the current point to coordinates `x` and `y` . Nothing is drawn if no current point was set.

(plt-plot x y)

Type: DE(graphenv.lsh)

See: `current-object`See: `object-size`

(OBSOLETE) Moves the current point to the actual coordinates `x` and `y` , and draws a small symbol like those used for showing the data points on a curve.

The shape of this symbol is defined by the argument `object` of `new-plot-port` or `setup-axes` . this argument must be a function that draws the object at the specified position. The following functions are predefined:

- `object-nil` does nothing. It is useful if you want cancel the effect of `plt-plot` .
- `open-square` , `closed-square` , `open-circle` , and `closed-circle` respectively draw a small hollow square, a small filled square, a small hollow circle or a small filled circle.
- `open-up-triangle` , `closed-up-triangle` , `open-down-triangle` , `closed-down-triangle` , draw various kind of small triangles.
- `straight-cross` and `oblique-cross` draw small crosses.

All these predefined object functions use the global variable `object-size` for defining the size of the object drawn by `plt-plot` . The default object may be overridden by setting the variable `current-object` to a new object function.

(plt-sd x y sd)

Type: DE(graphenv.lsh)

Draws a standard deviation bar given by **sd** around point (*x* , *y*). The horizontal size of the bar is defined by the global variable **object-size** .

(in-plot-port l1 ... ln)

Type: DMD(graphenv.lsh)

(OBSOLETE) Function **in-plot-port** unpacks the current plot port info and stores it in local variables while executing *l1 ... ln* . You may then call the **(plt-rbx x)** and **(plt-rby y)** functions to compute the pixel coordinates of a real point (*x* , *y*).

current-object

Type: VAR(graphenv.lsh)

See: **(plt-plot x y)**

(OBSOLETE) If **current-object** is equal to the empty list, the function **plt-plot** will draw the default object defined in the current port. However you can store a new object definition function in **current-object** . This function will then be called instead of the default one.

object-size

Type: VAR(graphenv.lsh)

See: **(plt-plot x y)**

See: **(plt-sd x y sd)**

(OBSOLETE) This global variable controls the size of the drawings performed by **plt-plot** and **plt-sd** .

4.17.8 Events.

See: Ogre.

Several functions are provided for handling mouse and keyboard events occurring on graphic windows. Lush has a centralized event queue (see a full description in the Events and Timers section). It is possible to obtain the events occurring on a window by creating an **EventLock** object. Messages **read-event** and **check-event** allow for testing and waiting for events.

This mechanism is used by a few handy functions **get-click** , **get-vector** and **get-rect** , as demonstrated by the incredibly shrinking paint program **lushpaint** below:

```
(de lushpaint()
  (let ((window ()))
    (new-window)
```

```

(draw-line 2 5 18 5)
(draw-rect 22 2 18 8)
(fill-rect 42 2 18 9)
(gprintf 102 10 "Quit")
(clip 0 12 1000 1000)
(let ((ok t)
      (mode draw-line)
      (echo get-vector))
  (while ok
    (let ((x y) (get-click))
      (cond
        ((> y 10)
         (apply mode (echo)) )
        ((point-in-rect x y '(0 0 20 10))
         (setq mode draw-line echo get-vector) )
        ((point-in-rect x y '(20 0 20 10))
         (setq mode draw-rect echo get-rect) )
        ((point-in-rect x y '(40 0 20 10))
         (setq mode fill-rect echo get-rect) )
        ((point-in-rect x y '(100 0 20 10))
         (setq ok ()) ) ) ) )
  (delete window) ) )

```

(NOTE: more serious paint programs should probably be written with Ogre).

Graphic Event Functions.

(new EventLock window) Type: CLASS(graphenv.lsh)

Creating an **EventLock** object sets up window **window** to record graphic events in the event-queue. Events can be retrieved with messages **read-event** and **check-event** .

The initial state of the window is restored when the **EventLock** object is destroyed.

(=> EventLock read-event) Type: MSG(graphenv.lsh)

See: Event Lists.

See: (eventinfo)

This function returns the first available event on the window associated to the **EventLock** object. Each event is represented by a list whose first element broadly describes the event type. Additional information can be obtained using function **eventinfo** .

If the event queue is empty, this function blocks until an event occurs.

(=> EventLock check-event) Type: MSG(graphenv.lsh)

See: Event Lists.

See: (eventinfo)

This function returns the first available event on the window associated to the **EventLock** object. Each event is represented by a list whose first element

broadly describes the event type. Additional information can be obtained using function `eventinfo` .

This function returns the empty list when no events are available.

(hilite mode x1 y1 x2 y2) Type: DX

The `hilite` function is used for displaying transient drawings when the mouse button is depressed. As soon as the mouse button is released, these transient drawings are cleared, and the window is refreshed.

Four kinds of drawings are supported, whose names are defined by "`graphenv.lsh`" :

- When `mode` is `hilite-invert` , the rectangle defined by the points (`x1` , `y1`) and (`x2` , `y2`) is inverted. The result however is poor on color displays.
- When `mode` is `hilite-rect` , the rectangle defined by the points (`x1` , `y1`) and (`x2` , `y2`) is outlined with a dashed line.
- When `mode` is `hilite-line` , a dashed line is drawn from point (`x1` , `y1`) to point (`x2` , `y2`).
- When `mode` is `hilite-none` , any transient graphics are cleared, and the window is refreshed.

Transient drawings created with `hilite` are very efficient, and totally disappear when the mouse button is released. This function is useful for providing an echo to the user during a mouse drag operation. For example, function `hilite` can display a rectangular outline while the user is dragging the mouse for selecting a part of an image.

(get-click [rect]) Type: DE(`graphenv.lsh`)

Waits for a mouse click on the current window. Returns a list (`x` `y`) giving the coordinates of the mouse during the mouse click. Mouse clicks will be ignored outside the rectangle specified by the optional argument `rect` .

(get-vector [rect]) Type: DE(`graphenv.lsh`)

Waits for a mouse down on the current window. Tracks the mouse moves with a dashed line, until the mouse button is released. Returns a list (`x1` `y1` `x2` `y2`) telling where the mouse button has been depressed and where it has been released.

The optional argument `rect` indicates a bounding rectangle which must entirely contain the resulting vector. The echoed dashed line reflects then this limitation.

(get-rect [rect]) Type: DE(`graphenv.lsh`)

Waits for a mouse down on the current window. Tracks the mouse moves with a dashed rectangle, until the mouse button is released. Returns a list (`x` `y` `w` `h`) telling the location of the top left corner of the rectangle, its width and its height.

The optional argument `rect` indicates a bounding rectangle which must entirely contain the resulting rectangle. The echoed dashed rectangle reflects then this limitation.

Event Lists.

See: (eventinfo)

Events returned by methods `read-event` and `check-event` are encoded as lists, according to the following templates:

- (`mouse-down x y`) : The mouse button has been depressed at position (`x` , `y`). You may use function `eventinfo` to query the identity of the button and the status of the shift and control keys.
- (`mouse-drag x_1 y_1 x_2 y_2`) : The mouse button has been depressed at position (`x_1` , `y_1`). The mouse has then been moved to position (`x_2` , `y_2`). The mouse button still is depressed. You may use function `eventinfo` to query the identity of the button and the status of the shift and control keys.
- (`mouse-up x_1 y_1 x_2 y_2`) : The mouse button has been depressed at position (`x_1`,`y_1`). The mouse has then been moved to position (`x_2`,`y_2`), where the button has been released. You may use function `eventinfo` to query the identity of the button and the status of the shift and control keys.
- (`"c" x y`) : The key associated to the single character string "c" has been hit, while the mouse pointer was located at position (`x` , `y`). These events often are referred as keypress events.
- (`arrow-left x y`) , (`arrow-right x y`) , (`arrow-up x y`) or (`arrow-down x y`) : The left (respectively right, up and down) arrow has been hit, while the mouse pointer was located at position (`x` , `y`). You may use function `eventinfo` to query the exact identity of the depressed key and the status of the shift and control keys.
- (`help x y`) : The keyboard dependent help key has been hit, while the mouse pointer was located at position (`x` , `y`). You may use function `eventinfo` to query the exact name of the depressed key and the status of the shift and control keys.
- (`fkey x y`) : A function key has been depressed. You must use function `eventinfo` to get more information about the identity of the function key and of the status of the shift and control keys.
- (`resize w h`) The window has been resized by the user, and is now `w` x `h` pixels large.
- (`delete`) : Windowing systems often provides a mean to delete a window. For example, X11 ICCM compliant window managers are able to send a `WM_DELETE` message to a client. If an event handler is attached to the window, a delete event is added to the event queue, and the window is not destroyed. A program thus gets a chance to display a confirmation dialog before deleting the window.

- (`sendevent x y`) : These event is generated by function `sendevent` . Integer values `x` and `y` are arbitrary.

These few events give enough information for building graphics interfaces. The class library "`libogre/ogre.lsh`" is designed for that purpose.

4.17.9 System Specific Graphic Functions

Lush provides also system specific functions. The name of these functions always begin with the name of the supported graphic system. For example , "x11" is the prefix for X-Windows specific functions and "win" is the prefix for Windows (95 and NT) specific functions.

Specific graphic functions for X11

The following functions are only implemented on Unix systems implementing the X Windows standard. This includes "Motif" and "OpenWindows" based systems.

(`x11-window` [[[`x y w h`] `name`]]) Type: DX

See: (`new-window` [[`x y w h`] [`name`]])

This is the low level function called by `new-window` on X Windows systems.

This function creates a window of size `w` and `h` at position `x` and `y` named `name` on the default X Windows screen. During the first call of function `x11-window` , Lush searches the variable `display` and the environment variable `DISPLAY` for the name of the X Windows server and display to use.

Unlike `new-window` , this function does not store the newly created window in the variable `window` , and thus does not make this window current.

(`x11-fontname s`) Type: DX

See: (`font` [`fontname`])

This function comes with the X11 driver. It converts a legal PostScript(TM) font name, into a X11 font name.

String `s` must be a legal PostScript font name, composed of a font family, an optional font style, and a font size, possibly separated by dashes, like "`Helvetica-18`" , "`Times-Roman24`" , "`Courier-Bold-12`" . This function returns the corresponding font name under the X11. conventions.

Example:

```
? (when x11-fontname
  (x11-fontname "Times-Roman18") )
= "-*-times-medium-r-normal-*-18-*
```

(`x11-depth`) Type: DX

Returns the depth of the current screen.

(`x11-lookup-color name`) Type: DX

Lookup the (R , G , B) components of a color named `name` .

(`x11-configure` [`raise x y w h`]) Type: DX

Changes the window location (*x* , *y*) and size (*w* , *h*). Giving the empty list as argument leaves the corresponding characteristic unchanged. When argument **raise** is true , the window is deiconified and raised.

This function works only with X11R4 or greater. When effective , it returns a list (**visible w h**) whose elements indicate the window visibility and size.

(**x11-text-to-clip text**) Type: DX

This function stores string **text** into the X11 copy/paste mechanism. This function will some day recognize the copy/paste protocol used and behave accordingly.

It implements today a cut buffer based mechanism suitable for exchanging text with emacs and xterm. We plan to support selection as soon as possible.

(**x11-clip-to-text**) Type: DX

This function retrieves a string from the X11 copy/paste mechanism. It returns the empty list () if no string is available. This function will some day recognize the copy/paste protocol used and behave accordingly.

It implements today a cut buffer based mechanism suitable for exchanging text with emacs and xterm. We plan to support selection as soon as possible.

Specific graphic functions for Windows

[note: this function is not available in the current version of Lush]

The following functions are only implemented on Windows

(**wbm-window** [[[*x y*] *w h*] *name*])

See: (new-window [[*x y*] *w h*] [*name*])

See: (wbm-toplevel-window [[[*x y*] *w h*] *name*])

This is the low level function called by **new-window** on Windows 95 and NT systems. This function is available under WinLush only.

Function **wbm-window** creates a window of size *w* and *h* at position *x* and *y* named *name* in the WinLush environment. The command line version of Lush cannot create graphic windows yet.

Unlike **new-window** , this function does not store the newly created window in the variable **window** , and thus does not make this window current.

(**wbm-toplevel-window** [[[*x y*] *w h*] *name*])

See: (new-window [[*x y*] *w h*] [*name*])

See: (wbm-window [[[*x y*] *w h*] *name*])

This function is very similar to function **wbm-window** described above. Instead of creating a child window of the main WinLush window (according to the Windows MDI scheme), function **wbm-toplevel-window** creates a toplevel window that is not attached to the main WinLush window.

(**wpr-window** [[*w h*] *name*])

See: (print-window *w h* [*destination*])

See: (wpr-printers)

This is the low level function called by **print-window** on Windows 95 and NT systems.

Function **wpr-window** creates a window descriptor which actually accesses the printer named *name* . The default printer is used if argument *name* is not

specified. Function `wpr-printers` should be used to obtain the list of valid printer names.

Unlike `new-window` or `print-window`, this function does not store the newly created window in the variable `window`, and thus does not make this window current.

Arguments `w` and `h` specify the width and height of the coordinate system. The drawings will be scaled to fit the size of the page returned by your printer.

(wpr-printers)

See: `(wpr-window [[w h] name])`

This function returns a list of printers names suitable for function `wpr-window`. The default printer name comes first.

(win-depth)

Returns the depth of the default display of your computer. This is the number of bits specified in the Display Properties dialog box.

(win-configure [raise x y w h])

Changes the window location (`x`, `y`) and size (`w`, `h`). Giving the empty list as argument leaves the corresponding characteristic unchanged. When argument `raise` is true, the window is deiconified and raised.

This function returns a list (`visible w h`) whose elements indicate the window visibility and size.

(win-text-to-clip text)

This function stores string `text` into the clipboard.

(win-clip-to-text)

This function retrieves a string from the clipboard. It returns `()` if the clipboard does not currently contain a string.

Knowing how a system dependent feature is supported.

The function `gdriver-feature` implements a small database about the features supported by the common graphic drivers. We suggest you use this function to implement portable programs.

(gdriver-feature feat opt) Type: DE

This function provides a portable mean to scan the current graphic system for a feature. It returns the function implementing the feature specified by symbol `feat` on the given window. The result depends on the parameter `opt`.

- When `opt` is symbol `loose`, this function always returns a function. It returns the feature function as soon as a partial implementation exists. It returns function `progn` if the feature is not implemented.

This option is usefull for calling minor services (e.g. putting text into the clipboard) for which failure has little consequence.

- When `opt` is symbol `strict`, this function returns a function if and only if the feature is fully implemented. It returns the empty list otherwise.

Features are the following.

- `depth` : This feature is a generalization of `x11-depth` .
- `text-to-clip` : This feature is a generalization of `x11-text-to-clip` .
- `clip-to-text` : This feature is a generalization of `x11-clip-to-text` .
- `configure` : This feature is a generalization of `x11-configure` .
- `fontname` : This feature is a generalization of `x11-fontname` .
- `lookup-color` : This feature is a generalization of `x11-lookup-color` .

Here is an example of the use which can be made of `gdriver-feature` .

```
? (new-window 400 400 400 400 "test")
= ::WINDOW:X11:8c440
? ((gdriver-feature 'configure 'loose) ()) 300 300 400 400)
= (t 400 400)
```

4.17.10 Producing Encapsulated PostScript Files.

Lush can output graphics commands into an Encapsulated PostScript File (EPS). You can then send these files to a PostScript printer or use adequate software packages to modify these files and integrate them into your documents. There is even a Lush command that translates these files into a sequence of Lush drawing commands.

Note: Many computers use incompatible conventions for marking the end of lines (Unix machines use the character NL , Macs use the character CR and PC use a sequence CR NL.) Several popular software exhibit a strange behavior when they attempt to read a sequence of characters which appears to them as a very long line. You must probably use a text filter like `dos2unix` to convert your EPS files.

(ps-window [[w h] name])

Type: DX

See: (print-window w h [destination])

Function `ps-window` is the low level function called by `print-window` on systems that provide no other alternative. All versions of Lush however support Encapsulated PostScript files.

Function `ps-window` returns a window descriptor which actually writes PostScript commands into a file named `name` . When argument `name` is omitted, the PostScript commands are written to file `"tlooutput.ps"` .

Unlike `new-window` or `print-window` , this function does not store the newly created window in the variable `window` , and thus does not make this window current.

Arguments `w` and `h` specify the width and height of the coordinate system. The target rectangle will be scaled to fit a standard A4 or Letter format page.

(gspecial string)

Type: DX

This function sends a system specific command represented by string **string** . This function is ignored if the graphics driver does not recognize the string **string** . In the case of a window descriptor created with **ps-window** , the string **string** is inserted verbatim into the Encapsulated PostScript file.

(ps-play filename [function])

Type: DE

See: (ps-window [[[x y] w h] name])

This function interprets PostScript files created with **ps-window** . Function **ps-play** parses the PostScript file **filename** , converts it into a sequence of SN calls, and applies function **function** to these calls. The default value for the optional argument **function** is **eval** . Function **ps-play** thus plays the contents of the PostScript file in the current window.

Examples:

To play a EPS file into a window, type:

```
? (let ((window (new-window)))
    (ps-play "your_file.ps") )
```

To see the contents of a EPS file, type:

```
? (ps-play "your_file.ps" pprint)
```

Note: This function only interprets PostScript files created by the Lush PostScript driver. It is by no means a complete PostScript interpreter.

(ps-plot filename x y w h [bb-override])

See: (ps-play filename [function])

This function plots EPS-files created with SN or Lush in the window referred to by the symbol **window** in the region determined by **x y w** and **h** . **x y** correspond to the coordinates of the region's upper-left corner, **w** to its width and **h** to its height. The plot is scaled to fill exactly this region using the EPS "%BoundingBox" directive which can be overruled by the optional **bb-override** parameter. If **w** or **h** are 0 or nil, scaling in their respective directions are chosen as to preserve proportions. If both **w** and **h** are 0 or nil, no scaling is performed.

4.17.11 Graphical Utilities.

A family of Lush functions are provided to perform elementary computations on rectangles. A rectangle is stored as a list (x y w h) . The first two elements **x** and **y** are the coordinates of its top left corner; **w** is its width , **h** its height.

(rect-2-ppbrect r)

Type: DE(graphenv.lsh)

Returns rectangle `r` under a format suitable for using with `new-plot-port`.**(window-rect)**

Type: DE(graphenv.lsh)

Returns the boundaries of the current window as a rectangle.

(point-in-rect x y r)

Type: DX

The `point-in-rect` function returns `t` if point `(x , y)` is located inside rectangle `r`.**(rect-in-rect r1 r2)**

Type: DX

The `rect-in-rect` function returns `t` if rectangle `r2` encloses rectangle `r1`.**(collide-rect r1 r2)**

Type: DX

Returns the intersection of `r1` and `r2`. If rectangle `r1` and `r2` do not intersect, `collide-rect` returns the empty list.**(bounding-rect r1 r2)**

Type: DX

Returns the smallest rectangle that encloses both rectangles `r1` and `r2`.**4.17.12 Editable Graphics with comdraw/ivtools**

Author(s): Secil Ugurel

A very useful tool for scientists (and many other people) is the ability to produce editable graphics. Lush provides this feature through the "comdraw" graphic driver. The editability allows one to manually change colors, add text, change line styles, and do all kinds of other things to Lush graphics generated this way.

Opening a `comdraw` window is done simply with:

```
(comdraw-window 10 10 500 500 "My Comdraw Window")
```

where the first four parameters are the desired positions and dimensions of the window. Any Lush graphic command executed while such a window is the current window will be drawn in this window and will be instantly editable. The `graphics-batch` construct has the effect of grouping the objects drawn within its scope.

Installing ivtools and comdraw

`comdraw` is part of the `ivtools` package, available at [<http://www.ivtools.org/ivtools/>]. As of Nov 2002, the version available at this site (1.0.6) needs a patch to work well with Lush. This patch will be part of the next version of `ivtools`. In the meantime, it is best to download RPMs of the patched version from the Lush download page.

```
(comdraw-window [[[x y] w h] name])
```

Type: DX

create a `comdraw` window. This function pops-up a `comdraw` window of size `w` and `h` at position `x` and `y` with name `name` on the screen. To keep the created `comdraw` window in the window variable and make it the current window the variable `window` should be set to the result of `comdraw-window`, as in:

```
(setq window (comdraw-window))
```

support for graphics commands

Almost all the graphics functions of Lush are supported by the `comdraw` driver. However, some are not supported, or have a different behavior than with the `x11` driver.

`comdraw` Graphics are considerably slower than `x11` graphics.

Functions `clip` and `addclip` are supported, but due to the lack of an internal clipping function in `comdraw`, only line clipping is performed in `comdraw` windows.

`graphics-batch` groups the object created by the evaluation of the graphics commands within its scope.

The graphics event functions `hilite`, `get-click`, `get-vector`, `get-rect` and the `EventLock` objects, `read-event` and `check-event` functions are not supported. The imaging functions `hinton-map` and `get-image` are not supported.

comdraw Specific Functions

Using the `gspecial` function in Lush, `Comdraw`-specific commands can be sent directly to the `Comdraw` interpreter. Below is a list of `Comdraw` commands (See `Comdraw` Manual for further information). Be aware that the origin of an

original Comdraw window is the bottom left corner. When using standard Lush graphic functions on a comdraw-window , appropriate conversions are performed to set the origin to the upper hand corner. However when `gspecial` is used directly , no such transformation is performed.

```

compview=rect(x0,y0,x1,y1) -- creates a rectangle
compview=rectangle(x0,y0,x1,y1) -- creates a rectangle
compview=line(x0,y0,x1,y1) -- creates a line
compview=arrowline(x0,y0,x1,y1) -- creates an arrow line
compview=ellipse(x0,y0,r1,r2) -- creates an ellipse
compview=text(x0,y0 textstr) -- creates a text string on pixel coordinates x0, y0
compview=multiline(x0,y0[,x1,y1,...]) -- creates a multiline
compview=arrowmultiline(x0,y0[,x1,y1,...]) -- creates a multiline with arrows
compview=openspline(x0,y0[,x1,y1,...]) -- creates an open spline
compview=arrowspline(x0,y0[,x1,y1,...]) -- creates an open spline with arrows
compview=polygon(x0,y0[,x1,y1,...]) -- creates a polygon
compview=closed spline(x0,y0[,x1,y1,...]) -- creates a closed spline
compview=raster(x0,x1,y0,y1) -- creates an empty raster
xylist=center(compview :xy :yx :x :y :scrn) -- returns the center of a compview (de
rectlist=mbr(compview :lbrt :lrbt :scrn) -- returns the minimum bounding rectangle
ptlist=points(compview :scrn) -- returns the point list from a compview graphic
font(fontnum) -- sets the current font from the menu
brush(brushnum) -- sets the current brush from menu
pattern(patternnum) -- sets the current pattern from menu
colors(fgcolornum bgcolornum) -- sets current fore ground and background colors from
nfonts() -- returns the size of font menu
nbrushes() -- returns the size of brush menu
npatterns() -- returns the size of pattern menu
ncolors() -- returns the size of color menus
select([compview ...] :all :clear) -- makes the listed graphics the current selecti
delete(compview [compview ...]) -- deletes the graphic(s)
move(dx dy) -- moves to the current selection
scale(xflt yflt) -- scales the current selection
rotate(degflt) -- rotates the current selection
newgroup=growgroup(groupview compview) -- adds graphic to an existing group graphic
newgroup=trimgroup(groupview compview) -- removes graphic from an existing group gr
compview=setattr(compview [:keyword value [:keyword value [...]]) -- sets attribut
attrlist(compview) -- returns attribute list of a component
compview=frame([index]) -- returns composite component for a frame, defaults to cur
val=at(list|attrlist|compview n :set val) -- returns (or set) the nth item in a lis
num=size(list|attrlist|compview) -- returns size of a list.
update() -- updates the viewer
error=save([pathstr]) -- saves the document (to the pathname)
compview=import(pathstr :popen :next) -- imports graphic file from pathname or URL,
(:next imports next in numeric series)

```

```

export(compview[,compview[,...compview]] [path] :host host_str :port port_int :socket :string
-- exports a graphic file from pathname or URL, or from a command
compview=paste(compview [xscale yscale xoff yoff | a00,a01,a10,a11,a20,a21]) -- pastes a graph
val=pastemode([val] :get) -- toggles or set paste mode, default is 0, always pastes new graph.
compview=addtool(pathname) -- adds button to toolbar based on zero-centered idraw drawing.
ncols() -- onscreen horizontal extent in pixels
nrows() -- onscreen vertical extent in pixels
dx,dy=stod(sx,sy) -- converts from screen to drawing coordinates
sx,sy=dtos(dx,dy) -- converts from drawing to screen coordinates
dx,dy=gtod(compview gx,gy) -- converts from graphic to drawing coordinates
gx,gy=dtog(compview dx,dy) -- converts from drawing to graphic coordinates
handles(flag) -- enable/disable current selection tic marks and/or highlighting
highlight(compview compviewgs) -- set the highlight graphic state for a graphic
zoom(zoomflt) -- zoom by factor
zoomin() -- zoom-in by 2
zoomout() -- zoom-out by 2
pan(px py) -- pan viewer
smallpanup() -- small pan up
smallpandown() -- small pan down
smallpanleft() -- small pan left
smallpanright() -- small pan right
largepanup() -- large pan up
largepandown() -- large pan down
largepanleft() -- large pan left
largepanright() -- large pan right
tilefile(inpath outpath [xsize] [ysiz]) -- tiles pgm or ppm image file
poke(compview x y val) -- pokes pixel value into raster
pcols(compview) -- returns number of columns in a raster
prows(compview) -- returns number of rows in a raster
pflush(compview) -- flushes pixels poked into a raster
pclip(compview x1,y1,x2,y2,x3,y3[,...,xn,yn]) -- clips raster with polygon
alpha(compview [alphaval]) -- sets/gets alpha transparency
-- requires plotmtv and pstoeedit --
barplot([var_str value_float] [...] :title title_str :xtitle xtitle_str
        :ytittle ytittle_str :valtitle valtitle_str :newview) -- displays a barplot
run(filename) -- runs commands from file
quit() -- quits this interpreter
exit() -- exits entire application
pause([msgstr]) -- pauses script execution until C/R

```

Internals

The interaction with the Comdraw process is performed via a pseudo-terminal. Lush commands are converted to Comdraw commands and sent through the terminal. Lush then waits for comdraw to write its answer. This text-based handshake make the whole process a bit slow.

What is different in the patched version of comdraw

The following commands have been added to comdraw to support certain Lush graphic functions (such as colors and raster images). Those functions are not present in v1.0.6 of ivtools. They are present in the patched version of ivtools available from the Lush download page , and (we are promised) from the next official version of comdraw.

- `colorrgb(fgcolornum bgcolornum)`: sets current fore ground and background colors by RGB name. The colorname format is "#RRBBGG" for 8 bits per component.
- `fontbyname(fontname)`: sets the current font by X font name. Fontname is a string.
- `pokeline(compview x y w vallist)`: pokes pixel values of a line listed in vallist into raster compview. w is the width of the line (length of the list).

4.18 Events and Timers**4.18.1 Event Queue.**

See: Event Lists.

See: Ogre.

The Lush interpreter provides a centralized mechanism for queuing and dispatching events. Events are represented by arbitrary non-empty lists. The event lists discussed in section "Event Lists" are merely the events generated by the window system.

Function `sendevent` is the simplest way to generate an event. Besides the event itself, function `sendevent` requires an "event handler", that is to say an arbitrary non null lisp object. The event handler can be merely used as a key to identify the target of the event. Yet it is strongly suggested that the event handler should be a lisp object that recognizes method `handle` .

Besides function `sendevent` , events can be generated by defining a timer using `create-timer` , or by associating an event handler with a graphic window using function `set-event-handler` .

Although events can be polled manually using functions `testevent` , `checkevent` , and `waitevent` , it is often more practical to let Lush dispatch the events automatically. Lush silently polls the event queue whenever it is waiting for input on the console. All available events are dispatched by calling the method `handle` of their event handler with the event itself as a single argument. These automatic call allow the implementation of asynchronous graphic interfaces, like those implemented by the Ogre library.

(set-event-handler w h)

Type: DX

See: `Ogre`.

Associates the event handler `h` with window `w` . The new event handler for window `w` replaces the old one. Providing the empty list as argument `h` simply removes the previous event handler without defining a new one. Once an event handler has been attached to a window, graphics events occurring on that window, like mouse or keyboard interaction, are posted on the lush event queue and dispatched as usual.

(create-timer handler delay [interval])

Type: DX

See: `Timers`

See: `(new Timer [delay [interval]] [callback])`

Creates a timer that generates an event `(timer <timerid >` for handler `handler` every `interval` milliseconds after an initial delay of `delay` milliseconds. Specifying an interval of zero milliseconds creates a one shot timer that fires only once after `delay` milliseconds. This function returns a timer identifier that can be used with function `kill-timer` . Section "Timers" describe a more convenient way to define a timer.

(kill-timer timerid)

Type: DX

See: `(create-timer handler delay [interval])`

Destroys the timer `timerid` .

(sendevent handler event)

Type: DX

Posts and event `event` with event handler `handler` . Argument `event` must be a non empty list. Argument `handler` must be a non null object. It is customary (but not mandatory) to make sure that the event handler recognizes method `handle` for processing automatically dispatched events.

(sendevent x y)

Type: DX

This obsolete form of `sendevent` takes two integer arguments and posts event `(sendevent <x y >` to the event handler associated with the current window.

(testevent [h])

Type: DX

Function `testevent` returns the first event available for handler `h` without removing the event from the queue. When argument `h` is omitted , function `testevent` implicitly uses the event handler associated with the current window.

Function `testevent` works differently when argument `h` is the empty list. It returns the event handler associated with the first pending event `e`, or the empty list if no events are pending.

When no suitable event is pending `e`, function `testevent` returns the empty list without waiting.

(`checkevent` [`h`])

Type: DX

Function `checkevent` returns the first event available for handler `h` and removes it from the queue. When argument `h` is omitted, function `checkevent` implicitly uses the event handler associated with the current window.

Function `checkevent` works differently when argument `h` is the empty list. It returns the event handler associated with the first pending event `e`, without modifying the queue.

When no suitable event is pending `e`, function `checkevent` returns the empty list without waiting.

(`waitevent`)

Type: DX

Function `waitevent` first tests the event queue and returns the event handler associated with the first pending event. Otherwise function `waitevent` waits until an event occurs and returns the associated event handler.

(`process-pending-events`)

Type: DX

Function `process-pending-events` dispatches all pending events. While there are pending events `e`, function `process-pending-events` removes the first event from the queue `q`, checks whether the associated event handler recognizes method `handle` `h`, and `h`, calls method `handle` of the event handler object with the event as single argument.

Function `process-pending-events` returns when no events are available on the event queue. Function `process-pending-events` is implicitly called whenever events become available while Lush is waiting for user input on the console. Long Lush programs can call this function from time to time in order to maintain event driven graphic interface active during the execution of the program.

(`==>` `eventhandler` `handle` `event`)

Type: MSG

See: (`process-pending-events`)

Method `handle` of event handler objects is automatically called when lush dispatches events, either because it is waiting for user input on the console, or because function `process-pending-events` has been called.

4.18.2 Timers.

```
(new Timer [delay [interval]] [callback])
```

Type: CLASS

Creates a new timer that fires every `interval` milliseconds after an initial delay of `delay` milliseconds. Specifying an interval of zero milliseconds creates a one shot timer that fires only once after `delay` milliseconds.

A timer event is posted into the lush event queue when the timer fires. This event is dispatched when lush waits for user input or when function `process-pending-events` is called. Dispatching the timer event causes function `callback` to be called with the timer as a single argument.

Example:

```
? (new timer 1000 (lambda(c) (printf "One second\n"))))
```

Timers only post events if the previously posted timer events have been dispatched timely. This feature prevents the event queue to grow very large when lush is not able to dispatch the events fast enough.

```
(=> timer set delay [interval])
```

Type: MSG

Sets the timer delay `delay` and periodicity `interval` expressed in milliseconds. Passing the empty list as argument `delay` cancels the timer.

```
(=> timer setcall callback)
```

Type: MSG

Sets the timer callback function to `callback` . This function is called whenever lush dispatches events generated by this timer.

4.19 Documentation and Help System

Lush has a convenient documentation system that allows users to:

- browse the Lush documentation using the GUI tool (`helptool`)
- quickly access the documentation of a particular function or topic from the Lush prompt by typing `^Atopic` or by calling `(apropos "topic")` .
- consult HTML , LaTeX , PostScript , or DjVu versions the Lush manual

The documentation system allows developers to easily include the code and its documentation in the same file. It also makes it easy to organize sections and entries into a complete manual.

4.19.1 Using the Online Manual

`^Atopic`

Typing `^A` (caret then A or control-A) followed by a topic at the Lush prompt is the easiest and quickest way to obtain help. This command displays a numbered list of those entries and sections of the manual that contain `topic` as a substring. The user is then asked to pick one of the choices by entering the corresponding number. Entering `q` or `enter` returns to the Lush prompt. If the list has only one match, the corresponding manual entry is shown right away. The first time this command is invoked, it parses and reads the manual. The `^A` macro-character calls the `(apropos "<topic ")>` function described below. Example:

`? ^Aregex-match`

```
(regex-match <r> <s>) [DX]
```

Returns `t` if regular expression `r` exactly matches the entire string `s`. Returns the empty list otherwise. Example:

```
? (regex_match "(+|-)?[0-9]+(\\|\\. [0-9]*)?" "-56")
= t
```

```
= ()
```

`(apropos topic-string [force])`

Type: DE

This command displays a numbered list of those entries and sections of the manual that contain `topic` as a substring. The user is then asked to pick one of the choices by entering the corresponding number. Entering `q` or `enter` returns to the Lush prompt. If the list has only one match, the corresponding manual entry is shown right away. The first time this command is invoked, it parses and reads the manual.

If the optional boolean parameter `force` is non-nil, the Lush manual is re-read.

`(helptool [book [title]])`

Type: DM(lsh/libogre/helptool.lsh)

Author(s): Yann LeCun Last Modified: 2002-10-17 17:51:38

Helptool is a GUI-based browser and search engine for the Lush manual. It is invoked by simply typing `(helptool)` at the Lush prompt.

Helptool provides an "explorer-like" GUI to browse through the sections and entries of the documentation. The optional argument `book` can be one of the following:

- if absent or nil, the default documentation book is read from `lsh/lush-manual.hlp` (if necessary) and displayed. This book is read only once in a session.
- if equal to `t`, the default document book is read (or re-read if it was previously read), and displayed.
- if equal to a string, a documentation book is read from the file passed in the string (which can be a `.hlp` or a `.lsh` file). The second optional argument is used as the title of the book. If the second argument is not present, the first argument will be used as the title. This is a convenient way of browsing the documentation of a Lush program that is not part of the standard package. It is also a convenient way to use Lush as a documentation browser for other projects.

The Helptool window is composed of an "explorer-like" area on the left that contains a hierarchical list of sections and entries, and a document display area on the right that shows the body of the selected entry. Entries can be opened and closed with a mouse click on the pink square icon. The content of an entry can be displayed by simply clicking on it.

Entering a string or a regular expression in the search tool at the top of the window will "filter" entries that contain that string or match that regular expression. Hitting the `enter` key or the `next` button (or the `space` key in the explorer area) will jump from matching entry to matching entry.

Most users will prefer to use the keyboard shortcuts that provide a very fast and convenient way to navigate the manual. The explorer area, the document area, and the search tool have different keyboard shortcuts. It should be noted that the various areas must be clicked on to get the keyboard focus. The most important key in the explorer area is the space bar, which allows for continuous sequential reading.

Here are the explorer area keyboard shortcuts:

- `down-arrow` or `n` : move to the next entry (without displaying the document)
- `up-arrow` or `p` : move to the previous entry at the same hierarchical level (without displaying the document)
- `right-arrow` : open the current entry (show the subentries) and display the body.
- `left-arrow` : close the current entry (hide the subentries), or move to the parent entry if the current entry is already closed (or if it has no subentries).

- **space** : This provides a simple way to read to manual sequentially. Hitting space will display the body of the current entry in the document area if it is not already being displayed. If it is already being displayed, the document area will be scrolled down. If the document area is already at the bottom, it will go to the next entry that matches the current search string (or to the next entry if the search string is empty).
- **backspace** : scroll up the document area.
- **enter** : go to the next entry matching the search string (or the next entry if the search string is empty) and display the corresponding body in the document area.
- **c** : close all open entries (collapse the entire tree) and go to the title entry.

Examples of how to invoke Helptool:

- `(helptool)` : normal call.
- `(helptool t)` : force re-reading of lush manual
- `(helptool "sdl/sdl.hlp")` : read all entries in file `sdl.hlp` and display as a book entitled "sdl/sdl.hlp".
- `(helptool "sdl/libSDL.lsh" "LibSDL")` : read all entries in `libSDL.lsh` and display as a book entitled "LibSDL".
- `(helptool "/home/yann/blah.hlp" "my blah")` : read all entries in `/home/yann/blah.hlp` and display as book entitled "my blah".

At startup, `helptool` only loads the title lines of each section or entry, but not the bodies. The bodies are read on demand and cached. They are re-read if the file where they reside has been modified since last time it was read.

Known bug: Reading the entire Lush manual takes several seconds.

4.19.2 Writing Documentation

Documentation is either extracted from comments in Lush files (with `.lsh` extension), or extracted from Lush documentation files (with `.hlp` extension).

Help entries can contain text, and/or can contain other help entries as subtopics. Entries can be organized hierarchically into books that can be browsed graphically using `(helptool)`. Entries in a file are added to the main documentation by inserting the file name at the desired location in one of the `.hlp` files read by `(helptool)`.

Looking at or at any `.lsh` or `.hlp` file in, , or is a good way to learn how to organize Lush documentation files or the comments of a Lush file in order to provide documentation.

Creating a new Section or Entry

Sections or entries are created in .lsh or .hlp files by a title line that starts with the compound macro-character `#?` . Title lines can create new sections , describe variables , functions , classes , and methods.

Here are a few examples of typical title lines:

- `#? *** All about Choucroute Garnie`
: Add a Section
- `#? ** << choucroute-garnie.lsh`
: Include .lsh file
- `#? << choucroute-garnie.hlp`
: Include .hlp file
- `#? choucroute-garnie`
: Describe Variable
- `#? (garni-la-choucroute <saucisse> <patates>)`
: Describe Function
- `#? * choucroute`
: Describe Class
- `#? * (new choucroute)`
: Describe Constructor
- `#? (==> <choucroute> garni <saucisse> <patates>)`
: Describe Method

As shown above , the `#?` of each entry can optionally be followed by one or more asterisks. The number of asterisks specifies the hierarchical level of the entry in the manual relative to the other entries in the same file. Entries with fewer asterisks than a previous entry in the same file will become subtopics of that previous entry:

```
#? *** My Section Title.
my section text goes here
#? ** My Subsection Title 1
my subsection 1 text goes here
#? ** My Subsection Title 2
my subsection 2 text goes here
```

Entries that document functions , classes , methods , and variables , should be written according to the following convention:

- **Function:** Function description entries indicate how to call the function. Arguments are indicated by angle brackets. Here are a few examples:

```
#? (car <list>)
#? (+ <n_1> ... <n_n>)
#? (new-window [<w> <h>] <name>)
```

- **Variable:** Description of variables should simply include the variable name on the title line. It is recommended that the text include a

```
{<type> VAR}
```

directive.

- **Class:** Class descriptions are usually associated with an entry for its constructor, indicating how to create an instance. Method documentation entries show how to call the method and should be subheadings of the constructor entry:

```
#? * (new helpinfo <record>)
#? (==> <helpinfo> tty-print)
```

A special directive in `.hlp` files indicate that entries at that location should be read from another file. The following example inserts all toplevel entries from file `"toto.hlp"` as a section marked with two stars (the stars of the entries in the file `"toto.hlp"` are only used to place the those entries relative to each other).

```
#? ** << toto.hlp
```

Document Body: Introduction

The body (text) of an entry or section can be written following the title line.

In a `.hlp` file, the body of an entry is whatever text is found until the end-of-file or the next entry title line (i.e. the next `#?` at the beginning of a line). An empty line is interpreted as a paragraph boundary (but a line with only spaces is not).

In a `.lsh` file, the body of an entry is whatever bloc of consecutive comment lines (i.e. lines that begin with one or more semicolons) that follow the title line. By convention, the lines should begin with two semicolons. A lines with only semicolons is interpreted as a paragraph boundary.

The body of an entry or section can be written according to two different syntaxes. The "normal" syntax comprises a set of curly-brace-enclosed tags that have the same name and more or less the same function as the corresponding HTML tags. The "old" syntax (present for backward compatibility) comprises TROFF/NROFF-like tags that must be at the beginning of a line and begin with a dot. Blank lines indicate paragraph boundaries in both syntaxes

Debugging a documentation file or a Lush file with documentation entries can be done easily by calling (`helptool "the-file-to-be-debugged"`) . The entry bodies are re-read whenever the file is modified , but the entries themselves (the title lines) are not.

Document Body: the Brace Syntax

The brace document format consists of segments of text enclosed between curly braces with a tag immediately following the opening brace. In this syntax , the following characters or character sequences have a special meaning:

- **two carriage returns (blank line)** : paragraph break
- **tag...** : a brace tag expression
- **,(...)** : read Lush expression in paren and insert it in the document. It will be evaluated at rendering time. If the result is a string , it will be inserted in the text.
- **\(...)** : same as above
- **,\$expr** : read Lush object following dollar sign and insert it in the document. Unlike with the comma/paren construct , this allows to insert symbols and string literals. The Lush object will be evaluated at rendering time. If the result is a string , it will be inserted in the text.
- **\\$expr** : same as above
- **„expr** : read and evaluate the Lush expression and substitute the result (preferably a string) in the text. Unlike with with the single comma constructs , the evaluation is performed at read time.
- **...** : display content of bracket as bold/highlighted (mainly used for short Lush code snippets or arguments of the function being described).

To insert text that contain the above combinations requires escaping. This can be done with the **,\$"any legal Lush string"** construct.

A brace expression is an expression of the form:

```
{<tag1> word1 {<tag2> word21 word22} word3 .... wordn}
```

A brace expression is read by the Lush reader just like a regular Lush expression. The reader translates the above example into a list like this:

```
(<tag1> -1 1 "word1" 1
  (<tag2> -1 1 "word21" 1 "word22")
  1 "word3" 1 "... " 1 "wordn")
```

Such brace lists produce the appropriate document when evaluated in the right context.

Here are the tags. Three dots indicate that text or other brace expression can be inserted.

Basic text formatting tags:

- **br** : line break
- **p** : paragraph break
- **li ...** : list item. Preferrably (though not necessarily) appears within a **ul** tag.
- **ul li ...** : unstructured list with list items
- **div ...** : does nothing , just groups its arguments
- **pre ... /pre** : preformatted text (no need to escape special characters)
- **code ... /code** : preformatted code displayed in blue (no need to escape special characters)
- **img ...** : insert image

Text attribute tags :

- **b ...** : turn to bold font
- **i ...** : set font to italic
- **c ...** : highlight , and turn bold (for code and arguments)

Metatags: The content of these tags is not displayed in the main document but is displayed in the header:

- **author ...** : author
- **symbol ...** : which Lush symbol does this entry describe.
- **location ...** : override the name of the file in which this function/entry is defined. By default , this is automatically set to file being read.
- **keywords ...** : keywords for searching this entry
- **date ...** : date of creation/modification of the object being described. By default , this is automatically set to the last modification time of the file in which the entry is defined (if this file is known).

- `title ...` : override for the name of the entry. By default this is set to the entry title line.
- `type ...` : type of object being described. By default this is guessed from the entry title whenever possible.

Special Tags:

- `ex lush-expression` : insert a demonstration of typing the lush-expression at the Lush prompt and getting the result.
- `hlink` : hyperlink to URL or other entry in the document
- `see` : hyperlink to URL or other entry in the document

Conditional tags:

- `if-html ...` : insert only if rendering to HTML
- `if-latex ...` : insert only if rendering to LaTeX
- `if-ogre ...` : insert only if rendering to helptool
- `if-text ...` : insert only if rendering to plain text

Tags that are legal but not yet implemented: these tags are legal but are not functional at the moment and are reserved for future implementations (they currently act as `div`):

- `desc ...` : short one-line description of the entry
- `u ...` : underlined text
- `tt ...` : typewriter font
- `font ...` : select font
- `center ...` : center text
- `h1 ...` : big title
- `h2 ...` : medium title
- `h3 ...` : small title

Document Body: the "Dot-Tag" Legacy syntax

Predecessors to Lush (SN and TL3) used a format for documentation inspired by TROFF/NROFF with formatting tags that begin with a dot and are placed at the beginning of the line. Much of the Lush documentation is still in this format, which is why we document it here.

Within the text, strings enclosed in angled brackets are displayed in boldface and highlighted in a color different from the regular text. This should be used for short segments of Lush code, file names, arguments of the function being documented, and other literals.

Tags are divided in three categories: tags with no argument, tags whose argument is on the same line as themselves, and tags whose "argument" (or scope) is the text until the next tag.

Formatting tags that take no arguments:

- **.P** or **.PP** : paragraph break (blank line plays the same role), or return to paragraph fill mode
- **.BR** : line break (without paragraph skip).

Formatting tags whose argument is on the same line:

- **.HLINK** : hyperlink to another entry of a URL.
- **.SEE** : reference/hyperlink to another entry.
- **.EX** : example, a Lush expression that will be evaluated and whose result is displayed in the document.
- **.IMG** or **.EPS** : the path to an image file (in any format that ImageMagick understands).

Formatting tags whose argument is whatever text is between it and the next tag:

- **.LI** or **.IP** : list item.
- **.PRE** or **.VP** : pre-formatted text.
- **.CODE** : displayed code.
- **.DIV** : switch to paragraph mode (without a paragraph skip like .P)

Conditional inclusion tags (active until the next tag):

- **.IFOGRE** or **.IFHELP** or **IFHLP** : include only in helptool output.
- **.IFTEX** : include only in LaTeX output.
- **.IFTXT** : include only in plain text output.
- **.IFHTML** : include only in HTML output

Metadata tags (whose argument is on the same line):

- **.TYPE** : type of entry (e.g. DX , DE , CLASS , MSG....)
- **.DESC** : short one-line description.
- **.FILE** : file where the definition for the entry resides.
- **.AUTHOR** or **.AUTH** : author(s)
- **.DATE** : date of last modification.

Directive **.EX** lets you run a lisp expression and display it as an example. It is preferable that the expression carries no side effect.

Directive **IMG** allows to include a picture (in any format):

```
.IMG toto.png
```

4.20 Obsolete functions

Functions described in this section are largely obsolete , but are included in the standard library for convenience , and bakward compatibility.

4.20.1 Operations on Lists of Numbers

These are functions that operate on list of numbers interpreted as vectors of real numbers. They are provided mostly for convenience (and for historical reasons) , but are largely obsolete. Serious vector computation should be performed using the much more powerful matrix/tensor engine , rather than with these list functions.

```
(add-lists 11 12)
```

Type: DX

Return the list of the sums term by term of the elements of lists 11 and 12 . This function is equivalent to:

```
(mapcar '+ 11 12)
```

```
(diff-lists 11 12)
```

Type: DX

Return the list of the differences term by term of the elements of lists 11 and 12 . This function is equivalent to:

```
(mapcar '- 11 12)
```

(abs-dist 11 [12])

Type: DX

Returns the absolute (L1) distance between the lists of numbers 11 and 12 . The absolute distance is the sum of the absolute differences between the elements of 11 and 12 .

If argument 12 is omitted, a lists of zeros is assumed.

Example:

```
? (abs-dist (range 1 10))
= 55
```

(sqr-dist 11 [12])

Type: DX

Returns the squared (L2) distance between the lists of numbers 11 and 12 . The squared distance is the sum of the squares of the differences between the elements of 11 and 12 .

If argument 12 is omitted , a lists of zeros is assumed.

(sup-dist 11 [12])

Type: DX

Returns the supremum (L-infinity) distance between the lists of numbers 11 and 12 . The supremum distance is the largest absolute difference between elements of 11 and 12 .

If argument 12 is omitted , a lists of zeros is assumed.

(mean-abs-dist 11 [12])

Type: DE(sysenv.lsh)

See: (abs-dist 11 [12])

Returns the average absolute distance between elements of 11 and 12 . This is defined as the absolute distance divided by the number of elements in lists 11 and 12 .

(mean-sqr-dist 11 [12])

Type: DE(sysenv.lsh)

See: (sqr-dist 11 [12])

Returns the average squared distance between elements of 11 and 12 . This is defined as the squared distance divided by the number of elements in lists 11 and 12 .

(mean-sup-dist 11 [12])

Type: DE(sysenv.lsh)

See: (sup-dist 11 [12])

Returns the average supremum distance between elements of 11 and 12 . This is defined as the supremum distance divided by the number of elements in lists 11 and 12 .

(hamming-dist [margin] 11 [12])

Type: DX

Returns the number of elements of 11 and 12 whose absolute difference is greater than margin .

If argument 12 is omitted , a list of zeroes is assumed.

The default value for margin is 0 . In this case hamming-dist returns the number of elements strictly different in 11 and 12 .

(quadrant-dist 11 12)

Type: DX

Returns the number of elements of 11 and 12 whose sign is different.

(mean-hamming-dist [margin] 11 [12])

Type: DE(sysenv.lsh)

See: (hamming-dist [margin] 11 [12])

Returns the averaged hamming distance between lists 11 and 12 . The maximal result thus is 1.

(mean-quadrant-dist 11 12)

Type: DE(sysenv.lsh)

See: (quadrant-dist 11 12)

Returns the averaged quadrant distance between lists 11 and 12 . The maximal result thus is 1.

4.20.2 Statistical Functions on Lists of Numbers.

These are functions that operate on list of numbers. They are provided mostly for convenience (and for historical reasons) and are largely obsolete. Serious statistical computation should be performed using the much more powerful matrix/tensor engine , rather than with these list functions.

(sup 1)

Type: DX

Return the largest element of the list 1 .

Example:

```
? (sup (range 1 10))
= 10
```

(inf 1)

Type: DX

Return the smallest element of the list `1` .

Example:

```
? (inf (range 1 10))
= 1
```

(rank list target [width])

Type: DX

Returns the list of the indices of the elements of `list` whose distance to `target` is smaller or equal than `width` . If argument `width` is omitted, 0 is assumed.

Example:

```
? (de rank-of-max(1)
    (car (rank 1 (sup 1)))) )
= rank-of-max
? (rank-of-max '(2 1 4 5 -4 -3 -2) )
= 3
```

(a-rank-of-min 1 [comp])

Type: DE(miscenv.lsh)

The function `a-rank-of-min` returns the first rank of the minimum value in a list `1` . The optional argument `comp` is a boolean function which takes two arguments and tests the strict superiority of its first argument on its second.

(a-rank-of-max 1 [comp])

Type: DE(miscenv.lsh)

The function `a-rank-of-min` returns the first rank of minimum value in list `1` . Optional argument `comp` is a boolean function which takes two arguments and tests the strict superiority of its first argument on its second.

(mean 1)

Type: DX

Returns the average of the elements of list `1` .

Example:

```
? (mean (range 1 10))  
= 5.5
```

(median 1)

Type: DX

Returns the median of the elements of list 1 .

Example:

```
? (median (append (range 5 30 2) (range 1 10)))  
= 9
```

(ksmallest 1 k)

Type: DX

Returns the k -th smallest elements of list 1 . Argument k must be an integer between 1 and the size of the list 1 .

Example:

```
? (ksmallest (append (range 5 30 2) (range 1 10)) 10)  
= 8
```

(quantile 1 f)

Type: DX

Returns the f -th quantile of the elements of list 1 . Argument f must be a number between 0 and 1.

Example:

```
? (quantile (append (range 5 30 2) (range 1 10)) 0.9)  
= 25
```

(sdev 1)

Type: DX

Returns the standard deviation of the elements of list 1 .

Example:

```
? (sdev (range 1 100))  
= 28.8661
```

(cov lx ly)

Type: DX

Returns the covariance of the elements of list `lx` versus the elements of list `ly` .

Example:

```
? (cov (range 1 10) (range 11 30 2))
= 16.5
```

(regression lx ly)

Type: DX

Simple monovariate linear regression.

Arguments `lx` and `ly` are two lists of numbers. Function `regression` returns a list of the form `(A B R)` .

- The function $y=Ax+B$ is then the best linear approximation of dependencies between lists `lx` and `ly` .
- The correlation coefficient `R` measures the fit of the approximation. The closer `R` is to 1, the better the approximation is.

Example:

```
? (regression (range 1 10) (range 11 30 2))
= (2 9 1)
```

(x-estimate lr y)

Type: DE(sysenv.lsh)

See: `(regression lx ly)`

Given the result list `lr` of the `regression` , returns a linear estimation of `x` given `y` .

(y-estimate lr x)

Type: DE(sysenv.lsh)

See: `(regression lx ly)`

Given the result list `lr` of the `regression` , returns a linear estimation of `y` given `x` .

```
(sum 1)
```

Type: DX

Returns the sum of the elements of the list `1 .`

Example:

```
? (sum (range 1 10))
= 55
```

4.20.3 DZ.

Lush contains a mechanism called “DZ” which allows to define very efficient numerical functions. This mechanism is based on a small stack machine, and on a DZ compiler. The DZ module is a remnant from an early attempt at compiling SN numerical to make it run faster. It was eventually superseded by the DH function (i.e functions compiled to C), but was kept around.

New DZ functions are easily defined with the functions `dz` , `dz-compile` and `zlambda` . They return an object from of class DZ which can be executed like a 'de' but which can only contain and return numerical expressions. Except for the overhead associated to small DZ functions, typical DZ functions execute 15 to 20 time faster than their DE equivalent.

Two important restrictions are enforced in `dz` function declaration:

- If the DZ refer a symbol defined outside the DZ function, it is evaluated only once at compile time and must return either another previously created DZ function (in which case it is "inlined") or a number (which is then considered a constant). In other word, inside a DZ function only static scoping is allowed.
- Only numerical expression can be used by the small machine, that is functions that return list, string or matrix are not allowed. Other functions like `if` , `for` , `let` are allowed provided they do not push a non numerical expression onto the stack.

Predefined DZs.

A couple of useful DZ functions are predefined.

```
(sgn n) Type: DZ
```

Returns +1 if `n` is equal or greater than 0 , -1 otherwise.

```
? (sgn -2)
= -1
```

```
? (sgn 1.2)
= 1
```

(abs n) Type: DZ
Returns the absolute value of *n* .

? (abs 123.3)
= 123.3

? (abs -23.3)
= 23.3

(int n) Type: DZ
Returns the integral part of *n* .

? (int -4.5)
= -5

? (int 4.5)
= 4

(sqrt n) Type: DZ
Returns the square root of *n*

? (sqrt 15)
= 3.873

(0-x-1 n) Type: DZ
This function implements a piecewise saturated linear function. It returns 0 if *n* is smaller than -1. It returns 1 if *n* is larger than +1. It returns *n* if *n* is in the -1 to +1 range.

? (0-x-1 -2)
= -1

? (0-x-1 0.7)
= 0.7

? (0-x-1 1.3)
= 1

(0-1-0 n) Type: DZ
This function implements the indicator function of the -1 to +1 range. It returns 1 if *n* is in the -1 to +1 range. It returns 0 otherwise.

```
? (0-1-0 -2)
= 0
```

```
? (0-1-0 0.7)
= 1
```

(sin n) Type: DZ
Returns the sine of n radians.

```
? (sin (/ 3.1415 3))
= 0.866
```

(cos n) Type: DZ
Returns the cosine of n radians.

```
? (cos (/ 3.1415 3))
= 0.5
```

(tan n) Type: DZ
Returns the tangent of n radians.

```
? (tan (/ 3.1415 3))
= 1.7319
```

(atan n) Type: DZ
Returns the arc tangent of n , in radians.

```
? (* 4 (atan 1))
= 3.1416
```

(exp n) Type: DZ
Returns the exponential of n .

```
? (exp 1)
= 2.7183
```

(exp-1 n) Type: DZ
Returns the exponential of n minus 1. This function gives an accurate value of $\exp(n)-1$ even for small values of n .

```
? (exp-1 0.5)
= 0.6487
```

(log n) Type: DZ

Returns the natural logarithm of n .

? (log 2)

= 0.6931

(log1+ n) Type: DZ

Returns the natural logarithm of n plus 1. This function gives an accurate value of $\log(1+n)$ even for small values of n .

? (log1+ 1)

= 0.6931

(sinh n) Type: DZ

Returns the hyperbolic sine of n .

? (sinh 1)

= 1.1752

(cosh n) Type: DZ

Returns the hyperbolic cosine of n .

? (cosh 1)

= 1.5431

(tanh n) Type: DZ

See: (qtanh n)

Returns the hyperbolic tangent of n .

? (tanh 1)

= 0.7616

(qtanh n) Type: DZ

See: (tanh n)

Returns a rational approximation of the hyperbolic tangent of n . This function is orders of magnitude faster than **tanh** . Its accuracy however is quite low (about 0.00001.)

? (qtanh 1)

= 0.7616

The formula of the approximation is: $\left[\tanh(x) \approx \frac{P(x)^{16-1}}{P(x)^{16+1}} \right]$ where $\left[P(x) = 1 + 0.125x + 0.0078125x^2 + 0.00032552083x^3 \right]$

(qdtanh n) Type: DZ

Returns a rational approximation of the derivative of the hyperbolic tangent of n . The accuracy of this function is quite low (about 0.00001), but its computation time is very small.

? (qdtanh 1)

= 0.42

The formula of the approximation is: $\left[\frac{d}{dx} (1 - \tanh(x)^2) \approx 1 - \left(\frac{P(x)^{16-1}}{P(x)^{16+1}} \right)^2 \right]$ where $\left[P(x) = 1 + 0.125x + 0.0078125x^2 + 0.00032552083x^3 \right]$

(qstdsigmoid n) Type: DZ

See: (qtanh n)

Returns the value of standard sigmoid at point n , computed with a rational approximation to the hyperbolic tangent. The standard sigmoid is defined as $\left[1.7159 \operatorname{tanh} \left(\frac{2}{3}x \right) \right]$

(qdstdsigmoid n) Type: DZ

See: (qdtanh n)

Returns the value of the derivative of standard sigmoid at point n , computed with a rational approximation to the hyperbolic tangent. The standard sigmoid is defined as $\left[1.7159 \operatorname{tanh} \left(\frac{2}{3}x \right) \right]$

(qexpmx n) Type: DZ

Returns a rational approximation of the exponential of minus the absolute value of n . The accuracy of this function is quite low (about 0.00001), but its computation time may be smaller on certain computers.

? (qexpmx 1)

= 0.3679

The formula of the approximation is: $\left[e^{-x} \approx \frac{1 - P(-x)^8}{P(-x)^8} \right]$ where $\left[P(x) = 1 + 0.125x + 0.0078125x^2 + 0.00032552083x^3 + 0.000010172526x^4 \right]$

(qdexpmx n) Type: DZ

Returns a rational approximation of the derivative of the exponential of minus the absolute value of n . The accuracy of this function is quite low (about 0.00001), but its computation time may be smaller on certain computers.

? (qdexpmx 1)

= 0.3679

The formula of the approximation is: $\left[\frac{d}{dx} e^{-x} \approx -\frac{1}{P(x)} e^{-x} \right]$ where $\left[P(x) = 1 + 0.125x + 0.0078125x^2 + 0.00032552083x^3 + 0.000010172526x^4 \right]$

(qexpmx2 n) Type: DZ

Returns a rational approximation of the exponential of minus the square of n . The accuracy of this function is quite low (about 0.00001), but its computation time may be smaller on certain computers.

```
? (qexpmx 1)
= 0.3679
```

The formula of the approximation is: $\left[e^{-x^2} \approx \frac{1}{P(x^2)} \right]$ where $\left[P(x) = 1 + 0.125x + 0.0078125x^2 + 0.00032552083x^3 + 0.000010172526x^4 \right]$

(qdexpmx2 n) Type: DZ

Returns a rational approximation of the exponential of minus the square of n . The accuracy of this function is quite low (about 0.00001), but its computation time may be smaller on certain computers.

```
? (qdexpmx 1)
= 0.3679
```

The formula of the approximation is: $\left[\frac{d}{dx} e^{-n^2} \approx -2x \frac{1}{P(n^2)} e^{-n^2} \right]$ where $\left[P(x) = 1 + 0.125x + 0.0078125x^2 + 0.00032552083x^3 + 0.000010172526x^4 \right]$

Compiling DZ.

The following functions are useful to define DZ functions.

(dz name args . body) Type: DM

Create a new function which evaluate its argument and which is similar to a **de** except that it works only on numerical expression and is much faster. Sets the symbol **name** to this function.

args is the argument list, made of symbol names.

body is the function itself. Each list in **body** will be evaluated when the function is called. The result of the last evaluation will be returned by the function call.

Example:

```
? (dz fact (n)
  (let ((res 1))
    (for (i 2 n)
      (setq res (* i res)))
    res))
```

```
= fact
? (fact 5)
= 120
?
```

(zlambda args . body) Type: DF

Returns a function which evaluates its argument and is similar to a `de` except that it works only on numerical expression, and is much faster than a `de`.

`args` is the argument list, made of symbol names.

`body` is the function itself. Each list in `body` will be evaluated when the function is called. The result of the last evaluation will be returned by the function call.

Example:

```
? ((zlambda (x)
    (* x x) ) (+ 4 5))
= 81
```

(dz-compile de-function) Type: DE

Returns a function which evaluates its argument and is similar to a `DE` except that it works only on numerical expression, and is much faster than a `DE`.

`de-function` is the `de` function to be compiled.

(dz-load n program) Type: DX

This function returns a new `DZ` which takes `n` arguments.

The list `program` describes which instruction will be executed by the stack machine. Each element in this list can be a label declarations or a stack machine instruction.

- A label declaration is a string. If this string is used as argument of a branch opcode, the branch will jump at the label location. If this string is used as argument of a subsequent stack relative opcode, this opcode will refer to the current top of the stack.
- A stack machine instruction is a list composed of an opcode string and of its arguments. A suffix in the opcode strings defines the opcode addressing mode, and thus defines what are the arguments.

Caution must be exercised when you are programming this machine: An incorrect program may hang SN3. Usually `DZ` programs are composed by the `DZ` compiler, which is assumed correct.

(dz-spline index) Type: DX

See: `(dz-dspline index)`

Generates a `DZ` function that computes a cubic spline interpolations of a set of `n` points specified by the `n` by 2 matrix `index`. Argument `index` thus must

be a n by 2 index. Its first column contains the abscissa, its second column contains the ordinates of the points.

The function `dz-spline` may be used for generating the derivative of this cubic spline interpolation.

(dz-dspline index) Type: DX

See: (dz-spline index)

Generates a DZ function that computes the derivative of a cubic spline interpolations of a set of n points specified by the n by 2 matrix `index` . Argument `index` thus must be a n by 2 index. Its first column contains the abscissa, its second column contains the ordinates of the points.

The function `dz-spline` may be used for generating the cubic spline itself.

(dz-def dz) Type: DX

See: (dz-load `n_1 n_2 program`)

See: DZ opcodes.

This function returns the definition of the DZ function `dz` . This definition is a list (`n . program`) , where n is the number of arguments of the DZ, and `program` is a list of stack machine instructions.

(dz-trace boolean) Type: DX

It function sets or reset the trace mode for the stack machine In trace mode , each instruction is disassembled , and the stack is displayed after the execution of each instruction. This function is mostly useful for designing and debugging the DZ compiler.

This function does nothing if the trace mode has not been compiled. This is the usual setup , for efficiency reasons.

DZ opcodes.

See: (dz-load `n program`)

A DZ program is a list which describes which instruction will be executed by the stack machine. Each element in this list can be a label declarations or a stack machine instruction.

A label declaration is a string. If this string is used as argument of a branch opcode, the branch will jump at the label location. If this string is used as argument of a subsequent stack relative opcode, this opcode will refer to the current top of the stack.

A stack machine instruction is a list composed of an opcode string and of its arguments. Opcodes fall into five broad categories:

- Opcodes whose name ends with a “#” are “immediate” opcodes. They take one numeric argument. For example (“ADD#” 2.2) adds 2.2 to the top of the stack.
- Opcodes whose name ends with a “@” are “stack relative” opcodes. They take one positive integer argument, which describes an element of the stack. For example (“ADD@” 2) adds the third topmost element of the stack to the top of the stack.

If a label is given as operand of a stack relative opcode, it refers to the position of the top of the stack when the label has been defined. In addition, labels "arg1" , "arg2" , etc... refer to the arguments of the DZ.

- Opcodes whose name begins with "BR" are "relative branches". If the branching condition is true, their integer argument is added to the program counter.

If a label is given as operand of a relative branch, it jumps immediately after the label position in the program.

- Two special opcodes, "SPLINE" and "DSPLINE" take a 1D matrix as operand. DZ using these opcodes usually are generated with the `dz-spline` and `dz-dspline` function.
- The other opcodes do not have arguments. For example "ADD" adds the two topmost elements of the stack, and replace them on top of the stack by their sum.

Two other rules are enforced by the `dz-load` function:

The arguments of a DZ are pushed on the stack when the execution starts. The execution stops when the last opcode has been executed. The result must be the only element on the stack. The opcode "POP@" is handy for removing the arguments of the DZ.

It is forbidden to make loops whose iterations would add or remove elements to the stack. The size of the stack at each position of the program thus is determined statically.

Here are a brief description of the opcodes.

Opcode	Stack Before -> Stack After	Notes
NOP	* -> *	
STACK	* -> *	prints the stack.
PRINT	* -> *	prints the top of the stack.
ERROR	* ->	produces a run time error.
POP	x * -> *	
DUP	x * -> x x *	
PUSH# C	* -> C *	
PUSH@ R	x a1..aR * -> aR x a1..aR *	
POP@ R	x a1..aR * -> x *	
SET@ R	x a1..aR * -> a1..aR-1 x *	
MINUS	x * -> -x *	
INVERT	x * -> 1/x *	
ADD1	x * -> x+1 *	
SUB1	x * -> x-1 *	
MUL2	x * -> x*2 *	
DIV2	x * -> x/2 *	

RAND		* -> r *	uniform random number in [0,1[
GAUSS		* -> r *	gaussian random number (m=0,s=1)
SPLINE N DATA	x *	-> f(x) *	computes a spline interpolation
DSPLINE N DATA	x *	-> f'(x) *	computes its derivative

ADD# C		x * -> x+C *
ADD@ R	x a1..aR *	-> x+aR a1..aR *
ADD	x y *	-> x+y *
SUB# C		x * -> x-C *
SUB@ R	x a1..aR *	-> x-aR a1..aR *
SUB	x y *	-> y-x *

Idem for

MUL, DIV, MIN, MAX,
 DIVI (integer division),
 MODI (remainder),
 POWER (exponentiation)

SGN	x *	-> sgn(x) *
ABS	x *	-> abs(x) *

Idem for the following functions, equivalent to their
 lowercase lisp counterpart:

INT SQRT PIECE RECT
 SIN COS TAN ASIN ACOS ATAN
 EXP EXPM1 LOG LOG1P TANH COSH SINH
 QTANH QDTANH QSTDSIGMOID QDSTDSIGMOID
 QEXPMX QDEXPMX QEXPMX2 QDEXPMX2

BR L		* -> *	inconditional relative branch
BREQ L	x *	-> *	branch if x==0
BRNEQ L	x *	-> *	branch if x!=0
BRGT L	x *	-> *	branch if x>0
BRLT L	x *	-> *	branch if x<0
BRGEQ L	x *	-> *	branch if x>=0
BRLEQ L	x *	-> *	branch if x<=0

The following opcodes are handy for programming <for> loops, where
 <i> is the variable, <n> the limit and <s> the increment.

BEGFOR L	i n s *	-> i n s *	check and branch
		i n s *	pop 3 if it does not branch
ENDFOR L	i n s *	-> i+s n s *	increment, check and branch
		i n s *	pop 3 under if it does not branch

Chapter 5

Ogre: Object Oriented GUI Toolkit

5.1 Introduction to Ogre.

See: Graphics.

See: Events.

Lush provides all the necessary graphic primitives for programming a graphics interface.

- Output primitives include functions for opening windows and functions for producing various elementary drawings. Lush supports multiple fonts, multiple windows and even multiple window systems.
- Input functionalities are also present. Lush gathers all events in a common event queue. There are functions for inspecting this event queue (`checkevent` and `waitevent`). Hook functions are called when events are occurring while Lush is waiting for user commands (`event-hook` and `idle-hook`).

It is rather easy to build simple graphic interfaces using these primitives. For instance, the function `snpaint` implements a simple drawing program for drawing lines and rectangles using low level primitives.

Graphic interfaces however become quickly complex. The class browser described in the last chapter of this document contains various features, like menus, buttons, requesters, selectors and scrollbars. A direct program would be overwhelmingly complex.

This is why people have developed various toolkits which handle automatically a large part of this complexity. The Ogre library is such a toolkit entirely written using the Lush object oriented language. This manual describes how to build programs with this library.

In the Ogre library, each component of a graphics interface, a button for example, is implemented as an object which is an instance of a subclass of class `VisualObject` .

The class of such an object defines both the appearance of the object and its response to user input. Defining a new type of object (e.g. a menu) is thus just a matter of defining a new subclass of `VisualObject` .

There is a wide variety of such object classes:

- Basic objects are the major components of an interface. For instance, buttons, strings or check boxes are visible and respond to user interaction.
- Container objects are rectangular areas containing other objects and maintaining a certain layout. For instance there is a class `Column` which maintain several objects aligned in a column.
- Window objects are special container object affected to a window on the screen. Every object inserted in a window object appears into the window and is eligible for receiving user events.

There are two programming levels with the Ogre library. At the simplest level, you just use the provided classes for defining your interface. At the second level, you define subclasses of the standard Ogre classes in order to create new kind of graphic components.

5.2 Ogre Tutorial.

Here is the transcript of a trivial call to the ogre library. You must first initialize the Ogre library by calling function `ogre` :

```
? (ogre)
[ogre.lsh] (autoload)
= idle-hook
```

Let us then open a window `ww` containing a message string and a button arranged in a column named `cc` :

```
? (setq ww
  (new windowobject 0 0 400 300 "Essai"
    (setq cc (new column (new string "Press button to beep")
      (new stdbutton "Beeper"
        (lambda(c) (beep))) ) ) ) )
= ::windowobject:e3100
```

When you click on the button, the callback function `(lambda(c) (beep))` is called and produces an audible beep. Let us add now a check box into the column:

```
? (==> cc insert (new checkbox "check me"
                  (lambda(c)
                    (printf "Check box state is \\%l\\n"
                            (==> c getdata))))))
= ()
```

The callback function takes one argument. This argument is the check box object that we can query using method `getdata` . Let us add now an editable string named `ee` into the column:

```
? (==> cc insert (setq ee (new editstring 18 "hello")))
;; Argument 18 is the field width.
= ()
? (==> ee setdata "hello people")
= ()
? (==> ee getdata)
= "hello people"
```

Again we can manipulate the state of the editable string using method `setdata` and `getdata` . You may have noticed that the button width has changed when we have inserted the editable string. The column indeed manage its contents in order to keep them properly aligned. We can indeed move and resize the column as a whole:

```
? (==> cc move 50 70)
= (50 20 129 84)
? (==> cc resize 250 120)
= (50 20 250 100)
```

We can even add into the column a small object which lets you move the column with the mouse:

```
? (==> cc insert (new dragarea 50 20))
= ()
```

Arguments 50 and 20 are the width and height of this object. Since the drag area is inserted into a column, its width is adjusted to the column width.

Let us add a row with two exclusive buttons:

```
? (==> cc insert (new row (new radiobutton "choice1")
                          (new radiobutton "choice2") ) )
= ()
```

We can get (or set) the complete state of the column with a single message `getdata` (or `setdata`):

```
? (==> cc getdata)
= (() "hello people" () t)
? (==> cc setdata '(t "goodbye" t ()))
= ()
```

The Ogre library provides many more complex object, as shown in the following example:

```
? (setq h (new filerequester ww))
= ::filerequester:f3ef0
? (==> h popup)
= ::filerequester:f3ef0
? (==> h getdata)
= "/home/leon"
```

Let us add now a button for destroying the window.

```
? (==> ww insert (new stdbutton "Bye Bye."
                          (lambda(c) (==> thiswindowobject delete)) ) )
= ()
```

This button ends the tutorial.

5.3 The Ogre Class Library.

All the objects must interact in order to be drawn at the proper location and with the proper ordering. Similarly, user events (e.g. a mouse click) must be dispatched towards the object located below the mouse.

The kernel of the Ogre library ensures these interaction by defining how messages are propagated between the objects in a graphic interface.

5.3.1 Ogre Methods.

Event dispatching is mostly performed by the `Container` class. A container dispatches the events messages to its children according to simple rules. Keyboard events are sent to the “active” object. All other events are sent to the object located under the mouse pointer.

Repainting is managed asynchronously. When you want to redraw an object, you must send a message `expose` to this object to tell the Ogre system to schedule a repainting operation. When the system becomes idle, the Ogre library sends a `repaint` message to all objects to repaint with a proper ordering which ensures that the topmost objects are repainted last.

Similarly, you never change the location of an object directly. You send message `move`, `resize` or `moveresize` which tell the Ogre system to relocate

an object. The Ogre library then enforces all the rules of the container objects, performs a global recomputation of the location of all objects and places the objects by sending them a message `geometry` .

In general, most operation on the Ogre objects are implemented using two messages.

- The request message (e.g. `expose`) is used by a program to tell the Ogre library that you wish to perform a certain operation.
- The implementation message (e.g. `repaint`) is only called by Ogre to perform the operation and update all the other objects accordingly.

When you want to perform a certain operation, you must call the request message and not the implementation message. The Ogre library will call the implementation message for you at a proper time and with the proper order.

When you define a new class of Ogre objects, you must define the implementation method only. The corresponding request message will be inherited from the superclass.

Here is a list of the most useful message pairs:

```
\begin center \begin tabular
-c-1-1- \hline Operation & Request & Implementation \\ \hline \hline
Drawing & \typo expose & \typo repaint \\ & & \typo repaint-bw \\ & & \typo repaint-color \\ & & \typo backpaint \\ & & \typo backpaint-bw \\ & & \typo backpaint-color \\ \hline Moving and resizing & \typo move & \typo geometry \\ & \typo moverel & \typo manage-geometry \\ & \typo resize & \typo compute-geometry \\ & \typo moveresize & \\ \hline Making an object visible & \typo insert & \typo realize \\ & \typo remove & \\ \hline \end tabular \end center
```

Operation	Request	Implementation
Drawing	<code>expose</code>	<code>repaint</code>
		<code>repaint-bw</code>
		<code>repaint-color</code>
		<code>backpaint</code>
		<code>backpaint-bw</code>
		<code>backpaint-color</code>
Moving and resizing	<code>move</code>	<code>geometry</code>
	<code>moverel</code>	<code>manage-geometry</code>
	<code>resize</code>	<code>compute-geometry</code>
	<code>moveresize</code>	
Making an object visible	<code>insert</code>	<code>realize</code>
	<code>remove</code>	

5.3.2 Ogre Class Hierarchy.

There are in fact two major kinds of classes:

- “User classes” define the most useful graphics components. For instance, class `StdButton` defines a standard push button, class `Menu` manages popup menus and class `Requester` manages subwindows which pop up to request information from the user.
- “Abstract classes” are located on top of the class hierarchy. They define slots and request methods for managing these interactions in behalf of instances of their subclasses. Abstract classes are provided as a choice of superclasses for defining new classes of Ogre objects.

Here is a display of the current class hierarchy in the Ogre library. Classes displayed with a star are the abstract classes.

```
visualobject *
  control *
    editstring
    editnumber
    editsecretstring
  button *
    stdbutton
    tinybutton
    checkbox
    radiobutton
    filereqbutton
  menuitem
  knob *
    dragarea
    sizearea
  slider *
    hslider
    vslider
  scrollbar *
    hscrollbar
    vscrollbar
  textpane
  container *
  form *
    windowobject *
    autowindowobject *
  requester
  warningrequester
  errorrequester
  yesnorequester
```

```

        filerequester
        printrequester
viewer
    selector
    edittext
column
    menupopup
row
grid
frame
    framesize
viewerhole
emptyspace
darkspace
string
icon
menu
    choicemenu

```

Here is the use of the main abstract classes:

- Class **VisualObject** is the root of the Ogre class tree. It defines how objects are located and repainted. It also defines the default event processing messages.
- Class **Control** is the subclass of the simplest interactive objects. Such objects may be enabled or disabled. They may call a callback function when the user requests it.
- Class **Button** defines common messages handled by an object responding to a single mouse click.
- Class **Knob** defines common messages handled by an object which requests a geometry change of its container. For instance, inserting a **DragArea** in a container allow the user to move the container by clicking on the dragarea object.
- Class **Slider** and **Scrollbar** define the common methods for objects allowing the user enter a numerical value proportional to the location of a visible cursor.
- Class **Container** is the superclass of all objects containing other objects. It defines methods used for dispatching the events to the proper objects as well as methods for enforcing a certain layout.
- Class **Form** is a special container used for defining new objects composed of other objects. For instance, a **Viewer** is a **Form** containing two scrollbars and a container whose contents is moved according to the scrollbar position.

- Class `WindowObject` is the most important class. It defines a container object attached to a graphic window on the screen. Every object inserted in a window object becomes visible and is eligible for receiving events.

5.4 Ogre Utility Functions.

5.4.1 Initializing the Ogre library.

(ogre)

Type: DE(ogre.lsh)

A single function named `Ogre` loads and initializes the library. You must call this function on top of each file involved with a graphics interface. Function `Ogre` is often defined as an autoload function which loads the Ogre library and calls the actual `Ogre` function.

This function initializes the library and sets up the event dispatcher. In particular, it creates the object `ogre-task` and defines the functions `event-hook` and `idle-hook`. If the library was already initialized, function `Ogre` returns immediately.

ogre-task

Type: VAR(ogre.lsh)

See: (event-hook)

See: (idle-hook)

See: (process-pending-events)

The object located in variable `ogre-task` is set up by function `ogre`. This object performs three central tasks in the Ogre library:

- It records all operations requested by the user and delayed by the library. This is the case of most repainting requests.
- It starts a repainting operation whenever function `idle-hook` is called.
- It sends event messages to the proper window object whenever function `event-hook` is called.

When Lush is waiting for user commands on the standard input, the Ogre library manages events asynchronously. The user can thus either type lisp commands or activate graphic interfaces.

When a Lush program is running, however, the Ogre library manages events when the function `process-pending-events` is called. It is a good practice to call this function a few times during long programs.

5.4.2 Error handling in the Ogre library.

See: Errors in Lush.

See: (debug-hook)

See: Interruptions in Lush.

See: (break-hook)

The Ogre library runs Lush functions whenever an event occurs.

An error condition occurs when these functions are incorrectly designed. The error message is printed as usual, but you are not prompted for a debug toplevel.

An interruption occurs if you type `Ctrl-C` in the Lush console. The break message is printed as usual, but you are not prompted for a break toplevel.

If you wish to be prompted for these debugging utilities, you must redefine functions `ogre-debug-hook` and `ogre-break-hook` . Here is a simple way to achieve this:

```
(setq ogre-debug-hook nice-debug-hook)
(setq ogre-break-hook nice-break-hook)
```

5.4.3 Ogre Color Palette.

Rendering a graphic interface strongly depends on the capabilities of your screen.

- A black and white screen is barely able to display a few levels of gray using special dithering patterns. Objects must be rendered using simple and visible graphics.
- On the other hand, a color screen is able to render three dimensional looking objects using several colors for rendering the various shadow level.

Ogre provides a way to test if you are using a black and white or a color screen. It also provides several utility functions for displaying three dimensional looking objects.

color-palette

Type: VAR(ogre.lsh)

See: (color [c])

If you are using a black and white display, this variable contains the empty list. You can then select three colors using function `color` .

- Color `color-fg` is the foreground color (usually black),
- Color `color-bg` is the background color (usually white)
- Color `color-gray` is a 50% dithered gray shade used for rendering disabled objects.

If you are using a color display, this variable contains an array containing the color numbers used by Ogre for displaying an object. Several variables are used to name these colors:

<code>palette-left</code>	(for rendering light)
<code>palette-right</code>	(for rendering shadow)
<code>palette-up</code>	(for rendering raised objects)
<code>palette-down</code>	(for rendering depressed objects)
<code>palette-disabled</code>	(for rendering disabled texts)
<code>palette-selected</code>	(for rendering selected objects)

(new-palette r g b)

Type: DE(ogre.lsh)

See: (allocolor r g b)

See: (==> windowobject palette p)

This function creates a new palette array whose dominant color is defined by its RGB component `r` , `g` and `b` .

You can then store the resulting array into variable `color-palette` for redefining the default palette. You can also pass the new palette to an existing window object using message `palette` .

(getcolor colname)

Type: DE(ogre.lsh)

See: (color [c])

This function returns the color number for the color named `n` in the current palette. Although any number can be passed as argument `n` , we suggest using one of the predefined symbols `palette-left` , `palette-right` , `palette-up` , `palette-down` , `palette-disabled` or `palette-selected` for specifying a color. The resulting color number can then be applied using function `color` .

This function is equivalent to:

```
(color-palette <colname>)
```

(setcolor colname)

Type: DE(ogre.lsh)

See: (color [c])

This function sets the current color to the color named `n` in the current palette. Although any number can be passed as argument `n` , we suggest using one of the predefined symbols `palette-left` , `palette-right` , `palette-up` , `palette-down` , `palette-disabled` or `palette-selected` for specifying a color.

This function is equivalent to:

```
(color (getcolor <colname>))
```

(fill-background x y w h)

Type: DE(ogre.lsh)

See: (fill-rect x y w h)

This function fills a rectangle with the background color `palette-up` of the current palette. The 3d rendering functions are designed for being called on such a background.

(draw-up-rect x y w h)

Type: DE(ogre.lsh)

See: (draw-rect x y w h)

This function draws the border of a rectangle using the colors defined in the current palette. The border only is drawn and appears as a slightly raised line.

(draw-up-round-rect x y w h)

Type: DE(ogre.lsh)

See: (draw-round-rect x y w h [r])

This function draws the border of a rectangle with round corners using the colors defined in the current palette. The border only is drawn and appears as a slightly raised line.

(fill-up-rect x y w h [c])

Type: DE(ogre.lsh)

See: (fill-rect x y w h)

This function fills a rectangle using color number `c` and draws a border in order to give the illusion of a rectangle located above the background.

The default value for `c` is the color number indicated by `palette-up` .

(fill-down-rect x y w h [c])

Type: DE(ogre.lsh)

See: (fill-rect x y w h)

This function fills a rectangle with the color number `c` and draws a border in order to give the illusion of a rectangle located below the background.

The default value for `c` is the color number indicated by `palette-down` .

(fill-up-round-rect x y w h [c])

Type: DE(ogre.lsh)

See: (fill-round-rect x y w h [r])

This function fills a rectangle with round corners using color number `c` and draws a border in order to give the illusion of a rectangle located above the background.

The default value for `c` is the color number indicated by `palette-up` .

(fill-down-round-rect x y w h [c])

Type: DE(ogre.lsh)

See: (fill-round-rect x y w h [r])

This function fills a rectangle with round corners using the color number `c` and draws a border in order to give the illusion of a rectangle located below the background.

The default value for `c` is the color number indicated by `palette-down` .

(fill-up-circle x y r [c])

Type: DE(ogre.lsh)

See: (fill-circle x y r)

This function fills a circle using color number `c` and draws a border in order to give the illusion of a circle located above the background.

The default value for `c` is the color number indicated by `palette-up` .

(fill-down-circle x y r [c])

Type: DE(ogre.lsh)

See: (fill-circle x y r)

This function fills a circle using color number `c` and draws a border in order to give the illusion of a circle located below the background.

The default value for `c` is the color number indicated by `palette-down` .

5.4.4 Ogre Fonts.

Font names used by the font function used to be device dependent. A couple of functions have been defined for selecting fonts in a device independent way.

(font-18)

Type: DE(ogre.lsh)

Selects a 18 points font useful for drawing titles.

(font-12b)

Type: DE(ogre.lsh)

Selects a 12 points bold font mostly used for button labels.

(font-12)

Type: DE(ogre.lsh)

Selects a all-purpose 12 points font.

(font-8f)

Type: DE(ogre.lsh)

Selects a fixed width 8 points font used for text input area.

(font-8)

Type: DE(ogre.lsh)

Selects a small 8 points font.

5.5 Visual Objects.

Class `VisualObject` is the root of the Ogre class tree and defines the default appearance (nothing) and the default response to user input (nothing too) of a graphics interface component.

Class `VisualObject` however defines most request messages which are inherited by the graphic components. It also defines the default implementation of most implementation messages.

Four slots are defined by class `VisualObject` :

```
(defclass visualobject object
  (rect (0 0 0 0))
  oldrect
  itscontainer
  window )
```

- Slot `window` is the descriptor of the window that contains the object. This slot is set as soon as a window is associated to the object.
- Slot `itscontainer` is set when the object is inserted into a container. It indicates which container (the father) manages this object.
- Slot `rect` is a list `(x y w h)` that indicates the rectangle associated to the object in the window coordinates.
- Slot `oldrect` is a temporary storage used by method `moveresize` to store the last acknowledged geometry of an object. This slot is non null whenever a geometry change has been requested.

5.5.1 (new VisualObject w h)

Type: CLASS(ogre.lsh)

Creates and return a new instance of class `VisualObject` . Arguments `w` and `h` are the width and the height of the object.

This constructor is seldom used for creating new instances. Derived classes (e.g. `Control`) usually call this constructor method within their own constructor method in order to initialize the `VisualObject` part of their instances.

5.5.2 VisualObject Request Messages.

Class `VisualObject` defines request messages inherited by most other objects for performing two central tasks:

- Repainting the graphic interface components
- Changing the geometry of graphic interface components.

Unless you have a thorough knowledge of the internals of Ogre, you should neither redefine nor override these messages when defining a subclass of `VisualObject`.

(==> `VisualObject expose [rect]`)

Type: MSG(ogre.lsh)

Message `expose` tells the Ogre library that a visual object needs to be repainted. Argument `rect` indicates with more accuracy which rectangle in the current window must be repainted. When argument `rect` is omitted, the entire object rectangle is considered.

Exposure messages are propagated down until they reaches the window object. The clipped rectangle are then added to the damaged area list. Repainting is usually performed when the event queue becomes empty.

(==> `VisualObject repair-damaged`)

Type: MSG(ogre.lsh)

Message `repair-damaged` starts the repainting operation without delay. When this message is sent to an Ogre object, the Ogre system immediately repaints the damaged areas.

(==> `VisualObject moveresize x y w h`)

Type: MSG(ogre.lsh)

Message `moveresize` ask the Ogre library to redefine the rectangle associated to an Ogre object. Arguments `x` and `y` are the coordinates of the topleft corner of the rectangle. Arguments `w` and `h` are respectively the width and the height of the rectangle.

This request is then signaled to the object's container which gets a chance to redefine its own geometry and to enforce a particular layout. When these recomputations are finished, the Ogre library effectively calls method `geometry` to relocate the objects.

(==> `VisualObject move x y`)

Type: MSG(ogre.lsh)

Message `move` is used for changing the location of the topleft corner of the object's rectangle to coordinates `x` and `y`. It actually calls method `moveresize`.

(==> VisualObject **moverel** *xr yr*)

Type: MSG(ogre.lsh)

Message **moverel** is used for adding *xr* and *yr* to the coordinates of the topleft corner of the object's rectangle. It actually calls method **moveresize** .

(==> VisualObject **resize** *w h*)

Type: MSG(ogre.lsh)

Message **resize** is used for changing the width and height of the object's rectangle to *w* and *h* . It actually calls method **moveresize** .

(==> VisualObject **geometry** *x y w h*)

Type: MSG(ogre.lsh)

See: (==> Container **manage-geometry**)

Method **geometry** performs a geometry change on a given object. Unlike method **moveresize** , the change takes place immediatly with all the overhead of recomputing the global position of all objects.

Method **geomerty** is normally called by the method **manage-geometry** defined by the container objects for enforcing a particular layout. You may call this method directly (this is a backward compatibility requirement), although it is more efficient to use method **moveresize** .

(==> VisualObject **front** *x y w h*)

Type: MSG(ogre.lsh)

See: (==> Container **manage-geometry**)

Although the repainting process becomes inefficient, object rectangles may overlap. Methods **front** performs an immediate geometry change and places the object on top of the object stack.

(==> VisualObject **back** *x y w h*)

Type: MSG(ogre.lsh)

See: (==> Container **manage-geometry**)

Although the repainting process becomes inefficient, object rectangles may overlap. Methods **back** performs an immediate geometry change and places the object at the bottom of the object stack.

5.5.3 VisualObject Implementation Methods.

Implementation methods are overridden by each subclass of **VisualObject** in order to define the behavior of each type of object. Class **VisualObject** provides a default implementation for four types of implementation methods:

- Method `realize` is called whenever an object is associated to or dissociated from a window. Redefining a realization method allows for precomputing font sizes and color numbers at once.
- The geometry computation method is called whenever the geometry requirement of an object may have changed. This method defines the minimum size of a graphic object.
- Repaint methods are called when the Ogre library detects that a part of the object rectangle has been damaged, either because method `expose` has been called or because the object has been moved or uncovered by another object.
- Events methods are called whenever the user performs certain mouse or keyboard actions. They define what actions are performed in response to user events. Event messages are documented in the next section.

```
(==> VisualObject realize window)
```

Type: MSG(ogre.lsh)

This method is called whenever an object is associated to a window (realized) or dissociated from a window (unrealized).

- An object is realized when it is inserted into a realized container. A window object is always realized. When an object becomes realized, method `realize` is called with a valid window descriptor as argument `window`. Method `realize` then sets the slot `window`, compute the geometry requirement of the object and calls `expose` in order to paint the object.
- An object is unrealized when it is removed from its container. Method `realize` is then called with an empty list as argument.

The default definition of method `realize` is given below.

```
(defmethod visualobject realize (w)
  (when (setq window w)
    (==> this compute-geometry)
    (==> itscontainer change-geometry) )
  (==> this expose rect) )
```

Subclasses of `VisualObject` may override this default definition. The new definition however must call the superclass method `realize` in order to perform the essential tasks described above.

Before overriding method `realize`, you should also consider overriding method `compute-geometry` instead of method `realize`.

(==> VisualObject **compute-geometry**)

Type: MSG(ogre.lsh)

Method **compute-geometry** is called when the object is realized and when a change of the object state changes the geometry requirements of the object.

This method must either return an empty list or compute the minimal size of the object's rectangle and enforce this minimal size by sending a message **resize** and return the new rectangle **rect** .

The default method **compute-geometry** just returns the empty list. Class **String** for instance overrides this method to compute the size of the string text and resizes the string object to the correct size.

(==> VisualObject **backpaint**)

Type: MSG(ogre.lsh)

Method **backpaint** is called by the Ogre library for repainting the background of an object. In the case of a container object, this method is called before repainting the contents of the object.

The default method **backpaint** tests whether you have a black and white or a color display and calls method **backpaint-bw** or **backpaint-color** respectively. It is therefore advisable to override methods **backpaint-bw** and **backpaint-color** instead of **backpaint** .

(==> VisualObject **backpaint-bw**)

Type: MSG(ogre.lsh)

Method **backpaint-bw** is called by the Ogre library for repainting the background of an object on a black and white display. In the case of a container object, this method is called before repainting the contents of the object.

The default **backpaint-bw** method just clears the object's rectangle with the background color **color-bg** . This is suitable for most cases.

(==> VisualObject **backpaint-color**)

Type: MSG(ogre.lsh)

Method **backpaint-color** is called by the Ogre library for repainting the background of an object on a color display. In the case of a container object, this method is called before repainting the contents of the container object.

The default **backpaint-color** method just clears the object's rectangle with the background color defined in the current palette. This is suitable for most cases.

(==> VisualObject **repaint**)

Type: MSG(ogre.lsh)

Method `repaint` is called by the Ogre library for repainting the foreground of an object. In the case of a container object, this method is called after repainting the contents of the container object.

The default method `repaint` tests whether you have a black and white or a color display and calls method `repaint-bw` or `repaint-color` respectively. It is therefore advisable to override methods `repaint-bw` and `repaint-color` instead of `repaint`.

(==> `VisualObject repaint-bw`)

Type: `MSG(ogre.lsh)`

Method `repaint` is called by the Ogre library for repainting the foreground of an object on a black and white display. In the case of a container object, this method is called after repainting the contents of the container object.

The default method `repaint-bw` does nothing.

(==> `VisualObject repaint-color`)

Type: `MSG(ogre.lsh)`

Method `repaint` is called by the Ogre library for repainting the foreground of an object on a color display. In the case of a container object, this method is called after repainting the contents of the container object.

The default method `repaint-color` does nothing.

5.5.4 Event Methods.

See: (==> `Control activate d`)

See: Events.

Whenever the user performs certain mouse or keyboard actions, the Lush kernel generates an event as explained in the preceding chapter.

When the Ogre library detects an event in a window, it builds an ordered list of objects eligible for handling the event. An event message is then sent to the object with highest priority. If the event method returns the symbol `ignored`, the library proceeds with the next object in the ordered list. This process stops as soon as an object accepts the event (i.e. until an event message returns a value different from symbol `ignored`.)

Starting with the highest priority, the objects eligible for receiving event messages are:

- The toplevel container,
- The chain of embedded containers (if any) managing the topmost object,
- The topmost object located below the mouse cursor,
- The current activated control object of the window. Activation is discussed later in section “Control”.

Class `VisualObject` defines a default method for the event messages. These default methods just reject the event by returning symbol `ignored` .

In practice, the default event rejection mechanism ensures that mouse events are handled by the object located below the mouse cursor and keyboard events are handled by the activated objects. It is however possible to change these settings by overriding the event methods of objects with a higher priority.

(==> `VisualObject mouse-down x1 y1`)

Type: `MSG(ogre.lsh)`

See: `(eventinfo)`

Method `mouse-down` is called when the mouse button is depressed. Arguments `x1` and `y1` indicate the location of the mouse cursor in the current window when the mouse button has been depressed.

You can use function `eventinfo` to find the name of the mouse button and the state of the shift and control keys. Names returned by function `eventinfo` are quite machine dependent however.

(==> `VisualObject mouse-drag x1 y1 x2 y2`)

Type: `MSG(ogre.lsh)`

Method `mouse-drag` is called whenever the mouse is moved while the mouse button is depressed. Arguments `x1` and `y1` indicate the location of the mouse cursor in the current window when the mouse button was first depressed. Arguments `x2` and `y2` indicate the coordinates of the mouse cursor after the move.

(==> `VisualObject mouse-up x1 y1 x2 y2`)

Type: `MSG(ogre.lsh)`

Method `mouse-up` is called when the mouse button is released. Arguments `x1` and `y1` indicate the location of the mouse cursor in the current window when the mouse button was first depressed. Arguments `x2` and `y2` indicate the coordinates of the mouse cursor when the button was released.

(==> `VisualObject keypress c x y`)

Type: `MSG(ogre.lsh)`

Method `keypress` is called whenever an ASCII key is hit. Argument `c` is a string which contains a single character for that key. Arguments `x` and `y` are the location of the mouse cursor when the key was hit.

(==> `VisualObject arrow-left x y`)

Type: `MSG(ogre.lsh)`

See: `(eventinfo)`

Method `arrow-left` is called whenever the user types the left arrow key. Arguments `x` and `y` are the location of the mouse cursor when the key was hit.

You can use function `eventinfo` to find the exact name of the key and the state of the shift and control keys. Key names returned by function `eventinfo` are quite machine dependent however.

(==> `VisualObject arrow-right x y`)

Type: `MSG(ogre.lsh)`

See: `(eventinfo)`

Method `arrow-right` is called whenever the user types the right arrow key. Arguments `x` and `y` are the location of the mouse cursor when the key was hit.

You can use function `eventinfo` to find the exact name of the key and the state of the shift and control keys. Key names returned by function `eventinfo` are quite machine dependent however.

(==> `VisualObject arrow-up x y`)

Type: `MSG(ogre.lsh)`

See: `(eventinfo)`

Method `arrow-up` is called whenever the user types the up arrow key. Arguments `x` and `y` are the location of the mouse cursor when the key was hit.

You can use function `eventinfo` to find the exact name of the key and the state of the shift and control keys. Key names returned by function `eventinfo` are quite machine dependent however.

(==> `VisualObject arrow-down x y`)

Type: `MSG(ogre.lsh)`

See: `(eventinfo)`

Method `arrow-down` is called whenever the user types the down arrow key. Arguments `x` and `y` are the location of the mouse cursor when the key was hit.

You can use function `eventinfo` to find the exact name of the key and the state of the shift and control keys. Key names returned by function `eventinfo` are quite machine dependent however.

(==> `VisualObject help x y`)

Type: `MSG(ogre.lsh)`

See: `(eventinfo)`

Method `help` is called whenever the user types the system dependent help key. Arguments `x` and `y` are the location of the mouse cursor when the key was hit.

The help key under Windows is function key F1. The help key under X11 can be configured using program `xmodmap` .

You can use function `eventinfo` to find the exact name of the key and the state of the shift and control keys. Key names returned by function `eventinfo` are quite machine dependent however.

(==> VisualObject **fkey** x y)

Type: MSG(ogre.lsh)

See: (eventinfo)

Method **fkey** is called whenever the user types a key which does not correspond to an ASCII character and yet cannot be processed as an arrow key or a help key. You must then use function **eventinfo** to analyze the keyname and process the event.

Since key names returned by function **eventinfo** are quite machine dependent, your event handling procedure should use this name as a hint rather than expecting well defined values.

Remark: All operating systems define certain hot keys for various purposes. Lush cannot override these assignments. Under Windows for instance, keys F9, F10 and CTRL-F4 are directly processed by the operating system and never passed to WinLush.

(==> VisualObject **size** w h)

Type: MSG(ogre.lsh)

Method **size** is called whenever the geometry of an object changes in response to a user interaction or to a change in a related object.

5.6 Control Objects.

See: Visual Objects.

See: Event Methods.

Class **Control** is the abstract class for interactive graphics objects like buttons, check boxes, menu items, etc...

Since class **Control** is a subclass of class **VisualObject** , all the properties and methods of class **VisualObject** also apply to class **Control** . In addition, class **Control** provides more support for defining the interactive graphics objects.

- Class **Control** provides standard methods for enabling or disabling a control object. Slot **disabled** contains a non nil value when a control is disabled.
- Class **Control** provides standard methods for changing the activation status of a control object.
- Class **Control** defines standard methods for changing the appearance of an object. Slot **text** contains whatever data represents the appearance of an object.
- Class **Control** defines standard methods for storing and handling the state resulting of user interaction. Slot **data** contains whatever data represent this state.

- Class `Control` stores a callback function in slot `callback` . This function is called when the user performs certain action defined in by implementation of the specific control object.

5.6.1 (new Control w h f)

Type: CLASS(ogre.lsh)

See: Ogre Callbacks.

Creates and return a new instance of class `Control` . Arguments `w` and `h` are the width and the height of the control object. Argument `f` is the callback function for this control object.

This constructor is seldom used for creating new instances. Derived classes (e.g. `Button`) usually call this constructor method within their own constructor method in order to initialize the `Control` part of their instances.

5.6.2 Enabling or Disabling a Control Object.

When disabled, the control object ignores all user interaction. This is signalled by having the control rendered with a specific appearance. A disabled button on a black and white display is rendered with a dithering pattern. The label of disabled button on a color display is rendered with a light gray color instead of black.

The enabled/disabled status of a control object is stored in the slot `disabled` defined by class `Control` .

- Slot `disabled` contains () if the object is enabled.
- Slot `disabled` contains a number (the disable count) is the object is disabled.

Implementation methods usually test slot `disabled` before performing their task. Repainting methods must change the rendering colors according to the enabled/disabled status of the object is disabled. Event methods must ignore event messages when the object is disabled.

Two request methods, `enable` and `disable` , are defined by class `Control` for changing the enabled/disabled status of an object:

(==> `Control disable`)

Type: MSG(ogre.lsh)

If message `disable` is sent to an enabled control object, the object is disabled. If message `disable` is send to an already disabled control object, the object's disable count is increased.

(==> `Control enable`)

Type: MSG(ogre.lsh)

When message `enable` is sent to a control object, the object's disable count is decreased by one. The object is enabled when this count reaches zero.

The disabled count feature proves useful for disabling an object either permanently or temporarily:

- For permanently enabling or disabling a control object whose state is unknown, enable it first and possibly disable it.
- For temporarily disabling an object, just disable it. It will retrieve its previous state when you will enable it again.

(new DisableLock c1...cn)

Type: CLASS(ogre.lsh)

It is often useful to disable a button during the execution of some Lush program and to restore its initial state. This task could be achieved by sending messages `disable` and `enable` .

If however an error occurs during the program execution, the button would remain disabled. The best way to solve this problem consists in creating a lock object defined by class `DisableLock` .

If you wish to temporarily disable object `c1` to `cn` , proceed as follows:

```
(let ((lock (new DisableLock c1...cn)))
  ;; call your Lush program here
  (.....) )
```

When created, the lock object sends a message `disable` to the objects `c1` to `cn` . The lock is destroyed when you leave the `let` instruction. The lock destructor then sends a messages `enable` to our objects `c1` to `cn` .

If an error occurs during the execution of the Lush program, the lock is destroyed by the garbage collector and the objects retrived to their initial state.

5.6.3 Activation of a Control Object.

The active control object receives all the event messages that were not processed by the objects located below the mouse cursor. Only one object might be activated in a given window at a given instant.

Activation is especially useful for redirecting keyboard events toward objects containing some editable text. Having the keyboard events processed by the object located under the mouse often seems unnatural because the keyboard does not move physically like a mouse.

The user activates such an editable object with a mouse click in the object rectangle. It is customary in Ogre that the editable objects ignore the keyboard events unless they are active. Most keyboard events are then unconditionnally directed to the active object without regard to the position of the mouse cursor.

On the other hand , it is often useful to implement accelerator keys for performing tasks usually achieved by menu items or buttons. Accelerator key

events must be defined by the containers managing the related objects. This definition ensures that accelerator keys events are not sent to the active object but handled directly.

The activation status of a control object is stored in slot `activated` (sic) defined by class `Control`. This slot contains a non nil value if the control object is active. Implementation methods of editable objects usually test slot `activated` before performing their task. Repainting methods must indicate the activation status of an object. Event methods must ignore event messages unless the object is activated.

Two request methods are defined by class `Control` for testing or changing the activation status of an object. These methods ensure that one object only is activated in a given window.

(==> `Control activate d`)

Type: MSG(ogre.lsh)

Changes the activation status of the target object to `d`. If argument `d` is `()`, the target object is deactivated. If argument `d` is `t`, the target object is activated and the previous active object is deactivated.

5.6.4 Appearance of a Control Object.

Class `Control` defines a slot named `text` which contains whatever data is useful for defining the appearance of a control object. In most cases so far, this slot contains a single string which is either a button label or a checkbox caption.

Two request methods are defined by class `Control` for obtaining and changing the contents of this slot:

(==> `Control settext d`)

Type: MSG(ogre.lsh)

This request methods changes the contents of the slot `text` of the target object to the value given by argument `d`.

Since this change causes a general change in the object appearance, the object is sent a message `compute-geometry` for redefining its geometry requirements and a message `expose` for updating the display. Sending message `settext` sometimes triggers a global relocation of all objects in the window.

(==> `Control gettext`)

Type: MSG(ogre.lsh)

This request method returns the contents of slot `text`. It is preferred to use this method rather than accessing directly the slot `text` because complex control objects may use other slots for controlling the appearance of the object.

5.6.5 State of a Control Object.

Class `Control` defines a slot named `data` which contains whatever data is useful for storing the state of a control object resulting from the user interaction. For instance, slot `data` of a check box contains a boolean value. Similarly, slot `data` of an editable string contains the current value of the string.

Three methods `hasdata`, `setdata` and `getdata` are used for obtaining and changing the state of a control object. These methods are more implementation method than request methods. Several subclasses of `Control` override these methods in order to pre-process or post-process the state information.

These methods are mostly useful because they provide an abstract way to save and restore the state of a collection of control objects. For instance, a container object can save or restore the state of all its descendants with a single message.

(==> Control setdata d)

Type: MSG(ogre.lsh)

This method changes the state of a control object to `d`.

The default method sets the slot `data` of the target object to the value given by argument `d`. Since this change usually changes the object appearance, the object is sent a message `expose` for updating the display.

(==> Control getdata)

Type: MSG(ogre.lsh)

This method returns the state of a control object.

The default method just returns the contents of slot `data`. It is preferred to use this method rather than accessing directly the slot `data` because complex control objects may use other slots for controlling the state of the object.

(==> Control hasdata)

Type: MSG(ogre.lsh)

If an object contains some state information, this method returns the object itself. If an object contains no state information, this method returns the empty list.

Method `hasdata` is used to test if an object provides an adequate implementation for methods `setdata` and `getdata`. The default method `hasdata` indicated that there is some state information in the control object. Certain subclasses, like push button, override this method in order to signal that they carry no state information.

5.6.6 Ogre Callbacks.

See: Window Objects.

See: Popup Requesters.

See: Forms.

See: Menus.

Each control object has a callback function. Callback functions are executed when certain event messages are received. For example, a push button executes the callback function when it is depressed.

Callbacks functions are called with one argument which is the caller object. They are executed within the caller scope and therefore directly access the slots of the caller object. In addition certain local variables are set:

- Variable `thiswindowobject` always refers to the toplevel container of the control object. See the discussion on class `WindowObject` for more information.
- Variable `thisrequester` possibly refers to the closest requester containing the control object. See the discussion on class `Requester` for more information.
- Variable `thisform` possibly refers to the closest form containing the control object. See the discussion on class `Form` for more information.
- Variable `thismenu` possibly refers to the menu object containing the menu item which has launched the callback function. See the discussion on class `Menu` for more information.

Three methods are implemented by class `Control` for handling callback functions.

(==> `Control setcall f`)

Type: MSG(ogre.lsh)

Message `setcall` redefines the callback function of a target object to be function `f`. Trouble will surely occur if the argument `f` is not a function with one argument or an empty list.

(==> `Control execute`)

Type: MSG(ogre.lsh)

Message `execute` starts the callback function with the object itself as argument. The object is disabled while the callback function is running.

(==> `Control trigger`)

Type: MSG(ogre.lsh)

Message `trigger` asks the object to simulate an user action. When a push button receives a trigger message, for example, it is drawn as depressed, executes the callback function and is redrawn as usual.

The default implementation defined by class `Control` just calls method `execute`. Subclasses of `Control` usually override method `trigger` by sending false event messages to the object.

5.7 Container Objects.

See: Visual Objects.

Class `Container` is the abstract class for container objects. A container object manages several objects called its sons. Sons are only visible through the rectangle of the container. The container dispatches event messages, propagates repainting messages, controls the geometry of its sons and performs a couple of other critical tasks.

Since class `Container` is a subclass of class `VisualObject`, all the properties and methods of class `VisualObject` also apply for class `Container`. Class `Container` however defines new slots and new methods for handling container objects.

- Slot `contents` contains a list whose elements (`son x y w h`) point to the managed objects `son` and cache the coordinates for their rectangles.
- Class `Container` defines private method for implementing the core of the Ogre library. These methods deal with the repainting process and the geometry management.
- Class `Container` defines public request methods for handling several control objects as a whole. Method `enable` and `disable` allow you to disable or enable all the control objects located in a container. Method `getdata` returns a list containing the states of all control objects located in the container. This list is built by sending messages `getdata` to the control objects. Method `setdata` takes such a list as argument and sets the state of all control objects located in the container.

Subclasses of `Container` are very frequently defined for defining structuring container whose sons are arranged according to a certain layout. Classes `Row` and `Column` are examples of such subclasses.

5.7.1 (new Container x y w h ...contents...)

Type: CLASS(ogre.lsh)

Creates a new container object. Arguments `x`, `y`, `w` and `h` describe the initial rectangle of the container object. The arguments denoted as `...contents...` are graphical objects initially managed by the container object.

This constructor is seldom used for creating new instances. Derived classes (e.g. `Row`) usually call this constructor method within their own constructor method in order to initialize the `Control` part of their instances.

5.7.2 Repainting.

When it is time to repaint the contents of a container, the container object receives a message `backpaint` for rendering its background. The objects managed by the container are then rendered on top of this background. Finally the

container object receives a message `repaint` for drawing above the managed objects.

During this operation, the clipping rectangle is the container object's rectangle. Only the portion of the sons which overlap the container's rectangle are rendered.

5.7.3 Inserting and Removing Objects.

Three request methods are implemented for inserting new objects into a container or removing inserted objects from a container.

(==> **Container insert what**)

Type: MSG(ogre.lsh)

Inserts object `what` on top of the objects located in the container.

When an object is inserted into a container, the coordinates of the topleft corner of the container's rectangle are added to the coordinates of the object's rectangle. Therefore, if the initial rectangle of an object is located at coordinates (0, 0), the object appears on the topleft corner of the container.

If there is a policy for the layout of the container, the new layout is computed and each object in the container is moved, resized and repainted.

(==> **Container remove what**)

Type: MSG(ogre.lsh)

Removes object `what` from the container.

When an object removed from a container, the coordinates of the topleft corner of the container's rectangle are subtracted from the coordinates of the object's rectangle. Inserting the object again thus inserts the object at the same location with respect to the container's topleft corner.

(==> **Container removall**)

Type: MSG(ogre.lsh)

Removes all objects managed by the container.

5.7.4 Geometry Management.

Class `Container` implements the core of the geometry management system of the Ogre library. This system aims at providing containers with the opportunity of enforcing certain geometry policies on their sons. For instance:

- Class `Container` enforces a very weak policy. If you move an instance of class `Container`, the managed objects move with their container.
- Class `Row` aligns the objects on a row and ensures that all objects have the same sufficient height,

- Class `Column` aligns the objects on a row and ensures that all objects have the same sufficient width,
- Class `Grid` aligns the objects on a grid with a predefined number of columns. Objects located in a same column have the same sufficient width and objects located in a same row have the same sufficient height.

These structuring containers make Ogre more attractive and easier to use than asking the programmer to precompute the position of each object for all the interface configurations. Yet, geometry management is the most difficult component of the Ogre library.

Containers define the geometry policy using only two implementation methods, `compute-geometry` and `manage-geometry`. Here are the steps involved in a geometry computation:

1. An object managed by a container receives a geometry change request (i.e. `moveresize`, `move` or `resize`).
2. The container of this object receives a message `compute-geometry` in order to recompute its size requirements by looking at the slots `rect` of the managed objects. It might then call method `resize` in order to enforce its new minimal size.
3. The final rectangle of the container is stored in slot `rect` of the container. The container then receives a message `manage-geometry` to get a chance to assign a final rectangle to all managed objects by sending them a message `geometry`.

The hidden part of the iceberg lies in the global computation of the object's location. If the container decides to change its geometry requirements during step 2, the container of the container then receives a message `compute-geometry` and gets a chance to participate to this geometry discussion. When this second container sets the first container's geometry, it calls the first container's method `geometry` which calls the first container's method `manage-geometry` and effectively performs step 3.

Important Note: The geometry management system has been significantly revamped in Lush. It is now faster and skinnier. Old programs still work correctly unless they redefine method `geometry` of class `WindowObject`.

(==> **Container compute-geometry**)

Type: MSG(ogre.lsh)

The implementation method `compute-geometry` is called when the container is realized and when a change in their object state changes the geometry requirements of the object.

This method must compute the minimal size of the container on the basis of the information stored in the slots `rect` of the managed objects or cached in the list of managed objects stored in slot `contents` of the container. This method may enforce a minimal size by sending a message `resize`.

The default method `compute-geometry` just returns the empty list.

Example:

Class `Row` defines the following method `compute-geometry` which computes the minimal size of a row:

```
(defmethod row compute-geometry ()
  (when window
    (let* ((height (sup (cons 0 (all (((son x y w h) contents)) h))))
          (width (sum (all (((son x y w h) contents)) w))) )
      (=> this resize (+ width (* hspace (length contents))) height)
      rect ) ) )
```

(=> Container `manage-geometry`)

Type: MSG(ogre.lsh)

The implementation method `manage-geometry` is called when it is time to locate the sons of the container on the basis of the final rectangle assigned to the container.

This method must compute the final rectangle of all the managed objects according to the container's rectangle (found in slot `rect` of the container) and according to the managed object initial rectangles (found in slot `rect` of the managed objects and cached in the list of managed objects located in slot `contents` of the container.) Method `manage-geometry` then sends a message `geometry` to the managed objects in order to assign a final rectangle to these objects.

The default method `manage-geometry` defined in class `Container` does nothing. A secondary mechanism just moves the managed objects in order to maintain their position with respect to the topleft corner of the container's rectangle.

Example:

Class `Row` defines the following method `manage-geometry` which arranges the managed objects in a row starting on the left of the container's rectangle.

```
(defmethod row manage-geometry ()
  (let ((x y w h) rect))
  (incr x (div hspace 2))
  (each ((i contents))
    (let ((w (nth 3 i))) ;; get the object's width
      (=> (car i) geometry x y w h)
      (incr x (+ w hspace)) ) ) ) )
```

(new `GeometryLock` `c`)

Type: CLASS(ogre.lsh)

A significant amount of time is spent in geometry management when many objects are inserted into a realized container. It is then advisable to postpone the geometry until all objects have been inserted.

Such a task is accomplished by creating an object instance of class `GeometryLock` as follows:

```
(let ((lock (new GeometryLock c)))
      ;; insert your objects here into container c
      (==> c insert ...))
```

When created, the lock sets a flag in container `c` which suspend the usual geometry management. When the lock is destroyed, the flag is cleared and the geometry management is started once. If an error occurs during the object insertion, the lock is destroyed by the garbage collector and the container is left with an acceptable state.

5.7.5 Control Management.

Five request methods are defined by class `Container` for handling as a whole several control objects managed by a container or by its sons.

(==> `Container disable`)

Type: `MSG(ogre.lsh)`

When a container receives a message `disable`, all the control objects managed by the container or by its sons receive a message `enable`. This is useful for controlling the enabled/disabled state of several control objects at once.

```
(defmethod container disable ()
  (each ((i contents))
        (==> (car i) disable) ) )
```

(==> `Container enable`)

Type: `MSG(ogre.lsh)`

When a container receives a message `enable`, all the control objects managed by the container or by its sons receive a message `enable`. This is useful for controlling the enabled/disabled state of several control objects at once.

```
(defmethod container enable ()
  (each ((i contents))
        (==> (car i) enable) ) )
```

(==> Container **hasdata**)

Type: MSG(ogre.lsh)

Method **hasdata** returns the list of stateful control objects managed by the container or by any of its sons. This is achieved by concatenating the result of running method **hasdata** on all the objects managed by the container.

```
(defmethod container hasdata ()
  (flatten (all ((i contents) (==> (car i) hasdata))) )
```

(==> Container **getdata**)

Type: MSG(ogre.lsh)

Method **getdata** returns the list of states returned by all control objects managed by the container or by any of its sons.

```
(defmethod container getdata ()
  (all ((i (==> this hasdata))
        (==> i getdata) ) )
```

(==> Container **setdata** d)

Type: MSG(ogre.lsh)

Method **setdata** sets the states of all control objects managed by the container or by any of its sons using list **d** which is usually returned by a call to method **getdata** .

```
(defmethod container setdata (d)
  (each ((i (==> this hasdata))
        (j d) )
        (==> i setdata j) ) )
```

5.8 Container Flavors.

This chapter presents various flavors of predefined containers. Class **Form** defines a container abstract class for implementing composite objects. Class **WindowObject** defines a toplevel container associated to a window on your screen. Structuring containers help organizing the layout of a graphic interface.

5.8.1 Forms.

See: Ogre Callbacks.

Class **Form** is an abstract class used to group several cooperating objects into a single Ogre component. This component can then be inserted into a window or a requester.

Class `Form` is a trivial subclass of class `Container` . Class `Form` adds three properties to standard containers.

- Method `compute-geometry` makes a form object large enough to hold all its sons while keeping their initial relative positions.
- Method `manage-geometry` places the sons of a form object at the center of the form object while keeping their relative position.
- When Ogre calls the callback function of a descendant of a form object, it sets the temporary variable `thisform` to the form object.

These properties are handy for designing complex graphics objects composed of several components as subclasses of `Form` . Several classes, like class `WindowObject` , class `Requester` and class `Viewer` , are implemented as subclasses of `Form` .

Moreover, a form object is sometimes considered as an extended control object. For instance, we can gather a slider and an editable numeric field into a form object and use this object like an usual control object.

This is achieved by defining a subclass of `Form` whose constructor builds the elements of the composite object and keeps a pointer to these object in some of its slots. This subclass then must define methods `hasdata` , `setdata` , `getdata` , `execute` and `setcall` like those of a control object. Method `hasdata` usually returns `this` . The other methods are usually forwarded to the components of the composite object.

(`new Form ...contents...`)

Type: CLASS(ogre.lsh)

Creates a new form object managing objects `...contents...` . There is no need to provide a default size for a form object since the form size is derived from the component sizes and location.

5.8.2 Window Objects.

See: Events.

Class `WindowObject` is a subclass of class `Form` which implements the link between a physical window and its managed ogre objects.

- The constructor of class `WindowObject` creates a window on the screen with the same size than the window object rectangle. It call then method `realize` to tell its sons that a physical window is now attached to the graphic interface.
- Class `WindowObject` define methods for handling the window events and redirecting event messages to the proper descendant objects.

- Class `WindowObject` define slots and methods to manage the repainting process. The window object remembers which portions of the window has been damaged or exposed (using message `expose`) and sends repainting messages when appropriate.
- Class `WindowObject` defines methods for synchronizing the window object rectangle and the physical window geometry.

A window object is a toplevel container. Inserting a window object into another container is an almost certain cause of trouble.

Whole graphics interface are often defined as subclasses of `WindowObject` . Such subclasses usually contain new slots for storing interface specific data. Callback functions then send messages to `thiswindowobject` . These messages then are executed within the interface scope.

(new WindowObject x y w h name ...contents...)

Type: CLASS(ogre.lsh)

Creates a window object managing objects `...contents...` .

When a window object is created, the constructor creates a window on the screen named `name` with width `w` and height `h` . If both arguments `x` and `y` are 0 , the position of this window is defined by the default rule of your computer. If both arguments are positive, they define the position of the topleft corner of the window.

Exemple:

```
;;; First of all, initialize Ogre !
? (ogre)
= idle-hook
;;; create a window with a button
? (setq win (new WindowObject 100 100 300 300 "Example"
                          (setq b (new stdbutton "Hello"
                                             ;; its name
                                             (lambda(c) (beep))) )))) ;; its callback
= ::WindowObject:06f00
;;; Now we remove the button b from win
? (==> win remove b)
```

Note: The event handler of the newly created window is the corresponding Ogre window object. You can thus interactively obtain the Ogre window object for a given window by typing `(setq w (waitevent))` and clicking into the corresponding window.

(new AutoWindowObject x y w h name contents)

Type: CLASS(ogre.lsh)

Lush provides an additional class named `AutoWindowObject`. Unlike class `WindowObject`, class `AutoWindowObject` manages a single object `contents` and dynamically resizes the window to the size of the managed object.

The constructor arguments `w` and `h` are only hints for sizing the window. They may be empty lists instead of numbers. The size of the window is eventually given by the size of its contents.

WindowObject Request Methods.

(`==>` `windowobject palette p`) Type: `MSG(ogre.lsh)`

See: `Ogre Color Palette`.

Message `palette` redefines the color palette used by `Ogre` for repainting the objects managed by the `windowobject`. This method sets slot `color-palette` of the window object and calls `expose` in order to repaint the window.

Argument `p` can be the empty list (for black and white display) or a palette returned by function `new-palette` (for color display). The color palette selection is enforced even if you do not use an adequate screen.

(`==>` `windowobject setmodal m`) Type: `MSG(ogre.lsh)`

Message `setmodal` affects the redirection of event messages by defining a modal object. When a modal object is defined, all events are redirected to this object. This is useful for implementing requesters that prevent the user to reach the other components of the interface.

- When argument `m` is the empty list, message `setmodal` undefines the current modal object.
- When argument `m` is one of the objects managed by the window object, message `setmodal` makes this object modal.

(`==>` `windowobject getmodal`) Type: `MSG(ogre.lsh)`

Returns the current modal object in the window object or the empty list if no modal object is defined.

(`==>` `windowobject read-event`) Type: `MSG(ogre.lsh)`

See: `Event Lists`.

Method `read-event` returns the next available event in the window associated to the window object. Events are returned as standard Lush event lists. If no events are pending, method `read-event` blocks.

(`==>` `windowobject manage-event event`) Type: `MSG(ogre.lsh)`

See: `Event Lists`.

Method `manage-event` processes a Lush event list `event` and sends a corresponding event message to the proper object in the window object.

WindowObject Implementation Methods.

When defining a graphic interface as a subclass of class `WindowObject` you may override three methods:

- Method `delete` is called when you destroy the window using a the window manager feature (i.e. with the mouse).
- Methods `compute-geometry` and `manage-geometry` work as usual. They are however called in special cases allowing a dynamical change of the window size.

(==> `windowobject delete`) Type: MSG(ogre.lsh)

Method `delete` is called when the user deletes the onscreen window using the window manager features. In the same spirit, we suggest you to call method `delete` whenever you destroy a window object.

The default method `delete` just deletes the window object. You may override this method, for instance, to pop up a confirmation dialog to the user.

(==> `windowobject manage-geometry`) Type: MSG(ogre.lsh)

Beside its normal use, method `manage-geometry` is also called when you resize the window with the mouse. The mouse action is actually converted into a `size` event which is processed using the usual geometry management functions.

The default method `manage-geometry` does nothing. Overriding the method `manage-geometry` in a window object is often useful for adjusting the size of the window object components to the onscreen window size.

(==> `windowobject compute-geometry`) Type: MSG

Method `compute-geometry` usually computes the minimal size of the object and enforces this geometry using method `resize` . In the case of a window object , the actual size of the onscreen window is also adjusted.

The default method `compute-geometry` does nothing. You may override this method to adjust automatically the size of the onscreen window to the mimial size of its contents.

5.8.3 Structuring Containers

See: Container Objects.

Structuring containers are used for designing the layout of a graphics interface. They arrange their sons according to a certain policy. Ogre provides three class of structuring containers: class `Row` , class `Column` and class `Grid` .

These containers inherit all the slots and methods of an ordinary container. They just define specific methods `compute-geometry` and `manage-geometry` to enforce their layout policy.

In addition, two classes are provided for padding a row or a column with some space, class `EmptySpace` and class `DarkSpace` .

(`new Row ...contents...`)

Type: CLASS(ogre.lsh)

See: (==> `Container insert what`)

See: (==> `Container remove what`)

See: Control Management.

Returns a new row object managing the objects specified by `...contents...`

Class `Row` defines a structuring container which aligns its contents horizontally from left to right. The height of a row is determined by the managed object whose required height is largest.

An empty space is inserted between the objects. This space is defined by slot `hspace` of the row object and defaults to 4 points.

(new Column ...contents...)

Type: CLASS(ogre.lsh)

See: (`==>` `Container` insert `what`)

See: (`==>` `Container` remove `what`)

See: Control Management.

Returns a new column object managing the objects specified by `...contents...`

Class `Column` defines a structuring container which aligns its contents vertically from top to bottom. The width of a column is determined by the managed object whose required width is largest.

An empty space is inserted between the objects. This space is defined by slot `vspace` of the column object and defaults to 4 points.

Example:

```
;; Creates a window with three buttons
? (setq win
  (new WindowObject 100 100 300 300 "Essai"
    (new Column
      (new StdButton "One" (lambda(d) (print 1)))
      (new StdButton "Two" (lambda(d) (print 2)))
      (new StdButton "Three" (lambda(d) (print 3))))))
= ::WindowObject:06f00
```

(new Grid cols ...contents...)

Type: CLASS(ogre.lsh)

See: (`==>` `Container` insert `what`)

See: (`==>` `Container` remove `what`)

See: Control Management.

Returns a new grid object with `cols` columns and managing the objects specified by `...contents...`

Class `Grid` defines a structuring container which aligns its contents in a grid with `cols` columns. The grid is filled left to right. The width of each column is determined by the object of the column whose required width is largest. The height of each row is determined by the object of the row whose required height is largest.

Empty spaces are inserted between rows and columns. These spaces are defined by slot `vspace` and `hspace` of the grid object and defaults to 4 points.

(new EmptySpace w [h])

Type: CLASS(ogre.lsh)

Extra spaces may be added into a structuring container by inserting an emptyspace object. Such an object has no response to user events and is uniformly filled with the background color.

Class `EmptySpace` is a class of emptyspace objects. The constructor expression above returns a new emptyspace object of minimal width `w` and minimal height `h`. When argument `h` is omitted, a square minimal space is assumed.

Of course, the minimal geometry specifications interact with the geometry management of structuring containers. It is seldom necessary to define all the sizes of an emptyspace object.

Exemple:

```
;;; Creates a window with three buttons
;;; with a 20 pixels space between button 2 and button 3
? (setq win
  (new WindowObject 100 100 300 300 "Essai"
    (new Column
      (new StdButton "One" (lambda(d) (print 1)))
      (new StdButton "Two" (lambda(d) (print 2)))
      (new EmptySpace 20)
      (new StdButton "Three" (lambda(d) (print 3))) )))
= ::WindowObject:06f00
```

(new DarkSpace w [h])

Type: CLASS(ogre.lsh)

Similarly you can insert a darkspace object into a structuring container. On a black and white display, a darkspace is displayed as a black area. On a color display, a darkspace is displayed as a raised area.

Class `DarkSpace` is a class of emptyspace objects. The constructor expression above returns a new emptyspace object of minimal width `w` and minimal height `h`. When argument `h` is omitted, a square minimal space is assumed.

Of course, the minimal geometry specifications interact with the geometry management of structuring containers. Specifying a single argument of 4, for instance, will define a minimal width and height of 4 points. If such a darkspace object is inserted in a column, the width of object will be increased up to the column width, effectively displaying a 4 points thick line.

Exemple:

```
;;; Creates a window with three buttons
```

```

;;; with a 3 pixels line between button 2 and button 3
? (setq win
  (new WindowObject 100 100 300 300 "Essai"
    (new Column
      (new StdButton "One" (lambda(d) (print 1)))
      (new StdButton "Two" (lambda(d) (print 2)))
      (new DarkSpace 3)
      (new StdButton "Three" (lambda(d) (print 3))))))
= ::WindowObject:06f00

```

5.9 Buttons.

Buttons are small interactive area which can be used for triggering an action or storing boolean states. Ogre provides four class of buttons:

- Class `StdButton` and `TinyButton` for push buttons,
- Class `CheckBox` for boolean state buttons,
- Class `RadioButton` for exclusive buttons.

All these classes are subclasses of `Control` . Buttons thus inherit all the slots and methods defined in class `Control` .

5.9.1 Push Buttons.

A push button invokes a specific action when the user clicks the mouse button while the pointer is located above the button image. Two kinds of push buttons are defined by the Ogre library.

```
(new StdButton label call)
```

Type: CLASS(ogre.lsh)

See: (font-12b)

See: (==> Control settext d)

See: (==> Control gettext)

See: (==> Control setcall f)

Class `StdButton` implements a standard push button. The label of a standard push button is displayed using the font set by the function stored in slot `textfont` , which default to a 12 points bold font `font-12b` .

The constructor expression returns a push button whose name is given by string `label` . This label may be changed by sending a message `settext` . When the button is depressed, the callback function `call` is executed. The button remains disabled while the callback function has not returned.

Example:

```

;;; A window with a 3 state label
? (setq win
  (new WindowObject 100 100 300 300 "Essai"
    (new StdButton "One"
      (lambda(caller)
        (==> caller settext
          (selectq (==> caller gettext)
            ("One" "Two")
            ("Two" "Three")
            (t "One")) ) ) ) ) ) )
= :WindowObject:06f00

```

```
(new TinyButton label call)
```

Type: CLASS(ogre.lsh)

See: (font-12b)

See: (==> Control settext d)

See: (==> Control gettext)

See: (==> Control setcall f)

Class `TinyButton` implements a push button with a reduced size. A tiny push button is slightly smaller than a standard push button because the label is written using a smaller font `font-12` .

The constructor expression returns a push button whose name is given by string `label` . This label may be changed by sending a message `settext` . When the button is depressed, the callback function `call` is executed. The button remains disabled while the callback function has not returned.

5.9.2 Check Boxes.

A check box is composed of a small square button followed by a descriptive text. When the user clicks on the square or on the text , the state of the small square changes.

Class `CheckBox` implements check box buttons. Since class `CheckBox` is a subclass of class `Control` , all the slots and methods defined for control objects are inherited.

```
(new CheckBox label call)
```

Type: CLASS(ogre.lsh)

See: (==> Control settext d)

See: (==> Control gettext)

See: (==> Control setdata d)

See: (==> Control getdata)

See: (==> Control setcall f)

Returns a new check box whose descriptive text is given by string `label` . When a check box is depressed, its state changes and the callback function `call` is executed.

You can query or change the state of a check box (`t` or `()`) by sending messages `getdata` and `setdata` . The descriptive text can be modified using `settext` .

Example:

```
;;; A window with a check box that controls
;;; the enable/disable state of a push button
? (setq win
  (new WindowObject 100 100 300 300 "Essai"
    (new column
      (setq thebutton
        (new StdButton "beep"
          (lambda(c) (beep)) ))
      (new CheckBox "Disable button"
        (lambda(c)
          (if (==> c getdata)
              (=> thebutton disable)
              (=> thebutton enable) ) ) ) ) ) ) )
```

(new **ImageButton** up-image down-image disabled-image call)

Type: CLASS(ogre.lsh)

See: (=> Control setcall f)

Class **ImageButton** implements a push button whose appearance is an image drawn with `rgb-draw-matrix` . Each of the first three arguments must be an `idx2` or an `idx3` of `ubytes` containing the pixel data (greyscale if an `idx2`, RGB if an `idx3` with at least 3 elements in the last dimension). `up-image` is the image shown when the button is up, `down-image` when it is pushed, `disabled-image` when it is disabled. The A (alpha) component, if present, is ignored. `call` is the callback function called when the button is clicked.

Here is an example that shows how to create down and disabled image of an **ImageButton** from an up image.

```
(libload "libimage/image-io")
(libload "libimage/rgbaimage")
(ogre)
(setq up (image-read-rgba "my-button-image.png"))
(setq down (idx-copy up))
(rgbaim-contbright up down -1 0)
(setq disabled (idx-copy up))
(rgbaim-contbright up disabled 0.5 40)
(setq btn (new ImageButton up down disabled (lambda (c) (printf "coucou\n"))))
```

```
(setq window (new windowobject 10 10 200 200 "test" btn))
```

```
(new StdButtonBg text up-image down-image disabled-image call)
```

Type: CLASS(ogre.lsh)

See: (`==>` Control setcall f)

This works just like an `StdButton` except that the image of the button is an RGBA or grayscale image passed as argument (instead of the boring blueish rounded rectangle of `StdButton`). The text is displayed on top of the image. The three images passed as argument will be used respectively as the up button image, the clicked (down) button image, and the disabled button image. The text is drawn in black over the background image, so the image colors should not be too dark. The background images are automatically sized/resized to fit the text (without necessarily preserving the aspect ratio). The images must be `idx2` (for grayscale images) or `idx3` with the last dimension equal to 4 (RGBA images). The A (alpha) component is ignored.

Here is an example that shows how to create an `StdButtonBg` from one of the standard icons:

```
(libload "libimage/image-io")
(libload "libimage/rgbaimage")
(setq icondir (concat-fname lushdir "lsh/libogre/icons"))
(ogre)
(setq btn
  (new StdButtonBg
    "Click Me"
    (image-read-rgba (concat-fname icondir "button-brushed-metal-02-up.png"))
    (image-read-rgba (concat-fname icondir "button-brushed-metal-02-down.png"))
    (image-read-rgba (concat-fname icondir "button-brushed-metal-02-disabled.png"))
    (lambda (c) (printf "coucou\n"))))
(setq window (new windowobject 10 10 200 200 "test" btn))
(==> btn disable)
(==> btn enable)
```

5.9.3 Exclusive Buttons.

Exclusive button are two-state buttons which cooperate with the other buttons in the same container in order to ensure that at most one button is in a positive state.

Class `RadioButton` implements exclusive buttons. Since class `RadioButton` is a subclass of class `CheckBox` , all the slots and methods defined for checkbox and control objects are inherited.

```
(new RadioButton label call)
```

Type: CLASS(ogre.lsh)

See: (`==> Control settext d`)
 See: (`==> Control gettext`)
 See: (`==> Control setdata d`)
 See: (`==> Control getdata`)
 See: (`==> Control setcall f`)

Returns a new check box whose descriptive text is given by string `label` . You can query or change the state of a check box (`t` or `()`) by sending messages `getdata` and `setdata` . The descriptive text can be modified using `settext` .

You must insert a radiobutton in a container containing only radiobuttons. Whenever the user clicks on a radiobutton, the state of all the other radiobuttons in the container becomes `()` , while the state of the clicked radiobutton becomes `t` and the callback function `call` is called.

Example:

```
;; A window with a three radio buttons
;; building a three state choice
? (setq win
  (new WindowObject 100 100 300 300 "Essai"
    ;; create a common callback function
    (let ((call (lambda(d)
                  (printf "\\%s has been selected\\n"
                    (==> d gettext) ) )))
      (new column
        (new radiobutton "choice 1" call)
        (new radiobutton "choice 2" call)
        (new radiobutton "choice 3" call) ) ) ) )
```

Remarks:

- All radiobuttons in a same container are exclusive.
- A container of radiobuttons should not contain objects which are not instances of class `RadioButton` .

5.9.4 File Requester Button.

A file requester button is composed of a small square button which contains an icon. When the button is depressed, a file requester is popped up.

A file requester button is usually associated with a string editor, which displays the file selected in the requester.

See: File Requester.

See: Editable Strings.

(**new FileReqButton** areqmsg aedstring)

Type: CLASS

Returns a new file requester button. Argument **areqmsg** is the introductory message of the requester associated with the button. Argument **aedstring** is the string editor associated with the button.

5.10 Character String Objects.

Ogre provides several classes for handling character string:

- Class **String** provides a way to display fixed text.
- Class **EditString** provides a way to display a text string that the user can edit using the mouse and the keyboard.
- Class **EditNumber** provides a way to display a numerical text string that the user can edit using the mouse and the keyboard.
- Class **EditSecretString** provides a way to display a secret text string that the user can edit using the mouse and the keyboard.
- Classes **EditText** and **TextPane** implement a multiline text editor. Class **TextPane** is the basic class for the editor or viewer. Class **EditText** manages a text editor and two scrollbars.

5.10.1 Fixed Strings.

Class **String** is used for creating a caption text which does not respond to user interaction. Such objects are useful for displaying explanatory text or short message. Class **String** is a subclass of class **VisualObject** . This is not a control object because it has no user interaction capabilities.

(**new String** text)

Type: CLASS(ogre.lsh)

See: (==> **Control** settext d)

See: (==> **Control** gettext)

Creates a new string object displaying text **text** . The text displayed by a string object may be changed using **settext** like the text of a control object.

- Argument **text** may be a string containing or not newline characters. When **text** contains newline characters, a multiline text is displayed.
- Argument **text** may also be a list of strings, one per line.

The text is displayed using a font set by the function stored in slot **textfont** of the string object. This font default to a standard 12 points font.

5.10.2 Editable Strings.

See: Regular Expressions.

See: Activation of a Control Object.

See: (`==> Control setdata d`)

See: (`==> Control getdata`)

See: (`==> Control setcall f`)

An editable string is a single line text editor.

An editable string object becomes activated when the user clicks the mouse button above its rectangle. Keypress and arrow events then are bound to editing functions according to a keymap stored in global variable `EditStringKeymap` .

The text edited by an editable string is controlled by the value of the slot `regex` of the editable string. This slot contains a regular expression which must match the text at all times. If an editing action results in a non-matching text, the action is discarded and a beep is emitted.

- Class `EditString` implements the editable strings. Since class `EditString` is a subclass of class `Control` , all the slots and methods defined for control objects are inherited.
- Class `EditNumber` implements an editable string intended for entering or displaying numbers.
- Class `EditSecretString` implements an editable string intended for entering secret texts by displaying any character as an 'x'.

(`new EditString minsize [defaulttext]`)

Type: CLASS(ogre.lsh)

Creates a new editable string displaying at least `minsize` characters. The initial value of the text is specified by the optional string argument `defaulttext` .

Example:

```
;;; A window with two editable strings
? (setq win
  (new WindowObject 100 100 300 300 "Essai"
    (new grid 2
      (new string "First string:")
      (new editstring 16 "number one")
      (new string "Second string:")
      (new editstring 16 "number two") ) ) )
```

EditStringKeymap

Type: VAR(ogre.lsh)

The global variable `EditStringKeymap` contains the keymap for `EditString` objects as a list of key bindings. Every key binding has the form `(char action)` where `char` is a one character string and `action` is a symbol identifying a specific editing action.

Every typed character is matched against the keymap.

- If a binding matches the typed character, message action with no argument is sent to the editable string.
- If no binding matches, the character is inserted into the data string at the current cursor location. The cursor is then advanced by one position.

The default keymap is loosely modeled after the Emacs text editor:

```
(setq EditStringKeymap
 '( ( "\\n"      execute)          ;; <LineFeed>
    ( "\\r"      execute)          ;; <Enter> or <Return>
    ( "\\b"      backspace)        ;; <ctrl-H> or <Backspace>
    ( "\\x7f"    backspace)        ;; <Delete>
    ( "\\^D"     delete-char)      ;; <Ctrl-D>
    ( "Delete"   delete-char)      ;; Key <Delete>
    ( "\\^A"     begin-of-line)    ;; <Ctrl-A>
    ( "Home"     begin-of-line)    ;; Key <Home>
    ( "\\^E"     end-of-line)      ;; <Ctrl-E>
    ( "End"      end-of-line)      ;; Key <End>
    ( "\\^F"     arrow-right)      ;; <Ctrl-F> and <Right>
    ( "\\^B"     arrow-left)       ;; <Ctrl-B> and <Left>
    ( "\\^P"     arrow-up)         ;; <Ctrl-P> and <Up>
    ( "\\^N"     arrow-down)       ;; <Ctrl-N> and <Down>
    ( "\\^K"     kill)             ;; <Ctrl-K>
    ( "\\^Y"     yank) ) )        ;; <Ctrl-Y>
```

The binding second elements, actions, are the names of methods defined by class `EditString` to perform the various editing tasks. The valid actions are:

- Action `execute` calls the callback function. You can install a callback function using message `setcall` because an editable string is a control object.
- Action `backspace` deletes the character preceding the cursor.
- Action `delete-char` deletes the character located right after the cursor.
- Action `begin-of-line` moves the to cursor the beginning of the line.
- Action `end-of-line` moves the cursor to the end of the line.
- Action `arrow-left` moves the cursor one character left.

- Action `arrow-right` moves the cursor one character right.
- Action `arrow-up` recalls the previous string in the history buffer. Editable strings maintain a buffer of the last 20 strings typed by the user.
- Action `arrow-down` recalls the next string in the history buffer.
- Action `kill` deletes the text between the cursor and the end of the string. This text is copied into the clipboard.
- Action `yank` inserts text from the clipboard. The kill and yank actions implement in fact a simple Copy/Paste mechanism between editable strings.

5.10.3 Editable Numbers.

See: Regular Expressions.

See: (`==> Control setdata d`)

See: (`==> Control getdata`)

See: (`==> Control setcall f`)

(`new EditNumber minsize [defaultvalue]`)

Type: CLASS(`ogre.lsh`)

Class `EditNumber` is a subclass of class `EditString` used for editing numeric fields. It checks that the text is a valid number using a regular expression stored in slot `regex` at construction time. In addition, the methods `setdata` and `getdata` have been modified to handle numbers instead of strings.

Argument `minsize` is the minimal number of characters displayed in the object. The initial value of the number is specified by the optional number argument `defaultvalue` .

When sent to an editable numeric string, message `setdata` requires a numerical argument. Similarly, message `getdata` returns a number or the empty list when no text has been entered.

5.10.4 EditSecretString

This subclass of `EditString` works exactly like `EditString` but always display the string character as a string of "x" .

5.10.5 Multiline Text Editor.

See: (`==> Control setdata d`)

See: (`==> Control getdata`)

See: (`==> Control setcall f`)

See: Forms.

(**new EditText** *w h* [default])

Type: CLASS(ogre.lsh)

Creates a new edittext object (i.e. a text editor with both a vertical and horizontal scrollbar). Argument *w* is the number of visible columns. Argument *h* is the number of visible lines.

Class **EditText** is a subclass of **Forms** that defines composite object containing a **TextPane** object and two scroll bars. The text editor is actually implemented by class **TextPane** described hereafter.

Class **EditText** forwards messages **setdata** , **getdata** , **read-file** and **write-file** to this underlying **TextPane** object.

See: (new **TextPane** *w h* [**def editp vs hs**])

See: (==> **TextPane** **getdata**)

See: (==> **TextPane** **setdata** *arg*)

See: (==> **TextPane** **read-file** *fname*)

See: (==> **TextPane** **write-file** *fname*)

(**new TextPane** *w h* [**def editp vs hs**])

Type: CLASS(ogre.lsh)

Creates a new **TextPane** object. Class **TextPane** is a subclass of **Control** . It implements a simple multi-line text editor.

- Argument *w* is the number of visible columns.
- Argument *h* is the number of visible lines.
- Argument **def** specifies a default contents as a string or a list of strings. Single strings are searched for TAB and NL characters and splitted into multiple strings (one per line).
- Flag **editp** tells if the string is editable or just viewable.
- If scrollbars are attached to the text pane, you must pass them as arguments **vs** and **hs** . The text pane will then synchronize the scrollbars with the actual contents of the text pane area.

TextPaneKeymap

Type: VAR(ogre.lsh)

This A-list contains the associations between key strokes and editing methods executed by a **TextPane** object. The default keymap is loosely based on the Emacs key bindings.

```
(setq TextPaneKeymap
 '( ("\\n"   execute)
    ("\\r"   execute)
    ("\\b"   backspace)
```

```

("\\x7f"    backspace)
("\\^A"    begin-of-line)
("\\^B"    arrow-left)
("\\^D"    delete-char)
("\\^E"    end-of-line)
("\\^F"    arrow-right)
("\\^K"    kill)
("\\^N"    arrow-down)
("\\^P"    arrow-up)
("\\^V"    page-down)
("\\^Y"    yank)
("Delete"  delete-char)
("C-Home"  begin-of-text)
("C-End"   end-of-text)
("Home"    begin-of-line)
("End"     end-of-line)
("Prior"   page-up)
("Next"    page-down)
("\\x1b"   metakey "\\x1b")
("\\x1b<"  begin-of-text)
("\\x1b>"  end-of-text)
("\\x1bv"  page-up)
) )

```

The first element of each association describes the keystroke or keystroke combination. This string can contain an ASCII character, a string containing a keystroke combination, or a function key name. The leading character of keystroke combinations must be associated to method `metakey`. The function key names may be preceded by "C-" and "S-" to indicate that the control or shift key must be depressed.

The rest of the association specify which action should be called when the corresponding keystroke is typed by the user. The following actions are supported:

- Action `execute` calls the callback function. You can install a callback function using message `setcall` because an editable string is a control object.
- Action `backspace` deletes the character preceding the cursor.
- Action `delete-char` deletes the character located right after the cursor.
- Action `begin-of-line` moves the cursor to the beginning of the line.
- Action `end-of-line` moves the cursor to the end of the line.
- Action `begin-of-text` moves the caret at the beginning of the text.
- Action `end-of-text` moves the caret at the end of the text.

- Action `arrow-left` moves the cursor one character left.
- Action `arrow-right` moves the cursor one character right.
- Action `arrow-up` moves the cursor one line up.
- Action `arrow-down` moves the cursor one line down.
- Action `page-up` moves the cursor one screen up.
- Action `page-down` moves the cursor one screen down.
- Action `kill` deletes the text between the cursor and the end of the string. This text is copied into the clipboard.
- Action `yank` inserts text from the clipboard. The kill and yank actions implement in fact a simple Copy/Paste mechanism between editable strings.
- Action `metakey` is used to declare the prefix of a keystroke combination. The action argument will be combined with the next keystroke and the combined string will be searched in the keymap.

(==> `TextPane setdata arg`)

Type: `MSG(ogre.lsh)`

Sets the text in a `TextPane`. Argument `arg` specifies a default contents as a string or a list of strings. Single strings are searched for TAB and NL characters and splitted into multiple strings (one per line).

(==> `TextPane getdata`)

Type: `MSG(ogre.lsh)`

Returns a list of strings representing all lines of the text pane.

(==> `TextPane read-file fname`)

Type: `MSG(ogre.lsh)`

Reads the contents of text file `fname` into the text pane.

(==> `TextPane write-file fname`)

Type: `MSG(ogre.lsh)`

Writes the contents of the text pane into file `fname` .

5.11 Icons.

An icon is a simple object which displays a greyscale or colormapped picture. The `Icon` class is somewhat superseded by the `Ogrimage` class.

5.11.1 (new icon mat &optional [sx [sy [map]]])

Type: CLASS

Returns a new icon.

Argument `mat` is a 2D matrix containing the data to be represented as rectangles using 64 gray levels. By default, the values lower than 0 will be black and the values beyond 1 will be white.

Arguments `sx` and `sy` are the width and height of each rectangle.

Argument `cmap` may be a 1D integer matrix defining 64 colors. When it is defined, its colors are used instead of gray levels.

See: (gray-draw-matrix x y mat minv maxv apartx aparty)

See: (color-draw-matrix x y mat minv maxv apartx aparty cmap)

5.12 OgrImage.

An `OgrImage` is a simple object which displays a color or greyscale picture. An `OgrImage` can be used advantageously in combination with the `ImageViewer` class.

5.12.1 (new OgrImage m)

Type: CLASS

return a new `OgrImage` that will display picture `m` . `m` must be an `IDX2` (greyscale image) or `IDX3` whose last dimension must be 1 (grayscale) , 3 (RGB) , or 4 (RGBA).

5.12.2 (==> OgrImage get-selected)

Type: MSG

return the rectangle last selected by the user by clicking and dragging the mouse on the image.

5.13 ImageViewer.

An `ImageViewer` is a scrollable viewing area for an `OgrImage`.

5.13.1 (new ImageViewer w h m [scroll])

Type: CLASS

Create a new image viewer which will display image `m` in a scrollable viewing area of size `w` , `h` . If the optional `scroll` argument is true, the image scrolling will follow scrollbar dragging, whereas if `scroll` is nil or absent, the scrolling is only performed when the mouse button is released. Here is an example:

```
(libload "libimage/image-io")
(setq m (image-read-rgb (concat-fname lushdir "lsh/libimage/demos/sample.jpg")))
(ogre)
(setq w (new WindowObject 10 10 400 400 "asd" (new ImageViewer 340 340 m t)))
```

5.13.2 (==> ImageViewer get-selected)

Type: MSG

return the rectangle last selected by the user by clicking and dragging the mouse on the image.

5.14 Menus.

A menu is a popup column associated to a menu button. When the mouse is depressed on the menu button, a popup column appears and the user may drag the mouse over an item in the column. When the mouse button is released, the item is selected, its callback function is executed and the popup column disappears.

Menus are implemented with three classes:

- Class `Menu` defines the menu button. A menu button appears as a bold string preceded by a small upside down triangle. When the menu button is depressed, an instance of class `MenuPopup` is inserted into the menu button's window.
- Class `MenuPopup` is a subclass of class `Column`. It arranges the menu items and manages event messages by selecting the appropriate menu item.
- Class `MenuItem` is a subclass of class `Control`. Menu items may receive `settext` messages for changing their label and `setdata` messages for setting slot `data`. When slot `data` is not the empty list, a check mark appears on the left of the menu item.

The constructor of class `Menu` sets up all these objects in a single call.

5.14.1 Standard Menus.

```
(new Menu menuname label1 call1 ... labelN callN )
```

Type: CLASS(ogre.lsh)

Creates a menu button labelled `menuname` linked to a menu popup composed of items defined by the pairs `label1`, `call1` to `labelN`, `callN`. Argument `labeli` is the name of the `i` th menu item. Argument `calli` is the callback associated to the `i` th menu item.

Example:

```

;;; A window with a menu composed of 3 items
? (setq win
  (new WindowObject 100 100 300 300 "Essai"
    (new menu "Test Menu"
      ;; the first item toggles its check mark
      "Toggle"
      (lambda(item) (==> item setdata (not (==> item getdata))))
      ;; the second item activates the bell
      "Bell"
      (lambda(item) (beep))
      ;; the third one destroy this interface
      ;; remember that <thiswindowobject> always
      ;; refer to the closest window object.
      "Quit"
      (lambda(item) (==> thiswindowobject delete)) ) ) )

```

(==> Menu **finditems** label)

Type: MSG(ogre.lsh)

Message **finditems** returns the menu items matching argument **label** .

- If **label** is the empty list () , message **finditems** returns a list containing all the menu items.
- If **label** is a number, message **finditems** returns a list containing the **label** -th menu item.
- If **label** is a string, message **finditems** returns a list of menu items whose label string is equal to **label** .

(==> Menu **disable** ...labels...)

Type: MSG(ogre.lsh)

See: (==> Menu **finditems** label)

When a menu receives a message **disable** , it sends a message **disable** to all menu items identified by the arguments ...labels... . Therefore, the corresponding menu items are disabled.

Each argument of message **disable** is interpreted like the argument **label** of message **finditems**. It can be the empty list, a number or a string.

(==> Menu **enable** ...labels...)

Type: MSG(ogre.lsh)

See: (==> Menu **finditems** label)

When a menu receives a message **enable** , it sends a message **enable** to all menu items identified by the arguments ...labels... . Therefore, the corresponding menu items are disabled.

Each argument of message `disable` is interpreted like the argument `label` of message `finditems`. It can be the empty list, a number or a string.

(new MenuItem label call)

Type: CLASS(ogre.lsh)

Creates a new menu item labelled `label` . Argument `call` indicates the callback function for this menu item. Menu items may be inserted into a menu popup by sending a message `insert` to the menu object.

(==> Menu insert menuitem)

Type: MSG(ogre.lsh)

Message `insert` inserts a new menu item `menuitem` in the menu popup.

(==> Menu remove menuitem)

Type: MSG(ogre.lsh)

Message `insert` removes new menu item `menuitem` from the menu popup.

5.14.2 Choice Menus.

This class implements a multiple choice item. This object display a menu mark and the name of the currently selected item. When the user clicks on it , a menu popup is displayed and the user can select another item.

Choice menus supports the methods of standard menus , in addition to their own methods.

(new ChoiceMenu [items [callback]])

Type: CLASS

Creates a new choice menu that lets the user choose among the items specified in list of strings `items` . When argument `items` is unspecified , selection is impossible.

(==> choicemenu setitems items)

Type: MSG

Changes the items that the user can select with the choice menu to be the strings of list `items` . As a side effect , the current selection is reset to the empty list.

(==> choicemenu setdata d)

Type: MSG

Selects string `d` in the choice menu.

(==> **choicemenu getdata**)

Type: MSG

Returns the string selected in the choice menu.

(==> **choicemenu setcall callback**)

Type: MSG

Sets the callback function

(==> **choicemenu disable**)

Type: MSG

Disables a choicemenu.

(==> **choicemenu enable**)

Type: MSG

Enables a choicemenu.

5.15 Popup Requesters.

A requester is a prepared form containing a collection of graphical objects. When the requester receives a message `popup`, the form is displayed in the middle of the window and receives all event messages until it receives a message `popdown`.

Requesters are useful for asking the user to provide additional information about an action triggered by a button or a menu item. The callback function of the button or menu item then just send a message `popup` to the requester which performs all the remaining tasks.

- Class `Requester` implements the basic mechanism of a requester.
- Class `ErrorRequester` and `WarningRequester` implement requesters specialized for displaying error messages.
- Class `YesNoRequester` defines a requester with a standard state and two standard buttons for accepting or cancelling an action. These requesters are useful for entering extra information about an action triggered by a button or a menu item.
- Class `PrintRequester` implements a standard requester for selecting a printer.
- Class `FileRequester` implements a standard requester for entering a filename for reading or writing.

5.15.1 Requester.

Class `Requester` is a subclass of `Form` which implements the basic mechanisms of a requester. When the requester receives a message `popup`, the requester is inserted into a predefined window object until reception of a message `popdown`.

Example:

```
;;; A window with a button that pops a requester up.
;;; The requester contains an editstring and two buttons.
;;; 1- Create the window
? (setq win
  (new WindowObject 100 100 300 300 "Essai"
    (setq thebutton
      (new StdButton "Pop the requester up"
        (lambda(d) (==> therequester popup)) ))))
= ::WindowObject:06f00
;;; 2- Create the requester
? (setq therequester
  (new Requester win
    (new Column
      ;; The first button pops the requester down
      ;; Variable <thisrequester> always refers to the closest
      (new StdButton "Pop the requester down"
        (lambda(d)
          (==> thisrequester popdown)) )
      ;; The second button changes the label of the popup button
      (new EditString 8 "Pop it up again")
      (new StdButton "Change label"
        (lambda(d)
          (==> thebutton settext
            ;; Collective <getdata> on the requester!
            (car (==> thisrequester getdata))))))))))
= ::Requester:0702c
```

```
(new Requester support ...contents...)
```

Type: CLASS(ogre.lsh)

Creates a new requester containing objects `...contents...`. This requester will be inserted into the window object containing object `support` upon reception of a message `popup`.

```
(==> Requester popup)
```

Type: MSG(ogre.lsh)

When a requester receives message `popup` , it is inserted in the middle of the window object defined at requester creation time.

All event messages are then directed to the requester. Since no other component of the interface `oplevel` is active, it is advisable to insert a button whose action consists in sending message `popdown` to the requester itself.

(==> Requester **popdown**)

Type: MSG(ogre.lsh)

A popped requester is removed from its window object when it receives a message `popdown` . The usual event dispatching is then restored.

(==> Requester **popuplock**)

Type: MSG(ogre.lsh)

Message `popuplock` behaves very much like message `popup` . Unlike message `popup` , message `popuplock` blocks until the requester receives a message `popdown` . When message `popuplock` returns, the window object state is normal again.

Example:

```
;;; Pop up the requester defined above in this section and block:
? (==> therequester popuplock)
;;; Lush only returns when you activate the 'popdown' button
= ()
```

(==> Requester **popuphard**)

Type: MSG(ogre.lsh)

Message `popuphard` is an even more drastic form of `popuplock` . It pops up the requester and blocks until the requester is popped down like message `popuplock` .

While message `popuphard` is active, no other window object accepts user events. If the user click in another Ogre window, the computer beeps and put the window containing the requester above all other window on the screen.

Message `popuphard` is useful for requesting an information which requires immediate attention from the user.

(==> Requester **setsupport win**)

Type: MSG(ogre.lsh)

Message `setsupport` changes the support window of a requester.

5.15.2 Warning Requester.

Class `WarningRequester` is a subclass of `Requester`. A warning requester contains only a `String` object. The actual text of the string is defined when you pop the warning requester up.

(new `WarningRequester` support)

Type: `CLASS(ogre.lsh)`

Creates a new warning requester for displaying text in window object `support`.

(==> `WarningRequester popup text`)

Type: `MSG(ogre.lsh)`

Message `popup` first sets the string's text to `text` and pops the requester up. The requester may be popped down by sending the usual message `popdown`.

Using a warning requester with messages `popup` and `popdown` allows for displaying messages indicating what is being computed.

(==> `WarningRequester popuplock seconds text`)

Type: `MSG(ogre.lsh)`

Message `popuplock` first sets the string's text to `text` and pops the requester up during `seconds` seconds. When this delay is elapsed, the requester is popped down and method `popuplock` returns.

Using a warning requester with `popuplock` allows for displaying a temporary warning message.

5.15.3 Error Requester.

Class `ErrorRequester` is a subclass of `Requester`. An error requester contains a `String` object and a button labelled "Ok" for popping the requester down. The actual text of the string is defined when you pop the error requester up.

Such a requester is very useful for signaling an error to the user. The user must depress the button "Ok" to remove the requester. This button actually sends a message `popdown` to the error requester.

Example:

```
;;; Assuming window object <win> already exists:
? (setq error-dialog (new ErrorRequester win))
= ::ErrorRequester:07120
? (==> error-dialog popup "Error message")
= ()
```

(new **ErrorRequester** support)

Type: CLASS(ogre.lsh)

Creates a new error requester for displaying error messages in window object **support** .

(==> **WarningRequester popup text**)

Type: MSG(ogre.lsh)

Message **popup** first sets the string's text to **text** and pops the requester up. The requester may be popped down by sending the usual message **popdown**.

Using a warning requester with messages **popup** and **popdown** allows for displaying messages indicating what is being computed.

(==> **WarningRequester popuplock text**)

Type: MSG(ogre.lsh)

Message **popuplock** first sets the string's text to **text** and pops the requester up. The requester is popped down and method **popuplock** returns when the user depresses button "Ok" .

5.15.4 **Yes/No Requester.**

Class **YesNoRequester** is a subclass of **Requester** for implementing yes/no requesters. A yes/no requester contains a user defined part, a positive button and a negative button.

Like usual requesters, yes/no requesters are popped up when they receive a message **popup** or **popuplock** . Both the positive and the negative button pop the yes/no requester down.

When you create the yes/no requester however, you specify a callback function.

- When the positive button is depressed, the yes/no requester is popped down and the callback function is executed.
- When the negative button is depressed, the yes/no requester object is popped down and the state of the user defined part of the requester is restored to the state recorded when the requester was popped up. The callback function is not executed.

If you have popped up the requester using message **popuplock** , you can also test which button has been depressed by looking at the value returned by message **popupdown** . This value is the empty list if the negative button has been depressed.

Slots **yesbutton** and **nobutton** of a yes/no requester contain the positive and negative buttons. These buttons may be programmatically acted by sending them a message **trigger** .

(new YesNoRequester support dialog yes no call)

Type: CLASS(ogre.lsh)

Creates a new yes/no requester for displaying on the window object of object support .

Argument `dialog` is the user defined part of the requester. String `yes` is the label of the positive button. String `no` is the label of the negative button. Argument `call` is a callback function or the empty list.

Example:

```
;;; A button that pops up a yesnorequester
;;; for changing its label...
? (setq win
  (new WindowObject 100 100 400 200 "Essai"
    (setq thebutton
      (new StdButton "Hello"
        (lambda(c) (==> theyesnoreq popup)) )))
= ::windowobject:07010
? (setq theyesnoreq
  (new YesNoRequester win
    ;; The user defined part
    (new Column
      (new String "Change button label")
      (new EditString 8 "new label") )
    ;; The yes and no labels
    " Ok " "Cancel"
    ;; The callback
    (lambda(caller)
      (==> thebutton settext
        (car (==> caller getdata)) ))))))
= ::yesnorequester:07040

(==> YesNoRequester settext yes no default)
```

Type: MSG(ogre.lsh)

The label of the positive and negative buttons can be changed by sending a message `settext` to the yes/no requester. String `yes` becomes the label of the positive button. String `no` becomes the label of the negative button.

There is usually a default button indicated by a wider outline. This button is triggered if the user hits the carriage return key. This default button is usually the positive one, but may be changed with argument `default` of message `settext` .

- If `default` is `'yes` , the positive button is the default,
- if `default` is `'no` , the negative button becomes the default,
- if `default` is the empty list, no default button is set.

```
(==> YesNoRequester setcall callback)
```

Type: MSG(ogre.lsh)

Sets the callback function for the requester. Argument `callback` can be a function with one argument or the empty list.

```
(==> YesNoRequester ask yes no default)
```

Type: MSG(ogre.lsh)

See: `(==> YesNoRequester settext yes no default)`

See: `(==> Requester popuphard)`

This message should be used with yes/no requesters whose user defined component is a single String object.

When such a requester receives message `ask`, it first sets the text of the string object to `text` and the button labels to `yes` and `no`. The default button is indicated by argument `default`.

The requester is then popped up using `popuphard`. The user can then press either the positive or the negative button. When the requester is popped down, message `ask` returns `t` or `()` according to the user answer.

Example:

```
;;; Creates a confirmation dialog in an exiting win <win>.
? (setq confirm-dialog
  (new YesNoRequester win
    (new string "msg")           ; dummy message
    "yes" "no"                 ; dummy labels
    () )                       ; no callback
  = ::yesnorequester:070c0
  ? (==> confirm-dialog ask
    "Should I really do that"   ; the message
    "Proceed" "Please don't"     ; the button labels
    'no )                       ; the default
  ;;; Waiting for your answer...
  = t or ())
```

5.15.5 Print Requester.

See: `(print-window w h [destination])`

A print requester lets the user choose a valid destination for printing a graphics. Function `getdata` will return a string that can be used verbatim as the `destination` argument of function `print-window`.

```
(new PrintRequester w [callback])
```

Type: CLASS(ogre.lsh)

See: Yes/No Requester.

Creates a new print requester for displaying in the window object `w` . The optional argument `callback` defines a callback function that will be called if the user presses button "Ok" (as usual with Yes/No Requesters.)

(==> `PrintRequester` `getdata`)

Type: `MSG(ogre.lsh)`

See: (`print-window w h [destination]`)

Returns a string that can be used directly as argument `destination` of function `print-window` .

5.15.6 File Requester.

A file requester lets the user choose a filename using a list of files existing in various directories. The file requester contains a message string, a file selector, an editable filename string and two buttons labelled "Ok" and "Cancel" .

Starting from the current filename, the selector displays the contents of the directory. A first click on a selector item selects a file or directory and copies its name in the editable filename string. Alternatively, the user can type a filename in the filename field. The button "Ok" is disabled if the user types an invalid filename.

This selection is validated by a second click, by pressing the enter key or by depressing the button "Ok" .

- If the selected item is a directory, the contents of this directory is displayed and the user can select again an item.
- If the selected item is a file, the requester is popped down and the callback function is executed.

Class `FileRequester` is a subclass of class `YesNoRequester` which implements the standard Ogre file requester.

(`new FileRequester w [message flag filter callback]`)

Type: `CLASS(ogre.lsh)`

See: Yes/No Requester.

Creates a new file requester for displaying in the window object `w` . The contents of the message string is indicated by the optional argument `message` . The callback function is indicated by the optional argument `callback` or may be set using message `setcall` .

Special behavior are selected by argument `flag` :

- When `flag` is 'no-newfile' , it is forbidden to enter the name of a non existent file.
- When `flag` is 'ask-newfile' , validating the name of a non existent file pops up a confirmation requester. This is useful before creating a new file.

- When `flag` is `'ask-oldfile'`, validating the name of an existent file pops up a confirmation requester. This is useful before overwriting an existing file.

Argument `filter` is a function with one argument for testing the file-names. It returns a non nil value if its argument is a valid file name for this file request. For instance, this function might test an extension or examine the file header. Files rejected by the function `filter` are not displayed in the selector. If the user types such an invalid filename in the filename field, validation is prevented by disabling the button "Ok" .

`(==> FileRequester setdata fname)`

Type: MSG(ogre.lsh)

Sets the current file name of a `FileRequester` object to file or directory `fname` .

`(==> FileRequester getdata)`

Type: MSG(ogre.lsh)

Gets the full current file name of a `FileRequester` object.

`(==> FileRequester getdir)`

Type: MSG(ogre.lsh)

Gets the name of the directory displayed in the selector of a `FileRequester` object.

`(==> FileRequester getbase)`

Type: MSG(ogre.lsh)

Returns the base name of the current filename selected in a `FileRequester` object. If the selected filename is a directory, this message returns the empty list.

`(==> FileRequester setparm message flag [filter])`

Type: MSG(ogre.lsh)

Changes the parameters of a `FileRequester` object.

`(==> FileRequester ask message flag [filter [fname]])`

Type: MSG(ogre.lsh)

Changes the parameters of a `FileRequester` object, sets the current filename to file or directory `fname` , pops up the requester and waits until the request is complete.

- If the user validates a file name, a non nil value is returned. The filename can be accessed by sending a `getdata` message.

- If the user selects the Cancel button, the empty list is returned.

(`ogre-ask-file winobj message flag [filter [fname]]`)

This function popups a default filerequester in window `win` .

Other arguments are the same as for method `ask` of class `FileRequester` .

See: (`==> FileRequester ask message flag [filter [fname]]`)

5.16 Movable and Resizable Objects.

Class `DragArea` and `SizeArea` define objects used for moving or resizing their container. Special container classes `Frame` and `FrameSize` define containers which may be moved and resized by the user.

5.16.1 (new DragArea w h)

Type: CLASS(`ogre.lsh`)

Creates a new dragging area with minimal width `w` and minimal height `h` .

A dragging area appears as a rectangle with a gray outline. When you depress the mouse inside a dragging area, you can drag its container to another place.

Two slots affect the behavior of a dragging area:

- The moved object always remains inside rectangle `constraintrect` . This rectangle defaults to the empty list `()` which means that the object remains confined within its container's rectangle.
- Slot `magnet` should contain a list `(mx my)` . The moved object keeps aligned on an invisible grid whose columns are `mx` pixels wide and whose rows are `my` pixels high. It defaults to the empty list `()` which means that objects may be moved at any pixel position.

5.16.2 (new SizeArea w h)

Type: CLASS(`ogre.lsh`)

Creates a new sizing area with minimal width `w` and minimal height `h` .

A sizing area appears as two overlapping gray squares. When you depress the mouse inside a sizing area, you can change the size of its container. Just drag the mouse until you touch a container boundary. This container boundary then follows the mouse until you release the mouse button.

Two slots affect the behavior of a size area.

- The resized object always remains confined within the rectangle `constraintrect` . This rectangle defaults to the empty list `()` which means that the object remains confined within its container's rectangle.

- Slot `formfactor` should contain a list `(fx fy)` . The width of resized object remains an integer multiple of `fx` , the height of the resized object remains an integer multiple of `fy` , and both coefficients are equal. Slot `formfactor` defaults to the empty list `()` which means that no form factor constraint apply.

5.16.3 (new Frame x y w h ...contents...)

Type: CLASS(ogre.lsh)

Creates a new frame at location `x` and `y` with minimal width `w` and minimal height `h` . The frame initially manage objects `contents` .

Class `Frame` is a subclass of class `Container` . A frame always contains a dragging area located in the background of the container. A frame thus is a container that the user can move.

Example:

```
? (setq win (new windowobject 0 0 400 400 "Essai"
             (new Frame 50 30 200 140
              (new stdbutton "Beep"
               (lambda(c) (beep)) ) ) ) )
```

5.16.4 (new FrameSize x y w h ...contents...)

Type: CLASS(ogre.lsh)

Creates a new sizable frame at location `x` and `y` with minimal width `w` and minimal height `h` . The frame initially manage objects `contents` .

Class `FrameSize` is a subclass of class `Frame` . A sizable frame always contains a dragging area located in the background of the container. A sizable frame also contains a small size area in its bottom right corner. A sizable frame thus is a container that the user can move and resize.

Example:

```
? (setq win (new windowobject 0 0 400 400 "Essai"
             (new FrameSize 50 30 200 140
              (new stdbutton "Beep"
               (lambda(c) (beep)) ) ) ) )
```

5.17 Sliders and Scrollbars.

Sliders and scrollbars are graphical objects for entering numerical values.

- Sliders are rendered as symbolized potentiometers. Sliders are mostly used as a way to enter numerical values.

- Scrollbars are rendered as a gray area with a white handle. Scrollbars are mostly used for controlling other objects.

The user can grab and move the handle with the mouse. The user might also maintain the mouse button depressed on either side of the handle. In this case, the handle moves slowly towards the mouse pointer.

The standard callback function is called whenever the user releases the mouse button. An additional callback might be set up with message `setdrag` . This callback is called whenever the user moves the handle.

Sliders and scrollbars are implemented by a set of specialized subclasses of class `Control` .

<code>Control</code>	abstract class for all controls
<code>Slider</code>	abstract class for all sliders and scrollbars
<code>HSlider</code>	horizontal sliders
<code>VSlider</code>	vertical sliders
<code>Scrollbar</code>	abstract class for all sliders and scrollbars
<code>HScrollbar</code>	horizontal scrollbars
<code>VScrollbar</code>	vertical scrollbars

Example: Example:

```
? (setq win (new windowobject 100 100 400 200 "Sliders & Scrollbars"
  (new row
    (new grid 2
      (new emptyspace 100 100)
      (new vslider 0 100 ())
      (new hslider 0 100 ())
      (new emptyspace 10 10) )
    (new grid 2
      (new emptyspace 100 100)
      (new vscrollbar 100 ())
      (new hscrollbar 100 ())
      (new emptyspace 10 10) ) ) ) )
```

5.17.1 Sliders.

Sliders are used for entering numerical values using an analog handle. Horizontal sliders are implemented by class `HSlider` . Vertical sliders are implemented by class `VSlider` . Both class `HSlider` and `VSlider` are subclasses of class `Slider` which implement a number of useful methods.

```
(new HSlider mini maxi [callback])
```

Type: CLASS(ogre.lsh)

Returns a horizontal slider for entering integer values in the range `mini` to `maxi` . The callback function `callback` is called whenever the user releases the mouse button.

(new VSlider mini maxi [callback])

Type: CLASS(ogre.lsh)

Returns a vertical slider for entering integer values in the range `mini` to `maxi` . The callback function `callback` is called whenever the user releases the mouse button.

(=> Slider setrange mini maxi)

Type: MSG(ogre.lsh)

See: **(=> Scrollbar setrange mini maxi [prop])**

Message `setrange` redefines the minimal and maximal values allowed in a slider or a scrollbar.

(=> Slider setstep step)

Type: MSG(ogre.lsh)

Message `setstep` redefines the possible increments of the values allowed in a slider or a scrollbar.

The initial increment is always 1. This initial increment ensures that the slider or scrollbar is limited to integer values. Specifying the empty list as an increment means that any value in the legal range are allowed.

(=> Slider setdrag call)

Type: MSG(ogre.lsh)

Message `setdrag` sets up a secondary callback function which is called whenever the handle moves. Installing such a callback function is useful for displaying echo to the user.

5.17.2 Scrollbars.

Scrollbars are a special kind of slider which are used for controlling which information is displayed in other objects.

Scrollbars are implemented by classes `HScrollbar` and `VScrollbar` which are indirect subclasses of class `Slider` . All the methods defined by class `Slider` are thus inherited by scrollbars.

(new HScrollbar maxi [callback])

Type: CLASS(ogre.lsh)

Returns a horizontal scrollbar for entering integer values in the range 0 to `maxi` . The callback function `callback` is called whenever the user releases the mouse button.

```
(new VScrollbar maxi [callback])
```

Type: CLASS(ogre.lsh)

Returns a vertical scrollbar for entering integer values in the range `mini` to `maxi`. The callback function `callback` is called whenever the user releases the mouse button.

```
(=> Scrollbar setrange mini maxi [prop])
```

Type: MSG(ogre.lsh)

See: (`=> Slider setrange mini maxi`)

Message `setrange` redefines the minimal and maximal values allowed in a scrollbar. When the optional argument `prop` is provided, the maximal value is reduced by `prop` and the value `prop` is used for defining the size of the knob. This is useful for scrolling lists of texts.

5.18 Composite Objects.

This chapter describes a few standard objects composed of several elementary objects. Although composite objects usually are subclasses of `Form`, they obey the protocols defined for class `Control`.

5.18.1 Viewers.

A viewer object display a rectangular portion of a particular object (the viewer's contents). The size of the visible part of this object depends on the viewer size. The user control which part of the object is visible by sending messages `setpos` or by using the optional scrollbars provided by the viewer.

Class `Viewer` implements viewer objects.

```
(new Viewer w h contents [hp vp])
```

Type: CLASS(ogre.lsh)

Creates a new viewer on object `contents`. The viewer is `w` points wide and `h` points high. The optional argument `hp` (resp. `vp`) is a boolean value (`t` or `()`) controlling whether an horizontal (resp. vertical) scrollbar is displayed or not.

Example:

```
? (setq win (new WindowObject 0 0 400 200 "Essai"
           (new Viewer 300 150
             (new column (new stdbutton "One" ())
                         (new emptyspace 60 60)
                         (new stdbutton "Two" ())
                         (new stdbutton "Three" ())
                         (new stdbutton "Four" ()))
```

```
t t) ) )
```

```
(==> Viewer setpos h v)
```

Type: MSG(ogre.lsh)

Message `setpos` defines the visible portion of the contents of a viewer. This portion is a rectangle whose top left corner is located `h` pixels to the right and `v` pixels below the top left corner of the contents of the Viewer object.

```
(==> Viewer sethpos h)
```

Type: MSG(ogre.lsh)

Message `sethpos` just change the horizontal coordinate of the visible portion of the contents of a viewer.

```
(==> Viewer setvpos v)
```

Type: MSG(ogre.lsh)

Message `setvpos` just change the vertical coordinate of the visible portion of the contents of a viewer.

```
(==> Viewer setcontenu object)
```

Type: MSG(ogre.lsh)

Message `setcontenu` changes the object viewed through a Viewer object. Argument `object` must be a valid graphic component.

5.18.2 Selectors.

A selector is an object which allow the user to select one or several strings within a list of strings. Only a few strings are visible at a given time. A scrollbar located along the right side of the Selector object controls which strings are visible. Selected strings are highlighted.

Class `Selector` implements a selector object. The complete behavior of a selector is controlled by three slots of the selector object: `multiple` , `call1` and `call2` .

- When slot `multiple` contains a non nil value, the user might select multiple items. In this case, message `getdata` sets and message `setdata` returns the list of the selected strings.
- If slot `multiple` contains the empty list, one item at most can be selected at a given time. Message `getdata` sets and message `setdata` returns the selected string or the empty list.

Two callback functions are called when the user selects items.

- Callback `call1` is called whenever the user selects an item by clicking the mouse button over an unselected string.

- Callback `call2` is called whenever the user selects an item a second time by clicking the mouse button over a highlighted string.

(**new Selector** `nvisible` [`callback` [`items`]])

Type: CLASS(ogre.lsh)

Creates a new selector able to display `nvisible` strings at once.

Argument `callback` of the constructor controls the values of these flags.

- If argument `callback` is `t`, the user might select multiple items.
- If argument `callback` is a function, one item at most can be selected at a given time. The specified callback function is called whenever the user selects a new item.
- If argument `callback` is a non-nil list, one item at most can be selected at a given time. The list is assumed to be made of two functions or nil values. The first one is used whenever the user selects a new item. The second one is used whenever the user selects again a selected item.

Argument `items` optionally gives a list of strings initially displayed in the selector.

(**==> Selector** `getdata`)

Type: MSG(ogre.lsh)

Message `getdata` returns the selected strings in a selector:

- If slot `multiple` contains a non nil value, the user can select multiple items. In this case, message `getdata` returns the list of the selected strings.
- If slot `multiple` contains the empty list, the user can select at most one item. Message `getdata` then returns the selected string or the empty list if no string is selected.

(**==> Selector** `getdatanum`)

Type: MSG(ogre.lsh)

Instead of returning the strings themselves, like message `getdata`, message `getdatanum` returns the order number of the selected strings in a selector:

- If slot `multiple` contains a non nil value, the user can select multiple items. In this case, message `getdatanum` returns the list of the order numbers of selected strings.
- If slot `multiple` contains the empty list, the user can select at most one item. Message `getdata` then returns the order number of the selected string or the empty list if no string is selected.

(==> Selector **setdata** data)

Type: MSG(ogre.lsh)

Message **setdata** sets which strings are selected in a selector. The selected strings are displayed on a highlighted background.

- If slot **multiple** contains a non nil value, argument **data** must be a list of strings or a list of numbers indicating which items must be selected.
- If slot **multiple** contains the empty list, argument **data** is the empty list (to desselect the selected item), a string or a number indicating which string must be selected.

(==> Selector **setpos** pos)

Type: MSG(ogre.lsh)

Message **setpos** makes the **pos** -th string appear on top of the selector and updates the scrollbar in accordance.

(==> Selector **setitems** items)

Type: MSG(ogre.lsh)

Message **setitems** redefines a new list of strings **items** displayed in a selector and clears the current selection.

Since the minimal size of a selector depends on the width of the largest string in the item list, sending message **setitems** can trigger a geometry negotiation and readjust the location of all interface components.

5.19 A Complete Example.

Whole graphics interface are often defined as subclasses of class **WindowObject** . Such subclasses usually contain new slots for storing interface specific data. Callback functions then send messages to **thiswindowobject** . These messages then are executed within the interface scope.

5.19.1 The Class Browser “classtool”.

This section describes a complete application of the Ogre library. This application is nothing but the standard Tlisp class browser. This class browser is invoked by typing the command (**classtool**) .

(**classtool** [c1])

Type: DE(classtool.lsh)

Invokes the Lush class browser on class **c1** . The class browser displays the subclasses, superclasses, slots and methods of a class. Class **object** is assumed if no argument is specified.

The class browser interface is composed of a menu and six selectors:

- Selector “Sub-Classes” displays the names of the subclasses of class `c1` . A mouse click on one of these subclasses changes the current class to the selected class.
- Selector “Slots” displays the names of all slots defined by class `c1` .
- Selector “Method” displays the names of all method selectors defined for class `c1` .
- Selector “Super-Classes” displays the names of the successive superclasses of `c1` . A mouse click on one of these subclasses changes the current class to the selected class.
- Selector “Inherited Slots” displays the names of all slots defined by the successive superclasses of class `c1` . A mouse click on a method item prints the definition of the selected method.
- Selector “Inherited Methods” displays the names of all methods inherited from the successive superclasses of class `c1` . A mouse click on a method item prints the definition of the selected method.

The top of the class browser contains a menu and an information string. The information string displays the number of slots, the number of inherited slots, the number of methods and the number of inherited methods for the current class. The name of the menu is always the name of the current class, displayed in large characters.

The menu itself contains four items:

- Selecting item "Show Class" prints the definition of class `c1` .
- Selecting item "Show Subtree" prints an indented list of the subclasses of class `c1` .
- Selecting item "Refresh" reads again the information for class `c1` and reflects possible changes on the display.
- Selecting item "Select" pops up a requester which lets the user enter the name of a class. The class browser jumps to the selected class when the user presses button "Ok" . If the typed class name is not a valid class name, a message is displayed.

5.19.2 The Program “classtool”.

Here is a review of the main components of the “classtool” program. You can look at the complete listing of this program in file:

```
<lshdir>/lib/classtool.lsh
```

\vspace 5mm \bf i) Initial definitions.

The first executable line of file "`classtool.lsh`" initializes the Ogre library by calling function `ogre` . This is necessary to ensure that the Ogre class library is properly loaded and initialized.

(ogre)

Then we define a subclass `c-classtool` of class `WindowObject` . This class contains several slots for referencing the major components of the interface.

```
(defclass c-classtool windowobject
  the-menu          ;; the menu
  the-string        ;; the information string
  the-i-classes     ;; the superclass selector
  the-i-slots       ;; the inherited slots selector
  the-i-methods     ;; the inherited methods selector
  the-classes       ;; the subclass selector
  the-slots         ;; the slots selector
  the-methods       ;; the method selector
  the-error-requester ;; a signaling requester
  the-class-requester ;; the requester
  cl )              ;; the current class
```

\vspace 5mm \bf ii) Constructor Method.

We define then the constructor of class `c-classtool` . This constructor first calls the constructor of its superclass `WindowObject` and defines the contents of the window object. This very long call sets up the major components of the `classtool` interface.

All the interface is a single column which contains:

- A row implements the menu bar. The menu bar contains only one menu object (stored in slot `the-menu`) and an information string (stored in slot `the-string`).

The menu defines the five items documented above. The callback functions of the menu items do not perform very much. They merely send an appropriate action message to the interface itself (acceded through variable `thiswindowobject`) or pop up a suitable the requester.

```
(setq the-menu
  (new Menu "object"
    "Show Class"
    (lambda(c) (==> thiswindowobject display-action))
    "Show Subtree"
    (lambda(c) (==> thiswindowobject subtree-action))
    "Refresh"
```

```
(lambda(c) (==> thiswindowobject refresh-action))
"Select"
(lambda(c) (==> the-class-requester popup))
"Quit"
(lambda(c) (==> thiswindowobject delete)) ) )
```

- A dark space object provides a clear separation between the menu bar and the rest of the interface. Although the minimal space allocated for the dark space 3 points wide, the column layout policy makes this space as wide as the column itself.
- A grid with 3 columns contains all the other objects. It includes the selector titles (string objects) and the selector themselves which are stored in appropriate slots of the classtool object. The callback functions of the selector do not perform much. They just send messages `classes-action` or `methods-action` to the classtool object itself.

The grid also include three `emptyspace` object which specify a minimal size for the grid columns. This technique avoids troublesome geometry changes because it ensures that the selectors are already wide enough for displaying most names.

The constructor of class `c-classtool` then adjust the font used in the menu title by directly poking into the object slot `textfont` . It calls then method `compute-geometry` to ensure that the object size is adjusted for the new font.

```
(setq :the-menu:textfont font-18)
(==> the-menu compute-geometry)
```

The constructor of class `c-classtool` then creates two requesters.

- The error requester `the-error-requester` is used later for displaying error messages.

```
(setq the-error-requester
      (new ErrorRequester this) )
```

- The class requester is popped up when you select menu item `"select"` for directly entering a class name. It defines a callback function which just sends a message `select-action` to the classtool object.

We poke a new regular expression in the slot `regex` of the editable string in order to ensure that the text typed by the user is a valid symbol.

```
(setq the-class-requester
      (new YesNoRequester this
```

```
(new column
  (new String "Type a class name")
  (new DarkSpace 3)
  (let ((x (new EditString 20)))
    (setq :x:regex "[A-Za-z]?[-_|A-Za-z0-9]*") x) )
" Ok " "Cancel"
(lambda(c) (==> thiswindowobject select-action)) ) )
```

\vspace 5mm \bf iii) Method “setclass”.

Method `setclass` is then defined. This method collects the class information for the selected class and updates the information displayed in the selectors.

It first checks that its argument is a valid class and pops up the error requester if this check is negative.

```
(defmethod c-classstool setclass(c)
  (if (not (and c (classp c)))
    (==> the-error-requester popup "This is not a valid class")
```

If the check is positive, method `setclass` displays a message “working” in the message string and force an immediate display update using message `repair-damaged` .

```
(==> the-string settext "<<working>>")
(==> this repair-damaged)
```

The menu title is then changed to the class name using method `settext` . Since the display update is delayed until all events are processed, this change becomes visible when all selectors are updated.

Method `setclass` then collects the class information into six lists: the subclass list (`cc`), the slot list (`cs`), the method list (`cm`), the superclass list (`ic`), the inherited slot list (`is`) and the inherited method list (`im`). We take a particular care of separating the various inherited classes by a dummy entry in the inherited lists.

This information is then loaded into the selectors using method `setitems` . The message string is then updated to the class statistics.

```
(==> the-classes setitems cc)
(==> the-slots setitems cs)
(==> the-methods setitems (sort-list cm >))
(==> the-i-classes setitems ic)
(==> the-i-slots setitems is)
(==> the-i-methods setitems im)
(==> the-string settext
  (sprintf "  \\%1:  \\%d+\\%d slots,  \\%d+\\%d methods"
    cn (length cs) isc (length cm) imc) ) ) ) )
```

\vspace 5mm \bf iv) Action methods.

We define then the various action methods which are called by the menu items, by the selectors, or by the class selection requester.

Method `classes-action` is called when the user selects a class in the subclass or superclass requester. This method just sends a message `setclass` to switch to the new class.

```
(defmethod c-classtool classes-action(c)
  (let ((cn (==> c getdata)))
    (==> this setclass (apply scope (list (named cn)))) ) )
```

Method `methods-action` is called when the user selects a method in the method selector or the inherited method selector. Method `methods-action` first makes sure that the user did not select a dummy entry used for separating the classes in the inherited method selector. It uses then function `pretty-method` to print the definition of the selected method.

```
(defmethod c-classtool methods-action(c)
  (let ((m (==> c getdata))
        (cl cl))
    (when (<> (left m 5) "=====")
      (setq m (named m))
      (while (and cl ~(member m (methods cl)))
        (setq cl (super cl)) )
      (when cl
        (==> this repair-damaged)
        (==> c setdata ())
        (print)
        (pretty-method cl m) ) ) ) )
```

Method `refresh-action` is called by the menu item "Refresh" . It just calls method `setclass` on the current class (from slot `cl`) in order to reload the class information.

```
(defmethod c-classtool refresh-action()
  (==> this setclass cl) )
```

Method `display-action` is called by the menu item "Show Class" . It just prints the class definition when adequate.

```
(defmethod c-classtool display-action()
  (if (not (super cl))
    (beep)
    (print)
    (pprint (nconc (list 'defclass (classname cl) (classname (super cl)))
                  (slots cl) ) ) ) )
```


Chapter 6

Dynamic Loader/Linker

One of the coolest features of Lush is its dynamic loader/linker. The dynamic loader/linker allows to load object files (.o, .so, or .a) and make their functions easily accessible from the interpreter. Functions in dynamically loaded object files can be written in any language (though C is preferred).

Object files can be dynamically loaded into Lush using the `(mod-load "mylibrary.o")` construct.

Let's say you have written a C file called `tt titi.c /tt` with the following content:

```
float sq(float x)
{
    return x*x;
}
```

You have compiled the file and produced the object file `titi.o`. Calling `sq` from the Lush interpreter is as simple as the following. First, dynamically load the object file into Lush

```
? (mod-load "titi.o")
```

then, write a compiled lisp function whose only purpose is to call `tt sq /tt`. To be able to call C from Lush:

```
? (de square (x) ((-float-) x)
    (cheader "extern float sq(float);")
    (float #{ sq( $x ) #} )
? (dhc-make "junk" square)
```

The function `square` can now be called:

```
? (square 5)
= 25
```

Dynamically loaded Modules can be `.o` object files, but also `.a` files (static libraries), or `.so` files (shared objects). Functions in dynamically loaded modules can call any external function or variable defined or used by Lush as well as external functions and variable defined in other modules. In particular, the C library functions are accessible. A function defined in a module however is not always executable. Indeed, its module might call an undefined function, or a function defined by another non executable module. In fact, four situations occur:

- Initialized and executable modules are the only accessible modules. All functions referenced by these modules have been found, and the initialization routine (e.g. `init_essai`) has been successfully executed, creating descriptors for the new lisp functions defined in the modules. At this point, all new lisp functions defined by such a module are accessible and work as expected.
- Uninitialized modules reference some undefined functions, or some functions defined by a non executable module. Therefore, the initialization function has not been executed, and the descriptors for the new lisp functions defined by such a module have not been created.
- Modules may be initialized but non executable. Such a situation occurs when a module has been initialized and executable, but is no longer executable, because it uses some function or global variable which is no longer defined, because its module has been unloaded.
- Finally, certain modules do not define an initialization function. Such modules just define C functions used by other modules. We say that such a module is in a unknown state.

6.1 (mod-load filename)

Type: DE(sysenv.lsh)

This function loads a piece of binary code into LUSH. It can be used to load various kind of files containing object code:

- Shared libraries (usually a file with extension `".so"`)
- Object files (usually a file with extension `".o"`)
- Library archives (usually a file with extension `".a"`) containing a collection of object files. `mod-load` only loads the components of a `.a` file that provide definitions for currently undefined symbols. You may want to use functions `mod-create-reference` to create fake undefined symbols and selectively load parts of the library before actually using them.

6.2 (mod-unload filename)

Type: DE(sysenv.lsh)

This function removes a piece of binary code previously loaded with `mod-load`. Only object files (extension ".o") and library archives (extension ".a") can be safely removed. It is currently not possible to unload a shared library module (usually a file with extension ".so") because these files are dealt with using operating system facilities that seldom provide unloading support.

This operations encompasses three steps:

- Destroying all lisp functions previously defined by the module. Calling such functions will cause an error in the future.
- Relinquishing the memory utilized by the module.
- Checking the executability of all loaded modules, and mark the lisp functions defined by a module as partially linked, if this modules is no longer executable.

6.3 (find-shared-library name [extlist])

Type: DE(sysenv.lsh)

Returns the pathname of a shared library named `name`. Shared libraries are searched in the directories specified by variable `shared-include-path` which is initialized by "`lushenv.lsh`".

The optional argument `extlist` is a list of possible filename extensions. The default value is either null (when `name` already contains an extension) or the system dependent filename extension for shared libraries.

```
? (find-shared-library "libm")
= "/usr/lib/libm.so"
```

6.4 (find-static-library name [extlist])

Type: DE(sysenv.lsh)

Returns the pathname of a static library named `name`. Static libraries are searched in the directories specified by variable `static-include-path` which is initialized by "`lushenv.lsh`".

The optional argument `extlist` is a list of possible filename extensions. The default value is either null (when `name` already contains an extension) or the system dependent filename extension for static libraries.

```
? (find-static-library "libm")
= "/usr/lib/libm.a"
```

6.5 (find-shared-or-static-library name [extlist])

Type: DE(sysenv.lsh)

Returns the pathname of a library named `name` . A shared library is first searched with the supplied extension list `extlist` . Otherwise a static library is searched.

```
? (find-shared-or-static-library "libm")
= "/usr/lib/libm.so"
```

6.6 (mod-list)

Type: DE(sysenv.lsh)

This function returns the list of the currently loaded modules.

Example:

```
? (mod-list)
= ("/home/leonb/lush/src/lush" "/home/leonb/test/essai.o")
```

6.7 (mod-undefined)

Type: DX

This function returns a list with the names of all undefined C functions and global variables in the current modules.

Example:

```
? (mod-undefined)
= ("compute_squares" "numbers_of_squares")
```

6.8 (mod-status)

Type: DE(sysenv.lsh)

Displays a summary of all loaded modules.

6.9 (mod-inquire filename)

Type: DE(sysenv.lsh)

This function returns a list describing the status of a loaded module defined by the object file `filename` . The first element of this list is a string describing the states of a module. When the initialization function has been called, the names of the new lisp functions are provided in the remaining part of this list.

6.10 (mod-create-reference string1 ... stringN)

Type: DX

Creates a fake undefined symbol that will be considered when loading libraries (archive files like "foo.a"). The loader indeed only loads the library components which define symbols currently undefined. You may want to use functions `mod-create-reference` to create fake undefined symbols and load certain parts of the library before actually using them.

6.11 (mod-compatibility-flag boolean)

Type: DX

The old loader was based on the DLD-3.2.3 library. This library had significant bugs in the code checking the executability of a module. The new loader implements these checks properly. This rightful code may prevent you to load your old files. You can use function `mod-compatibility-flag` with a non nil argument to loosen the checks until the new system is almost as buggy as the old one.

6.12 Low level module functions.

Most of the above functions are in fact written in Lisp using a set of lower level functions. LUSH modules are represented by lisp object of class |MODULE| . The following functions manipulate these objects.

6.12.1 (module-list)

Type: DX

Returns the list of all currently loaded modules. Unlike `mod-list` this function returns the module objects instead of the module filenames.

6.12.2 (module-filename m)

Type: DX

Returns the filename associated with module object `m` .

6.12.3 (module-executable-p m)

Type: DX

Test if the code for module `m` is executable.

6.12.4 (module-unloadable-p m)

Type: DX

Test if module `m` can be removed from memory.

6.12.5 (module-initname m)

Type: DX

Returns the name of the C initialization function for module `m` .**6.12.6 (module-depends m)**

Type: DX

Returns the list of all the initialized modules that depends on module `m` . This function is useful to evaluate the consequences of a call to `module-unload` .**6.12.7 (module-never-unload m)**

Type: DX

Make sure that the module `m` will not be unloaded. Attempts to unload the module will cause an error.**6.12.8 (module-defs m)**

Type: DX

Returns an alist describing all the primitives defined by module `m` . This alist is populated the first time the module becomes executable. It is the user's responsibility to make this primitives available for general use by defining conveniently named variables.**6.12.9 (module-load filename [hookfunction])**

Type: DX

Loads the binary code file `filename` and returns a module object.Function `hookfunction` is called with two arguments when the state of the module changes. The first argument is a selector symbol representing the nature of the state change. The second arguments is the module object itself.

The selector symbol can take the following values:

- `init` The hook function is called with selector `init` just after calling the initialization function of the module. Most useful hook functions will scan the list of definitions returned by `module-defs` and define global symbols to access the newly defined primitives.
- `exec` The hook function is called with selector `exec` whenever the executability of the module changes because of loading/unloading another module.
- `unlink` The hook function is called with selector `unlink` just before unloading the module. This is a good time to revert the changes made during the module initialization.

6.12.10 (module-unload m)

Type: DX

Unloads the binary module `m` .**6.13 Extending the Interpreter**

While most users will prefer limit themselves to writing lisp functions (possibly with inline C code) and compiling them, some adventurous users may need to extend the interpreter more directly. This section describes how to do that by making use of the dynamic loader/linker.

Here is an example of a file, named `"essai.c"` , which defines a new lisp function written in C for computing the square of the hypotenuse of triangles.

```

/* ----- Beginning of File "essai.c" ----- */
#include "header.h"
/* This is the function that does the work */
double hypotenuse(x,y)
double x,y;
{
    double z = x*x + y*y;
    printf("hypot(\\%f,\\%f)=\\%f\\n",x,y,z);
    return z;
}

/* This is the interpreter interface function */

DX(xhypotenuse) {
    ARG_NUMBER(2);
    ALL_ARGS_EVAL;
    return NEW_NUMBER( hypotenuse(AREAL(1),AREAL(2)));
}

/* This is the initialization routine.
 * Its name is formed by prepending "init_" to the file name
 */
void init_essai()
{
    dx_define("hypotenuse",xhypotenuse);
}
/* These two (optional) definitions to guarantee that
 * this module will only be loaded by compatible versions of LUSH.
 */
int majver_essai = 40; /* LUSH_MAJOR */

```

```
int minver_essai = 10; /* LUSH_MINOR */
/* ----- End of File "essai.c" ----- */
```

Integrating this function into the LUSH interpreter can be achieved in two ways:

- Copying this file in the "src" directory of Lush, adding "essai.o" to the list of objects in the "Makefile" , calling "init_essai()" from file "toplevel.c" and recompiling everything. This will build a new version of Lush with the new function.
- Compiling this file separately and loading the resulting object file into Lush at run-time. This solution is named dynamic linking of an external module. Modules are loaded with the function `mod-load` , and unloaded with the function `mod-unload` . These functions allocate the necessary memory, relocate the machine code, and resolve the external references.

The compilation is performed by the command:

```
globina\% gcc -c -I/home/leonb/lush/include essai.c
```

The resulting "essai.o" file can be loaded into LUSH with the following command.

```
? (mod-load "essai.o")
= "/home/user/subdir/essai.o"
? (hypotenuse 3 4)
hypot(3.000000,4.000000)=25.000000
= 25
```

6.14 Debugging modules with gdb.

Although GDB is unaware of Lush's dynamic linking capabilities, it contains decent support to debug code contained in shared libraries (i.e. files with extension ".so"). The best way to debug modules therefore consists in building them as shared libraries.

Example:

```
globina\% gcc -shared -o essai.so -I/home/leonb/lush/include essai.c
globina\% gdb lush
... GDB starts
(gdb) run
... LUSH starts
? (mod-load "essai.so")
= "/home/leonb/essai.so"
```

```
<CTRL-C>
(gdb) br xhypotenuse
Break in xhypotenuse
(gdb) cont
? (hypotenuse 3 4)
Breakpoint 1, 0x4001a8c1 in xhypotenuse () from /home/leonb/essai.so
(gdb) ....
```

There are two limitations:

- Shared libraries must be completely linked. They can only reference symbols contained in the main lush executable and its libraries (not in other dynamically loaded modules)
- Shared libraries cannot be unloaded and replaced by another version.

Chapter 7

CLush: Compiled Lush

Author(s): Leon Bottou , Yann LeCun

[under construction]

One of Lush's most interesting features is its Lisp-to-C compiler and the ability to freely intermix Lisp and C code in a single function. Unlike many other interpreted/compiled languages , and contrary to many well established principles of language design , the two dialects that the Lush interpreter and the Lush compiler implement are two very different languages , though they share the same syntax. The compiled dialect is strongly typed , lexically bound , and has no garbage collector , while the interpreted dialect has loose typing , is dynamically bound , and is garbage collected.

In fact , the Lush compiler is not designed to replace the interpreter , but rather to complement it. Lush applications are generally a combination of compiled and interpreted code. The compiled part is often itself a combination of Lush and C code. The compiled code will generally contain the "expensive" and heavily numerical parts of an application , where performance is the main requirement , while the interpreted part will contain the high-level logic , the user interface , the memory management , etc , where flexibility is more important than pure performance.

7.1 A Simple Example

[under construction]

Numerous examples of compiled functions with and without in-line C-code are available in the libraries (in directories `lsh` and `packages`).

7.2 Using `libload` and `dhc-make`.

Lush contains convenient functions that implement capabilities similar to `make`. Properly using these functions ensures that the Lush interpreter loads the most

recent version of the functions as necessary , and that all functions are judiciously recompiled when functions they call have been redefined.

Users should just follow the following four rules:

- Every Lush file should use the function `libload` to load all required files (i.e. files defining functions or classes required by this lush file).
- The compilation of Lush function and classes should be performed by inserting one or several calls to `dhc-make` at the end of the lush file that defines those functions and classes.
- Users should identify the main Lush file of the project , that is to say the file that directly or indirectly loads all the Lush files of your project. This is usually the file with the higher level functions. It is sometimes convenient to create a Lush file for the sole purpose of being the main Lush file.
- When you want to load the most recent version of all your project files , you just need to load the main Lush file using either `libload` or the macro-character `!^L!` . This will load all modified files , reload all files that depends on modified files , and recompile all functions and classes that need recompilation , either because they have been modified , or because they depend on modified classes or functions.

See: (`libload libname [opt]`)

See: (`dhc-make fname fspec1 ... fspecN`)

7.3 Compilation Functions

7.3.1 High-Level Compilation Functions

(`dhc-make fname fspec1...fspecN`)

Type: DM(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Function `dhc-make` controls the compilation of classes and functions specified by the `fspecs` arguments.

Although this function can be used from the command line , it mostly is used within a Lush file to control the compilation of classes and functions defined in this file. We will assume now that `dhc-make` is invoked from a Lush file.

The first argument `fname` is either the empty list or a string representing filenames of the C and object files output by the Lush compiler.

- Passing the empty list instructs `dhc-make` to generate a suitable filename derived from the name of the Lush file containing the `dhc-make` command. Generated C files will be created in a subdirectory "C" of the directory containing the Lush file. Generated object files will be created in a subdirectory of "C" whose platform dependent name is provided by expression (`getconf "host"`) .

- Passing a string specifies a particular base name. Function `dhc-make` adds the suitable suffixes for C files or object files, replacing any other suffix possibly present in `fname`. Argument `fname` may specify a directory for both the C and object files. Otherwise the C and object files will be generated as explained above.

Each of the following arguments `fspec` is either a function name or a list composed of a class name and method names. The following example, for instance, first compiles function `foo`, then compiles methods `rectangle` and `area` of class `rectangle`, and finally compiles function `bar`.

```
(dhc-make ()
      foo
      (rectangle rectangle area)
      bar )
```

The order of the `fspec` arguments is important because no function, class or method can be compiled before compiling all functions, classes, or methods it references. This order can be different from the function and class definition order in the Lush file.

Before starting the compilation, function `dhc-make` tests whether the target files already exist and are up-to-date:

The C file will be generated if any of the following conditions is true.

- no such C file exists,
- the Lush file is more recent than the existing C file,
- the Lush file depends on another Lush file which is more recent than the existing C file. These dependencies are those collected by function `libload`.

The object file will be generated if any of the following conditions is true:

- no such object file exists,
- the C file is more recent than the existing object file, as happen, for instance, when the C file was regenerated during this invocation of `dhc-make`.

Finally `dhc-make` performs a `mod-load` of the object file and checks that the compiled functions are now executable. It returns the name of the object file.

Note: Function `dhc-make` can also be invoked from the command line. When this is the case, it always generates fresh C and object files, and derives their names from the last `fspec` argument.

(dhc-make-with-libs fname libs fspec1...fspecN)

Type: DM(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

See: (dhc-make fname fspec1 ... fspecN)

This function is similar to `dhc-make` but takes an additional argument `libs` to handle cases where the compiled functions depend on external C libraries. This argument can take two values:

- Argument `libs` may be a list of library names. Function `dhc-make-with-libs` works like `dhc-make` but also loads all the specified libraries after loading the compiled functions, and before checking whether the compiled functions are executable.
- Argument `libs` may be the symbol `t`. Function `dhc-make-with-libs` then works like `dhc-make` but does not cause an error if the compiled functions are not executable because they contain unresolved references. The caller is expected to load additional modules using `mod-load` in order to satisfy these unresolved references.

(dhc-make-with-c++ fname libs fspec1...fspecN)

Type: DM(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

See: (dhc-make fname fspec1 ... fspecN)

See: (cinline format var1 ... varn)

This function is similar to `dhc-make-with-libs` but uses the C++ compiler instead of the C compiler to generate the object file. This makes no difference except that one can use C++ constructs when inlining C code with `cinline`. This is particularly handy when interfacing external C++ libraries.

(dhc-make-all fname fspeclist libs)

Type: DE(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

See: (dhc-make fname fspec1 ... fspecN)

See: (dhc-make-with-libs fname libs fspec1 ... fspecN)

This is the elementary function called by `dhc-make`, `dhc-make-with-libs` and `dhc-make-with-c++`.

Argument `fspeclist` is a list containing the `fspec1 ... fspecN` arguments of function `dhc-make`. Argument `fname` is similar to the `fname` argument of functions `dhc-make` or `dhc-make-with-libs`. Argument `libs` is similar to the `libs` argument of function `dhc-make-with-libs`.

7.3.2 Customizing the Behavior of dhc-make.

The following variables control the behavior of `dhc-make`.

dhc-make-force

Type: VAR(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Setting this variable to a non null value forces the recompilation of all files, even when they have not been modified.

The following example, for instance, forces the compilation of functions `foo` and `bar`.

```
(let ((dhc-make-force t))
  (dhc-make () (foo bar)) )
```

dhc-make-lushflags

Type: VAR(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

This variable contains a string containing the compilation flags for invoking the C or C++ compiler. The following example, for instance, insert additional options `"-w -DMMX"` for the C compiler:

```
(let ((dhc-make-lushflags (concat dhc-make-lushflags " -w -DMMX")))
  (dhc-make () (foo bar)) )
```

The initial value is derived from the Makefile variable `LUSHFLAGS`. It might reference other predefined variables prefixed with character `$`. These variables are expanded using `dhc-substitute-env`.

dhc-make-command

Type: VAR(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

This variable is a string containing the command for generating object files. Variables prefixed with `$` will be expanded using `dhc-substitute-env`. Four additional variables are defined:

- `$LUSHFLAGS` Compilation flags defined by `dhc-make-lushflags` or by the optional argument of `dhc-make-o`.
- `$INCS` The flag specifying the location of the Lush include files.
- `$SRC` The pathname of the source file.
- `$OBJ` The pathname of the object file.

dhc-make-overrides

Type: VAR(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

This variable contains an a-list of additional variable definitions for interpreting `dhc-make-lushflags` and `dhc-make-commands`. Definitions provided in `dhc-make-overrides` take precedence over the defaults provided by Lush.

The following example, for instance, invokes the C++ compiler instead of the C compiler:

```
(let ((dhc-make-overrides
      (cons (cons "CC" (dhc-substitute-env "$CXX")) dhc-make-overrides) ))
      (dhc-make () (foo bar)) )
```

7.3.3 Querying File Dependencies

The following functions can be used to determine whether `dhc-make` will invoke the compiler.

(dhc-make-get-dependencies [lushfile])

Returns the list of Lush files on which file `lushfile` is dependent. These are the files directly or indirectly loaded using `libload` . When argument `lushfile` is not provided , `file-being-loaded` is assumed.

(dhc-make-test fname)

Type: DM

This function tests whether a subsequent invocation of `dhc-make` with the same `fname` argument needs to generate a new version of the C file. When there is no need to generate a new C file , function `dhc-make-test` makes sure that the object file is up-to-date and returns the name of the object file. Otherwise it returns the empty list.

This function is useful to avoid processing long files when we are only interested in precompiled functions. See "`packages/lapack/lapack-s.lsh`" for an example.

7.3.4 Low-Level Compilation Functions.

Most users will rarely use the functions described in this section , but they come in handy in certain cases.

(dhc-make-o-filename src)

Returns a suitable filename for the object file

(dhc-generate-c filename '([func1 [funcn]]))

Translate lisp functions and classes `func1` ... `funcn` to C code and produces a source file suitable for a file named `filename` . Argument `filename` should be provided without the ".c" suffix and must be a legal C identifier as it is used for the initialization function in the C code.

(dhc-substitute-env str [htable])

Returns a copy of string **str** after substituting all environment variables. Substitution values are searched in **htable** when available, then passed to functions **getconf** and **getenv**.

(dhc-make-c fname fsymblist)

Compile functions or classes **fsymblist** into a new source file **fname**. Argument **fname** must be provided without the ".c" suffix.

(dhc-make-o src-file [obj-file [lushflags]])

Compile C source file **src-file** generated by the dh compiler into object file **obj-file**. Argument **lushflags** is an optional string containing compiler options.

(dhc-make-c-maybe sname fname fsymblist)

Compile functions or classes **fsymblist** producing the file **fname**. Argument **fname** must be provided without the suffix ".c". Compilation will only occur if the existing **fname** was created before the file **sname** or any file loaded from **sname** using **libload**.

(dhc-make-o-maybe src-file [obj-file [cflags]])

Same as **dhc-make-o** but only recompiles if source file is newer than object file

7.4 Design Philosophy

The design philosophy of the Lush compiler is somewhat unusual, and, admittedly, in stark violation of many commonly accepted principles of good programming language design. First, The Lush compiler is not designed to replace the interpreter, but to complement it. Lush applications are often a combination of compiled code for things where performance takes precedence over flexibility (e.g. "expensive" and heavily numerical functions), and interpreted code for things where flexibility takes precedence over performance (high-level logic, user interface, memory management).

Therefore, the Lush compiler is designed to generate very efficient code (practically as good as its C equivalent), but has limitations as to what it can and cannot compile. What that means for the programmer is that, although the compiled and interpreted Lush dialects use the same syntax, they are really two different languages. From now on, the compiled Lush dialect will be called CLush, while the interpreted Lush dialect will be called simply Lush.

This deliberate difference between the two dialects allows all kinds of really juicy hacks (with no performance penalty), that would be difficult to achieve if the two dialects were identical by design. Within a single program you can

simultaneously have in-line C code in the compiled part, run-time construction of code in the interpreted part (including run-time self modifying expressions like DMD splicing macros) with no hidden performance penalty.

The main differences between Lush and CLush are as follows:

- CLush is strongly typed (the type of all the variables must be declared, and is immutable), while Lush is not typed (e.g. a single variable can be assigned values of different types at different times in the same program).
- CLush uses lexical scoping, while Lush uses dynamic variable binding (if you don't know the difference, don't worry).
- The CLush compiler refuses to compile code that dynamically allocates objects that would require garbage collection. This allows the compiler to allocate all temporaries, local variables, and return values on the stack. While this provides clear performance advantages, it also puts strong limitations on what can and cannot be compiled (in particular most list manipulation functions are not usable in CLush)
- Dynamic allocation can be done in CLush, but the memory management must be done "by hand", and allocated objects must be deallocated explicitly, as with C. A `pool` class is provided in `libstd/dynamic.lsh` to facilitate dynamic allocation/deallocation of objects in compiled code. A good example of its use is the graph library in `lsh/libgraph`.
-

7.5 What is Compilable, and What isn't

[under construction]

There is a simple test to determine if a particular function is compilable or not: `(compilablep func)`. This may return `'yes`, `'maybe` or `()`. The value `'maybe` is returned when the argument is a macro whose compilability cannot be determined without expansion.

Basically the kinds of things that can be compiled by the CLush compiler are the kinds of things that you expect to be compilable in a "conventional" language like C. All the really "Lispish" stuff like `eval`, `lambda`, etc is not compilable (those things don't exist in C).

There are three types of things that cannot be compiled:

- Things that construct functions and code at run time (e.g. `lambda`, `mlambda`), or things that evaluate dynamically generated code (e.g. `eval`).
- Things that allocate dynamic structures that would require garbage collection
- Things that call non-compilable functions.

7.5.1 (compilablep function)

Type: DM

Tests if a function is compilable. Returns `t`, `maybe` or `()`. The value `'maybe` is returned when the argument is a macro whose compilability cannot be determined without expansion.

7.6 CLush-specific functions and constructs

7.6.1 Numerical Type Declarations

```
((-double-) var1 ... varn)
```

declare double precision floating point variables.

```
((-float-) var1 ... varn)
```

declare single precision floating point variables.

```
((-int-) var1 ... varn)
```

declare integer variables (32 bits).

```
((-short-) var1 ... varn)
```

declare short integer variables (16 bits).

```
((-byte-) var1 ... varn)
```

declare signed byte variables (8 bits).

```
((-ubyte-) var1 ... varn)
```

declare unsigned byte variables (8 bits).

```
((-gptr- [spec]) var1 ... varn)
```

This declares the symbols `var1 ... varn` as C-style pointers. In the absence of a `spec` the pointers will be equivalent to universal pointers (`void *`). An optional string `spec` can be specified to explicitly declare the C type of the pointer. Examples:

```
((-gptr- "int *") integer-pointer)
((-gptr- "FILE *") file-pointer)
(cpheader "<my_header_file.h>")
((-gptr- "MyObjectType *") myobject-pointer)
```

Argument `spec` may also be a list containing a Lisp class name. This will define the type of the variable as a pointer to the C version of that class. Example:

```
((-gptr- (myclass)) myclass-pointer)
```

The C type of `myclass-pointer` will be `"struct CClass_myclass *"` .

```
((-obj- (class)) var1...varN)
```

This syntax declares the variables `var1 ... varN` as being instances of the class `class` .

```
((-idx0- (type)) var1...varN)
```

This syntax declares the variables `var1 ... varN` as `idx0` (scalars) of type `type` , where `type` is one of the Lush numerical types (`-ubyte-` , `-double-` , `-float-` , etc).

```
((-idx1- (type)) var1...varN)
```

This syntax declares the variables `var1 ... varN` as `idx1` (vector) of type `type` , where `type` is one of the Lush numerical types (`-ubyte-` , `-double-` , `-float-` , etc).

```
((-idx2- (type)) var1...varN)
```

This syntax declares the variables `var1 ... varN` as `idx2` (matrix) of type `type` , where `type` is one of the Lush numerical types (`-ubyte-` , `-double-` , `-float-` , etc).

```
((-idx3- (type)) var1...varN)
```

This syntax declares the variables `var1 ... varN` as `idx3` (3-tensor) of type `type` , where `type` is one of the Lush numerical types (`-ubyte-` , `-double-` , `-float-` , etc).

```
((-idx4- (type)) var1...varN)
```

This syntax declares the variables `var1 ... varN` as `idx4` (4-tensor) of type `type` , where `type` is one of the Lush numerical types (`-ubyte-` , `-double-` , `-float-` , etc).

```
((-idx5- (type)) var1...varN)
```

This syntax declares the variables `var1 ... varN` as `idx5` (5-tensor) of type `type` , where `type` is one of the Lush numerical types (`-ubyte-` , `-double-` , `-float-` , etc).

```
((-idx6- (type)) var1...varN)
```

This syntax declares the variables `var1 ... varN` as `idx6` (6-tensor) of type `type`, where `type` is one of the Lush numerical types (`-ubyte-`, `-double-`, `-float-`, etc).

```
((-idx7- (type)) var1...varN)
```

This syntax declares the variables `var1 ... varN` as `idx7` (7-tensor) of type `type`, where `type` is one of the Lush numerical types (`-ubyte-`, `-double-`, `-float-`, etc).

7.6.2 Class Declaration and Compilation

Classes can be compiled with some restrictions. In short:

- The types of all the slots of a compiled class must be declared.
- It is preferable to include the constructor in the list of methods being compiled.
- The constructor of a compiled class **must** set initialize all the slots with an object of the type declared in the class definition.
- The parent class of a compiled class must be compiled prior to the class
- Methods defined in a subclass that overload a method of the parent class must have the **same prototype** as the method they overload (i.e. the same arguments with the same type, the same return type, and the same temporary variables).

Here is an example of compiled class declaration syntax (taken from `packages/gblearning/modules.lsh`):

```
(defclass mle-cost sn-module
  ((-obj- (logadd-layer)) logadder)
  ((-obj- (idx1-ddstate)) logadded-dist)
  ((-idx1- (-int-)) label2classindex))
```

Essentially, the type of each slot must be declared using the same syntax as parameters of a function.

Compiling a class with `dhc-make` and its variants is done by passing a list whose first element is the class name and the remaining elements are the methods to be compiled. Here is an example:

```
(dhc-make ()
  (sn-module)
  (mle-cost mle-cost fprop bprop bbprop))
```

It is not required to compile all the methods of a class. A compiled class can have a mixture of compiled and non-compiled methods.

7.6.3 Type Conversions

Functions are provided to convert objects from one type to another. Their behavior is somewhat different in interpreted and compiled mode.

Numerical Type Conversions

These functions convert one numerical or object type to another. In compiled mode, they behave like C casts. In interpreted mode, the argument is casted to the corresponding C type, and then converted back to `double` since `double` is the only numerical type manipulated by the interpreter.

(to-int arg) Type: DX

Casts `arg` into an integer.

(to-float arg) Type: DX

Casts `arg` into a single precision number.

(to-double arg) Type: DX

Casts `arg` into a double precision number.

(to-bool arg) Type: DX

Casts `arg` into a boolean.

(to-flt arg) Type: DX

Casts `arg` into a single precision number (deprecated name for `to-float`).

(to-real arg) Type: DX

Casts `arg` into a double precision number (deprecated name for `to-double`).

Non-Numerical Type Conversions

(to-gptr arg) Type: DX

Casts `arg` into a GPTR.

In interpreted mode `to-gptr` simply returns a `gptr` containing the address the C object associated with the object passed as argument (which can be an index, storage, string, a compiled function, or any object).

This function may not be used to convert an integer (or other numerical types) to a `gptr`, but merely to obtain the address of an object.

In compiled mode, it is useful to know that a `gptr` does not carry enough information to allow the compiler to determine when the corresponding object must be released. The following function is a sure loser:

```
(de never-do-this()
  (gptr (new myclass)) )
```

The compiler cannot understand that the object created by `new` is returned as a `Gptr` and may be used outside this function. The object is freed before returning, making the `gptr` useless. To prevent this, the object should be allocated in a pool using `in-pool` from `lsh/libstd/dynamic.lsh`.

(to-obj [class] arg) Type: DX

Casts `arg` into an object of class `class`. Argument `arg` must be a `gptr` or an object. Argument `class` is mandatory in compiled mode.

- When argument `arg` is an object, this function checks that the object class is a subclass of `class` and returns the unmodified object.
- When this function is called in interpreted mode with a `gptr` argument, the lisp-to-c interface retrieves and returns the lisp object associated with the C object pointed to by the GPTR. It also checks that the object class is indeed a subclass of `class`.
- When this function is called in compiled mode with a GPTR argument, the compiler converts the GPTR into a pointer to an object of the specified class. Class membership is checked at run time.

7.6.4 inline C code

CLush allows to freely intermix Lisp and C code within a single function. This can be done in several ways, the simplest of which is the "hash-brace" construct.

Hash-Brace Construct

The macro character pair `# c-code #` allows easy embedding of C code inside a Lisp function. The embedded C code is delimited by `# and #`. Functions containing hash-brace segments cannot be executed by the interpreter and must be compiled.

Any C code can be written within a hash-brace construct. Lisp expressions and variables can be evaluated and referred to in the Lisp code by preceding them with a dollar sign.

- `$legal-c-identifier` refers to the lisp variable of the same name
- `$legal-lisp-variable` refers to the lisp variable (the braces are required if the lisp variable name is not a legal C name, e.g. if it contains dashes).
- `($lisp expression)` refers to the result returned by the lisp expression.

Examples:

```
[under construction]
[insert juicy hash-brace examples here]
```

The hash-brace macro calls the lower-level compiler macro `cinline`.

Example:

```
(de add-to-integer( intg arg )
  ((-int-) intg arg)
  #{ $intg += $arg; #}
  intg )
```

expands to:

```
(de add-to-integer (intg arg)
  ((-int-) intg arg)
  (cinline " \\%s += \\%s; " intg arg)
  intg )

See: (cinline format [ arg1 [ arg2 ... [ argn ]]])
```

C Directives and Macros

(cphheader string) Type: DF

This directive can be used to insert a line in the C file generated by the CLush compiler. This is primarily used to include header files or to define macros. Examples:

```
(cphheader "#include <stdio.h>")
```

The lines are inserted at the beginning of the C file , before all the Lush header files.

(cheader string) Type: DF

Same as `cphheader` , but inserts the lines after the Lush header files instead of before.

(find-c-include name) Type: DE(sysenv.lsh)

Returns the pathname of a C include file named `name` . Include files are searched in the directories specified by variable `c-include-path` which is initialized by "lushenv.lsh" .

```
? (find-c-include "stdio.h")
= "/usr/include/stdio.h"
```

Here is an example of how to use `find-c-include` to precompute an include file:

```
(libload "libc/constants")
(defconstant qq (sprintf "#include \\\"\\%s\\\"" (find-c-include "stdio.h")))
= "qqq"
? (de asd () (cphheader @qq) #{ printf("asd\\n") #} ())
= asd
? (dhc-make "junk" asd)
```

(ccall string) Type: DE

call a C function. This is somewhat obsolete. Most users will prefer the hash-brace construct.

(cinline format var1...varn) Type: DF

Insert C code in the C file generated by the Lush compiler. The `format` is an printf-like format where instances of "%s" are substituted by the C name of the lisp variables `var1 ... varn` . Example:

```
;; set element 2 of v to x.
(de choucroute (v x)
  ((-idx1- (-int-)) v)
  ((-double-) x)
  (cinline "(IDX_PTR(\\%s,int))[2] = (int)(\\%s);" v x) ())
```

The C function generated by the Lush compiler looks like this:

```
extern_c char
C_choucroute (struct idx *L1_v, real L1_x)
{
  TRACE_PUSH ("C_choucroute");
  {
    (IDX_PTR (L1_v, int))[2] = (int) (L1_x);
    TRACE_POP ("C_choucroute");
    return 0;
  }
}
```

`cinline` is rarely used, as most users will prefer to use the hash-brace construct.

7.6.5 Compiler Directives

(ifcompiled compiled-expr interpreted-expr)

Type: DM

This allows to execute different codes in interpreted mode and in compiled mode.

(ifdef symb expr1 expr2)

Type: DM

[under construction]

#@expr

Type: DMC

Expressions prefixed by `#@` are not evaluated unless `dhc-debug-flag` is set to `t`. This can be used to insert debugging code.

dhc-debug-flag

Type: VAR

Turns on execution of code prefixed with `—#@—`

7.6.6 Dynamic Allocation in Compiled Code

[under construction] This will eventually become a tutorial on how to use the `pool` class defined in `libstd/dynamic.lsh`. The reference manual for `dynamic.lsh` is available in the standard library section of this manual.

See: Dynamic Allocation with Pools.

7.7 Interfacing Existing C/C++/FORTRAN Libraries to Lush

[under construction]

7.7.1 Symbolic constants.

(lsh/libc/constants.lsh) Last Modified: 2002-11-18 03:28:41

header files of C libraries are often filled with large amounts of `#define` directives that define numerical macros, flags and other constants. Interfacing these libraries to Lush requires to turn those constants into Lisp symbols. Lush contains a mechanism for defining such constants without filling up the global variable name space and risking name clashes. Constants can be defined with the `(defconstant "constant-name" value)` directive. Constants can later be accessed using the `@constant-name` macro or the `@@constant-name` macro. The `@` macro expands into a macro that evaluates to the value, while `@@` directly expands to the value at read time.

See: `@ identifier`

See: `(defconstant identifier val)`

(defconstant identifier val)

(lsh/libc/constants.lsh) Last Modified: 2002-11-18 03:28:41

Defines a symbolic constant named `identifier` and equal to `val`. Argument `identifier` must be a string or a symbol. Argument `val` must be a number or a string. These constant values are not accessible as variables but are accessible using the `—@—` macro-character.

@identifier @@identifier

(lsh/libc/constants.lsh) Last Modified: 2002-11-18 03:28:41

These macro-character expand to a expressions that retrieve the value of the symbolic constant named `identifier`. The `@` form causes the retrieval to occur at macro-expansion time. while the `@@` form causes the retrieval to occur at read time. These macros are particularly useful for using symbolic constants in compiled code.

Argument `identifier` must be a legal C identifier name.

Compatibility notes:

- If `identifier` has not been defined using `defconstant`, the global value of the symbol named `identifier` is used.
- If the macro-character `|@|` cannot read a legal C identifier name, it reads a lisp expression and arranges for that expression to be evaluated during the macro expansion time.
- The expansion of the `@@` form depends on the execution of the previous expressions, which may cause discrepancies when loading tokenized files. However, the `@@` is more efficient in interpreted mode since the expansion occurs in advance of execution.

7.7.2 Interfacing C code

[under construction]

7.7.3 Interfacing C++ code

[under construction]

7.7.4 Interfacing FORTRAN code

[under construction]

7.8 More on the Lisp-C Interface

[under construction]

7.8.1 DH: the Compiled Function Class

`(dhinfo-t dhfunc)`

Type: DX

[under construction]

`(dhinfo-c dhfunc)`

Type: DX

[under construction]

7.8.2 Classinfo

`(classinfo-t class)`

Type: DX

[under construction]

Similar to the `SN3` function but takes a class as argument instead of now obsolete `cclass` objects.

(classinfo-c class)

Type: DX

[under construction]

Similar to the SN3 function but takes a class as argument instead of now obsolete cclass objects.

7.8.3 Controlling the Lisp/C interface.

(lisp-c-map [arg])

Type: DX

This function displays the internal data structure of the interface between Lisp and Compiled code. The LISP-C interface maintains a sorted avl tree of all objects currently allocated. This table relates the address of the compiled structure, the type of the structure, the way the object was created and possibly the Lisp representation of this object.

- Calling this function without arguments returns the number of entries in the table.
- Calling this function with `arg` equal to 0 or `()` prints all the table. You better check the size of the table before doing this.
- Calling this function with `arg` equal to 1 prints all the table entries that needs to be synchronized before calling a compiled function.
- Calling this function with `arg` equal to 2 prints all the table entries that represent temporary objects.
- Calling this function with `arg` equal to an object, a storage, an index, or a string prints the table entry (if any) associated with this object.

You can redirect the output of this command using function `writing` .

Example:

```
? (lisp-c-map)
= 12
? (lisp-c-map ())
= L 0x1e29b0 str "A string for class c1"
< L 0x2f10a0 idx ::INDEX1:<3>
< L 0x2f12e0 obj:c2 ::c2:2f1330
= L 0x2f1370 idx ::INDEX1:<3>
= L 0x2f1c98 srg ::FSTORAGE:static@2fe820
= L 0x2f1e18 str "A second string for class c1"
= L 0x2ff050 str "A third string for class c1"
= L 0x2ff068 obj:c1 ::c1:2f1b08
= L 0x2ff080 idx ::INDEX1:<3>
```

```

=          L 0x2ff0a8 srg          ::FSTORAGE:static@2fe82c
=          L 0x31cb20 obj:c1       ::c1:1e2980
=          L 0x31cb38 srg          ::FSTORAGE:ram@1e2968:<3>
= 12

```

The first character describes the position of the entry in the balanced tree. The second character tells whether the object was created by LISP (L) or by compiled code (C) (e.g. Pool). Each entry displays then the address of the compiled object and the lisp representation of the object. An object created by C code may have no lisp representation until it is referenced by a represented object or until the proper GPTR is casted by function <obj>.

(lisp-c-dont-track-cside)

Type: DX

Restrict the synchronization of the lisp and C data structures in a way similar to previous versions of the system (SN3.x). C objects managed by compiled code will appear as gptrs in Lisp code instead of being translated into equivalent Lisp objects.

(lisp-c-no-warnings ..exprs..)

Type: DY

Evaluates expressions ..exprs.. like function progn without displaying the usual diagnostic about the synchronization of the lisp and C data structures.

7.9 Compiler Internals

7.9.1 (dhc-generate-c filename '([func1 [funcn]]))

(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Translate lisp functions and classes `func1 ... funcn` to C code and produces a source file suitable for a file named `filename`. Argument `filename` should be provided without the ".c" suffix and must be a legal C identifier as it is used for the initialization function in the C code.

7.9.2 (compilablep function)

Type: DM(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Tests if a function is compilable. Returns `t`, `maybe` or `()`. The value `'maybe` is returned when the argument is a macro whose compilability cannot be determined without expansion.

7.9.3 (dhc-substitute-env str [htable])

(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Returns a copy of string `str` after substituting all environment variables. Substitution values are searched in `htable` when available, then passed to functions `getconf` and `getenv`.

7.9.4 `varlushdir`

Type: VAR(lushenv.lsh)

This variable is defined by "`lushenv.sn`" when the owner of the lush directory is different from the owner of the lush process. Compiled versions of the system lush files will be created in a shadow directory tree located under the `varlushdir` directory.

7.9.5 (`dhc-make-cdir cdir [create]`)

(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

This function returns `cdir` when `varlushdir` is the empty list. When `varlushdir` is the name of an existing directory, this function checks if `cdir` is a subdirectory of the Lush directory `lushdir`. If this is the case, it returns the name of a subdirectory under `varlushdir` with the same name. When argument `create` is `t`, function `dhc-make-cdir` creates all the required directories under `varlushdir`.

7.9.6 (`dhc-make-c-filename src`)

(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Returns a suitable filename for the C file generated from the LSH file `src`

7.9.7 (`dhc-make-o-filename src`)

(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Returns a suitable filename for the object file generated from the C or LSH file `src`

7.9.8 (`dhc-make-get-dependencies [sname]`)

(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Returns a list of files that are loaded directly or indirectly when loading file `sname`. This function takes advantage of the information collected by function `libload`. Argument `sname` defaults to the currently loaded file.

7.9.9 `dhc-make-lushflags`

Type: VAR(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Compilation flags used by `dhc-make-o`. This is similar to the variable `LUSHFLAGS` in the makefiles. Variables prefixed with `$` will be expanded using `dhc-substitute-env`.

7.9.10 `dhc-make-command`

Type: VAR(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Compilation command used by `dhc-make-o`. Variables prefixed with `$` will be expanded using `dhc-substitute-env`. Four additional variables are defined:

- `$LUSHFLAGS` Compilation flags defined by `dhc-make-lushflags` or by the optional argument of `dhc-make-o`.
- `$INCS` The flag specifying the location of the Lush include files.
- `$SRC` The pathname of the source file.
- `$OBJ` The pathname of the object file.

7.9.11 `dhc-make-overrides`

Type: VAR(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

This variable contains an a-list of variable definitions for overriding those usually derived from `getconf` and `getenv`.

7.9.12 `dhc-make-force`

Type: VAR(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Setting this variable to a non null value forces the recompilation of all files, even when they have not been modified.

7.9.13 `dhc-make-essential-libs`

Type: VAR(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

This variable contains a list of archive libraries that should be reloaded when undefined symbols remain. This is useful on some BSD systems.

7.9.14 `(dhc-make-o src-file [obj-file [lushflags]])`

(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Compile C source file `src-file` generated by the dh compiler into object file `obj-file`. Argument `lushflags` is an optional string containing compiler options.

7.9.15 `(dhc-make-o-maybe src-file [obj-file [cflags]])`

(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Same as `dhc-make-o` but only recompiles if source file is newer than object file

7.9.16 (dhc-make-c fname fsymblist)

(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Compile functions or classes `fsymblist` into a new source file `fname` . Argument `fname` must be provided without the ".c" suffix.

7.9.17 (dhc-make-c-maybe sname fname fsymblist)

(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Compile functions or classes `fsymblist` producing the file `fname` . Argument `fname` must be provided without the suffix ".c". Compilation will only occur if the existing `fname` was created before the file `sname` or any file loaded from `sname` using `libload` .

7.9.18 (dhc-make-test fname)

Type: DM(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Test if both the C file named `fname` and the associated object file are up-to-date relative to the currently loaded file and its dependencies. Returns either () or the name of the up-to-date object file. When argument `fname` is the empty list , a suitable filename is constructed from the currently loaded file using the same rules as `dhc-make` .

7.9.19 (dhc-make-all fname fsymblist liblist)

(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Document this.

7.9.20 (dhc-make fname f1 [f2 ...[fn]])

Type: DM(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Document this.

7.9.21 (dhc-make-with-libs fname library-list f1 [f2 ...[fn]])

Type: DM(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Document this.

7.9.22 (dhc-make-with-c++ fname library-list f1 [f2 ...[fn]])

Type: DM(lsh/compiler/dh-compile.lsh) Last Modified: 2002-12-23 09:57:09

Same as `dhc-make-with-libs` but compiles with the c++ compiler instead of the c compiler.

Note that the dynamic loader is not currently able to execute static initializers possibly present in the generated code. The lush compiler obviously does not generate such construct. On the other hand everything can happen when using inlined C++ code.

7.9.23 (dhc-insert-op 1 op)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Concatenate all strings in `1` separated with `op` .

```
? (insert-op '( "3" "4" "5" "*" ) "*" )
= "3*4*5"
```

7.9.24 (dhc-nolast 1)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Remove last element of list.

7.9.25 (dhc-remove-dup 1)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Remove duplicates in a list.

7.9.26 (dhc-remove-eqdup 1)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Remove pointer duplicates from a list.

7.9.27 (dhc-remove-nth n 1)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Remove `n` th element from list `1`**7.9.28 (dhc-split-list 1 test)**

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Returns two list containing the element of list `1` . The first list contains the element of `1` that pass the test `test` , the second list the remaining elements.

```
? (dhc-split-list '(1 4 2 6 4 3 2 3 2 1 3) (lambda (x) (= 3 x)))
= ((3 3 3) (1 4 2 6 4 2 2 1))
```

7.9.29 (dhc-postincr symb)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Increments variable `symb` by one and returns the value of the variable `*before*` incrementing.

7.9.30 `dhc-debug-flag`

Type: VAR(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Turns on execution of code marked with `—#@—`

7.9.31 `—#@—`

Type: DMC(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Expressions prefixed by `#@` are not evaluated unless `dhc-debug-flag` is set to `t` .

7.9.32 (`dhc-add-c-declarations str`)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

To be used in a `dhm-c` macro. Adds string `str` into the declaration part of the current block of C code.

7.9.33 (`dhc-add-c-statements str`)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

To be used in a `dhm-c` macro. Adds string `str` to the statement part of the current block of C code.

7.9.34 (`dhc-add-c-epilog str`)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

To be used in a `dhm-c` macro. Adds string `str` to the epilog part of the current block of C code.

7.9.35 (`dhc-add-c-externs str`)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

To be used in a `dhm-c` macro. Adds string `str` to the C code segment declaring external symbols.

7.9.36 (`dhc-add-c-metaexterns str`)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

To be used in a `dhm-c` macro. Adds string `str` to the C code segment declaring external symbols used in meta information.

7.9.37 (`dhc-add-c-header str`)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

To be used in a `dhm-c` macro. Adds string `str` to the C code segment declaring files included after the standard LUSH include files.

7.9.38 (dhc-add-c-pheader str)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

To be used in a `dhm-c` macro. Adds string `str` to the C code segment declaring files included `*before*` the standard LUSH include files.

7.9.39 (dhc-class-to-struct-decl type)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Returns a string containing the C structure declaration for objects of class `type` .

7.9.40 (dhc-class-to-vtable-decl type)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Returns a string containing the C structure declaration for the virtual table of class `type` .

7.9.41 (dhc-type-to-c-decl type)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Returns a short string with the C equivalent of type `type` .

7.9.42 (dhc-declare-var name type)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

generate C code to declare variable of name `name` of type `type` in the current block , initialize it , and free it on exit.

7.9.43 (dhc-declare-temp-var type [clue])

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Uses `tmpnames-seed` to generate a unique temp variable of type `type` . Returns the C string representation.

7.9.44 dhc-debug-stack

Type: VAR(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Variable containing the source causing the error

7.9.45 (dhc-error string [arg])

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Called when we get a syntax error in the source

7.9.46 (dhc-check-symbol source)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Take an argument `source` and check if it is a symbol. If it is return the corresponding object of class `dhc-symbol`.

7.9.47 (dhc-internal-error str)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Called when we detect inconsistent internal state

7.9.48 dhc-type

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

A `dhc-type` is a structure which contains almost all the information which is available at compile time of an expression. Information which is not meant to be used in the construction of the DHDOC should be kept away of `dhc-type` (that info can be used in the `t-node` structure instead). For instance `u-bump` and `u-access` have an impact on the DHDOC (whether a variable is writable , whether an argument is a temporary) but the flag `ignore-return-value` in a `progn` is not type information and has nothing to do in the `dhc-type` construct.

7.9.49 (==> dhc-type print)

Type: MSG(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

A type printer , for debugging purposes

7.9.50 (==> dhc-type hashcode)

Type: MSG(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

returns a hashcode for a type. only to be called after the type is fully unified.

7.9.51 (==> dhc-type access [v])

Type: MSG(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

set or return the access type of a type. Possible values: `'read` or `'write`

7.9.52 (new dhc-type class [a1] [a2])

Type: CLASS(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

- Declaration of simple types: `(new dhc-type type)`. Where `type` is a symbol from like, `'dht-flt` `'dht-bool`, `'dht-nil`, etc...
- Declaration of arrays: `(new dhc-type 'dht-out ndim [type])` Where `type` can be any type (default is `(new dhc-type 'dht-flt)`). `ndim` is an integer.

- Declaration of object: (new dhc-type 'dht-obj *type-list*) where *type-list* is a list of cons of the form (name . type)
- Declaration of functions: (new dhc-type 'dht-func *type-list*) where *type-list* is a list of valid types corresponding to the arguments

7.9.53 (dhc-desc-to-type desc)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Convert a type description to a type. Argument *desc* is a type description similar to those returned by *dhinfo-t* and *classinfo-t* .

7.9.54 (new dhc-symbol name lex [fmt])

Type: CLASS(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

return an symbol object which knows the lisp name *name* , the C name , the scope level (passed in *lex* . *lex* is zero if the symbol is an argument). *fmt* is by default "L%d.%s" where %d refers to the scope level (*lex*) and %s to the C name of the symbol (computed from the lisp name with *dhc-lisp-to-c-name*).

7.9.55 (dhc-search-symtable symbolname table)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Search a symbol table

7.9.56 (dhc-add-to-symtable table symbolobject)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Return a new symbol table with one more element

7.9.57 (dhc-add-symbol-table symb lex)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Add a symbol named *symb* at lexical level *lex* into the current symbol table.

7.9.58 (dhc-add-global-table symb)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Add a symbol to the global table

7.9.59 (new t-node t-node-list type [source] [symbol])

Type: CLASS(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Returns a t-node (type node). T-nodes contains all the type information of a LISP expression. If *source* is a symbol, then the slot *symbol* should contain an object of class *dhc-symbol* .

A t-node contains 5 slots:

```
<tn-list> <type> <source> <symbol> <ignore>
```

The slot `type` holds the return type of the expression which is parsed. `is` a t-node. If the expression is terminal, the slot `tn-list` should be the empty list and the slot `source` should contain the corresponding source.

Examples:

Parsing the expression `(+ a b)`

```
returns the t-node <<+: (dhm)>, <a: (flt)>, <b: (flt)>: (flt)>
```

In this t-node the slots are:

```
tn-list : a list of 3 t-nodes = (<+: (dhm)> <a: (flt)> <b: (flt)>)
```

```
type    : a return dhc-type node = (flt)
```

```
source  : the original source = (+ a b)
```

For the first element of the above t-node

```
i.e. (car :t_node_above:tn-list)
```

```
tn-list : a list of t-nodes = ()
```

```
type    : a return dhc-type node = (dhm)
```

```
source  : the original source = +
```

7.9.60 (dhc-make-t-node expr)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Return a treetype (a t-node) for `expr` but leave all the types unknown.

7.9.61 (dhc-copy-source-tree source)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Utility for `dhc-get-treetype`. Copies a source tree list (except leaves). Returns the copy.

7.9.62 (dhc-get-treetype source)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

This function returns the treetype of a expression.

```
? (dhc-get-treetype '(+ 3 4))
= <<+: (dhm)>, <3: (number)>, <4: (number)>: (number)>
? (dhc-get-treetype '(lambda (a) (-flt- a) a))
= <<lambda: (dhm)>
  , <<a: (flt)>: (unk)>
  , <<-flt-: (dhm)>, <a: (flt)>: (flt)>
  , <a: (flt)>
  : (func ((flt)) () (flt))>
```

7.9.63 (dhc-get-type source)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

This function returns the dhc-type of expression `source`

7.9.64 (dhm-t func (source) . body)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Installs a new DHM-T for the function named `func` .

The DHM-T function will be called when parsing a list expression starting with function `func` . This list expression is passed as argument to the DHM-T function. It should construct and return a t-node for the provided source code.

7.9.65 (dhm-t-declare model f1 ... fn)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Associates the dhm-t for function `model` with all functions `f1` to `fn` .

7.9.66 (dhm-c symb (source treetype replace) . body)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Installs a new DHM-T for the function named `func` .

The DHM-C function will be called when generating C code for a list expression starting with function `func` . This list expression is passed as argument `source` of the DHM-C function. Argument `treetype` is the t-node prepared by the DHM-T function. The function should call the `dhc-add-c-xxx` functions to compose the C code and return a string representing a C expression for the return value. When argument `replace` is non nil , it contains the name of a variable where the return value must be stored. The dhm-c function should then return this name.

7.9.67 (dhm-c-declare model f1 ... fn)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Associates the dhm-c for function `model` with all functions `f1` to `fn` .

7.9.68 (get-dhm-t symb)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Return the dhm-t for function `symb`

7.9.69 (get-dhm-c symb)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Return the dhm-c for function `symb`

7.9.70 (get-dhm-target source)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

This function can be called within a dhm to return the function object for which the dhm applies

7.9.71 (dhm-p func (source) . body)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Installs a new DHM-P for the function named `func`

The DHM-P function will be called when preprocessing a list expression starting with function `func` . The list expression is passed as `source` . It should construct a copy of `source` , using function `dhc-pp` on all parts of `source` that can be subject to macro-expansion

There are relatively few DHM-P functions because the SN3 compiler lacked this facilities. In the absence of a DHM-P function , all arguments of a function call are considered for macro-expansion. This is obviously wrong for functions like `let` or `quote` .

7.9.72 (get-dhm-p symb)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Return the dhm-p for function `symb`

7.9.73 (dhc-pp body)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

Macro preprocessor for the dhc compiler. return a version of lisp expression `body` with all macros (both `dm` and `dmd`) expanded.

7.9.74 (dhc-parse-replacement-source-t source newsource)

(lsh/compiler/dh-util.lsh) Last Modified: 2002-11-06 20:30:05

A dangerous hack to allow source replacement from within a dhm-t. Modifies the source tree being parsed and calls the dhm-t for the "newsource" The "source" argument should be the input argument of the calling dhm-t. See the dhm-t for `incr` for usage. No corresponding dhm-c is needed.

7.9.75 (process-numerical-args-t arglist rettype)

(lsh/compiler/dh-macro.lsh) Last Modified: 2002-10-14 03:30:40

Process the arguments of a numerical expression , and determine the type of the returned value. The arguments are passed in `arglist` , and the type of the return value in `rettype` .

7.9.76 u-nodes

[under construction]

Chapter 8

Standard Libraries

This section describes all the libraries that are distributed with Lush in the `lsh` directory , but are not loaded by default into the interpreter environment.

8.1 Shell Commands

(`lsh/libc/shell.lsh`) Last Modified: 2002-10-09 02:43:39

This set of functions is designed to facilitate the writing of "shell scripts" in Lush. It includes functions that are more or less equivalent to common shell commands such as `ls` , `cd` , `cp` , `mv` , `rm` etc. Those commands manipulate and return lists of strings. Other functions are provided to manipulate those lists of strings.

8.1.1 String List Utilities

(`lsh/libc/shell.lsh`) Last Modified: 2002-10-09 02:43:39

(**glob** `regex` `l`)

Type: DE(`lsh/libc/shell.lsh`) Last Modified: 2002-10-09 02:43:39

Return a list of the elements of `l` that match the regular expression `regex` . `l` must be a list of strings , and `regex` a string containing a valid Lush regular expression.

(**split-words** `s`)

Type: DE(`lsh/libc/shell.lsh`) Last Modified: 2002-10-09 02:43:39

split a string of words (space-separated strings) into a list of those words.

(**merge-words** `l` [`c`])

Type: DE(`lsh/libc/shell.lsh`) Last Modified: 2002-10-09 02:43:39

turns a list of strings into a single string composed of the concatenation of the elements of `l` interspersed with spaces. if a string is provided in `c`, it is interspersed instead of spaces.

(read-lines f)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

read file `f` and return a list of strings, each of which is a line of the file.

(write-lines l)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

write each string in list `l` as a line.

#Pexpr

(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

equivalent to (write-lines `expr`)

8.1.2 File Information

(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

(file-type f)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

return the type of file `f` as an atom equal to `'reg` for a regular file, `'dir` for a directory, `'chr` for character device, etc.

(file-size f)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

return the size of file `f` in bytes

(file-mode f)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

return the permissions of file `f` as an integer.

(file-uid f)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

return the userid of the owner of file `f`

(file-gid f)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

return the groupid of file `f`

(file-atime f)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

return the last access time of file **f** as a string of the form "YYYY MM DD hh mm ss". The string can be turned into a list with split-words.

(file-mtime f)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

return the last modification time of file **f** as a string of the form "YYYY MM DD hh mm ss". The string can be turned into a list with split-words.

(file-ctime f)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

return the creation time of file **f** as a string of the form "YYYY MM DD hh mm ss". The string can be turned into a list with split-words.

8.1.3 Variables

(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

shell-home

(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

user home directory

shell-dirstack

(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

directory stack manipulated by pushd and popd

shell-tmpdir

(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

directory used by the function sh to store temporary files.

8.1.4 Directories

(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

(pwd)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

return current directory (equivalent to (chdir)).

(cd [d])

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

change current directory to **d** , or to user home if **p** is not present.

(pushd d)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

Temporarily change current directory to **d** . Returning to the previous directory can be done with `popd`.

(popd)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

return to the current directory before the last `pushd`.

8.1.5 Shell Commands

(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

This set of functions is designed to facilitate the writing of "shell scripts" in Lush. It includes functions that are more or less equivalent to common shell commands such as `ls` , `cd` , `cp` , `mv` , `rm` etc. Those commands manipulate and return lists of strings. Other functions are provided to manipulate those lists of strings.

(rm f)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

remove file **f** (be careful).

(cp from to)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

copy file **from** into file **to**

(mv from to)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

move file from **from** to **to** .

(ls d1 d2 ... dn)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

return a list of all the files in **d1** (in case-insensitive lexicographic order) , followed by all the files in **d2** , etc... Invisible files ("`.`" , "`..`" , "`.xxx`" , etc) are not included.

(ls-a d1 d2 ... dn)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

return a list of all the files in **d1** (in case-insensitive lexicographic order) , followed by all the files in **d2** , etc... Invisible files (".", "..", ".xxx", etc) are included.

(sh cmd [1])

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

run shell command **cmd** and return the standard output as a list of strings (one string per line). The optional argument **1** is a list of strings that will be written to a temporary file (one line per string) and fed to the standard input of command **cmd**

(sh-find dir pattern)

Type: DE(lsh/libc/shell.lsh) Last Modified: 2002-10-09 02:43:39

calls the standard unix command find to find all the files whose name fit the pattern **pattern** (this is a unix-style shell regex , not a Lush regex). Example:

```
(sh-find "." "*.lsh")
```

8.2 Tensor/Matrix/Vector/Scalar Libraries

8.2.1 Generic IDX Macros and Functions

(lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

Resizing Macros and Functions

(lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

A set of macros and functions to resize vectors , matrices and tensors.

(midx-m1resize m n0) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

a macro redims **idx1 m** in-place with the new size **n0** . Unlike **redim**, it does not create a temporary variable. Unlike **idx-changedim**, this macro increases the size of the storage if necessary. very efficient. In interpreted mode, the function **idx1-resize** can be called.

See: (**idx1-resize m n0**)

(midx-m2resize m n0 n1) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

a macro that redims **idx2 m** in-place

(midx-m3resize m n0 n1 n2) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

a macro that redims **idx3 m** in-place

(midx-m4resize m n0 n1 n2 n3) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

a macro that redims idx4 m in-place

(idx-f1resize m n0) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx1 of float m in-place with the new size n0 . Unlike redim, it does not create a temporary variable. Unlike idx-changedim, this macro increases the size of the storage if necessary. very efficient. exists also in macro form.

See: (midx-m1resize m n0)

(idx-f2resize m n0 n1) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx2 of float m in-place

(idx-f3resize m n0 n1 n2) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx3 of float m in-place

(idx-f4resize m n0 n1 n2 n3)) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx4 of float m in-place

(idx-i1resize m n0) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx1 of int m in-place with the new size n0 . Unlike redim, it does not create a temporary variable. Unlike idx-changedim, this macro increases the size of the storage if necessary. very efficient. exists also in macro form.

See: (midx-m1resize m n0)

(idx-i2resize m n0 n1) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx2 of int m in-place

(idx-i3resize m n0 n1 n2) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx3 of int m in-place

(idx-i4resize m n0 n1 n2 n3)) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx4 of int m in-place

(idx-s1resize m n0) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx1 of short m in-place with the new size n0 . Unlike redim, it does not create a temporary variable. Unlike idx-changedim, this macro increases the size of the storage if necessary. very efficient. exists also in macro form.

See: (midx-m1resize m n0)

(idx-s2resize m n0 n1) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx2 of short m in-place

(idx-s3resize m n0 n1 n2) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx3 of short m in-place

(idx-s4resize m n0 n1 n2 n3) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx4 of short m in-place

(idx-b1resize m n0) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx1 of byte m in-place with the new size n0 . Unlike redim, it does not create a temporary variable. Unlike idx-changedim, this macro increases the size of the storage if necessary. very efficient. exists also in macro form.

See: (midx-m1resize m n0)

(idx-b2resize m n0 n1) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx2 of byte m in-place

(idx-b3resize m n0 n1 n2) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx3 of byte m in-place

(idx-b4resize m n0 n1 n2 n3) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx4 of byte m in-place

(idx-u1resize m n0) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx1 of ubyte m in-place with the new size n0 . Unlike redim, it does not create a temporary variable. Unlike idx-changedim, this macro increases the size of the storage if necessary. very efficient. exists also in macro form.

See: (midx-m1resize m n0)

(idx-u2resize m n0 n1) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx2 of ubyte m in-place

(idx-u3resize m n0 n1 n2) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx3 of ubyte m in-place

(idx-u4resize m n0 n1 n2 n3) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx4 of ubyte m in-place

(idx-g1resize m n0) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx1 of gptr m in-place with the new size n0 . Unlike redim, it does not create a temporary variable. Unlike idx-changedim, this macro increases the size of the storage if necessary. very efficient. exists also in macro form.

See: (midx-m1resize m n0)

(idx-g2resize m n0 n1) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx2 of gptr m in-place

(idx-g3resize m n0 n1 n2) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx3 of gptr m in-place

(idx-g4resize m n0 n1 n2 n3) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx4 of gp tr m in-place

(idx-d1resize m n0) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx1 of double m in-place with the new size n0 . Unlike redim, it does not create a temporary variable. Unlike idx-changedim, this macro increases the size of the storage if necessary. very efficient. exists also in macro form.

See: (midx-m1resize m n0)

(idx-d2resize m n0 n1) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx2 of double m in-place

(idx-d3resize m n0 n1 n2) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx3 of double m in-place

(idx-d4resize m n0 n1 n2 n3) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

redims idx4 of double m in-place

Macros for common operations

(lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

(midx-m2oversample input nlin ncol output) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

macro for 2D oversampling

Obsolete Iterators for Inline C code

(lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

These function are kept for backward compatibility , but a largely superseded by cidx-bloop described in the Core section.

(cinline-idx1loop m type c j p1 [p2...]) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

man, that's a tough one to explain. don't try this at home anyway. this thing is like a loop for cinline. It executes the cinline instructions in p1 ... for each element of matrix m (must be an idx1). c , and j must be strings containing valid C variable identifiers, and type is a string that must contain a valid C type (like "float"). during the loop, C variable c will point to the current element of m , and C variable j will contain the index of that element. So a macro that fills the elements of a vector with its index plus a constant n can be written:

```
(cinline-idx1loop mat "flt" "p" "i" (cinline "*p = i+\\%s;" n))
```

(cinline-idxloop2 m1 type1 m2 type2 c1 c2 j p1 [p2...]) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

same as `cinline-idx1loop` , but loops simultaneously on two vectors

(`cinline-idxloop3` *m1 type1 m2 type2 m3 type3 c1 c2 c3 j p1 [p2...]*)
(lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

same as `cinline-idx1loop` , but loops simultaneously on three vectors

(`cinline-idx2loop` *m type c i j [11...]*) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

same as `cinline-idx1loop` but loops over all elements of an `idx2`. `i` and `j` are strings containing the C name of the row and column indices.

(`cinline-idx2loop2` *m1 type1 m2 type2 c1 c2 i j [1...]*) (lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

same as `cinline-idx1loop2` but loops over all elements of an `idx2`. `i` and `j` are strings containing the C name of the row and column indices.

(`cinline-idx2loop3` *m1 type1 m2 type2 m3 type3 c1 c2 c3 i j [1...]*)
(lsh/libidx/idx-macros.lsh) Last Modified: 2002-11-26 12:24:35

same as `cinline-idx1loop3` but loops over all elements of an `idx2`. `i` and `j` are strings containing the C name of the row and column indices.

8.2.2 IDX of Single Precision Floating Point Numbers

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

Basic operations on vectors , matrices , and tensors of floats

(`idx-f2timesf2` *A B C*)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

2d x 2d matrix multiply: $C = A * B$

(`idx-flreverse` *vector*)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

Reverses in-place the order of the elements of a `vector` . not very efficient , but I didn't have time to re-write a faster one.

(`idx-f1fill` *m f*)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

multiply elements of `m` by float `f` , put result in `q`

(`idx-f2fill` *m f*)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

fill `idx2` `m` with value `f`

(`idx-f3fill` *m f*)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

fill `idx3` `m` with value `f`

(idx-f1dotc m f q)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
multiply elements of m by float f , put result in q

(idx-f2dotc m f q)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
multiply elements of m by float f , put result in q

(idx-f3dotc m f q)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
multiply elements of m by float f , put result in q

(idx-f1dotcacc m f q)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
multiply elements of m by float f , accumulate result in q

(idx-f2dotcacc m f q)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
multiply elements of m by float f , accumulate result in q

(idx-f3dotcacc m f q)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
multiply elements of m by float f , accumulate result in q

(idx-f1addc m f q)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
add elements of m to float f , put result in q

(idx-f2addc m f q)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
add elements of m by float f , put result in q

(idx-f3addc m f q)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
add elements of m by float f , put result in q

(idx-fllincomb x cx y cy z)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

performs a linear combination of vectors **x** and **y** using coefficients **cx** and **cy** . Put result in **z** .

(idx-fltanh x1 y1)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

Computes the regular tanh of **x1** element by element , result in **y1** .

(idx-flinv x1 y1)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

puts inverse of each element of **x1** into corresponding elements in **y1**

(idx-flsign x y)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

Elements of **y** are set to the sgn function of elements of **x** .

(idx-flclip x y)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

(idx-flsup m)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

returns largest element in **m**

(idx-flmax m [r])

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

if **r** is present , put largest element of idx1 **m** into idx0 **r** otherwise returns largest element.

(idx-f2max m [r])

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

if **r** is present , put largest element of idx1 **m** into idx0 **r** otherwise returns largest element.

(idx-flinf m)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

returns smallest element in **m**

(idx-f1min m [r])

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

if **r** is present , put smallest element of **idx1 m** into **idx0 r** otherwise returns smallest element.

(idx-f2min m [r])

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

if **r** is present , put smallest element of **idx1 m** into **idx0 r** otherwise returns smallest element.

(idx-f1indexmax m)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

returns index (int) of largest element in **m**

(idx-f1indexmin m)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

returns index (int) of smallest element in **m**

(index-of-max m)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

return index of largest element in **idx1 m**

(idx-f1fill-with-index m)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

fill the elements of **idx1 m** with their index.

(idx-f1avg vector)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

Returns the average of the numbers in the input **vector** .

(idx-f1prod vec)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

returns the product of the elements of the input vector **vec** ,

(idx-f1logsum m)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

computes a "log-add" of the elements in **m** , i.e. $-\log(1/n \sum_i (\exp(-m_i)))$ where **n** is the dimension of **m** . This is the wrong logsum advocated by Yann. Use **idx-f1logadd** instead.

(idx-f2logsum m)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

same as `idx-f1logsum` , but in 2 dimensions $-\log(1/n \sum_{ij} (\exp(-m_{ij})))$ where n is the number of elements in m This is the wrong logsum advocated by Yann. Use `idx-f2logadd` instead.

(idx-f1logadd m)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

computes a "log-add" of the elements in m , i.e. $-\log(\sum_i (\exp(-m_i)))$

(idx-f2logadd m)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

same as `idx-f1logadd` , but in 2 dimensions $-\log(\sum_{ij} (\exp(-m_{ij})))$

(idx-f1logdotf1 m p)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

computes a "log-add" of term-by-term products of elements in m and p , i.e. $-\log(\sum_i (p_i \sum_i (\exp(-m_i))))$ p should be a normalized probability vector for this to be meaningful

(idx-f2logdotf2 m)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

same as `idx-f1logdotf1` , but in 2 dimensions

(idx-f1blogdotf1 m p out)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

This is the "bprop" corresponding to `idx-f1logdotf1 m` is the input (preferably positive costs) , p is the prior vector , and on output `out` is the vector of partial derivatives of `(idx-f1logdotf1 m p)` with respect to each element of m .

(idx-f1subextf1 x y z)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

Computes the external subtraction $z[i, j] = x[i] - y[j]$

(idx-f1subf0 x1 y0 z1)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

Subtracts `(y0)` from elements of vector `x1` . Results in vector `z1` .

(idx-f1mulacc x1 y1 z1)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
 $z1[i] += x1[i] * y1[i]$. not efficient. Could be improved a lot.

(idx-f1entropy m c)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
 computes the entropy of a vector ($-\sum P_i \log P_i$) this does a normalization and adds a constant c to all the elements (so it does not crash if they are 0)

(idx-f1subsample-fast in nlin)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
 subsamples vector (considered as a signal) by integer factor $nlin$ the returned output is a vector whose size is $\text{int}(N/nlin)$ where N is the size of in . there is a `idx-f1subsample` defined in `libconvol.sn`, it is a little slower and takes 3 parameters.

(idx-f4dotf3 x y z)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
 $z[i] = \text{sum}_{jkl} x[i, j, k, l] * y[j, k, l]$

(idx-f3extf1 x y z)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
 External (or cross-) product of 3d x with 1d y to yield 4d z . $z[i, j, k, l] = x[i, j, k] * y[l]$

(idx-f1extf3 x y z)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
 External (or cross-) product of 1d x with 3d y to yield 4d z . $z[i, j, k, l] = x[i] * y[j, k, l]$

(idx-f3dotf3 a b c)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
 c is an `idx0` which will contain the sum of all the products of elements from a and b . (this redundant with `idx-dot`).

(idx-f3dotf2 x y z)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34
 $z[i] = \text{sum}_{jk} x[i, j, k] * y[j, k]$

(idx-f4dotf1 x y z)

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

 $z[j, k, l] = \text{sum}_i x[i, j, k, l] * y[i]$ **(idx-f4dotflacc x y z)**

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

 $z[j, k, l] += \text{sum}_i x[i, j, k, l] * y[i]$ **(idx-flextf2 x y z)**

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

External (or cross-) product of 1d x with 2d y to yield 3d z . $z[i, j, k] = x[i] * y[j, k]$ **(idx-f3dotf1 x y z)**

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

 $z[j, k] = \text{sum}_i x[i, j, k] * y[i]$ **(idx-f3dotflacc x y z)**

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

 $z[j, k] += \text{sum}_i x[i, j, k] * y[i]$ **(make-f2-place-target n)**

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

return an $n \times n$ matrix with 1s on the diagonal and -1s everywhere else.**(idx-flrange n1 n2 delta)**

(lsh/libidx/idx-float.lsh) Last Modified: 2002-11-26 12:24:34

Returns a vector of all the numbers between $n1$ and $n2$, stepping by $delta$. Not compilable under current version.**8.2.3 IDX of Double Precision Floating Point Numbers**

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

Basic operations on vectors , matrices , and tensors of doubles

(idx-d2timesd2 A B C)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

2d x 2d matrix multiply: $C = A * B$

(idx-d1reverse vector)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

Reverses in-place the order of the elements of a **vector** . not very efficient , but I didn't have time to re-write a faster one for SN3.1

(idx-d1fill m f)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

multiply elements of **m** by float **f** , put result in **q**

(idx-d2fill m f)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

fill idx2 **m** with value **f**

(idx-d3fill m f)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

fill idx3 **m** with value **f**

(idx-d1dotc m f q)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

multiply elements of **m** by float **f** , put result in **q**

(idx-d2dotc m f q)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

multiply elements of **m** by float **f** , put result in **q**

(idx-d3dotc m f q)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

multiply elements of **m** by float **f** , put result in **q**

(idx-d1dotcacc m f q)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

multiply elements of **m** by float **f** , accumulate result in **q**

(idx-d2dotcacc m f q)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

multiply elements of **m** by float **f** , accumulate result in **q**

(idx-d3dotcacc m f q)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02
multiply elements of m by float f , accumulate result in q

(idx-d1addc m f q)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02
add elements of m to float f , put result in q

(idx-d2addc m f q)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02
add elements of m by float f , put result in q

(idx-d3addc m f q)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02
add elements of m by float f , put result in q

(idx-d1lincomb x cx y cy z)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02
performs a linear combination of vectors x and y using coefficients cx and cy . Put result in z .

(idx-d1tanh x1 y1)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02
Computes the regular tanh of $x1$ element by element , result in $y1$.

(idx-d1inv x1 y1)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02
puts inverse of each element of $x1$ into corresponding elements in $y1$

(idx-d1sign x y)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02
Elements of y are set to the sgn function of elements of x .

(idx-d1clip x y)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

(idx-d1sup m)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02
returns largest element in m

(idx-d1max m [r])

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

if **r** is present , put largest element of idx1 **m** into idx0 **r** otherwise returns largest element.

(idx-d2max m [r])

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

if **r** is present , put largest element of idx1 **m** into idx0 **r** otherwise returns largest element.

(idx-d1inf m)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

returns smallest element in **m**

(idx-d1min m [r])

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

if **r** is present , put smallest element of idx1 **m** into idx0 **r** otherwise returns smallest element.

(idx-d2min m [r])

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

if **r** is present , put smallest element of idx1 **m** into idx0 **r** otherwise returns smallest element.

(idx-d1indexmax m)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

returns index (int) of largest element in **m**

(idx-d1indexmin m)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

returns index (int) of smallest element in **m**

(index-of-max m)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

return index of largest element in idx1 **m**

(idx-d1fill-with-index m)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

fill the elements of idx1 **m** with their index.

(idx-d1avg vector)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

Returns the average of the numbers in the input `vector` .**(idx-d1prod vec)**

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

returns the product of the elements of the input vector `vec` ,**(idx-d1logsum m)**

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

computes a "log-add" of the elements in `m` , i.e. $-\log(1/n \sum_i (\exp(-m_i)))$ where `n` is the dimension of `m` . This is the wrong logsum advocated by Yann. Use `idx-d1logadd` instead.**(idx-d2logsum m)**

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

same as `idx-d1logsum` , but in 2 dimensions $-\log(1/n \sum_{ij} (\exp(-m_{ij})))$ where `n` is the number of elements in `m` This is the wrong logsum advocated by Yann. Use `idx-d2logadd` instead.**(idx-d1logadd m)**

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

computes a "log-add" of the elements in `m` , i.e. $-\log(\sum_i (\exp(-m_i)))$ **(idx-d2logadd m)**

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

same as `idx-d1logadd` , but in 2 dimensions $-\log(\sum_{ij} (\exp(-m_{ij})))$ **(idx-d1logdotd1 m p)**

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

computes a "log-add" of term-by-term products of elements in `m` and `p` , i.e. $-\log(\sum_i (p_i \sum_i (\exp(-m_i))))$ `p` should be a normalized probability vector for this to be meaningful**(idx-d2logdotd2 m)**

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

same as `idx-d1logdotd1` , but in 2 dimensions

(idx-d1blogdotd1 m p out)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

This is the "bprop" corresponding to `idx-d1logdotd1 m` is the input (preferably positive costs), `p` is the prior vector, and on output `out` is the vector of partial derivatives of `(idx-d1logdotd1 m p)` with respect to each element of `m`.

(idx-d1subextd1 x y z)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

Computes the external subtraction $z[i, j] = x[i] - y[j]$

(idx-d1subd0 x1 y0 z1)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

Subtracts (`y0`) from elements of vector `x1`. Results in vector `z1`.

(idx-d1mulacc x1 y1 z1)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

$z1[i] += x1[i] * y1[i]$. not efficient. Could be improved a lot.

(idx-d1entropy m c)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

computes the entropy of a vector ($-\sum P_i \log P_i$) this does a normalization and adds a constant `c` to all the elements (so it does not crash if they are 0)

(idx-d1subsample-fast in nlin)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

subsamples vector (considered as a signal) by integer factor `nlin` the returned output is a vector whose size is $\text{int}(N / \text{nlin})$ where `N` is the size of `in`. there is a `idx-d1subsample` defined in `libconvol.sn`, it is a little slower and takes 3 parameters.

(idx-f4dotd3 x y z)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

$z[i] = \sum_{jkl} x[i, j, k, l] * y[j, k, l]$

(idx-d3extd1 x y z)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

External (or cross-) product of 3d `x` with 1d `y` to yield 4d `z`. $z[i, j, k, l] = x[i, j, k] * y[l]$

(idx-d1extd3 x y z)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

External (or cross-) product of 1d x with 3d y to yield 4d z . $z[i, j, k, l]$
 $= x[i]*y[j, k, l]$

(idx-d3dotd3 a b c)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

c is an idx0 which will contain the sum of all the products of elements from
 a and b .

(idx-d3dotd2 x y z)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

 $z[i] = \text{sum}_{jk} x[i, j, k]*y[j, k]$
(idx-d4dotd1 x y z)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

 $z[j, k, l] = \text{sum}_i x[i, j, k, l]*y[i]$
(idx-d4dotd1acc x y z)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

 $z[j, k, l] += \text{sum}_i x[i, j, k, l]*y[i]$
(idx-d1extd2 x y z)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

External (or cross-) product of 1d x with 2d y to yield 3d z . $z[i, j, k] =$
 $x[i]*y[j, k]$

(idx-d3dotd1 x y z)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

 $z[j, k] = \text{sum}_i x[i, j, k]*y[i]$
(idx-d3dotd1acc x y z)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

 $z[j, k] += \text{sum}_i x[i, j, k]*y[i]$
(make-d2-place-target n)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

return an $n \times n$ matrix with 1s on the diagonal and -1s everywhere else.

(idx-d1range n1 n2 delta)

(lsh/libidx/idx-double.lsh) Last Modified: 2002-11-06 18:37:02

Returns a vector of all the numbers between **n1** and **n2** , stepping by **delta** . Not compilable under current version.

8.2.4 IDX of Integers

(lsh/libidx/idx-int.lsh) Last Modified: 2002-08-19 03:01:28

Basic operations on vectors , matrices , and tensors of integers.

(idx-i2addc m1 c m2)

(lsh/libidx/idx-int.lsh) Last Modified: 2002-08-19 03:01:28

(idx-i1concat int-vec1 int-vec2 ...)

(lsh/libidx/idx-int.lsh) Last Modified: 2002-08-19 03:01:28

Concatenate the **int-vectors** in argument list and return new big vector.

(random-permute perm-vector)

(lsh/libidx/idx-int.lsh) Last Modified: 2002-08-19 03:01:28

randomly permute the elements of the given **permutation-vector** (an integer vector)

(multinomial-sample probabilities)

(lsh/libidx/idx-int.lsh) Last Modified: 2002-08-19 03:01:28

Given a vector of **probabilities** of length N summing to one for events 0 to N-1 , return an integer i associated to these events with probability **probabilities** [i]. The algorithm samples a uniform number between 0 and 1 and loops through the **probabilities** vector , so its running time is O(N).

N.B.: if the sum S of the elements is 1, then the value N-1 will be sampled more frequently with extra probability 1 - S. if the sum S of the elements is 1 , then only the first M elements will be sampled , with M being such that it is the largest integer such that the sum of the **probabilities** from 0 to M-1 is above 1.

(fast-multinomial-sample cumulative-probabilities)

(lsh/libidx/idx-int.lsh) Last Modified: 2002-08-19 03:01:28

Given a vector of **cumulative-probabilities** of length N, monotonically increasing between 0 and 1, return an integer between 0 and N-1 sampled from the corresponding distribution. The algorithm samples a uniform number between 0 and 1 and does a binary search in the cumulative probabilities, so its running time O(log(N)) whereas the running time of (multinomial-sample **probabilities**) is O(N). For this call to be really fast, the computation of

`cumulative-probabilities` from `probabilities` must be done ahead of time (because it takes $O(N)$ time). It can be done with `(distribution2cumulative probabilities cumulative-probabilities)`.

See: `(multinomial-sample probabilities)`

See: `(distribution2cumulative probabilities cumulative-probabilities)`

`(distribution2cumulative probabilities cumulative-probabilities)`

`(lsh/libidx/idx-int.lsh)` Last Modified: 2002-08-19 03:01:28

Fill `cumulative-probabilities` by putting at the i -th position the sum the numbers in `probabilities` from the beginning to the $(i-1)$ -th position. Therefore, `cp[0]=0` and `cp[N-1]` may be 1. This is $P(\text{value} < i)$.

`(element-in-set int-element int-set)`

`(lsh/libidx/idx-int.lsh)` Last Modified: 2002-08-19 03:01:28

Return true iff `int-element` is in the `int-set` vector.

`(idx-ilmax m)`

`(lsh/libidx/idx-int.lsh)` Last Modified: 2002-08-19 03:01:28

returns largest element in `m`

`(idx-ilmin m)`

`(lsh/libidx/idx-int.lsh)` Last Modified: 2002-08-19 03:01:28

same as `idx-m1min` for vectors of integers

`(idx-ilfill-with-index m)`

`(lsh/libidx/idx-int.lsh)` Last Modified: 2002-08-19 03:01:28

fill the elements of `idx1 m` with their index.

`(idx-ilfill m c)`

`(lsh/libidx/idx-int.lsh)` Last Modified: 2002-08-19 03:01:28

`((-idx1- (-int-)) m):`

`((-int-) c):`

CREATED: P.Haffner, April 97

DESCRIPTION: Fills 1D int matrix with constant `<c>`

(idx-i2fill m c)

(lsh/libidx/idx-int.lsh) Last Modified: 2002-08-19 03:01:28

```
((-idx2- (-int-)) m):
((-int-) c):
CREATED: P.Haffner, April 97
DESCRIPTION: Fills 2D int matrix with constant <c>
```

8.2.5 IDX of unsigned bytes

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

Utilities for matrices and tensors of unsigned bytes [partially documented]

(idx-ub1add m c)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

```
((-idx1- (-ubyte-)) m):
((-int-) c):
CREATED: P.Haffner, April 97
DESCRIPTION: Adds 1D ubyte matrix with constant <c>
```

(idx-ub2add m c)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

```
((-idx2- (-ubyte-)) m):
((-int-) c):
CREATED: P.Haffner, August 97
DESCRIPTION: Adds 2D ubyte matrix with constant <c>
```

(idx-ub2mul m c)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

```
((-idx2- (-ubyte-)) m):
((-int-) c):
CREATED: P.Haffner, August 97
DESCRIPTION: Multiplies 2D ubyte matrix with constant <c>
```

(idx-ub2-i2mul i2 u2 c)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

((-idx2- (-ubyte-)) i2):
((-idx2- (-ubyte-)) u2):
((-int-) c):
CREATED: P.Haffner, Sept 97
DESCRIPTION:

(idx-ub3mat-add m1 m2)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

((-idx3- (-ubyte-)) m1 m2):
CREATED: Oct 97
DESCRIPTION:

(idx-ub1fill m c)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

((-idx1- (-ubyte-)) m):
((-int-) c):
CREATED: P.Haffner, April 97
DESCRIPTION: Fills 1D ubyte matrix with constant <c>

(idx-ub2fill m c)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

((-idx2- (-ubyte-)) m):
((-int-) c):
CREATED: P.Haffner, April 97
DESCRIPTION: Fills 2D ubyte matrix with constant <c>

(idx-ub-to-bits ub)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

(idx-ub-from-bits ub bits)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

(idx-ilindexmax m)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

(idx-m2opp m)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

((-idx2- (-flt-)) m):
 CREATED: P.Haffner, June 97
 TEMPORARIES: ?
 DESCRIPTION: invert sign of elements

(idx-u2sub ubim v)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

(idx-u2rev ubim)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

(idx-copy-inv u1 u2)

(lsh/libidx/idx-ubyte.lsh) Last Modified: 2002-08-19 03:01:28

((-idx2- (-ubyte-)) u1 u2):
 CREATED: P.Haffner, August 97
 DESCRIPTION: same as idx-copy, with byte inversion.

8.2.6 Arrays of vectors and matrices

(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36

arrays of idx1 of floats , and idx2 of floats. This class is convenient way of handling sets of vectors or matrices of disparate sizes.

(to-idx1flt gptr)

(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36

Casts a gptr into a (-idx1- (-flt-)) No checks are performed in compiled mode.

(to-idx2flt gptr)

(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36

Casts a gptr into a (-idx2- (-flt-)) No checks are performed in compiled mode.

idx1-array

(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36

array of idx1 of floats.

(new idx1-array n) Type: CLASS(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36

Create an array of idx1 of float. Each entry can have a different sizes. The methods of this object are:

```
(==> <idx1-array> set <i> <mat>)
<i> is an int, which must be less than <n>, <mat> is an idx1 of float.
This method set the <i>-th entry of the array to <mat>.
(==> <idx1-array> get <i>)
<i> is an int, which must be less than <n>.
This method returns the <i>-th entry of the array.
(==> <idx1-array> set-array <ar>)
Take an array of idx1 and do a set for each element (not compiled).
```

Example:

```
(setq ar {[3 4 5] [1] [1 2 3]})
(setq n (idx-dim ar 0))
(setq c-ar (new idx1-array n))
(for (i 0 (1- n)) ; equivalent to: (==> c-ar set-array ar)
  (==> c-ar set i (ar i)))
```

In a compiled or interpreted function

```
(de foo (c-ar ...)
  ((-obj- (idx1-array)) c-ar)
  (for (i 0 (1- n))
    ... (==> c-ar get i)... ; returns an idx1 of float
```

See: (new idx2-array n)

(==> idx1-array set i mat) Type: MSG(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36**(==> idx1-array get i)** Type: MSG(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36**(==> idx1-array set-array ar)** Type: MSG(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36**idx2-array**

(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36

array of idx2 of floats

(new idx2-array n) Type: CLASS(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36

Create an array of idx2 of float. Each entry can have a different sizes. The methods of this object are:

```
(==> <idx2-array> set <i> <mat>)
<i> is an int, which must be less than <n>, <mat> is an idx2 of float.
This method set the <i>-th entry of the array to <mat>.
(==> <idx2-array> get <i>)
<i> is an int, which must be less than <n>.
This method returns the <i>-th entry of the array.
(==> <idx2-array> set-array <ar>)
Take an array of idx2 and do a set for each element.
```

Example:

```
(setq ar {[[3 4 5][2 2 2]] [[1]] [[1] [2] [3]]})
(setq n (idx-dim ar 0))
(setq c-ar (new idx2-array n))
(for (i 0 (1- n)) ;; equivalent to: (==> c-ar set-array ar)
  (==> c-ar set i (ar i)))
```

In a compiled or interpreted function

```
(de foo (c-ar ...)
  ((-obj- (idx1-array)) c-ar)
  (for (i 0 (1- n))
    ... (==> c-ar get i)... ;; returns an idx1 of float
```

See: (new idx1-array n)

(==> idx2-array set i mat) Type: MSG(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36

(==> idx2-array get i) Type: MSG(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36

(==> idx2-array set-array ar) Type: MSG(lsh/libidx/idx-array.lsh) Last Modified: 2002-08-19 13:44:36

8.2.7 IDX reading and writing

(lsh/libidx/idx-io.lsh) Author(s): Pascal Vincent Last Modified: 2002-12-23 09:57:11

Compilable routines for reading/writing vector/matrix/tensors from/to an open file.

These routines are based on the C-like Compilable file I/O routines. They use their own matrix file format, different from the traditional Lush matrix formats. Nevertheless, the traditional functions `load-matrix` and `save-matrix` are able to handle this format.

IDX file format for compilable I/O

(lsh/libidx/idx-io.lsh) Author(s): Pascal Vincent Last Modified: 2002-12-23 09:57:11

the Lush interpreter has several function to read and write vectors/matrices/tensors in various portable ways, unfortunately, these functions cannot be called in compiled code at the moment. To alleviate the problem, a number of compilable

functions are provided to read/write IDX. unfortunately again, the formats are not compatible with the those of the core read/write functions. This discrepancy exists for largely historical reasons and will eventually be fixed.

The format used by the compilable functions is as follows:

- A long word (4 bytes) coding the matrix type (see detail below)
- A long word (4 bytes) for EACH dimension of the matrix (thus for an idx2 it will be 2 long-words, the first one for the height, the second one for the width of the matrix)
- A (possibly reversed) memory-dump of the elements of the matrix (short: 2 bytes, int: 4 bytes, f1t: 4 bytes real: 8 bytes)

The matrix types are encoded in the following way (hexadecimal representation):.PRE - idx1 of ubyte: 00 00 08 01 - idx2 of ubyte: 00 00 08 02 - idx3 of ubyte: 00 00 08 03 - idx1 of byte: 00 00 09 01 - idx2 of byte: 00 00 09 02 - idx3 of byte: 00 00 09 03 - idx1 of short: 00 00 0B 01 - idx2 of short: 00 00 0B 02 - idx3 of short: 00 00 0B 03 - idx1 of int: 00 00 0C 01 - idx2 of int: 00 00 0C 02 - idx3 of int: 00 00 0C 03 - idx1 of f1t: 00 00 0D 01 - idx2 of f1t: 00 00 0D 02 - idx3 of f1t: 00 00 0D 03 - idx1 of real: 00 00 0E 01 - idx2 of real: 00 00 0E 02 - idx3 of real: 00 00 0E 03

The long-words and the elements of the matrix MUST correspond to the byte order and coding of the Sparc stations (classical big-endian, and IEEE standard floating point representation) If the system is not a Sparc, then the functions that write or load these matrixes reverse the bytes if necessary to keep the file format consistant between systems.

(fread-idx1-ubyte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-gp1tr-) file) ; the file pointer returned by a fopen
((-idx1- (-ubyte-)) m) ; the matrix to be read from the file
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: reads matrix <m> from file <file-pointer>
             <m> is "resized" to fit the size of the matrix in the file
```

(fwrite-idx1-ubyte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-gp1tr-) file) ; the file pointer returned by a fopen
((-idx1- (-ubyte-)) m) ; the matrix to be written to the file
```

RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx2-ubyte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx2- (-ubyte-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx2-ubyte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx2- (-ubyte-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx3-ubyte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx3- (-ubyte-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx3-ubyte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
((-idx3- (-ubyte-)) m) ; the matrix to be written to the file
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: writes matrix <m> to file <file-pointer>

(fwrite-idx4-ubyte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
((-idx4- (-ubyte-)) m) ; the matrix to be written to the file
RETURNS: ()
CREATED: Yoshua Bengio 05 Feb 98
COMPILABLE: Yes
DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx4-ubyte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
((-idx4- (-ubyte-)) m) ; the matrix to be read from the file
RETURNS: ()
CREATED: Yoshua Bengio 05 Feb 98
COMPILABLE: Yes
DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fread-idx1-byte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
((-idx1- (-byte-)) m) ; the matrix to be read from the file
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx1-byte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx1- (-byte-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx2-byte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx2- (-byte-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx2-byte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx2- (-byte-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx3-byte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx3- (-byte-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx3-byte file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx3- (-byte-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx1-short file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx1- (-short-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx1-short file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx1- (-short-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx2-short file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx2- (-short-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx2-short file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx2- (-short-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx3-short file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx3- (-short-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx3-short file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx3- (-short-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx1-int file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx1- (-int-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx1-int file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx1- (-int-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx2-int file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx2- (-int-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx2-int file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx2- (-int-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx3-int file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx3- (-int-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx3-int file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx3- (-int-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx1-flt file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx1- (-flt-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx1-flt file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx1- (-flt-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx2-flt file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx2- (-flt-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx2-flt file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
((-idx2- (-flt-)) m) ; the matrix to be written to the file
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx3-flt file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
((-idx3- (-flt-)) m) ; the matrix to be read from the file
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: reads matrix <m> from file <file-pointer>
<m> is "resized" to fit the size of the matrix in the file

(fwrite-idx3-flt file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
((-idx3- (-flt-)) m) ; the matrix to be written to the file
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx1-real file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
((-idx1- (-real-)) m) ; the matrix to be read from the file
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: reads matrix <m> from file <file-pointer>
<m> is "resized" to fit the size of the matrix in the file

(fwrite-idx1-real file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx1- (-real-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx2-real file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx2- (-real-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx2-real file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx2- (-real-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(fread-idx3-real file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx3- (-real-)) m) ; the matrix to be read from the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: reads matrix <m> from file <file-pointer>
 <m> is "resized" to fit the size of the matrix in the file

(fwrite-idx3-real file-pointer m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-gptr-) file) ; the file pointer returned by a fopen
 ((-idx3- (-real-)) m) ; the matrix to be written to the file
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: writes matrix <m> to file <file-pointer>

(save-idx1-ubyte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-str-) filename):
 ((-idx1- (-ubyte-)) m)
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: saves matrix <m> to file <filename>

(load-idx1-ubyte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-str-) filename):
 ((-idx1- (-ubyte-)) m):
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: loads matrix <m> from file <filename>
 <m> is "resized" to fit the size of the matrix in the file

(save-idx2-ubyte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

((-str-) filename):
 ((-idx2- (-ubyte-)) m):
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: saves matrix <m> to file <filename>

(load-idx2-ubyte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):  
((-idx2- (-ubyte-)) m):  
RETURNS: ()  
CREATED: Pascal Vincent 04/05/96  
COMPILABLE: Yes  
DESCRIPTION: loads matrix <m> from file <filename>  
              <m> is "resized" to fit the size of the matrix in the file
```

(save-idx3-ubyte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):  
((-idx3- (-ubyte-)) m):  
RETURNS: ()  
CREATED: Pascal Vincent 04/05/96  
COMPILABLE: Yes  
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx3-ubyte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):  
((-idx3- (-ubyte-)) m):  
RETURNS: ()  
CREATED: Pascal Vincent 04/05/96  
COMPILABLE: Yes  
DESCRIPTION: loads matrix <m> from file <filename>  
              <m> is "resized" to fit the size of the matrix in the file
```

(save-idx4-ubyte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):  
((-idx4- (-ubyte-)) m):  
RETURNS: ()  
CREATED: Yoshua Bengio, 5 Feb 98  
COMPILABLE: Yes  
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx4-ubyte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx4- (-ubyte-)) m):
RETURNS: ()
CREATED: Yoshua Bengio, 5 Feb 98
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
              <m> is "resized" to fit the size of the matrix in the file
```

(save-idx1-byte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx1- (-byte-)) m)
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx1-byte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx1- (-byte-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
              <m> is "resized" to fit the size of the matrix in the file
```

(save-idx2-byte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx2- (-byte-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx2-byte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx2- (-byte-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
             <m> is "resized" to fit the size of the matrix in the file
```

(save-idx3-byte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx3- (-byte-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx3-byte filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx3- (-byte-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
             <m> is "resized" to fit the size of the matrix in the file
```

(save-idx1-short filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx1- (-short-)) m)
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx1-short filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx1- (-short-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
             <m> is "resized" to fit the size of the matrix in the file
```

(save-idx2-short filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx2- (-short-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx2-short filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx2- (-short-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
             <m> is "resized" to fit the size of the matrix in the file
```

(save-idx3-short filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx3- (-short-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx3-short filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx3- (-short-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
             <m> is "resized" to fit the size of the matrix in the file
```

(save-idx1-int filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx1- (-int-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx1-int filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx1- (-int-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
             <m> is "resized" to fit the size of the matrix in the file
```

(save-idx2-int filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx2- (-int-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx2-int filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):  
((-idx2- (-int-)) m):  
RETURNS: ()  
CREATED: Pascal Vincent 04/05/96  
COMPILABLE: Yes  
DESCRIPTION: loads matrix <m> from file <filename>  
              <m> is "resized" to fit the size of the matrix in the file
```

(save-idx3-int filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):  
((-idx3- (-int-)) m):  
RETURNS: ()  
CREATED: Pascal Vincent 04/05/96  
COMPILABLE: Yes  
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx3-int filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):  
((-idx3- (-int-)) m):  
RETURNS: ()  
CREATED: Pascal Vincent 04/05/96  
COMPILABLE: Yes  
DESCRIPTION: loads matrix <m> from file <filename>  
              <m> is "resized" to fit the size of the matrix in the file
```

(save-idx1-fft filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):  
((-idx1- (-flt-)) m):  
RETURNS: ()  
CREATED: Pascal Vincent 04/05/96  
COMPILABLE: Yes  
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx1-fft filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx1- (-flt-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
               <m> is "resized" to fit the size of the matrix in the file
```

(save-idx2-fft filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx2- (-flt-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx2-fft filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx2- (-flt-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
               <m> is "resized" to fit the size of the matrix in the file
```

(save-idx3-fft filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx3- (-flt-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx3-flt filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx3- (-flt-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
             <m> is "resized" to fit the size of the matrix in the file
```

(save-idx1-real filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx1- (-real-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx1-real filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx1- (-real-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
             <m> is "resized" to fit the size of the matrix in the file
```

(save-idx2-real filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx2- (-real-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx2-real filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx2- (-real-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
              <m> is "resized" to fit the size of the matrix in the file
```

(save-idx3-real filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx3- (-real-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: saves matrix <m> to file <filename>
```

(load-idx3-real filename m)

(lsh/libidx/idx-io.lsh) Last Modified: 2002-12-23 09:57:11

```
((-str-) filename):
((-idx3- (-real-)) m):
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: loads matrix <m> from file <filename>
              <m> is "resized" to fit the size of the matrix in the file
```

8.2.8 IDX mapping from files

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

functions to map files into virtual memory and to access the data as a matrix.

(mmap-idx s stor ndim magic)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

low-level function to map a file as an IDX.

(mmap-idx1-ubyte s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28
map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx2-ubyte s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28
map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx3-ubyte s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28
map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx4-ubyte s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28
map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx1-byte s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28
map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx2-byte s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28
map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx3-byte s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28
map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx4-byte s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28
map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx1-short s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28
map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx2-short s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28
map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx3-short s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx4-short s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx1-int s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx2-int s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx3-int s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx4-int s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx1-float s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx2-float s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx3-float s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx4-float s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it

(mmap-idx1-double s)

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it**(mmap-idx2-double s)**

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it**(mmap-idx3-double s)**

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it**(mmap-idx4-double s)**

(lsh/libidx/idx-map.lsh) Last Modified: 2002-08-19 03:01:28

map the IDX file **s** into virtual memory and returns an IDX to access it**8.2.9 Convolution, Subsampling, Oversampling.**

(lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

This library contains functions for performing 1D and 2D convolutions , backconvolutions , subsampling with local averaging , and oversampling by sample replication . Each functionality is provided as a type-independent splicing macro (DMD) with names of the form `midx-m1xxx` and `midx-m2xxx` , as well as type specific functions for floats , doubles , and bytes (with names of the form `idx-f1xxx` , `idx-d1xxx` , `idx-u1xxx` . Two kinds of convolution functions are provided: generic 1D and 2D convolutions for any size kernel , and optimized 1D convolutions with kernels of size 2 , 3 , 4 and 5. The optimized versions are faster than the regular ones by about a factor of 2 through the use of loop unrolling and circular register stacks.

Fast 1D Convolutions

(lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

a library of 1D convolutions that uses automatically generated code with unrolled loops and circular register stacks to limit memory accesses and speed up execution.

(midx-m1fastconvolacc in kernel out ksize type) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

ultrafast 1D convolution with accumulation of result into the output. **in** must be a CONTIGUOUS `idx1` , **kernel** an `idx1` , **out** a CONTIGUOUS `idx1` , **ksize** must be the size of the kernel , and **type** a string containing a valid C type , e.g. "float" (preferably the type of the objects in the 3 vectors). This macro is expanded at compile time into a bunch of "inline" instructions. It

unrolls loops with a circular register stack to minimize index calculations and memory access. Though this code is particularly fast on processors that have lots of floating point registers (MIPS , etc) , it manages to to beat the pedestrian version by about a factor of 2 on pentiums (reaching approx 200MFLOPS per GHz for kernel of size 5). THIS CODE DOES NOT DO ANY RUN-TIME CHECK FOR ARRAY SIZE , SO BE CAREFULL.

(idx-f1fastconvol2acc input kernel output) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

super-fast 1D convolution with size 2 kernel (with unrolled loops and optimal register assignment) **input** and **output** must be contiguous idx1 , **kernel** must be an idx1 all containing floats.

(idx-d1fastconvol2acc input kernel output) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

super-fast 1D convolution with size 2 kernel (with unrolled loops and optimal register assignment) **input** and **output** must be contiguous idx1 , **kernel** must be an idx1 all containing doubles.

(idx-u1fastconvol2acc input kernel output) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

super-fast 1D convolution with size 2 kernel (with unrolled loops and optimal register assignment) **input** and **output** must be contiguous idx1 , **kernel** must be an idx1 all containing ubytes.

(idx-f1fastconvol3acc input kernel output) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

super-fast 1D convolution with size 3 kernel. result is ACCUMULATED in the output (with unrolled loops and optimal register assignment) **input** and **output** must be contiguous idx1 , **kernel** must be an idx1 all containing floats.

(idx-d1fastconvol3acc input kernel output) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

super-fast 1D convolution with size 3 kernel. result is ACCUMULATED in the output (with unrolled loops and optimal register assignment) **input** and **output** must be contiguous idx1 , **kernel** must be an idx1 all containing doubles.

(idx-u1fastconvol3acc input kernel output) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

super-fast 1D convolution with size 3 kernel. result is ACCUMULATED in the output (with unrolled loops and optimal register assignment) **input** and **output** must be contiguous idx1 , **kernel** must be an idx1 all containing ubytes.

(idx-f1fastconvol4acc input kernel output) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

super-fast 1D convolution with size 4 kernel. result is ACCUMULATED in the output (with unrolled loops and optimal register assignment) **input** and **output** must be contiguous idx1 , **kernel** must be an idx1 all containing floats.

(idx-d1fastconvol4acc input kernel output) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

super-fast 1D convolution with size 4 kernel. result is ACCUMULATED in the output (with unrolled loops and optimal register assignment) `input` and `output` must be contiguous `idx1` , `kernel` must be an `idx1` all containing doubles.

(`idx-ulfastconvol4acc input kernel output`) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

super-fast 1D convolution with size 4 kernel. result is ACCUMULATED in the output (with unrolled loops and optimal register assignment) `input` and `output` must be contiguous `idx1` , `kernel` must be an `idx1` all containing ubytes.

(`idx-flfastconvol5acc input kernel output`) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

super-fast 1D convolution with size 5 kernel. result is ACCUMULATED in the output (with unrolled loops and optimal register assignment) `input` and `output` must be contiguous `idx1` , `kernel` must be an `idx1` all containing floats.

(`idx-dlfastconvol5acc input kernel output`) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

super-fast 1D convolution with size 5 kernel. result is ACCUMULATED in the output (with unrolled loops and optimal register assignment) `input` and `output` must be contiguous `idx1` , `kernel` must be an `idx1` all containing doubles.

(`idx-ulfastconvol5acc input kernel output`) (lsh/libidx/idx-convol.lsh)
Last Modified: 2002-10-09 02:43:39

super-fast 1D convolution with size 5 kernel. result is ACCUMULATED in the output (with unrolled loops and optimal register assignment) `input` and `output` must be contiguous `idx1` , `kernel` must be an `idx1` all containing ubytes.

Generic Convolutions

(lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

Regular 1D and 2D convolution and backconvolutions for `idx` of doubles , floats , and ubytes. Each function take 3 arguments: `input` , `kernel` , `output`. The size of the output plus the size of the kernel minus 1 must equal to the size of the input in all dimensions (to prevent annoying boundary conditions). The backconvolution can be seen the transposed operator of the convolution.

(`midx-m1convol input kernel output`) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

macro for 1D convolution. all arguments are `idx1` (of any numerical type).

(`idx-flconvol input kernel output`) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

1D convolution on `idx1` of floats

(`idx-m1convol input kernel output`) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

1D convolution on `idx1` of floats (alias of `idx-flconvol`)

(idx-d1convol input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

1D convolution on idx1 of double

(idx-u1convol input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

1D convolution on idx1 of ubytes

(midx-m1backconvol input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

macro for 1D back convolution

(idx-f1backconvol input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

1D back convolution on floats

(idx-m1backconvol input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

1D back convolution on floats (alias of idx-f1backconvol.

(idx-d1backconvol input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

1D back convolution on doubles

(idx-u1backconvol input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

1D back convolution on ubytes

(midx-m2convol input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

macro for 2D convolution. all arguments are idx2.

(idx-f2convol input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D convolution. all arguments are idx2 of floats

(idx-m2convol input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D convolution. all arguments are idx2 of floats. (alias for idx-f2convol)

(idx-d2convol input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D convolution. all arguments are idx2 of doubles.

(idx-u2convol input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D convolution. all arguments are idx2 of ubytes

(midx-m2convolacc input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

macro for 2D convolution with accumulation . all arguments are idx2.

(idx-f2convolacc input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D convolution with accumulation . all arguments are idx2 of floats

(idx-m2convolacc input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D convolution with accumulation . all arguments are idx2.

(**idx-d2convolacc** input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D convolution with accumulation . all arguments are idx2 of double.

(**idx-u2convolacc** input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D convolution with accumulation . all arguments are idx2 of ubyte.

(**midx-m2backconvol** input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

macro for 2D back-convolution. all arguments are idx2.

(**idx-f2backconvol** input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D back-convolution. all arguments are idx2 of floats.

(**idx-m2backconvol** input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D back-convolution. all arguments are idx2 of floats (alias for idx-f2baconvol)

(**idx-d2backconvol** input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D back-convolution. all arguments are idx2 of floats.

(**idx-u2backconvol** input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D back-convolution. all arguments are idx2 of floats.

(**midx-m2backconvolacc** input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

macro for 2D back-convolution. all arguments are idx2.

(**idx-f2backconvolacc** input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D back-convolution. all arguments are idx2 of floats.

(**idx-m2backconvolacc** input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D back-convolution. all arguments are idx2 of floats.

(**idx-d2backconvolacc** input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D back-convolution. all arguments are idx2 of floats.

(**idx-u2backconvolacc** input kernel output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

2D back-convolution. all arguments are idx2 of floats.

1D and 2D Subsampling and Oversampling

(lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

functions that subsample 1D and 2D idx where each output sample is the average of **n** adjacent input samples (with no overlap). The oversampling functions simply replicate the samples **n** times.

(**midx-m1subsample** input **n** output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

macro for subsampling and summing of size **n**

(idx-f1subsample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

subsampling and summing of size *n*

(idx-m1subsample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

subsampling and summing of size *n*

(idx-d1subsample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

subsampling and summing of size *n*

(idx-u1subsample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

subsampling and summing of size *n*

(midx-m1oversample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

macro for oversampling of size *n*

(idx-f1oversample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

oversampling of size *n*

(idx-m1oversample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

oversampling of size *n*

(idx-d1oversample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

oversampling of size *n*

(idx-u1oversample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

oversampling of size *n*

(midx-m2subsample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

macro for subsampling and summing of size *n*

(idx-f2subsample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

subsampling and summing of size *n*

(idx-m2subsample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

subsampling and summing of size *n*

(idx-d2subsample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

subsampling and summing of size *n*

(idx-u2subsample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

subsampling and summing of size *n*

(midx-m2oversample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

macro for oversampling of size *n*

(idx-f2oversample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

oversampling of size *n*

(idx-m2oversample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

oversampling of size *n*

(idx-d2oversample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

oversampling of size *n*

(idx-u2oversample input n output) (lsh/libidx/idx-convol.lsh) Last Modified: 2002-10-09 02:43:39

oversampling of size *n*

8.2.10 IDX sorting and binary search functions

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

sorting and bsearch functions for matrices of floats and ints

(idx-f1bsearch am v)

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

returns index of element in *idx1* of float *am* whose value is less than or equal to *v* using a binary search. this assumes the elements in *am* are sorted in ascending order.

(idx-d1bsearch am v)

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

returns index of element in *idx1* of double *am* whose value is less than or equal to *v* using a binary search. this assumes the elements in *am* are sorted in ascending order.

(idx-i1bsearch am v)

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

returns index of element in *idx1* of int *am* whose value is less than or equal to *v* using a binary search. this assumes the elements in *am* are sorted in ascending order.

(idx-f1sortup m)

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

sorts element in float vector *m* in ascending order

(idx-d1sortup m)

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

sorts element in float vector *m* in ascending order

(idx-flsortdown m)

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

sorts element in float vector m in descending order**(idx-d1sortdown m)**

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

sorts element in double vector m in descending order**(idx-i1sortup m)**

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

sorts element in integer vector m in ascending order**(idx-i1sortdown m)**

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

sorts element in integer vector m in descending order**(idx-fli1sortup m p)**

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

m is a vector of float , p is a vector of integers (same dimension as m). on output , m is sorted in ascending order , and p is sorted with the same permutation table.

(idx-fli1sortdown m p)

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

m is a vector of float , p is a vector of integers (same dimension as m). on output , m is sorted in descending order , and p is sorted with the same permutation table.

(idx-i1i1sortup m p)

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

m is a vector of int , p is a vector of integers (same dimension as m). on output , m is sorted in ascending order , and p is sorted with the same permutation table.

(idx-i1i1sortdown m p)

(lsh/libidx/idx-sort.lsh) Last Modified: 2002-11-06 18:37:02

m is a vector of int , p is a vector of integers (same dimension as m). on output , m is sorted in descending order , and p is sorted with the same permutation table.

8.2.11 IDX operations on squares of elements

(lsh/libidx/idx-squops.lsh) Last Modified: 2002-08-19 03:01:28

matrix and tensor operations that use the squares of the elements. This is used primarily for second derivative backpropagations in gradient-based learning algorithms.

(idx-m1squextm1 m1 m2 m3)

(lsh/libidx/idx-squops.lsh) Last Modified: 2002-08-19 03:01:28

square outer product of **m1** and **m2** . $M3ij = M1i * M2j^2$

(idx-m2squextm2 m1 m2 m3)

(lsh/libidx/idx-squops.lsh) Last Modified: 2002-08-19 03:01:28

square outer product of **m1** and **m2** . $M3ijkl = M1ij * M2kl^2$

(idx-m1squextm1acc m1 m2 m3)

(lsh/libidx/idx-squops.lsh) Last Modified: 2002-08-19 03:01:28

square outer product of **m1** and **m2** . $M3ij += M1i * M2j^2$

(idx-m2squextm2acc m1 m2 m3)

(lsh/libidx/idx-squops.lsh) Last Modified: 2002-08-19 03:01:28

square outer product of **m1** and **m2** . $M3ijkl += M1ij * M2kl^2$

(idx-m2squdotm1 m1 m2 m3)

(lsh/libidx/idx-squops.lsh) Last Modified: 2002-08-19 03:01:28

multiply vector **m2** by matrix **m1** using square of **m1** elements $M3i = \sum_j M1ij^2 * M2j$

(idx-m4squdotm2 m1 m2 m3)

(lsh/libidx/idx-squops.lsh) Last Modified: 2002-08-19 03:01:28

multiply matrix **m2** by tensor **m1** using square of **m1** elements $M3ij = \sum_{kl} M1ijkl^2 * M2kl$

(idx-m2squdotm1acc m1 m2 m3)

(lsh/libidx/idx-squops.lsh) Last Modified: 2002-08-19 03:01:28

multiply vector **m2** by matrix **m1** using square of **m1** elements $M3i += \sum_j M1ij^2 * M2j$

(idx-m4squadotm2acc m1 m2 m3)

(lsh/libidx/idx-squops.lsh) Last Modified: 2002-08-19 03:01:28

multiply matrix **m2** by tensor **m1** using square of **m1** elements $M3_{ij} += \sum_{kl} M1_{ijkl}^2 * M2_{kl}$ **(idx-m1squadotm1acc m1 m2 m3)**

(lsh/libidx/idx-squops.lsh) Last Modified: 2002-08-19 03:01:28

dot product between **m1** and **m2**, except square of terms of **m1** are used: $M3 += \sum_i M1_i^2 * M2_i$ **(idx-m2squadotm2acc m1 m2 m3)**

(lsh/libidx/idx-squops.lsh) Last Modified: 2002-08-19 03:01:28

dot product between matrices **m1** and **m2**, except square of terms of **m1** are used: $M3 += \sum_{ij} M1_{ij}^2 * M2_{ij}$

8.3 C Interface Utilities

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

8.3.1 Libc interface

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

this provides an interface to several commonly used functions from the standard C library.

(malloc n)

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

malloc **n** bytes of memory and return a **gp**tr to it.**(free p)**

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

free malloc'ed memory area pointed to by **gp**tr **p**.**(memcpy dest src size)**

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

free malloc'ed memory area pointed to by **gp**tr **p**.**(little-endianp)**

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

return **t** if the CPU is little endian (e.g. x86). and **nil** if it is not little endian (like sparcs).

(testbit v b)

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

test if bit **b** of integer **v** is on. **b** must be between 0 and 31.

8.3.2 Control Structures

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

A few macros that can be used as control structures in compiled code.

(icfor (i begin end step) 11 [12...])

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

a C-like for loop with integer loop variable. **i** is a lisp symbol used as a loop variable, **begin end** and **step** must be numerical expression. They are all converted to int. This is equivalent to the C construct:

```
for (int i = begin; i<end; i++) { .... }
```

This is macro that generates inline C, and can therefore only be called in compiled code.

8.3.3 Pointer Manipulation/Dereferencing

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

(gptr+ p n)

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

add **n** to generic pointer **p** . The returned pointer points **n** bytes higher than **p** .

(gptr- p1 p2)

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

returns the distance between two pointers (in bytes).

(str-ptr s)

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

return a generic pointer to the data of string **s** (i.e. a pointer to a C-style null-terminated string). This must be used with caution because the data pointed to by the pointer may be garbage collected.

(idx-ptr idx)

Type: DX(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

return a **gptr** that points to the first element of **m** . This must be used with caution because the data pointed to by the pointer may be garbage collected (when **m** is destroyed) , or even displaced (when **m** is resized). **idx-ptr** is primarily used when calling C function that require a pointer to numerical data.

Peek/Poke/Sizeof

(lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

A number of functions are provided to dereference pointers to various numerical types and either get the value they point to, or set the value at the address they point to. They are of the form:

```
(peek-TYPE <ptr>)
(poke-TYPE <ptr> <value>)
```

where TYPE can be any Lush numerical type (ubyte, short, int, float, double, gptr). The first form returns the value pointed to by **ptr** , while the second form sets the value. Another set of functions, (**sizeof-TYPE**) returns the size of the corresponding type.

These function are useful for accessing values in a pre-existing C structure (allocated by a C function). Creating a numerical variable from Lush is better done by first creating an **idx0** of the appropriate type, and then getting a pointer to the data through **idx-ptr** . Example:

```
(setq z (double-matrix))
(my-c-function-that-takes-a-ptr-to-a-double (idx-ptr z))
```

```
(peek-ubyte s) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(poke-ubyte s v) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(sizeof-ubyte) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(peek-short s) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(poke-short s v) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(sizeof-short) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(peek-int s) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(poke-int s v) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(sizeof-int) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(peek-float s) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(poke-float s v) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(sizeof-float) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(peek-double s) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(poke-double s v) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
(sizeof-double) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
```

(**peek-gptr s**) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
 (**poke-gptr s v**) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54
 (**sizeof-gptr**) (lsh/libc/libc.lsh) Last Modified: 2002-09-18 11:19:54

8.4 Plotting

(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

This class library allows graph plotting. It contains the `Plotter` and the `PlotterCurve` classes providing functions for drawing multiple curves and modifying their appearance. Here is a sample code to plot various curves:

```
(libload "libplot/plotter")
;; create a new plotter with default size (open a window if necessary)
(setq p (new Plotter))
;; plot log function in blue from 0 to 10.
(==> p PlotFunc "log" log 0 10 0.1 (alloccolor 0 0 1))
;; show it on the screen
(==> p redisplay)

;; add a piece of a red ellipse
(==> p PlotXY "lips" sin cos 0 5 0.1 (alloccolor 1 0 0))
(==> p redisplay)

;; plot a green polygonal sinusoid
(setq x (range 0 10 0.5))
(setq y (all ((v x)) (sin v)))
(==> p plotlists "sine" x y (alloccolor 0 1 0) closed-circle)
(==> p Redisplay)

;; add a grid
(==> p SetGrid 1)
(==> p Redisplay)

;; move the plotter to 100,0 and set its size to 300,200
(==> p move 100 0)
(==> p setsize 300 200)
(cls)
(==> p redisplay)

;; plot log function in a postscript file
(let ((window (ps-window "/tmp/curve.ps")))
  (setq psp (new Plotter))
  (==> psp PlotFunc "log" log 0 10 0.1 (alloccolor 0 0 1))
  (==> psp redisplay))
```

```
;; plot log function in an editable graphic window using comdraw
(setq window (comdraw-window))
(setq cp (new Plotter))
(==> cp PlotFunc "log" log 0 10 0.1 (alloccolor 0 0 1))
(==> cp setgrid 1)
(==> cp redisplay)
```

8.4.1 Plotter

(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Class **Plotter** This is the main class to define **Plotter** objects in Lush. **Plotter** objects have slots to define properties of the axes such as scales , labels , legend , title , ticks and logscale. It also stores the curves , lines and texts that will be drawn on the **Plotter**. **rect** slot of the **Plotter** specifies the rectangular area of the **Plotter** on the current window. Slots of the **Plotter** object can be modified by methods of the class. They are computed automatically or set to their default values if not specified by the use.

(new Plotter [[*x y w h*])

Type: CLASS(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

The constructor **Plotter** creates a new **Plotter** object at pixel positions (*x* , *y*) , with width *w* and height *h* on the current window. If the **window** object is empty , it pops up a new window. *x* , *y* , *w* and *h* are optional. If *x* and *y* are omitted , they are both set to 0. If *w* and *h* are omitted , the size of the **Plotter** is equal to the size of the current window. For example:

```
(setq myplotter (new Plotter 600 600))
```

(==> Plotter move *x y*)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

move the upper left-hand corner of the plotter to screen position *x y* . No screen update is performed until the next call to **redisplay**.

(==> Plotter setsize *w h*)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

set the width and height of the plotter object to *w* and *h* . No screen update is performed until the next call to **redisplay**.

(==> Plotter Set *attribute value*)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Sets the attribute of the **Plotter** object to the given value. The **attribute** argument is a string and can be "xticks" , "yticks" , "xscale" , "yscale" , "xlabel"

, "ylabel" , "legend" , "title" , "grid" , "xscale2" , "yscale2" , "xticks2" , "yticks2" , "xlabel2" , "ylabel2". If an invalid attribute name is given error is generated. The **value** argument should match the attribute's data type. For "xticks" , "yticks" , "xticks2" and "yticks2" the value should be a list of numbers , for "xscale" and "yscale" it should be a list of two numbers , for "legend" , "title" , "xlabel" , "ylabel" , "xlabel2" and "ylabel2" , it should be a string. Example:

```
(==> myplotter Set "xticks" (range 0 10))
(==> myplotter Set "xscale" (list 0 10))
(==> myplotter Set "xlabel" "x")
```

Individual methods are also provided to set and change the attributes of the Plotter object.

```
(==> Plotter SetRectNth i m n)
```

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Divides the current window into **m** by **n** rectangles and sets the rectangular area of the **Plotter** to the **i** th one. The first rectangle is the one on the upper left corner. When the rectangle of a **Plotter** object is set , the graphics plot is restricted to that rectangular area.

For example:

```
;;divides the current window vertically and
;;plots the first object on the left and the second object on the right.
(==> plotter1 SetRectNth 1 1 2)
(==> plotter2 SetRectNth 2 2 1)
```

```
(==> Plotter SetXScale xmin xmax)
```

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Sets the **xscale** slot of the **Plotter** . **xmin** argument is the minimum axis limit and **xmax** is the maximum axis limit on x. If a scale for x is not set , LUSH selects the axis limits based on the curve points.

```
(==> Plotter SetYScale ymin ymax)
```

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Sets the **yscale** slot of the **Plotter** . **ymin** argument is the minimum axis limit and **ymax** is the maximum axis limit on y. If a scale for y is not set , LUSH selects the axis limits based on the curve points.

(==> Plotter SetXTicks x)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Sets the `xticks` slot of the `Plotter` . LUSH computes the ticks marks on the x axis based on the x scale. However the tick marks on the x axis can be specified with the `SetXTicks` method. Argument `x` is a list of numbers

(==> Plotter SetYTicks y)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Sets the `yticks` slot of the `Plotter` . LUSH computes the ticks marks on the y axis based on the y scale. However the tick marks on the y axis can be specified with the `SetYTicks` method. Argument `y` is a list of numbers.

(==> Plotter SetXLabel x)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Adds axis label to the x axis. Argument `x` is a string.

(==> Plotter SetYLabel y)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Adds axis label to the y axis. Argument `y` is a string.

(==> Plotter SetLegend l)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Adds a legend text to the up right corner of the `Plotter` rectangle. Argument `l` is a string.

(==> Plotter SetTitle t)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Adds a title text to the top of the `Plotter` rectangle. Argument `t` is a string.

(==> Plotter SetGrid x)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Specifies whether to draw grid lines or not. Argument `x` is either 0 or 1. When `x` is 1 grid lines are drawn , when `x` is 0 grids are not drawn.

(==> Plotter SetXScale2 xmin xmax)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Using the `Plotter` library it is possible to draw graphs with double x axes. `SetXScale2` method sets the scale of the second x axis. `xmin` argument is the minimum axis limit and `xmax` is the maximum axis limit on the second x axis. If no scale is specified for the second x axis , it is not drawn.

(==> **Plotter SetYScale2** ymin ymax)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Using the **Plotter** library it is possible to draw graphs with double y axes. **SetYScale2** method sets the scale of the second y axis. **ymin** argument is the minimum axis limit and **ymax** is the maximum axis limit on the second y axis. If no scale is specified for the second y axis, it is not drawn.

(==> **Plotter SetXTicks2** x)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Using the **Plotter** library it is possible to draw graphs with double x axes. **SetXTicks2** method sets the tick marks of the second x axis. Argument **x** is a list of numbers indicating the tick marks.

(==> **Plotter SetYTicks2** x)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Using the **Plotter** library it is possible to draw graphs with double y axes. **SetYTicks2** method sets the tick marks of the second y axis. Argument **y** is a list of numbers indicating the tick marks.

(==> **Plotter SetXLabel2** x)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Using the **Plotter** library it is possible to draw graphs with double x axes. **SetXLabel2** method sets the label of the second x axis. Argument **x** is a string.

(==> **Plotter SetYLabel2** y)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Using the **Plotter** library it is possible to draw graphs with double y axes. **SetYLabel2** method sets the label of the second y axis. Argument **y** is a string.

(==> **Plotter SetXLogScale**)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Using the **Plotter** library it is possible to draw logplots. **SetXLogScale** method converts the scale of the x axis to log.

(==> **Plotter SetYLogScale**)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Using the **Plotter** library it is possible to draw logplots. **SetYLogScale** method converts the scale of the y axis to log.

(==> Plotter ResetXLogScale)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Using the `Plotter` library it is possible to draw logplots. `ResetYLogScale` method converts the scale of the x axis from log to regular.

(==> Plotter ResetXLogScale)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Using the `Plotter` library it is possible to draw logplots. `ResetYLogScale` method converts the scale of the x axis from log to regular.

(==> Plotter NewCurve name)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Adds a new `PlotterCurve` object to the `Plotter`. `name` is the name of the curve and of type string. Curves are stored as alists in the `Plotter` objects. Each `PlotterCurve` is associated with a name and with the `NewCurve` method added to the list of curves in the `Plotter`.

(==> Plotter ClearCurve curvename)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Deletes the `PlotterCurve` object named `curvename` form the `Plotter`.

(==> Plotter AddPoint curvename x y)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Adds the data point (x, y) to the `PlotterCurve` object named `curvename` to the `Plotter`. If a curve named `curvename` does not exist, error is returned. For example:

```
(=> myplotter AddPoint "S" 0 10)
```

(==> Plotter SetCurveColor curvename c)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Sets the color of the `PlotterCurve` object named `curvename` to color `c`. Argument `c` is a color number (as returned by `alloc-color`). If a curve named `curvename` does not exist, error is returned.

(==> Plotter SetCurveSymbol curvename symbolfunc)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Sets the marker symbol of the `PlotterCurve` object named `curvename` to symbol `symbolname`. Argument `symbolfunc` must be a two-argument function whose role is to draw a symbol at the screen coordinates passed as arguments. A set of pre-defined such functions are provided: `open-square`, `closed-square`,

open-circle , closed-circle , open-up-triangle , open-down-triangle , closed-up-triangle , closed-down-triangle , straight-cross and oblique-cross .

(==> Plotter **SetCurveSymbolSize** curvename symbolsize)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Sets the size of the marker of the **PlotterCurve** named **curvename** to **symbolsize** . **symbolsize** is an integer normally between 1 and 5. If **symbolsize** is set to 0 no markers are drawn at data points. For example:

```
(==> myplotter SetCurveSymbolSize "s" 3)
```

(==> Plotter **SetCurveLine** curvename line)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Specifies whether to draw lines between markers of the **PlotterCurve** object or not. Argument **line** is 0 or 1. If **line** is 0 only markers at data points are drawn , if **line** is 1 , markers are connected. By default , **line** is 1.

(==> Plotter **SetCurveXAxis** curvename xaxis)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

In case of use of double x axes , associates the **PlotterCurve** object to the x axis specified by the argument **xaxis** . If **xaxis** argument is 0 the x axis at the bottom is used (default) , if it is 1 , x axis at the top is used to draw the curve. The scales for the second (top) axis should be set.

(==> Plotter **SetCurveYAxis** curvename yaxis)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

In case of use of double y axes , associates the **PlotterCurve** object to the y axis specified by the argument **yaxis** . If **yaxis** argument is 0 the y axis on the left is used (default) , if it is 1 , x axis on the right is used to draw the curve. The scales for the second (right) axis should be set.

(==> Plotter **PlotFunc** curvename func xmin xmax xstep [color] [symbol])

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Plot a scalar real function. Samples are computed for are all values between **xmin** and **xmax** by step of **xstep** . **color** is a color identifier (as returned by **alloccolor**) , and **symbol** is a function of two arguments (x and y coordinates) that is called each time a data point is to be plotted. Predefined such functions include nil (no symbol is plotted) , open-square , closed-square , open-circle , closed-circle , open-up-triangle , open-down-triangle , closed-up-triangle , closed-down-triangle , straight-cross and oblique-cross , or any user-defined functions

Example:

```
(libload "libplot/plotter")
(setq p (new Plotter))
(==> p PlotFunc "log" log 0 10 0.1 (alloccolor 1 0 0))
(==> p redisplay)
```

```
(==> Plotter PlotXY curvename funcx funcy tmin tmax tstep [color]
[symbol])
```

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Plot a parameterized curve $(X(t), Y(t))$. Samples are computed for all values of t between $tmin$ and $tmax$ by step of $tstep$. $color$ is a color identifier (as returned by `alloccolor`), and $symbol$ is a function of two arguments (x and y coordinates) that is called each time a data point is to be plotted. Predefined such functions include `nil` (no symbol is plotted), `open-square`, `closed-square`, `open-circle`, `closed-circle`, `open-up-triangle`, `open-down-triangle`, `closed-up-triangle`, `closed-down-triangle`, `straight-cross` and `oblique-cross`, or any user-defined functions

Example:

```
(libload "libplot/plotter")
(setq p (new Plotter))
(==> p PlotXY "log" log 0 10 0.1 (alloccolor 1 0 0))
(==> p redisplay)
```

```
(==> Plotter PlotLists curvename xlist ylist [color] [symbol])
```

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Plot a list of points whose x and y coordinates are corresponding elements in $xlist$ and $ylist$. In other words $xlist$ is the list of x coordinates of the data, $ylist$ is the list of y coordinates of the datapoints. $curvename$ is a curve identifier (generally a string). If the curve $curvename$ already exists, it is replaced. $color$ is a color identifier (as returned by `alloc-color`), and $symbol$ is a function of two arguments (x and y coordinates) that is called each time a data point is to be plotted. Predefined such functions include `open-square`, `closed-square`, `open-circle`, `closed-circle`, `open-up-triangle`, `open-down-triangle`, `closed-up-triangle`, `closed-down-triangle`, `straight-cross` and `oblique-cross`. Example:

```
(==> myplotter PlotLists "log" (range 1 100) (mapcar log x))
```

```
(==> Plotter AddText x y txtstring)
```

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Adds a new text to the **Plotter** object. Texts are stored as list of lists in the **Plotter** . Each element of the text list is a list of the coordinates of the corresponding text and its string. Arguments **x** and **y** are real coordinates (coordinates on the plotter)

(==> **Plotter AddLine** x1 y1 x2 y2)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Adds a new line to the **Plotter** object. Lines are stored as list of lists in the **Plotter** . Each element of the line list is a list of the coordinates of the line. Arguments **x1** , **x2** , **y1** and **y2** are real coordinates (coordinates on the Plotter).

(==> **Plotter GetClick**)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Returns the real coordinates of the point clicked on the **Plotter** . This method is not supported in a Comdraw window

(==> **Plotter GetDistance**)

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Computes the the length of the vector drawn by dragging and dropping the mouse pointer on the **Plotter** . The distance is in real coordinates (coordinates of the plotter). This method is not supported in a Comdraw window.

(==> **Plotter Redisplay** [location])

Type: MSG(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

This is fundamental method of a **Plotter** object. Every curve , line , text , etc is stored in the **Plotter** object and is drawn after a **Redisplay** command. To see the modifications on the window , **Redisplay** should be called. The argument **location** is optional. It is a list of x and y pixel coordinates , width and height of the **Plotter** . If specified as (list x y w h) , it displays the **Plotter** in the rectangular area (x y w h).

For example:

```
;;locates and reprints the plotter in the rectangle
;;specified by (100,100,500,500) on current window
(==> myplotter Redisplay 100 100 500 500)
```

8.4.2 low level utility functions

(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

(compute-xscale pobject)

(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Computes the scale of the x axis for a given `Plotter` object from the data points in each curve.

(compute-yscale pobject)

(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Computes the scale of the y axis for a given `Plotter` object from the data points in each curve.

(set-axes axes xscale yscale xticks yticks xlabel ylabel legend title grid [xlogscale] ylogscale])

(lsh/libplot/plotter.lsh) Last Modified: 2003-01-06 23:21:30

Plots the axes given the scales of the x and y axis , axis ticks , axis labels , legend and the title. Labels , legend and title can be empty strings. The `axes` argument is a list of `xaxismin` , `yaxismin` , `xaxismax` and `yaxismax`. `xaxismin` and `yaxismin` are the pixel coordinates of the origin of the axes , and `xaxismax` and `yaxismax` are the pixel coordinates of the top right corner of the axes. The `grid` argument is either 1 or 0 , specifies whether to have grid on the plotter or not. `xlogscale` and `ylogscale` are optional and are used if a logscale is desired on one of the axes.

For example:

```
(set-axes (list 100 600 800 100)
  (list 0 10) (list 0 10) (range 0 10) (range 0 10) "x" "y" "" "graph" 1)
```

8.5 C-Style Input / Output

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

functions this section provide a convenient interface to the C stdio library. Functions are provided to open and close C-style file descriptors and pipes. The most commonly used stdio functions are also provided , including `ftell` , `fseek` , `fprintf` , `rewind` , `fputc` , `fgetc` , and several others. C-style file descriptors manipulated by these functions are merely generic pointers (`-gptr-`). As a consequence , many of these function as "unsafe" , i.e. calling them with an invalid file pointer may cause the interpreter to crash.

8.5.1 (fprintf file-pointer args...)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a `fopen`

MACRO

DESCRIPTION: same as printf, but writes to a file rather than to standard output

RETURNS: ()

CREATED: Pascal Vincent 04/05/96

8.5.2 (fwrite-str file-pointer s)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ;; the file pointer returned by a fopen

((-str-) s) ;; the string to be written to the file

RETURNS: ()

CREATED: Pascal Vincent 04/05/96

COMPILABLE: Yes

DESCRIPTION: writes string <s> to file <file-pointer>

See: Compilable file I/O

8.5.3 (stdout)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

RETURNS: (-gptr-) the file pointer associated with the standard output

CREATED: Pascal Vincent 04/05/96

COMPILABLE: Yes

See: Compilable file I/O

8.5.4 (fopen filename type)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-str-) filename) ; the name (or path) of the file to be opened:

((-str-) type) ; the opening type: "rb" for reading, "wb" for creating and writing, "a" for

RETURNS: (-gptr-) a file pointer to the open file (the pointer is null in case of failiure)

CREATED: Pascal Vincent 04/05/96

COMPILABLE: Yes

DESCRIPTION: same as the C function fopen

See: Compilable file I/O

8.5.5 (fclose file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a fopen
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: same as the C function fclose

See: Compilable file I/O

8.5.6 (popen filename type)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-str-) filename) ; the name (or path) of the pipe to be opened:
((-str-) type) ; the opening type: "rb" for reading, "wb" for creating and writing
RETURNS: (-gptr-) a file pointer to the open pipe (the pointer is null in case of failure)
CREATED: Yann LeCun 08/29/96
COMPILABLE: Yes
DESCRIPTION: same as the C function popen

See: Compilable file I/O

8.5.7 (pclose file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a popen
RETURNS: ()
CREATED: Yann LeCun 08/29/96
COMPILABLE: Yes
DESCRIPTION: same as the C function pclose

See: Compilable file I/O

8.5.8 (ftell file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a fopen
RETURNS: (-int-) the current position in the file, relative to the beginning
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: same as C function ftell

See: Compilable file I/O

8.5.9 (fseek file-pointer pos)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

```
((-gptr-) file) ; the file pointer returned by a fopen
((-int-) pos) ; an absolute position in the file:
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: moves to the given position in the file
```

See: Compilable file I/O

8.5.10 (fseek-from-end file-pointer pos)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

```
((-gptr-) file) ; the file pointer returned by a fopen
((-int-) pos) ; a position in the file relative to the end
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: moves to the given position in the file
```

See: Compilable file I/O

8.5.11 (fseek-from-current file-pointer pos)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

```
((-gptr-) file) ; the file pointer returned by a fopen
((-int-) pos) ; a position in the file relative to the current position:
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: moves to the given position in the file
```

See: Compilable file I/O

8.5.12 (fgetc file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a fopen

RETURNS: (-int-) The next byte read from the file

CREATED: Pascal Vincent 04/05/96

COMPILABLE: Yes

DESCRIPTION: Reads a byte (character) from the file (same as C function fgetc)

See: Compilable file I/O

8.5.13 (fputc file-pointer val)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a fopen

((-int-) val) ; contains the byte to be written to the file

RETURNS: ()

CREATED: Pascal Vincent 04/05/96

COMPILABLE: Yes

DESCRIPTION: Writes a byte (character) to the file (same as C function fputc)

See: Compilable file I/O

8.5.14 (fread-ubyte file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a fopen

RETURNS: (-ubyte-) the next ubyte read from the file

CREATED: Pascal Vincent 04/05/96

COMPILABLE: Yes

See: Compilable file I/O

8.5.15 (fwrite-ubyte file-pointer val)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a fopen

((-ubyte-) val) ; contains the ubyte to be written to the file

RETURNS: ()

CREATED: Pascal Vincent 04/05/96

COMPILABLE: Yes

DESCRIPTION: Writes a ubyte to the given file

See: Compilable file I/O

8.5.16 (fread-byte file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

```
((-gptr-) file) ; the file pointer returned by a fopen
RETURNS: (-byte-) the next byte read from the file
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
```

See: Compilable file I/O

8.5.17 (fwrite-byte file-pointer val)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

```
((-gptr-) file) ; the file pointer returned by a fopen
((-byte-) val) ; contains the byte to be written to the file
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: Writes an byte (4 bytes on Sun-OS or Solaris) to the given file
```

See: Compilable file I/O

8.5.18 (fread-short file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

```
((-gptr-) file) ; the file pointer returned by a fopen
RETURNS: (-short-) the next short (2 bytes) read from the file
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
```

See: Compilable file I/O

8.5.19 (fwrite-short file-pointer val)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

```
((-gptr-) file) ; the file pointer returned by a fopen
((-short-) val)
RETURNS: ()
```

CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: Writes a short (2 bytes) to the given file

See: Compilable file I/O

8.5.20 (fread-int file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a fopen
RETURNS: (-int-) the next int (4 bytes on Sun-OS or Solaris) read from the file
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes

See: Compilable file I/O

8.5.21 (fwrite-int file-pointer val)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a fopen
((-int-) val) ; contains the int to be written to the file
RETURNS: ()
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes
DESCRIPTION: Writes an int (4 bytes on Sun-OS or Solaris) to the given file

See: Compilable file I/O

8.5.22 (fread-flt file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a fopen
RETURNS: (-float-) the next float (4 bytes) read from the file
CREATED: Pascal Vincent 04/05/96
COMPILABLE: Yes

See: Compilable file I/O

8.5.23 (fwrite-flt file-pointer val)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a fopen
 ((-float-) val):
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: Writes a float (4 bytes) to the given file

See: Compilable file I/O

8.5.24 (fread-real file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a fopen
 RETURNS: (-real-) the next real (8 bytes) read from the file
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes

See: Compilable file I/O

8.5.25 (fwrite-real file-pointer val)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) file) ; the file pointer returned by a fopen
 ((-real-) val):
 RETURNS: ()
 CREATED: Pascal Vincent 04/05/96
 COMPILABLE: Yes
 DESCRIPTION: Writes a real (8 bytes) to the given file

See: Compilable file I/O

8.5.26 (reverse_n ptr sz n)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) ptr) ;; pointer to the block of memory that must be reversed
 ((-int-) sz) ;; size of each item to reverse
 ((-int-) n) ;; number of items to reverse

RETURNS: ()

CREATED: Pascal Vincent 04/17/96

COMPILABLE: Yes

DESCRIPTION: Reverses order of bytes in <n> items of size <sz>
starting at memory location <ptr>

This is a tool for writing/reading file formats that are portable
across systems with processors that represent long-words
differently in memory (Sparc vs Intel-Pentium for ex.)
It can be called from inline-C as C_reverse8(ptr,n);

8.5.27 (fscan-int file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

(-gptr- file-pointer) ;; the file pointer returned by fopen

DESCRIPTION: Same as fscanf(file_pointer,"\\%d",&result) in C.

RETURNS: int

CREATED: Yoshua Bengio 23 August 1996

8.5.28 (fscan-flt file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

(-gptr- file-pointer) ;; the file pointer returned by fopen

DESCRIPTION: Same as fscanf(file_pointer,"\\%d",&result) in C.

RETURNS: float

CREATED: Yoshua Bengio 23 August 1996

8.5.29 (fscan-str file-pointer)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

(-gptr- file-pointer) ;; the file pointer returned by fopen

DESCRIPTION: same as fscanf(file_pointer,"\\%s",result).

The result must hold in a string of size < 1024 bytes.

RETURNS: string

CREATED: Yoshua Bengio 23 August 1996

8.5.30 (fgets file-pointer max-size)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

(-gptr- file-pointer) ;; the file pointer returned by fopen

(-int- max-size) ;; maximum allowed size of line in bytes
 DESCRIPTION: same as C's fgets(result,max_size,file_pointer).
 RETURNS: string
 CREATED: Yoshua Bengio 23 August 1996

8.5.31 (file-size file-name)

Type: DE(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10
 returns the size in byte of file file-name author: Leon Bottou.

8.5.32 (rewind f)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-gptr-) f):
 CREATED: Oct 97
 DESCRIPTION:

8.5.33 (skip-comments start f)

(lsh/libc/stdio.lsh) Last Modified: 2002-12-23 09:57:10

((-char-) start):
 ((-gptr-) f):
 CREATED: Oct 97
 DESCRIPTION:
 skip comments lines starting with <start>

8.6 Graph Library

8.6.1 Building and Manipulating Grammars

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50
 functions for easily building grammars. The basic process consists in first building a grammar graph with null transitions (transitions matching no symbols) and then creating a graph without null transitions.

(NULL-PTR s)

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50
 RETURNS: t if s is (). CREATED: P.Haffner

Macros for typing labels

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

(LAB-ASCII str) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Label for an ASCII character

(LAB-NULL) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Label for the null transition

(LAB-MARKER string) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Label for marker links (e.g. field markers , junk markers , etc..) Argument **str** must be a 2 character string.

(LAB-WILD) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Label for matching any data with a reasonable score This is a "wildcard" label used for matching meaningful punctuations.

(LAB-JUNK) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Label for matching any data with a quite poor score This is a "junk" label for matching meaningless subimages before or after the interesting point.

(LAB-SKIP) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Label for matching any character , without calling the recognizer The default score will be 0

Gset: multiple sets of integers

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

GMultiIntSet (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

This class represents efficiently a collection of sets of integers. This is used by the new version of **gram-cleanup** . Extra documentation will be added later...

Grammar cleanup and optimization

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

(gram-cleanup graph) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Returns a new graph containing no null transition and implementing the same grammar than **graph** .

Functions for building grammars

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

gram-or, gram-cat, gram-some, gram-many, gram-maybe (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

These functions allow for easily building grammars from regular expression style operators.

All these function take an arbitrary number of arguments which may be integer labels (using LAB macros), one character strings or complete grammars.

- **gram-or** returns a grammar which matches any of the alternatives represented by the labels or the grammars.
- **gram-cat** returns a grammar which matches the concatenation of the strings of labels represented by the labels or the grammars.
- **gram-some** returns a grammar which matches one or more occurrence of the concatenation of the strings of labels represented by its arguments.
- **gram-many** returns a grammar which matches zero or more occurrence of the concatenation of the strings of labels represented by its arguments.
- **gram-maybe** returns a grammar which matches zero or one occurrence of the concatenation of the strings of labels represented by its arguments.

The resulting grammar may contain null transitions whose label is provided by macro `LAB-NULL` . These transitions must be removed using function `gram-cleanup` .

```
(setq gr-digstring
      (gram-some (gram-or (LAB-ASCII "0")
                        (LAB-ASCII "1"))))
(setq gr-number
      (gram-cat gr-digstring
                (gram-maybe "." gr-digstring) ) )
(setq grammar-for-numbers
      (gram-cleanup gr-number) )
```

gram2-or, **gram2-some** (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

These are optimizations of `gram-or` and `gram-some` devised by Patrick. Do not use before inquiring PH about limitations and potential headaches.

Setting the penalties/cost of edges.

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

(gram-penalty cost grammar) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Sets a cost penalty to all paths traversing a particular grammar. The cost will be added to the costs of all links arriving on the endnode. We assume that grammar `grammar` contains no null paths (i.e. paths between the startnode and the endnode with null transitions only).

converting grammar paths to strings

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

(gram-meaning meaning grammar) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Override the meaning fields in grammar `grammar` . Only one link in every path in grammar `grammar` will have a meaning equal to `meaning` . All other links will have a null meaning.

Note: The resulting grammar may be cleaned up using `gram-cleanup` .

conversion from regular grammar to matrix representation

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

(grammar-to-matrix mat) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

This function encodes a grammar into a matrix.

(matrix-to-grammar mat) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

This function decodes a grammar encoded with `grammar-to-matrix` .

(gram-static lispgrammar) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Use this macro in compiled code to refer to a grammar defined by a lisp expression `lispgrammar` evaluated at compile time. This macro expand into a call to `matrix-to-grammar` on a static matrix.

converting strings to label sequences

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

(string-to-labels s) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Returns a matrix of integers containing the ascii codes of string `s` .

(labels-to-string im) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Returns a string made with characters whose ascii codes are in matrix `im` . Illegal ascii codes are marked with string "[?]" .

testing label sequences against a grammar

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

(gram-match grammar labels) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Returns `t` if labels in `idx1 labels` match grammar `grammar` .

handling trivgraphs

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

(count-labels graphin) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

count the number of links with labels in a trivial graph

(count-meanings graphin) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

count the number of links with meanings in a trivial graph

(trivgraph-to-labels graphin) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

This function returns an integer vector containing value of slots **label** for each link in the single path of trivial graph **graphin** . This function returns a matrix with a single element equal to -1 when the graph contains no path.

(trivgraph-to-meanings graphin) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

This function returns an integer vector containing value of slots **meaning** for each link in the single path of trivial graph **graphin** . Links whose slot **meaning** contains (LAB-NULL) are discarded. This function returns a matrix with a single element equal to -1 when the graph contains no meaningful path.

(trivgraph-to-data graphin) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

This function returns a gptr vector containing value of slots **data** for each link in the single path of trivial graph **graphin** .

(trivgraph-to-costs graphin scores) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

((-obj- (GGraph)) graphin):

((-idx1- (-flt-)) scores):

CREATED: P.Haffner

COMPILABLE: Yes

DESCRIPTION: This function just copies the scores of the meaningful links of trivial graph <graph> into matrix <scores>. It does not apply the squashing function applied by <trivgraph-to-scores>.

(trivgraph-to-meaningful-costs graph scores) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

This function just copies the scores of the meaningful links of trivial graph **graph** into matrix **scores** . It does not apply the squashing function applied by **trivgraph-to-meaningful-costs** .

(trivgraph-to-scores graph scores) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

This function copies the squashed scores of the links of trivial graph **graph** into matrix **scores** .

(trivgraph-to-meaningful-scores graph scores) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

This function copies the squashed scores of the meaningful links of trivial graph **graph** into matrix **scores** .

Selecting path that produce a particular label sequence

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

(g-restrict-to-meanings graph meanings) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Creates a new graph which implements the subset of paths of graph **graph** whose links with non null meanings implement the meanings given in the integer matrix **meaning** . The links of the new graph have the same cost , label , meaning and data than the original links.

(This code is derived from **g-compose**)

(g-restrict-to-meanings-2 graph labels) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

See: (**g-restrict-to-meanings graph labels**)

This function is similat to **g-restrict-to-meanings** . However the data fields of the restricted graph point to the original link instead of to the data pointed to by the original link.

(g-restrict-to-labels graph labels) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Creates a new graph which implements the subset of paths of graph **graph** whose links implement the labels given in the integer matrix **meaning** . The links of the new graph have the same cost , label , meaning and data than the original links.

(This code is derived from **g-compose**)

(g-restrict-to-labels-2 graph labels) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

See: (**g-restrict-to-labels graph labels**)

This function is similat to **g-restrict-to-labels** . However the data fields of the restricted graph point to the original link instead of to the data pointed to by the original link.

additional grammar functions

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

(lab-to-str lab) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

((-int-) lab):

RETURNS: a string of 1 or 4 characters

CREATED: P.Haffner

COMPILABLE: Yes

DESCRIPTION: transcodes an integer into a string of one or 2 characters. (2 characters are returned between brackets)

See: (**str-to-lab s**)

(str-to-lab s) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

((-str-) s):

RETURNS: integer

CREATED: P.Haffner

COMPILABLE: Yes

DESCRIPTION: Transcodes a 1 or 2 character string into an integer

See: (lab-to-str lab)

(**gram-match-meaning** grammar meanings) (lsh/libgraph/grammar.lsh)

Last Modified: 2002-10-14 03:30:50

Returns **t** if meanings in `idx1` meanings match grammar `grammar` .

(**gram-match-string** grammar <list of strings>) (lsh/libgraph/grammar.lsh)

Last Modified: 2002-10-14 03:30:50

for each string in the list , check it it can be generated as meanings of the graph `grammar`

(**string-to-labels2** s) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Returns a matrix of integers containing the ascii codes of string `s` . Unlike `string-to-labels`, brackets `[xx]` means double ascii codes.

See: (string-to-labels s)

(**labels-to-string2** im) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

Returns a string made with characters whose ascii codes are in matrix `im` . Unlike `labels-to-string`, double ascii codes are quoted in brackets `[xx]`.

See: (labels-to-string im)

Graph Analysis

(lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

(**g-complexity** graph) (lsh/libgraph/grammar.lsh) Last Modified: 2002-10-14 03:30:50

```
((-obj- (GGraph)) graph):
CREATED: P.Haffner
COMPILABLE: Yes
DESCRIPTION: display useful information about a graph
- number of nodes
- number of links
- number of links for different labels
- perplexity.
```

8.6.2 Manipulating Graphs.

(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Gbag class

(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

a useful Lisp class to handle dynamically allocated data structures.

(**new Gbag**) Type: CLASS(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Allocates a new collection of GPTRs. The housekeeping data structures will be allocated on pool specified with method `add` . The following messages are supported:

See: (`==>` `gbag add pool gptr`)

See: (`==>` `gbag rem gptr`)

See: (`==>` `gbag test gptr`)

See: (`==>` `gbag size`)

(`==>` `gbag size`) Type: MSG(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Returns the number of elements in the bag `gbag` .

(`==>` `gbag add pool gptr`) Type: MSG(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Adds a new element GPTR in the bag `gbag` . Memory will be allocated from pool `pool` . It is safer to always use the same pool with a given bag. This pool , as well as the data pointed to by `gptr` *must* outlive the bag.

(`==>` `gbag test gptr`) Type: MSG(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Tests if element `gptr` belongs to bag `gbag` .

(`==>` `gbag rem gptr`) Type: MSG(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Remove all occurrences of element `gptr` in the bag `gbag` . This function also resets the iterator.

(`==>` `gbag empty`) Type: MSG(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Removes all elements in a bag

(`gbagelt-peek-data gptr`) (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Peek data in a gbag element.

(`gbagelt-peek-next gptr`) (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Peek next gbag element.

(`all-bag (var bag) ...prog...`) (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Iterate `prog` with `var` taking all values in `bag` .

GGraph class

(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

(`new GGraph`) Type: CLASS(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Class `GGraph` defines the functions for handling graphs. Graphs are handled via two other classes of objects.

- Class `GNode` represent nodes of the graph.
- Class `GLink` represent arcs of the graph.

- :graph:apool** (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56
 Slot **apool** in a graph contains the pool used for storing all objects depending on the graph. This pool is created when you create the graph object. You may use this pool to allocate objects stored in the **data** slots of the nodes and the links.
- :graph:nnodes** (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56
 Slot **nnodes** in a graph contains the number of nodes allocated in a graph.
- :graph:startnode** (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56
 Slot **startnode** in a graph contains the start node of the graph. This node is created when the graph is created.
- :graph:endnode** (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56
 Slot **endnode** in a graph contains the final node of the graph. This node is created when the graph is created.
- :node:id** (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56
 Slot **id** in a node contains the ordinal number of the node. This number was allocated when you created the node. You should not change it.
- :node:data** (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56
 Slot **data** of a node contains an arbitrary GPTR. This GPTR is not used by the graph library and is available for you. You must of course take the usual precautions with GPTRs.
- :link:upnode** (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56
 Slot **downnode** of a link contains the upstream node of the link.
- :link:downnode** (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56
 Slot **downnode** of a link contains the downstream node of the link.
- :link:cost** (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56
 Slot **cost** of a link is a floating point number for storing the cost associated with a link. Costs are real number (both positive and negative). The smaller the cost the better the link. You can view these costs as $-\log(\text{probability})$.
- :link:dcost** (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56
 Slot **dcost** of a link is a floating point number intended for storing the derivative of the global cost with respect to the contents of the slot **cost**. This slot is set by functions **g-clear-gradients** and **g-forward-backward**.
- :link:label** (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56
 Slot **label** of a link contains an arbitrary integer. This integer is not used by the graph library and is available for you.
- :link:data** (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56
 Slot **data** of a link contains an arbitrary GPTR. This GPTR is not used by the graph library and is available for you. You must of course take the usual precautions with GPTRs.

GNode class

(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

GLink class

(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

(all-nodes (var graph) ...prog...) (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Iterate **prog** with **var** being all nodes in graph **graph** .

(all-downlink (var node) ...prog...) (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Iterate **prog** with **var** being all links downstream of **node** .

(all-uplink (var node) ...prog...) (lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Iterate **prog** with **var** being all links upstream of **node** .

(==> graph newnode data) Type: MSG(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Creates a new node in the graph. The new node points to data GPTR **data** . All data pointed to by **data** MUST outlive the graph. You can for instance allocate this data on the pool **:graph:apool** .

(==> graph link node1 node2 cost label data) Type: MSG(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Creates a new arc in the graph between nodes **node1** and **node2** . The new arc is assigned a cost **cost** , an integer interpretation **label** and the data pointed to by the GPTR **data** . All data pointed to by **data** MUST outlive the graph. You can for instance allocate this data on the pool **:graph:apool** .

(==> graph linkm node1 node2 cost label meaning data) Type: MSG(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

See: **(==> graph link node1 node2 cost label data)**

A link creation function which also sets the **meaning** field. The default creation function store the same value in both the **label** and **meaning** fields

(==> graph cut link) Type: MSG(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Removes arc **link** from the graph. No memory is reclaimed. Links may be cut in a frozen graph.

(==> graph freeze) Type: MSG(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Forbids any further alteration of the graph.

(==> graph dump) Type: MSG(lsh/libgraph/graph.lsh) Last Modified: 2002-08-31 00:52:56

Dumps a graph on the standard output.

8.6.3 Searching and Pruning Graphs

(lsh/libgraph/gsearch.lsh) Last Modified: 2002-08-31 00:52:56

searching the shortest path in a simple graph or a composed graph using Viterbi and such.

(g-viterbi graph)

(lsh/libgraph/gsearch.lsh) Last Modified: 2002-08-31 00:52:56

Applies viterbi algorithm to compute a graph whose single path is a copy of the best path of graph **graph** . Graph **graph** should not have cycles.

(g-compose-viterbi graph1 graph2 trans)

(lsh/libgraph/gsearch.lsh) Last Modified: 2002-08-31 00:52:56

Applies viterbi algorithm to compute the cheapest path in the graph obtained by composing **graph1** and **graph2** by transformer **trans** . This routine returns a new graph containing only this path.

(g-forward graph)

(lsh/libgraph/gsearch.lsh) Last Modified: 2002-08-31 00:52:56

Returns the log-added costs of all paths in lattice **graph** .

REMARK: The graph should not have cycles.

(g-compose-forward agraph bgraph trans)

(lsh/libgraph/gsearch.lsh) Last Modified: 2002-08-31 00:52:56

This function returns the log-added costs of all paths in the lattice obtained by composing **agraph** and **bgraph** with transformer **gtrans** .

REMARK: The composed graph should not have cycles.

(g-clear-gradients graph)

(lsh/libgraph/gsearch.lsh) Last Modified: 2002-08-31 00:52:56

This function clears the gradients **dcost** and **dacost** in all reachable nodes and links of the graph **graph** .

(g-backward gradin graph)

(lsh/libgraph/gsearch.lsh) Last Modified: 2002-08-31 00:52:56

When called immediatly after **g-forward** , this function credits a change **gradin** of the result of forward to the cost of each link. The result is added into the slots **dcost** of each reachable link.

REMARK: The graph should not have cycles.

8.6.4 Transformation Operators on Graphs

(lsh/libgraph/gtransform.lsh) Last Modified: 2002-08-31 00:52:56

(new GTrans)

Type: CLASS(lsh/libgraph/gtransform.lsh) Last Modified: 2002-08-31 00:52:56

A transformer is the object which embodies the rules used for creating a new graph from one or two graphs using `g-transform` or `g-compose` . Class `GTrans` is the base class for all transformers.

- Method `check` is used for selecting links in the initial graph(s) for which structures will be created between two nodes of the destination graph.
- Method `vcost` and `fcost` compute the cost of these structures without actually building them. Method `vcost` (Viterbi's cost) computes the cost of the best path. Method `fcost` (Forward's cost) computes the logadded cost of all paths.
- Method `build` build a structure in the resulting graph on the basis of the information provided by the corresponding links in the parent graphs. This involves computing cost, label and data.
- Method `bprop` is used during training to update the gradient with respect to the cost of the initial graph. The implementation of this method requires that the transformer remembers the association between the structures created by `build` and the initial links. This memory is cleaned by method `reset` .

The base class `GTrans` accepts all links of the product graph. The cost in the link of the product graph will be the sum of the costs of the corresponding links in both graphs. The label and data component are copied from the second and first graph respectively.

```
See: (==> gtrans check link1 link2 )
See: (==> gtrans vcost link1 link2 )
See: (==> gtrans fcost link1 link2 )
See: (==> gtrans build ngraph nnodefrom nnode to link1 link2 )
See: (==> gtrans remember rememberp )
See: (==> gtrans reset)
See: (==> gtrans bprop)
```

(==> gtrans check link1 link2)

Type: MSG(lsh/libgraph/gtransform.lsh) Last Modified: 2002-08-31 00:52:56

Tells `g-transform` or `g-compose` if they must create a structure in the new graph corresponding to the given links of the parent graphs. A structure may be as simple as a single link or as complex as a subgraph.

- In the case of `g-transform` , arguments `link1` and `link2` are equal and identify a link of the initial graph. This function returns `t` if a new structure corresponding to this link must be created in the new graph

- In the case of `g-compose`, argument `link1` and `link2` identify two links in the two initial graph. This function returns `t` if a new structure corresponding to these links must be created in the new graph

`(==> gtrans vcost link1 link2)`

Type: MSG(lsh/libgraph/gtransform.lsh) Last Modified: 2002-08-31 00:52:56

Given two links `link1` and `link2` in two graphs, this method returns a floating point number indicating the cost of the best path in the structure corresponding to links `link1` and `link2`.

The default function just adds the costs of both links.

`(==> gtrans fcost link1 link2)`

Type: MSG(lsh/libgraph/gtransform.lsh) Last Modified: 2002-08-31 00:52:56

Given two links `link1` and `link2` in two graphs, this method returns a floating point number indicating the loadadded cost of all paths in the structure corresponding to links `link1` and `link2`.

The default function just calls `vcost`.

`(==> gtrans build ngraph nnodefrom nnode to link1 link2)`

Type: MSG(lsh/libgraph/gtransform.lsh) Last Modified: 2002-08-31 00:52:56

Given to links `link1` and `link2` in two graphs (for which method `check` indicates that we should build something in the resulting graph), method `build` creates a new structure between nodes `nnodefrom` and `nnode to` in graph `ngraph`. This task involves computing cost, label and data for the new links.

The default function just builds a single new link whose label comes from `link2` and whose data slot comes from `link1`.

`(==> gtrans bprop)`

Type: MSG(lsh/libgraph/gtransform.lsh) Last Modified: 2002-08-31 00:52:56

Propagates the gradients of the links of the destination graph into the the corresponding links of the initial graph. The resulting gradients are accumulated into the fields `dcost` of the corresponding links.

If running the transformer involves parametric operations (e.g. running a network), the gradients are also propagated along these operations and finally accumulated into the parameter vector.

To achieve this task (optional), the transformer must remember all structures build since its creation.

`(==> gtrans remember rememberp)`

Type: MSG(lsh/libgraph/gtransform.lsh) Last Modified: 2002-08-31 00:52:56

Sets or resets memoryless mode. Argument **rememberp** is a boolean specifying if it is useful to remember the result of the next calls to the method **build** of the transformer.

(==> **gtrans reset**)

Type: MSG(lsh/libgraph/gtransform.lsh) Last Modified: 2002-08-31 00:52:56

Resets the data remembered by the transformer for a possible call of method **bprop** .

(**g-copy graph ngraph nstart nend**)

(lsh/libgraph/gtransform.lsh) Last Modified: 2002-08-31 00:52:56

Copies graph **graph** between nodes **nstart** and **nend** of a new graph **ngraph** . This function only considers the nodes and links that we can reach from both the starting node and the end node. The other nodes and link are not copied.

WARNING: This function merely copies the **data** pointers. The corresponding data are not copied into the new pool. You can change these data pointer within a new pass.

(**g-transform graph trans**)

(lsh/libgraph/gtransform.lsh) Last Modified: 2002-08-31 00:52:56

Transforms graph **graph** according to transformer **trans** . The resulting graph is build by calling method **build** of the transformer for all links in the initial graph for which method **check** returns true.

This function ideally should only builds the nodes and the links which are reachable from both the startnode and the endnode. Yet function **g-transform** may build extra nodes and links if the final graph has cycle(s) or if the transformer method **build** builds unnecessary nodes.

See: (new GTrans)

(**g-compose graph1 graph2 trans**)

(lsh/libgraph/gtransform.lsh) Last Modified: 2002-08-31 00:52:56

Composes graphs **graph1** and **graph2** according to transformer **trans** . The resulting graph is build by calling method **build** of the transformer between all nodes of the product graph for which method **check** returns true.

This function ideally should only builds the nodes and the links which are reachable from both the startnode and the endnode. Yet function **g-compose** may build extra nodes and links if the final graph has cycle(s) or if the transformer method **build** builds unnecessary nodes.

See: (new GTrans)

8.6.5 Other Graph Search Functions

(lsh/libgraph/rsearch.lsh) Last Modified: 2002-08-31 00:52:56

algorithms for searching a graph. the difference with the g-viterbi and g-forward is that the seqrch is carried from the initial (upper) nodes. besides , a startnode must be specified.

(r-viterbi graph startnode)

(lsh/libgraph/rsearch.lsh) Last Modified: 2002-08-31 00:52:56

Applies viterbi algorithm to compute a graph whose single path is a copy of the best path of graph **graph** . Graph **graph** should not have cycles. different from g-viterbi in 2 respects - reverse order - a startnode can be specified

(r-forward graph startnode)

(lsh/libgraph/rsearch.lsh) Last Modified: 2002-08-31 00:52:56

Returns the lor-added costs of all paths in lattice **graph** . different from g-forward in 2 respects - reverse order - a startnode can be specified

REMARK: The graph should not have cycles.

(r-restrict-to-meanings graph meanings startnode)

(lsh/libgraph/rsearch.lsh) Last Modified: 2002-08-31 00:52:56

Creates a new graph which implements the subset of paths of graph **graph** whose links with non null meanings implement the meanings given in the integer matrix **meaning** . The links of the new graph have the same cost , label and meaning than the original links. The data field points to the original link.

(This code is derived from **g-compose**)

8.7 Image Processing Libraries

8.7.1 Demos and examples

demo that loads and displays a sample image

(lsh/libimage/demos/show-image.lsh) Last Modified: 2002-08-16 18:44:07

8.7.2 Reading/Writing Image Files to/from IDX.

(lsh/libimage/image-io.lsh) Last Modified: 2003-01-07 03:35:12

Functions to read/write image files into/from an IDX. Image files can be in any format that ImageMagick's convert command can handle. The type of the input image is determined automatically.

(image-read-rgb f)

(lsh/libimage/image-io.lsh) Last Modified: 2003-01-07 03:35:12

read an image file (in any format), and return an rgbimage with the image in it (IDX3 of ubytes with last dimension=3). Accepted format are whatever ImageMagick's convert command can convert to PPM. The type of the input image is determined automatically. In case it is not, the the type can be prepended to the filename with a colon. Command line options prepended to the argument are passed to convert.

Example:

```
(setq img (image-read-rgb "myimage.png"))
(setq img (image-read-rgb "TIFF:myimage.dunno"))
(setq img (image-read-rgb "-geometry 50x50\\% myimage.jpg"))
```

(image-read-rgba f)

(lsh/libimage/image-io.lsh) Last Modified: 2003-01-07 03:35:12

read an image file (in any format), and return an rgbaimage with the image in it (IDX3 of ubytes with last dimension=4). The alpha channel is filled if the image being read has one. An alpha value of 0 means opaque and 255 means fully transparent (this is backward to many conventions in other places in the Lush library). Accepted format are whatever ImageMagick's convert command can convert to PPM. The type of the input image is determined automatically. In case it is not, the the type can be prepended to the filename with a colon. Command line options prepended to the argument are passed to convert.

Example:

```
(setq img (image-read-rgba "-geometry 50x50\\% myimage.png"))
(setq img (image-read-rgba "-geometry 50x50\\% PNG:myimage"))
```

(image-read-ubim f)

(lsh/libimage/image-io.lsh) Last Modified: 2003-01-07 03:35:12

read an image file (in any format) , and return a ubimage with the image in it (IDX2 of ubytes) Accepted format are whatever ImageMagick's convert command can convert to PPM. Conversion to grayscale is performed by Lush , since convert doesn't convert to PGM.

(image-write-rgb f im)

(lsh/libimage/image-io.lsh) Last Modified: 2003-01-07 03:35:12

Writes an image stored in an IDX3 of ubytes into an image file (in any format). The type of the output image file is determined by a string prepended to the filename with a colon. Command line options prepended to the argument are passed to the convert command

Example:

```
(image-write-rgb "TIFF:myimage.tiff"))
```

8.7.3 Reading/Writing PPM/PGM/PBM Image Files

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

Routines for reading and writing PBM/PGM/PPM image files into/from idx. Another set of functions built on top of these is used to read/write image files in other formats than PPM/PGM/PBM.

(pnm-header f)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

```
((-gpPtr- "FILE *") f):
```

Return the descriptive string of a PNM image file (P6, P4, or P5), and skip the file descriptor past the comments.

(mmap-idx3-storage f u d0 d1 d2)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

```
((-gpPtr- "FILE *") f):
```

```
((-idx3- (-ubyte-)) u):
```

```
((-int-) d0 d1 d2):
```

Maps file f (at the current position given by FTELL) into idx3 u does not check that f size is OK.

(mmap-idx2-storage f u d0 d1)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

```
((-gpPtr- "FILE *") f):
```

```
((-idx2- (-ubyte-)) u):
```

```
((-int-) d0 d1):
```

Maps file f (at the current position given by FTELL) into idx2 u does not check that f size is OK.

(pnm-fread-into-rgbx f out)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

((-gp_{ptr}- "FILE *") f):

read an image from a PBM/PGM/PPM file into an `idx3` of ubytes (RGB or RGBA). `f` must be a valid file descriptor (C pointer). `out` is appropriately resized if required. The last dimension is left unchanged but must be at least 3. Appropriate conversion is performed. extra color components (beyond the first 3) are left untouched.

See: `pnm`

(pnm-read-into-rgbx f out)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

read a PBM/PGM/PPM file into an `idx3` of ubytes. The output `idx` is appropriately resized if necessary. The last dimension must be at least 3.

(pnm-fread-rgb f)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

((-gp_{ptr}- "FILE *") f) `f` must be an open file descriptor. read a PBM/PGM/PPM file and return an `idx3` of ubytes with the image in it. Conversion to RGB is performed if necessary.

(pnm-read-rgb s)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

read a PBM/PGM/PPM file and return an `idx3` of ubytes with the image in it. Conversion to RGB is performed if necessary.

(pnm-fread-rgba f)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

((-gp_{ptr}- "FILE *") f) `f` must be an open file descriptor. read a PBM/PGM/PPM file and return an `idx3` of ubytes with the image in it. Conversion to RGB is performed if necessary. The alpha channel is set to 0.

(pnm-read-rgba s)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

read a PBM/PGM/PPM file and return an `idx3` of ubytes with the image in it. Conversion to RGBA is performed if necessary. The alpha channel is set to 0.

(pnm-fread-into-ubim f out)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

```
((-gptr- "FILE *") f):
((-idx2- (-ubyte-)) out):
```

read an image from a PBM/PGM/PPM file into an idx2 of ubytes (one byte per pixel). **f** must be a valid file descriptor (C pointer). **out** is appropriately resized if required. Appropriate conversion is performed (e.g. if the file contains an RGB image).

(pnm-read-into-ubim f out)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

read a PBM/PGM/PPM file into an idx2 of ubytes (greyscale image). The output idx is appropriately resized if necessary. Appropriate conversions to greyscale are performed

(pnm-fread-ubim f)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

((-gptr- "FILE *") f) **f** must be an open file descriptor. read a PBM/PGM/PPM file and return an idx2 of ubytes with the image in it. Conversion to greyscale is performed if necessary.

(pnm-read-ubim f)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

read a PBM/PGM/PPM file and return an idx2 of ubytes with the image in it. Conversion to greyscale is performed if necessary.

(ppm-mmap f)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

```
((-gptr- "FILE *") f):
maps a ppm RGB image (file descriptor <f>)
into memory and return an idx3 of ubytes
with the data.
```

(pgm-mmap f)

(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12

```
((-gptr- "FILE *") f):
maps a pgm image (file descriptor <f>)
into memory and return an idx2 of ubytes
with the data.
```

(ppm-fwrite-rgb f m)

```
(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12
writes RGB image m in file descriptor f as a PPM/PPM file. m is and idx3
of ubytes. f must be a file descriptor.
```

(ppm-write-rgb f m)

```
(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12
writes RGB image m in file f as a PPM/PPM file. m is and idx3 of ubytes.
```

(pgm-fwrite-ubim f m)

```
(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12
writes an ubimage m in file f as a PGM file. m is and idx2 of ubytes , and
f is an open file descriptor.
```

(pgm-write-ubim f m)

```
(lsh/libimage/pnm.lsh) Last Modified: 2002-12-23 09:57:12
writes greyscale ubimage m in file f as a PGM file. m is and idx2 of ubytes.
```

8.7.4 Reading/Writing PBM Image files

```
(lsh/libimage/pbm.lsh) Last Modified: 2002-08-19 03:01:28
reading and writing PBM image files (one bit per pixel , bitonal images)
from and into lists of runs , and idx.
```

(pbm-write-runs f runs nlin ncol)

```
(lsh/libimage/pbm.lsh) Last Modified: 2002-08-19 03:01:28
```

```
((-str-) f):
((-idx2- (-int-)) runs):
((-int-) nlin ncol):
CREATED: P. Haffner, Oct 97
DESCRIPTION:
writes <runs> in file <f> as a PBM file.
```

(pbm-write-ubim f m)

(lsh/libimage/pbm.lsh) Last Modified: 2002-08-19 03:01:28
 writes binary image *m* in file *f* as a PBM file. *m* is and *idx2* of ubytes.

(pbm-write-ubim-val f m v)

(lsh/libimage/pbm.lsh) Last Modified: 2002-08-19 03:01:28
 writes binary image *m* in file *f* as a PBM file. *m* is and *idx2* of ubytes. *mask*
 (value on) all values *v*

(pbm-read-ubim f white black)

(lsh/libimage/pbm.lsh) Last Modified: 2002-08-19 03:01:28
 read binary image *m* from PBM file *f* . code black as *black* and white as
white

8.7.5 Reading/Writing Run-Length-Encoded Image Files

(lsh/libimage/rle.lsh) Last Modified: 2002-09-06 18:27:25
 read and write RLE images. RLE images are a compact way of representing
 binary images by a list of "runs" of black pixels.

(rle-read-ubim f white black)

(lsh/libimage/rle.lsh) Last Modified: 2002-09-06 18:27:25
 read binary image *m* from RLE file *f* . code black as *black* and white as
white

(rle-read-runs f runs img-dim)

(lsh/libimage/rle.lsh) Last Modified: 2002-09-06 18:27:25

```
((-str-) f):
((-idx2- (-int-)) runs):
((-obj- (img-dim-C)) img-dim):
CREATED: P. Haffner, Feb 98
MODIFIES: <runs> and <img-dim>
RETURNS: ()
DESCRIPTION:
read runs IDX <runs> from RLE file <f>.
dimension 0 of <runs> is the number of runs
dimension 1 of <runs> must be >= 3 with
0: Y position
1: X1 position
2: X2 position
populates object img-dim with both dimensions
```

See: `img-dim-C`

EXAMPLE: `(setq runs-dim (new img-dim-C 0 0)) (setq runs (int-matrix 1 4))` (a 4th optional dimension for connected componanet analysis)
`(rle-read-runs filename runs runs-dim)`

(rle-write-runs f runs nlin ncol)

(lsh/libimage/rle.lsh) Last Modified: 2002-09-06 18:27:25

```
((-str-) f):
((-idx2- (-int-)) runs):
((-int-) nlin ncol):
CREATED: P. Haffner, Oct 97
DESCRIPTION:
writes <runs> in file <f> as a RLE file.
```

8.7.6 RGBA Images of ubytes

(lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

basic image processing on RGBA images where each pixel component is an unsigned byte.

rgbaimage

(lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

a `rgbaimage` is an `idx3` of ubytes intended to store images. The last dimension is 4, and contains the R, G, B, and Alpha channels respectively.

Geometric Transforms

(lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

A few functions to enlarge, reduce, resize, and rotate images.

(rgbaim-resize im width height mode) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

resize an `rgba` image to any size using bilinear interpolation. Appropriate local averaging (smoothing) is performed for scaling factors smaller than 0.5. If one of the desired dimensions is 0, an aspect-ratio-preserving scaling is performed on the basis of the other dimension. When both `width` and `height` are non zero, the last parameter, `mode` determines how they are interpreted.

- if either `width` or `height` is zero, `mode` is ignored.
- `mode=0`: fit the image into a `width` by `height` frame while preserving the aspect ratio
- `mode=1`: scale the image to `width` by `height` possibly changing the aspect ratio

- `mode=2`: `width` and `height` are interpreted as scaling ratios

The sizes of the output image are rounded to nearest integers smaller than the computed sizes, or to 1, whichever is largest.

(`rgbaim-enlarge in nlin ncol`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

enlarges image `in` with integer ratios `nlin` (vertical) `ncol` (horizontal). the enlarged image is returned. The horizontal (resp. vertical) size of the returned image is equal to the integer part of the horizontal (resp vertical) size of `in` divided by `ncol` (resp `nlin`). returns (copy-matrix `in`) when enlarge rate is 1

(`rgbaim-enlarge-into in out nlin ncol`) (lsh/libimage/rgbaimage.lsh)
Last Modified: 2002-10-14 03:30:51

enlarges image `in` with integer ratios `nlin` (vertical) `ncol` (horizontal) and write result in `out` . The horizontal (resp. vertical) size of the output image is equal to the integer part of the horizontal (resp vertical) size of `in` divided by `ncol` (resp `nlin`).

(`rgbaim-subsample in nlin ncol`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

subsamples image `in` with integer ratios `nlin` (vertical) `ncol` (horizontal). the subsampled image is returned. The horizontal (resp. vertical) size of the returned image is equal to the integer part of the horizontal (resp vertical) size of `in` divided by `ncol` (resp `nlin`). returns (copy-matrix `in`) when subsample rate is 1

(`rgbaim-subsample-into in out nlin ncol`) (lsh/libimage/rgbaimage.lsh)
Last Modified: 2002-10-14 03:30:51

subsamples image `in` with integer ratios `nlin` (vertical) `ncol` (horizontal) and write result into `out` . The horizontal (resp. vertical) size of the output image is equal to the integer part of the horizontal (resp vertical) size of `in` divided by `ncol` (resp `nlin`).

(`rgbaim-rot90-left inp`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

rotate image `inp` by 90 to the left (counter-clockwise) and return the result.

(`rgbaim-rot90-right inp`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

rotate image `inp` by 90 to the right (clockwise) and return the result.

(`rgbaim-rot180 inp`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

rotate image `inp` by 180 degrees and return the result.

(`rgbaim-rot180-inplace in`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

in-place 180 degree rotation of an image (caution: input image is destroyed).

(`rgbaim-crop in x y w h`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

crop rectangle (`x` , `y` , `w` , `h`) from image `in` and return the result (a copy).

(rgbaim-warp-quad in out background mode x1 y1 x2 y2 x3 y3 x4 y4 p1 q1 p3 q3) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

```
((-flt-) x1 y1 x2 y2 x3 y3 x4 y4 p1 q1 p3 q3):
((-int-) background mode):
((-idx2- (-ubyte-)) in out):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
warp rgbaimage <in> through a geometric transformation that maps
a quadrilateral into a rectangle (bilinear transform).
The quadrilateral is specified by <x1> <y1> <x2> <y2> <x3> <y3> <x4> <y4>
(points are numbered clockwise starting from upper left).
The rectangle is specified by the upper left and lower right points
which are respectively <p1><q1> and <p3><q3>.
The result is put in <out>. Clipping is automatically performed.
pixels outside of the bounds of <in> are assumed to have the value
<background> (which must be an idx1 of ubyte of size 4).
<mode>=0: no antialiasing, <mode>=1: antialiasing with
bilinear interpolation (2 times slower).
```

(rgbaim-warp in out background pi pj) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

```
((-idx1- (-ubyte)) background):
((-idx3- (-ubyte-)) in out):
((-idx2- (-int-)) pi pj):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
Warpes an image using <pi> and <pj> as tabulated coordinate transforms.
<in> and <out> are idx2 of ubytes. <background> is the value assumed outside
of the input image. <pi> and <pj> are integers idx whose dimensions must
be identical to the first two dimensions of <out>.
An output pixel at coordinate (x,y) takes the value of the input pixel at
coordinate <pj[x,y]>,<pi[x,y]>. If the values in <pi> and <pj> are non
integers the pixels are computed by bilinearly interpolating the input
image. The <pi> and <pj> matrices contain 32-bit integers which are
interpreted as 16 bit of integer part and 16 bits of fractinal part.
Integers values are assumed to fall in the center of each pixel,
so the upper left-hand corner of an image is at coordinate (-0.5, -0.5).
```

(*rgbaim-warp-fast* in out background pi pj) (lsh/libimage/rgbaimage.lsh)
 Last Modified: 2002-10-14 03:30:51

```
((-int-) background):
((-idx2- (-ubyte-)) in out):
((-idx2- (-int-)) pi pj):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: Y. LeCun
Warps an image using <pi> and <pj> as tabulated coordinate transforms.
<in> and <out> are idx2 of ubytes. <background> is the value assumed outside
of the input image. <pi> and <pj> are tabulated coordinates
This is essentially identical to rgbaim-warp, except no bilinear
interpolation is performed (only a nearest neighbor rule is used).
```

(*rgbaim-rotscale* src sx sy dst dx dy angle coeff bg) (lsh/libimage/rgbaimage.lsh)
 Last Modified: 2002-10-14 03:30:51

rotate, scale, and translate image *src* and put result in *dst* . point *sx* ,
sy in *src* will be mapped to point *dx* , *dy* in *dst* . Image will be rotated
 clockwise by *angle* degrees and scaled by *coeff* . Pixels that fall off the
 boundary are clipped and pixels in the destination that are not determined by a
 source pixel are set to color *bg* (which must be an idx1 of ubytes of size 4). It
 is generally preferable to call *rgbaim-rotscale-rect* before hand to get appropriate
 values for *dx* , *dy* and for the size of *dst* so that no pixel is clipped. Here is
 an example:

```
(let* ((w (idx-dim m 1))
      (h (idx-dim m 0))
      (wh (int-matrix 2))
      (cxcy (float-matrix 2))
      (bg (ubyte-matrix 4)))
  (rgbaim-rotscale-rect w h hotx-src hoty-src angle coeff wh cxcy)
  (let ((z (ubyte-matrix (wh 1) (wh 0) 4)))
    (rgbaim-rotscale m hotx-src hoty-src z (cxcy 0) (cxcy 1) angle coeff bg)))
```

(*rgbaim-rotscale-rect* w h cx cy angle coeff wh cxcy) (lsh/libimage/rgbaimage.lsh)
 Last Modified: 2002-10-14 03:30:51

Given an input image of width *w* , height *h* , with a "hot" point at coordi-
 nate *cx* and *cy* , this function computes the width , height , and hot point
 coordinates of the same image rotated by *angle* and scaled by *coeff* so as to
 ensure that no pixel in the rotated/scaled image will have negative coordinates
 (so the image will not be clipped). *wh* and *cxcy* which must be idx1 of floats
 with two elements. This function should be called before *rgbaim-rotscale*.

low-level geometric transform functions (lsh/libimage/rgbaimage.lsh)
 Last Modified: 2002-10-14 03:30:51

These are unlikely to be used directly by most users.

(*rgbaim-interpolate-bilin* background pin indimi indimj inmodi inmodj ppi ppj rez) (lsh/libimage/rgbaimage.lsh)

Author(s): Y. LeCun Last Modified: 2002-10-14 03:30:51

```
(-gptr- pin rez background):
(-int- indimi indimj inmodi inmodj ppi ppj):
SIDE EFFECTS: <rez>
DESCRIPTION:
returns a bilinearly interpolated RGBA pixel value for coordinate
<ppi> <ppj>. The image data is pointed to by <pin>, with
<indimi> <indimj> <inmodi> <inmodj> being the dimensions and modulus.
This function clips automatically if <ppi> <ppj> are outside of the
bounds by assuming the outside values are equal to <background>.
pixel values are ubytes, while coordinates are 32 bit fixed point
with 16 bit integer part and 16 bit fractional part.
The function does not use floating point arithmetics.
```

Color Processing

(lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

(*rgbaim-contbriht* in out c b) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

correct contrast and brightness of image *in* and put result in *out* (*in* and *out* must be the same size). *c* and *b* are floats that will affect the brightness and contrast respectively. *c* =1 and *b* =0 leave the image unchanged. Each output pixel is computed as follows: $\text{output} = (\text{input} - 128) * c + b + 128$ (clipped between 0 and 255). An image can be conveniently inverted by setting *c* to -1 and *b* to 0.

(*rgbaim-luminance* in out) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

compute luminance of each pixel and put result in image *out* the formula used is $\text{lum} = 0.299 * R + 0.587 * G + 0.114 * B$

(*rgbaim-lum2rgba* in out) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

transform a grey-level image *in* (idx2 of flts) into an RGBA image. *in* and *out* must have identical sizes , though no check of that is done

(*rgbaim-uvw* in out) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

transform an RGBA image into an UVWA image , where the V component contains the luminance information , and the U and W components collectively contain the chrominance information. In the UVW color coding scheme the euclidean distance is meant to reflect the subjective "perceptual" distance. See Digital Image Processing by W. Pratt , page 66.

(`rgbaim-saturate in min max out`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

Saturate pixel component values in image `in` to the range [`min` , `max`]. Result is put in `out` .

Blitting

(lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

(`rgbaim-blit x y in out`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

blit `in` into `out` at position `x` , `y` with alpha blending. an alpha value of 0 in `in` means opaque , 255 means totally transparent.

(`rgbaim-blitcolor x y in r g b out`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

blit color defined by `r` , `g` , `b` using `idx2 in` as a stencil into RGBA image `out` at position `x y` The values in `in` must be between 0 and 255. if the value is 255 , no paint is put in; if the value is 0 , the pixel takes the value specified by `r` , `g` , `b` .

Histograms

(lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

(`rgbaim-histo32 rgbaim ppal count`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

compute histogram of image `rgbaim` . `ppal` must be a 32768 by 3 matrix of floats which will contain the list of pixel values found. `count` must be a 32768 matrix which will contain the corresponding pixel counts. pixel colors are sorted by luminance. This function really computes a 32768 color palette by assigning each pixel to the cube it belongs to in the 32x32x32 RGB color cube. The color prototype assigned to each cube is the mean of the pixels belonging to the cube.

low-level histogram functions (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

(`rgbaim-init-lum2rgb`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

build a table of 32x32x32 elements containing RGB cell indices in ascending luminosity.

(`rgbaim-lum2rgb n`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

converts a luminosity index to a 5-bit RGB cell number.

Various unusual subsampling functions

(lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

(`rgbaim-subsample+ in nlin ncol`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

same as `rgbaim-subsample` , but adds one row and one column for lazy pixels

(`rgbaim-subsample-med3 rgbaim`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

subsample images by 3x3 using a median filter.

(`rgbaim-med3-subopt rgbaim`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

3x3 subsampling using the median smart median optimised in C for 3x3 squares: 1.3 sec

(`rgbaim-med3-slow rgbaim`) (lsh/libimage/rgbaimage.lsh) Last Modified: 2002-10-14 03:30:51

3x3 subsampling using the median brute force median , with a 9 element quicksort

8.7.7 RGBA Images of Floats

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

basic image processing on RGBA images where each pixel component is a float.

rgbafimage

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

a `rgbafimage` is an `idx3` of `flts` intended to store images. The last dimension is 4 , and contains the R , G , B , and Alpha channels respectively.

(`rgbafim-subsample in nlin ncol`)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

subsamples image `in` with integer ratios `nlin` (vertical) `ncol` (horizontal). the subsampled image is returned. The horizontal (resp. vertical) size of the returned image is equal to the integer part of the horizontal (resp vertical) size of `in` divided by `ncol` (resp `nlin`).

(`rgbafim-luminance in out`)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

compute luminance of each pixel and put result in image `out` the formula used is $\text{lum} = 0.299 * R + 0.587 * G + 0.114 * B$

(`rgbafim-lum2rgba in out`)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

transform a grey-level image `in` (`idx2` of `flts`) into an RGBA image. `in` and `out` must have identical sizes , though no check of that is done

(rgbafim-chromin-broken in z k out)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

attempts to compute the chrominance info in **in** (normalized colors with almost identical luminance) and put the result in **out** . Both **in** and **out** are idx3 of flts. **z** is a multiplier for the average pixel component value. **k** is a fudge factor that prevents the normalisation from blowing up if the color is black or almost black.

(rgbafim-chrominance in z k out)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

attempts to compute the chrominance info in **in** (normalized colors with almost identical luminance) and put the result in **out** . Both **in** and **out** are idx3 of flts. **z** is a multiplier for the average pixel component value. **k** is a fudge factor that prevents the normalisation from blowing up if the color is black or almost black.

(rgbafim-uvw in out)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

transform an RGBA image into an UVWA image , where the V component contains the luminance information , and the U and W components collectively contain the chrominance information. In the UVW color coding scheme the euclidean distance is meant to reflect the subjective "perceptual" distance. See Digital Image Processing by W. Pratt , page 66.

(rgbafim-fromuvw in out)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

converts back from UVWA to RGBA. It is advasable to saturate the resulting image , as numerical errors may cause the pixel values to exceed the range 0-255.

(rgbafim-saturate in min max out)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

Saturate pixel component values in image **in** to the range [**min** , **max**]. Result is put in **out** .

(rgbafim-blit x y in out)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

blit **in** into **out** at position **x** , **y** with alpha blending. an alpha value of 0 in **in** means opaque , 255 means totally transparent.

(rgbafim-blitcolor x y in r g b out)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

blit color defined by *r* , *g* , *b* using *idx2 in* as a stencil into RGBA image *out* at position *x y* The values in *in* must be between 0 and 255. if the value is 255 , no paint is put in; if the value is 0 , the pixel takes the value specified by *r* , *g* , *b* .

(rgbafim-init-lum2rgb)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

build a table of 32x32x32 elements containing RGB cell indices in ascending luminosity.

(rgbafim-lum2rgb n)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

converts a luminosity index to a 5-bit RGB cell number.

(rgbafim-histo32 rgbafim ppa1 count)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

compute histogram of image *rgbafim* . *ppa1* must be a 32768 by 3 matrix which will contain the list of pixel values found. *count* must be a 32768 matrix which will contain the corresponding pixel counts. pixel colors are sorted by luminance. This function really computes a 32768 color palette by assigning each pixel to the cube it belongs to in the 32x32x32 RGB color cube. The color prototype assigned to each cube is the mean of the pixels belonging to the cube.

(rgbafim-interpolate-bilin background pin indimi indimj inmodi inmodj ppi ppj rez)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

(-gptr- pin rez background):

(-int- indimi indimj inmodi inmodj ppi ppj):

SIDE EFFECTS: <rez>

AUTHOR: Y. LeCun

COMPILABLE: Yes

DESCRIPTION:

returns a bilinearly interpolated RGBA pixel value for coordinate <ppi> <ppj>. The image data is pointed to by <pin>, with <indimi> <indimj> <inmodi> <inmodj> being the dimensions and modulus. This function clips automatically if <ppi> <ppj> are outside of the bounds by assuming the outside values are equal to <background>. pixel values are flts, while coordinates are 32 bit fixed point with 16 bit integer part and 16 bit fractional part.

(rgbafim-warp in out background pi pj)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

```
((-idx1- (-flt)) background):
((-idx3- (-flt-)) in out):
((-idx2- (-int-)) pi pj):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
Warpes an image using <pi> and <pj> as tabulated coordinate transforms.
<in> and <out> are idx2 of flts. <background> is the value assumed outside
of the input image. <pi> and <pj> are tabulated coordinates which can
be filled up using compute-bilin-transform or similar functions.
Pixel values are antialiased using bilinear interpolation.
```

(rgbafim-warp-fast in out background pi pj)

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

```
((-int-) background):
((-idx2- (-flt-)) in out):
((-idx2- (-int-)) pi pj):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
Warpes an image using <pi> and <pj> as tabulated coordinate transforms.
<in> and <out> are idx2 of flts. <background> is the value assumed outside
of the input image. <pi> and <pj> are tabulated coordinates which can
be filled up using compute-bilin-transform or similar functions.
This is essentially identical to warp-rgbafimage, except no antialiasing
is performed (it goes a lot faster, but is not nearly as nice).
```

**(rgbafim-warp-quad in out background mode x1 y1 x2 y2 x3 y3 x4 y4 p1
q1 p3 q3)**

(lsh/libimage/rgbafimage.lsh) Last Modified: 2002-10-14 03:30:51

```
((-flt-) x1 y1 x2 y2 x3 y3 x4 y4 p1 q1 p3 q3):
((-int-) background mode):
```

((-idx2- (-flt-)) in out):
 RETURNS: Null
 SIDE EFFECTS: <out>
 AUTHOR: Y. LeCun
 COMPILABLE: Yes
 DESCRIPTION:
 transforms rgbafimage <in> (idx2 of flt) mapping quadrilateral <x1> <y1> <x2>
 <y2> <x3> <y3> <x4> (points are numbered clockwise starting from upper left)
 to rectangle whose upper left and lower right points are <p1><q1>, <p3><q3>.
 result is put in <out> (idx2 of flt). Clipping is automatically performed.
 pixels outside of the bounds of <in> are assumed to have the value
 <background>. <mode>=0: no antialiasing, <mode>=1: antialiasing with
 bilinear interpolation (2 times slower). execution time on sparc 10 is about
 5 ms in mode 0 and 10 ms in mode 1 for a 32x32 target image.

8.7.8 Greyscale Images of ubytes.

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11
 linear coordinate transform using bilinear interpolation.

ubimage

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11
 a ubimage is an idx2 of ubytes intended to store image.

(ubimage2flt in out)

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11

((-idx2- (-ubyte-)) in):
 ((-idx2- (-flt-)) out):
 RETURNS: Null
 SIDE EFFECTS: <out>
 AUTHOR: Y. LeCun
 COMPILABLE: Yes
 DESCRIPTION:
 Converts ubimage <in> into idx2 of flt <out>.
 Values are scaled down by 1/255.

(flt2ubimage in out)

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11

((-idx2- (-ubyte-)) out):
 ((-idx2- (-flt-)) in):

RETURNS: Null
 SIDE EFFECTS: None
 AUTHOR: Y. LeCun
 COMPILABLE: Yes
 DESCRIPTION:
 converts flt idx2 <in> to an ubimage <out>.
 a pixel values are scaled up by 255.

(ubimage2fltimage in out bkgd ink)

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11

((-idx2- (-ubyte-)) in):
 ((-idx2- (-flt-)) out):
 ((-flt-) bkgd ink):
 RETURNS: Null
 SIDE EFFECTS: <out>
 AUTHOR: C. Burges
 COMPILABLE: Yes
 DESCRIPTION:
 Generalization of Yann's ubimage2flt for arbitrary ink and background.
 Converts ubimage <in> into idx2 of flt <out>.

(fltimage2ubimage in out bkgd ink)

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11

((-idx2- (-flt-)) in):
 ((-idx2- (-ubyte-)) out):
 ((-flt-) bkgd ink):
 RETURNS: Null
 SIDE EFFECTS: <out>
 AUTHOR: C. Burges
 COMPILABLE: Yes
 DESCRIPTION:
 Generalization of Yann's flt2ubimage for arbitrary ink and background.
 Converts flt idx2 <in> to a ubimage <out>.

(ubim-subsample in nlin ncol)

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11

subsamples image *in* with integer ratios *nlin* (vertical) *ncol* (horizontal).
 the subsampled image is returned. The horizontal (resp. vertical) size of the
 returned image is equal to the integer part of the horizontal (resp vertical) size
 of *in* divided by *ncol* (resp *nlin*).

returns (copy-matrix *in*) when subsample rate is 1

(ubim-zoom ubim zoomx zoomy)

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11
Zoom an ubyte image **ubim** by the given integer zoom factors.

(ubim-invert im)

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11
reverse polarity of an ubimage. each pixel is replaced by 255 minus itself

(ubim-interpolate-bilin background pin indimi indimj inmodi inmodj ppi ppj)

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11

(-gptr- pin):
(-int- background indimi indimj inmodi inmodj ppi ppj):
RETURNS: (-ubyte-)
SIDE EFFECTS: None
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
returns a bilinearly interpolated pixel value for coordinate
<ppi> <ppj>. The image data is pointed to by <pin>, with
<indimi> <indimj> <inmodi> <inmodj> being the dimensions and modulus.
This function clips automatically if <ppi> <ppj> are outside of the
bounds by assuming the outside values are equal to <background>.
pixel values are ubytes, while coordinates are 32 bit fixed point
with 16 bit integer part and 16 bit fractional part.
The function does not use floating point arithmetics.

(ubim-warp in out background pi pj)

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11

((-int-) background):
((-idx2- (-ubyte-)) in out):
((-idx2- (-int-)) pi pj):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
Warpes an image using <pi> and <pj> as tabulated coordinate transforms.
<in> and <out> are idx2 of ubytes. <background> is the value assumed outside

of the input image. <pi> and <pj> are tabulated coordinates which can be filled up using `compute-bilin-transform` or similar functions. Pixel values are antialiased using bilinear interpolation.

(ubim-warp-fast in out background pi pj)

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11

```
((-int-) background):
((-idx2- (-ubyte-)) in out):
((-idx2- (-int-)) pi pj):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
Warps an image using <pi> and <pj> as tabulated coordinate transforms.
<in> and <out> are idx2 of ubytes. <background> is the value assumed outside
of the input image. <pi> and <pj> are tabulated coordinates which can
be filled up using compute-bilin-transform or similar functions.
This is essentially identical to warp-ubimage, except no antialiasing
is performed (it goes a lot faster, but is not nearly as nice).
```

(ubim-warp-quad in out background mode x1 y1 x2 y2 x3 y3 x4 y4 p1 q1 p3 q3)

(lsh/libimage/ubimage.lsh) Last Modified: 2002-08-30 03:50:11

```
((-flt-) x1 y1 x2 y2 x3 y3 x4 y4 p1 q1 p3 q3):
((-int-) background mode):
((-idx2- (-ubyte-)) in out):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
transforms ubimage <in> (idx2 of ubyte) mapping quadrilateral <x1> <y1> <x2>
<y2> <x3> <y3> <x4> (points are numbered clockwise starting from upper left)
to rectangle whose upper left and lower right points are <p1><q1>, <p3><q3>.
result is put in <out> (idx2 of ubyte). Clipping is automatically performed.
pixels outside of the bounds of <in> are assumed to have the value
<background>. <mode>=0: no antialiasing, <mode>=1: antialiasing with
bilinear interpolation (2 times slower). execution time on sparc 10 is about
5 ms in mode 0 and 10 ms in mode 1 for a 32x32 target image.
```

8.7.9 Greyscale Images of shorts

(lsh/libimage/shimage.lsh) Last Modified: 2002-08-30 03:50:11
 functions to manipulate images where each pixel is a 16-bit short.

shimage

(lsh/libimage/shimage.lsh) Last Modified: 2002-08-30 03:50:11
 a shimage is an idx2 of shorts intended to store image. the pixel values are interpreted as fixed point integers , with the upper byte being the integer part , and the lower byte being the fractional part.

(shimage2flt in out)

(lsh/libimage/shimage.lsh) Last Modified: 2002-08-30 03:50:11

```
((-idx2- (-short-)) in):
((-idx2- (-flt-)) out):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
Converts shimage <in> into idx2 of flt <out>.
Values are scaled down by 1/256.
```

(flt2shimage in out)

(lsh/libimage/shimage.lsh) Last Modified: 2002-08-30 03:50:11

```
((-idx2- (-short-)) out):
((-idx2- (-flt-)) in):
RETURNS: Null
SIDE EFFECTS: None
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
converts flt idx2 <in> to an shimage <out>.
a pixel values are scaled up by 256.
```

(shimage2fltimage in out bkgd ink)

(lsh/libimage/shimage.lsh) Last Modified: 2002-08-30 03:50:11

```
((-idx2- (-short-)) in):
```

```
((-idx2- (-flt-)) out):
((-flt-) bkgd ink):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: C. Burges
COMPILABLE: Yes
DESCRIPTION:
Generalization of Yann's shimage2flt for arbitrary ink and background.
Converts shimage <in> into idx2 of flt <out>.
```

(fltimage2shimage in out bkgd ink)

(lsh/libimage/shimage.lsh) Last Modified: 2002-08-30 03:50:11

```
((-idx2- (-flt-)) in):
((-idx2- (-short-)) out):
((-flt-) bkgd ink):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: C. Burges
COMPILABLE: Yes
DESCRIPTION:
Generalization of Yann's flt2shimage for arbitrary ink and background.
Converts flt idx2 <in> to a shimage <out>.
```

**(shim-interpolate-bilin background pin indimi indimj inmodi inmodj
ppi ppj)**

(lsh/libimage/shimage.lsh) Last Modified: 2002-08-30 03:50:11

```
(-gptr- pin):
(-int- background indimi indimj inmodi inmodj ppi ppj):
RETURNS: (-short-)
SIDE EFFECTS: None
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
returns a bilinearly interpolated pixel value for coordinate
<ppi> <ppj>. The image data is pointed to by <pin>, with
<indimi> <indimj> <inmodi> <inmodj> being the dimensions and modulus.
This function clips automatically if <ppi> <ppj> are outside of the
bounds by assuming the outside values are equal to <background>.
pixel values are shorts, while coordinates are 32 bit fixed point
with 16 bit integer part and 16 bit fractional part.
The function does not use floating point arithmetics.
```

(shim-warp in out background pi pj)

(lsh/libimage/shimage.lsh) Last Modified: 2002-08-30 03:50:11

```
((-int-) background):
((-idx2- (-short-)) in out):
((-idx2- (-int-)) pi pj):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
Warpes an image using <pi> and <pj> as tabulated coordinate transforms.
<in> and <out> are idx2 of shorts. <background> is the value assumed outside
of the input image. <pi> and <pj> are tabulated coordinates which can
be filled up using compute-bilin-transform or similar functions.
Pixel values are antialiased using bilinear interpolation.
```

(shim-warp-fast in out background pi pj)

(lsh/libimage/shimage.lsh) Last Modified: 2002-08-30 03:50:11

```
((-int-) background):
((-idx2- (-short-)) in out):
((-idx2- (-int-)) pi pj):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
Warpes an image using <pi> and <pj> as tabulated coordinate transforms.
<in> and <out> are idx2 of shorts. <background> is the value assumed outside
of the input image. <pi> and <pj> are tabulated coordinates which can
be filled up using compute-bilin-transform or similar functions.
This is essentially identical to warp-shimage, except no antialiasing
is performed (it goes a lot faster, but is not nearly as nice).
```

(shim-warp-quad in out background mode x1 y1 x2 y2 x3 y3 x4 y4 p1 q1 p3 q3)

(lsh/libimage/shimage.lsh) Last Modified: 2002-08-30 03:50:11

```
((-flt-) x1 y1 x2 y2 x3 y3 x4 y4 p1 q1 p3 q3):
((-int-) background mode):
```

```
((-idx2- (-short-)) in out):
RETURNS: Null
SIDE EFFECTS: <out>
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
transforms shimage <in> (idx2 of short) mapping quadrilateral <x1> <y1> <x2>
<y2> <x3> <y3> <x4> (points are numbered clockwise starting from upper left)
to rectangle whose upper left and lower right points are <p1><q1>, <p3><q3>.
result is put in <out> (idx2 of short). Clipping is automatically performed.
pixels outside of the bounds of <in> are assumed to have the value
<background>. <mode>=0: no antialiasing, <mode>=1: antialiasing with
bilinear interpolation (2 times slower). execution time on sparc 10 is about
5 ms in mode 0 and 10 ms in mode 1 for a 32x32 target image.
```

8.7.10 Greyscale Images of Floats

(lsh/libimage/fimage.lsh) Last Modified: 2003-01-07 03:35:12
 basic image processing on greyscale images where each pixel is a float.

(fim-subsample in nlin ncol)

(lsh/libimage/fimage.lsh) Last Modified: 2003-01-07 03:35:12
 subsample image in by integer ratios nlin vertically and ncol horizontally
 (with averaging).

8.7.11 Computing Image Warping Maps

(lsh/libimage/image-transform.lsh) Last Modified: 2002-08-19 03:01:28
 compute displacement tables for 2D geometric transformations of images ,
 such as bilinear warps that transform any quadrilateral into any rectangle.

**(compute-bilin-transform dispi dispj x1 y1 x2 y2 x3 y3 x4 y4 p1 q1 p3
 q3)**

(lsh/libimage/image-transform.lsh) Last Modified: 2002-08-19 03:01:28

```
((-idx2- (-int-)) dispi):
((-idx2- (-int-)) dispj):
((-flt-) x1 y1 x2 y2 x3 y3 x4 y4 p1 q1 p3 q3):
RETURNS: Null
SIDE EFFECTS: <dispi> <dispj>
AUTHOR: Y. LeCun
COMPILABLE: Yes
DESCRIPTION:
Tabulates a bilinear transformation that maps the quadrilateral defined
```

by the `<xi>` `<yi>` to a rectangle whose upper left point is `<p1>` `<q1>` and lower right point is `<p3>` `<q3>`.

`<x1>` `<y1>` is the upper left point on the quadrilateral, and the points are numbered clockwise.

`<dispi>` and `<dispj>` must be `idx2` of `int`. On output, element `(i,j)` of `<dispi>` and `<dispj>` will be filled respectively with the abscissa and ordinate of the point in the quadrilateral that maps to point `(i,j)`. Values in these `idx2` are interpreted as 32bit fixed point numbers with 16 bit integer part and 16 bit fractional part.

`<dispi>` and `<dispj>` can subsequently be used as input to `warp-shimage`, or `warp-shimage-fast`.

This function makes minimal use of floating point arithmetics.

8.7.12 Connected Component Analysis

(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

elementary routines for performing a connected component (CC) analysis on images. The basic procedure consists in:

- Creating a new `CCAnalyzer` object for a given image
- Extracting the runs using method `run-analysis` .
- Extracting the connected components using method `cc-analysis` .

There are a few other methods for performing various cleanup on the runs and on the CCs. There is also a method for reconstructing the image from the cleaned CCs.

Another trick: There is a matrix named `ubcut` in the `CCAnalyzer` object. Every pixel set in this matrix will force the RUN and CC code to ignore connectivity on certain pixels. For instance, you can force arbitrary cuts by setting the pixels of a 4-connex line in this matrix. This has to be done before the call to `run-analysis` .

– Leon Bottou 9/96

CCAnalyzer

(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

Class for connected component analysis.

(new CCAnalyzer grayimage)

Type: CLASS(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

Class `CCAnalyzer` holds all the information necessary for the connected component analysis of an image `grayimage` .

(==> **ccanalyzer run-analysis threshold**)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

Computes a run length representation of the binary image obtained by thresholding the gray image **ubmatrix2d** at threshold **thres** . The run information is stored into slot **runs** of the **CCAnalyzer** object.

Each line of matrix **runs** is a 4-vector containing the following information:

- At index (**RUN-Y**) : The Y coordinate of the run.
- At index (**RUN-X1**) : The X coordinate of the first pixel of the run.
- At index (**RUN-X2**) : The X coordinate of the last pixel of the run.
- At index (**RUN-ID**) : The ID of the blob this run belongs to. The blob information is filled by the CA extraction code.

This function breaks runs on each non zero pixel of the image located in slot **ubcut** of **ccanalyzer** . The default value is all zeroes.

This function returns **t** on success.

(==> **ccanalyzer cc-analysis**)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

Perform the connected component analysis of the image represented by the runs matrix contained in the **ccanalyzer** object. This algorithm usually considers 8-connectivity. This is altered when the cut image **ubcuts** contains a non zero value indicating a forced cut.

The results are stored in the cc descriptor matrix **ccdesc** of the object. Each line of this matrix represent a connected component (CC).

- At index (**CC-NRUN**) : The number of runs in the CC.
- At index (**CC-FRUN**) : The index of the first run of the CC in the CC ordered run matrix **ccruns** stored in the object.
- At index (**CC-NPIX**) : The number of non background pixels in the CC.
- At index (**CC-WPIX**) : The sum of the pixel values (gray level) of the CC. These values are initially zero. Use function **cc-measure-gray** to compute these values.
- At indices (**CC-LEFT**) , (**CC-TOP**) , (**CC-RIGHT**) and (**CC-BOTTOM**) : The bounding box of the CC.

This function returns **t** on success. You may postprocess the returned CC with **cc-measure-gray** , **cc-remove-specs** , and other functions in this file.

(==> ccanalyzer remove-long-runs maxlen)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

This function removes the runs whose length is longer than `maxlen` . Long runs are sometimes good indicator for underlines. Note that nothing is changed in the image. You must use `reconstruct-from-runs` or `reconstruct-from-cc` for that.

(==> ccanalyzer reconstruct-from-runs)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

Reconstruct matrix `ubimg` from the run representation. This code saves the gray level information already present in `ubimg` . It just sets all pixels not included in a run to zero.

(==> ccanalyzer cc-measure-gray)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

Compute the average gray level of each CC and store it at index `(CC-WPIX)` of the CC descriptor matrix `ccdesc` .

(==> ccanalyzer remove-small-cc minpix mingray minweight)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

Remove small connected components

- whose size is less than `minpix` pixels,
- or whose average gray level is less than `mingray` ,
- or whose summed gray level is less than `minweight` .

Note that nothing is changed in the image or the run matrix. You must use `reconstruct-from-cc` for that.

(==> ccanalyzer reconstruct-from-cc)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

Reconstruct matrix `ubimg` from the connected components in `ccdesc` . This code saves the gray level information already present in `ubimg` . It just sets all pixels not included in a run to zero.

(==> ccanalyzer overlay-cc ccid ccimg originx originy)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

Overlays connected component `ccid` into ubyte matrix `ccimg` . Matrix `ccimg` is assumed to have its top left corner at position `originx` `originy` in the initial image. Connected components that do not fit matrix `ccimg` are silently clipped. Returns `t` when something was drawn.

(==> **ccanalyzer draw-cc x y**)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

Draws image at location **x y** by painting each connected component with a random color. Unfortunately, random colors are sometimes too close. [INTERPRETED ONLY]

(==> **ccanalyzer pick-cc x y**)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

returns the index of the connected component that contains the point **x, y**, or -1 if no component contains the point.

(==> **ccanalyzer image-height**)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

returns height of image on which CCA is performed

(==> **ccanalyzer image-width**)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

returns width of image on which CCA is performed

(==> **ccanalyzer run-histo histo hosto offset**)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

compute histogram of runs. **histo** is an N by P matrix, and **hosto** an N vector. this method will accumulate the histogram of run lengths for each scan line between **offset** and **offset +N-1**. on output, elements **i, j** of **histo** will be incremented by the number of runs on line **i - offset** that are of length **j**. element **i** of **hosto** will be incremented by the number of runs on scanline **i - offset** that are larger or equal to **P**.

(==> **ccanalyzer bbox-histo histo**)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

computes 2D histogram of bounding box heights and width and increment 2D matrix **histo** by it. Element **i, j** of **histo** will be incremented by the number of CC whose bounding box is of height **i** and of width **j**.

(==> **ccanalyzer ccbottoms thres**)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

return a matrix with **x** and **y** of horizontal center of bottom of each CC. CCs whose number of black pixels are smaller than **thres** are ignored.

(==> **CCAnalyzer get-cc-fltim cc-id cc-fltim cc-com cc-off**)

Type: MSG(lsh/libimage/cca.lsh) Last Modified: 2002-10-29 18:11:55

```

((-int-) cc-id) : index of the cc
((-idx2- (-flt-)) cc-fltim) : image in float
((-idx1- (-flt-)) cc-com) : center of mass, in float [X,Y]
((-idx1- (-int-)) cc-off) : offset [X,Y]
ADDED: Patrick Haffner, July 97

```

Given the index of a CC, returns:

- 1) the corresponding gray level float image, dilated (WITH OFFSET)
 - 2) the center of mass of the bitmap (WITHOUT OFFSET)
 - 3) the offset used to align the float image to the upper right corner.
- WARNING: fltim expands all the edges of the bounding box by 1 pixel, UNLESS this cannot be done because we hit the global image border.

8.7.13 Morphological Operations

(lsh/libimage/morpho.lsh) Author(s): Leon Bottou 09/96 Last Modified: 2002-09-22 22:06:59

Erosion, dilation, distance transforms, thresholding, and masking on images of ubytes.

(**ubim-internal-disttrans im dist**)

(lsh/libimage/morpho.lsh) Last Modified: 2002-09-22 22:06:59

Returns the internal distance transformation matrix of binary image `im`. The elements of the output ubyte matrix `dist` contains the distance to the closest background pixel in the image (background level is zero for this purpose). Distances are clipped to 255. The dimensions of `dist` must be 2 larger than that of `im`. `dist` will be resized to the appropriate size if necessary.

(**ubim-external-disttrans im dist**)

(lsh/libimage/morpho.lsh) Last Modified: 2002-09-22 22:06:59

computes the external distance transformation matrix of binary image `im`. The elements of the output ubyte matrix `dist` contains the distance to the closest non-background pixel in the image (background level is zero for this purpose). Distances are clipped to 255. The dimensions of `dist` must be 2 larger than that of `im`. `dist` will be resized to the appropriate size if necessary.

(ubim-positive-threshold im thres)

(lsh/libimage/morpho.lsh) Last Modified: 2002-09-22 22:06:59

Applies a threshold on ubyte image **im** . All values smaller than or equal to threshold are zeroed. Other values are set to 255.

(ubim-negative-threshold im thres)

(lsh/libimage/morpho.lsh) Last Modified: 2002-09-22 22:06:59

Applies a threshold on ubyte image **im** . All values smaller than or equal to threshold are set to 255. Other values are set to 0

(ubim-erosion im factor)

(lsh/libimage/morpho.lsh) Last Modified: 2002-09-22 22:06:59

Perform in place erosion of ubyte image **im** at distance **factor** .

(ubim-dilation im factor)

(lsh/libimage/morpho.lsh) Last Modified: 2002-09-22 22:06:59

Perform in place dilation of ubyte image **im** at distance **factor** .

(ubim-mask im mask)

(lsh/libimage/morpho.lsh) Last Modified: 2002-09-22 22:06:59

Sets to zero all bytes of ubyte image **im** whose corresponding pixel in **mask** is zero. This is a AND operation.

8.7.14 Morphological Operation on Images of Shorts.

(lsh/libimage/morpho-short.lsh) Last Modified: 2002-10-14 03:30:51

Erosion/dilation morphological operations on images of shorts.

CAUTION with the so-called "margin" parameter. a margin of 1 is sufficient in most cases, however the margin pixels MUST be filled with 0 (background/white) for the functions to work properly. (**dilero** does it automatically).

(ubim-fillborder shim c)

(lsh/libimage/morpho-short.lsh) Last Modified: 2002-10-14 03:30:51

fill a one-pixel border of the ubimage **shim** with value **c** . This function is useful for various morphological operations.

(shim-fillborder shim c)

(lsh/libimage/morpho-short.lsh) Last Modified: 2002-10-14 03:30:51

fill a one-pixel border of the shimage **shim** with value **c** . This function is useful for various morphological operations.

(shim-dilation im margin)

(lsh/libimage/morpho-short.lsh) Last Modified: 2002-10-14 03:30:51

```
((-idx2- (-short-)) im) ;; image to be dilated:
((-int-) margin):
RETURNS:
CREATED: Y. LeCun, modified P.Haffner
COMPILABLE: Yes
DESCRIPTION:
```

(shim-erosion im oldim mode margin)

(lsh/libimage/morpho-short.lsh) Last Modified: 2002-10-14 03:30:51

```
((-idx2- (-short-)) im) ;; image to be eroded:
((-idx2- (-short-)) orig-im) ;; original image, only used in mode 2:
((-int-) mode):
((-int-) margin) ;; erosion margin, typical values are 1 or 2:
RETURNS:
CREATED: Y. LeCun, modified P.Haffner
COMPILABLE: Yes
DESCRIPTION:
erosions with hexagonal tessellation
<dist> is an idx2 of shorts with dimensions identical to <im>+ border.
It will contain the distance transform on output.
<mode> should be
    1 ==> erosion
    2 ==> erosion and hole restoration
```

(idx-s2xors2 m1 m2)

(lsh/libimage/morpho-short.lsh) Last Modified: 2002-10-14 03:30:51

```
((-idx2- (-short-)) m1 m2):
RETURNS: (-idx2- (-short-)) m1x2
CREATED: P.Haffner
COMPILABLE: Yes
DESCRIPTION: performs a logical XOR between m1 and m2.
m1x2 pixel is 0 when m1 and m2 pixels are equal, 0 otherwise.
```

(dilero runs margin)

(lsh/libimage/morpho-short.lsh) Last Modified: 2002-10-14 03:30:51

```
(-idx2- runs):
(-int- margin):
RETURNS: an -idx2- matrix of runs
CREATED: P.Haffner
COMPILABLE: Yes
DESCRIPTION:
morphological closing operator.
successively applies dilation and erosion functions
<margin> represents the extend of the initial dilation
- 0 nothing happens
- 1 using an hexagonal tessellation, all pixel next to a black pixel becomes black.
- 2 distance is extended to two pixels.
```

8.7.15 Miscellaneous Image Macros and Conversions

(lsh/libimage/img-util.lsh) Last Modified: 2002-10-29 18:11:56
 basic macros and functions to manipulate dimensions and internal formats of images

(**rgbaim2rgbafim** in)

(lsh/libimage/img-util.lsh) Last Modified: 2002-10-29 18:11:56
 convert RGBA image of ubytes into RGBA image of floats.

(**rgbafim2rgbaim** in)

(lsh/libimage/img-util.lsh) Last Modified: 2002-10-29 18:11:56
 convert RGBA image of floats into RGBA image of ubytes.

img-dim-C

(lsh/libimage/img-util.lsh) Last Modified: 2002-10-29 18:11:56
 class that holds the two dimensions of an image d0: height d1: width
 (**new img-dim-C** d0 d1) Type: CLASS(lsh/libimage/img-util.lsh) Last Modified: 2002-10-29 18:11:56
 (**equal-dim** dimA dimB) (lsh/libimage/img-util.lsh) Last Modified: 2002-10-29 18:11:56

```
((-obj- (img-dim-C)) dimA dimB):
CREATED: Oct 97
DESCRIPTION:
t if dimA and dimB are equal
```

(**align-dimensions** img-dim align-block) (lsh/libimage/img-util.lsh) Last Modified: 2002-10-29 18:11:56

```

((-obj- (img-dim-C)) img-dim):
(-int- align-block):
CREATED: Oct 97
DESCRIPTION:
return new dimension object
- dimensions larger than img-dim
- multiple of align-block

```

8.7.16 Constants used in run-length encoded images.

(lsh/libimage/run-macros.lsh) Last Modified: 2002-09-06 18:27:25
 Indexes to access tables of runs

```

RUN-Y (f) 0
RUN-X1 (f) 1
RUN-X2 (f) 2
RUN-ID (f) 3
CC-NRUN (f) 0
CC-FRUN (f) 1
CC-WPIX (f) 2
CC-NPIX (f) 3
CC-LEFT (f) 4
CC-TOP (f) 5)1
CC-RIGHT (f) 6
CC-BOTTOM (f) 7

```

8.7.17 Converting RLE Image to greyscale pixelmap.

(lsh/libimage/runs2ubim.lsh) Last Modified: 2002-08-19 03:01:28

(nt-runs2ubim runs ubim)

(lsh/libimage/runs2ubim.lsh) Last Modified: 2002-08-19 03:01:28

```

((-idx2- (-int-)) runs):
((-idx2- (-ubyte-)) ubim):
CREATED: P.Haffner, August 97
TEMPORARIES: ?
DESCRIPTION:
maps runs into a ubyte image

```

(static-runs2ubim runs ubim white black)

(lsh/libimage/runs2ubim.lsh) Last Modified: 2002-08-19 03:01:28

```
((-idx2- (-int-)) runs):
((-idx2- (-ubyte-)) ubim):
((-int-) white black):
CREATED: P.Haffner, August 97
TEMPORARIES: ?
DESCRIPTION:
```

(nt-runs2ubim-dil runs ubim)

(lsh/libimage/runs2ubim.lsh) Last Modified: 2002-08-19 03:01:28

```
((-idx2- (-int-)) runs):
((-idx2- (-int-)) ubim):
CREATED: P.Haffner, August 97
TEMPORARIES: ?
DESCRIPTION:
1-dilattion while applying <nt-run2ubim>
```

(static-runs2ubim-dil runs ubim)

(lsh/libimage/runs2ubim.lsh) Last Modified: 2002-08-19 03:01:28

```
((-idx2- (-int-)) runs):
((-idx2- (-ubyte-)) ubim):
CREATED: P.Haffner, August 97
TEMPORARIES: ?
DESCRIPTION:
```

(static-runs2ubim-subsample runs ubim x-sub y-sub)

(lsh/libimage/runs2ubim.lsh) Last Modified: 2002-08-19 03:01:28

```
((-idx2- (-int-)) runs):
((-idx2- (-ubyte-)) ubim):
((-int-) x-sub y-sub):
CREATED: P.Haffner, August 97
TEMPORARIES: ?
DESCRIPTION:
<nt-runs2ubim> with subsampling
```

See: (nt-runs2ubim runs ubim)

(runs2ubim runs white black)

(lsh/libimage/runs2ubim.lsh) Last Modified: 2002-08-19 03:01:28

```
((-idx2- (-int-)) runs):
((-int-) white black):
CREATED: P.Haffner, August 97
TEMPORARIES: ?
DESCRIPTION:
```

(runs2ubim-subsample runs x-sub y-sub)

(lsh/libimage/runs2ubim.lsh) Last Modified: 2002-08-19 03:01:28

```
((-idx2- (-int-)) runs):
((-int-) x-sub y-sub):
CREATED: P.Haffner, August 97
DESCRIPTION:
Subsample on the fly --> grey levels
```

8.7.18 Color Clustering and Quantization

Color Clustering

(lsh/libimage/color-quantize.lsh) Last Modified: 2002-10-29 18:11:55

Functions to find representative colors from a precomputed color histogram. these functions just find the colors, they do not actually quantize the images. This is done in `rgbaimage`, and `rgbafimage`.

(cqu-init-proto ppa1 proto) (lsh/libimage/color-quantize.lsh) Last Modified: 2002-10-29 18:11:55

initialize prototypes for K-means. assumes that `ppa1` contains the color values sorted by luminosity (as returned by `rgbaim-histo32`). puts the prototypes regularly spaced along a line between the darkest and lightest color. `ppa1` : Nx3 matrix of floats containing a list of RGB color cells `proto` : Px3 matrix of floats, contains the initialized prototypes on output

(cqu-kmeans ppa1 count assign split proto label weight variance)
(lsh/libimage/color-quantize.lsh) Last Modified: 2002-10-29 18:11:55

performs a Kmeans algorithms on the color values contained in `ppa1`.
INPUT: `ppa1` : Nx3 matrix of floats containing a list of RGB color present in the picture `count` : N matrix of floats, number of pixels with the corresponding RGB color `assign` : N matrix of floats containing the label of the prototype to which the color cell is currently assigned. `split` : an int: only those color cells whose initial assigned label is equal to `split` will be taken into account. `proto` : Px3 matrix of floats, contains initialized RGB prototypes `label` : P matrix of ints, contains labels for the prototypes OUTPUT: `proto` : contains the updated

prototypes on output **assign** : the color cells whose assigned value was **split** will be assigned the value in **label** corresponding to their closest prototype.
weight : P matrix of float , contains the number of pixels assigned to each prototype
variance : P matrix of float , contains the variance associated with the prototypes

Color Segmentation and Quantization on RGBA images

(lsh/libimage/rgbaim-greypage.lsh) Last Modified: 2002-10-14 03:30:51

a set of function for detecting the background and foreground colors in a document image and for transforming them into grey-level images suitable for recognition , segmentation , etc... essentially sets the background to 0 , the foreground to 256 , the rest in the middle.

(rgbaim-project-greys rgbaim ubim polarity bsatur fsatur colors)
 (lsh/libimage/rgbaim-greypage.lsh) Last Modified: 2002-10-14 03:30:51

transforms an rgb image into a grey image where the foreground is 256 , the background is 0 , and the greys in between are between 256 and 0 the resulting image can be used for segmentation and recognition. **rgbaim** an idx3 of ubytes containing the input RGBA image. **ubim** an idx2 of ubytes that will contain the result it should be the same size as **rgbaim** **polarity** determines the polarity of the image: 0 means light background and dark foreground , 1 means dark background and light foreground , -1 should be passed if the polarity is unknown. In that case , the color cluster with the most numerous pixels is assigned to the background. pixel colors in **rgbaim** are projected on a straight segment between two color prototypes computed with K-means. The output pixels are given a grey value that depends on their position along that segment. The function that maps positions on the line to grey-levels is piece-wise linear saturation-type function. It is controlled by two points A and B. below A , points are assigned the value 0 , above B they are assigned 256. in between they are mapped linearly. the position of point A is determined by the argument **bsatur** a value of 0 puts A on the background prototype , a value of 0.5 puts it half way between the 2 prototypes. Point B is controlled similarly: 0 puts it at the foreground prototype and 0.5 puts it halfway. **colors** is a 2x3 matrix of rgb color clusters for the background and foreground respectively.

(rgbaim-project-colors rgbaim rgbaim2 polarity bsatur fsatur colors)
 (lsh/libimage/rgbaim-greypage.lsh) Last Modified: 2002-10-14 03:30:51

project the colors of an image onto a line joining two center clusters. **rgbaim** an idx3 of ubytes containing the input RGBA image. **rgbaim2** contains the result image on output. it should be the same size as **<rgbaim polarity** determines the polarity of the image: 0 means light background and dark foreground , 1 means dark background and light foreground , -1 should be passed if the polarity is unknown. In that case , the color cluster with the most numerous pixels is assigned to the background. pixel colors in **rgbaim** are projected on a straight segment between two color prototypes computed with K-means. The output pixels are given a grey value that depends on their position along that segment. The function that maps positions on the line to grey-levels is

piece-wise linear saturation-type function. It is controlled by two points A and B. below A , points are assigned the value 0 , above B they are assigned 256. in between they are mapped linearly. the position of point A is determined by the argument **bsatur** a value of 0 puts A on the background prototype , a value of 0.5 puts it half way between the 2 prototypes. Point B is controlled similarly: 0 puts it at the foreground prototype and 0.5 puts it halfway. **colors** is a 2x3 matrix of rgb color clusters for the background and foreground respectively.

(rgbaim-cluster-colors rgbaim ubim polarity bsatur fsatur proto weight variance) (lsh/libimage/rgbaim-greypage.lsh) Last Modified: 2002-10-14 03:30:51

See **rgbaim-project-greys** for an explanation of the arguments. This function performs color clustering, using the K-means algorithm. Results in **ubim** .

proto : n.colors x 3 matrix = colors of the clusters **weight** : n.colors vector = weights of the Gaussian mixture for each cluster **variance** : n.colors vector = variance parameter of each cluster Gaussian distribution

See: (rgbaim-greypage **rgbaim ubim polarity bsatur fsatur**)

(rgbaim-quickquant rgbaim ubim ppal assign) (lsh/libimage/rgbaim-greypage.lsh) Last Modified: 2002-10-14 03:30:51

quick color quantization of the pixels in **rgbaim** to the colors in palette **ppal** . Each pixel in the output image is a short whose value is the "label" of the palette color that is closest to the input pixel. **assign** is a vector of ints with the same first dimension as **ppal** which contains the "label" for each palette color. This function is relatively fast , but is approximate (the palette color associated with an input color is not always the closest). In particular , this function makes little sense if two palette colors are in the same cell in the 32x32x32 RGB cube (which means **ppal** should probably have less than 2^{15} colors).

low-level color quantization functions. (lsh/libimage/rgbaim-greypage.lsh) Last Modified: 2002-10-14 03:30:51

(rgbaim-greyquant rgbaim ubim proto polarity bsatur fsatur) (lsh/libimage/rgbaim-greypage.lsh) Last Modified: 2002-10-14 03:30:51

transforms and RGB image into a grey image with uniform foreground and background. **proto** is a 2x3 matrix of floats containing two color prototypes for background and foreground colors. **rgbaim** an idx3 of ubytes containing the input RGBA image. **ubim** an idx2 of ubytes that will contain the result it should be the same size as **rgbaim** **polarity** determines the polarity of the image: 0 means light background and dark foreground , 1 means dark background and light foreground , -1 should be passed if the polarity is unknown. In that case , the color cluster with the most numerous pixels is assigned to the background. pixel colors in **rgbaim** are projected on a straight segment between two color prototypes computed with K-means. The output pixels are given a grey value that depends on their position along that segment. The function that maps positions on the line to grey-levels is piece-wise linear saturation-type function. It is controlled by two points A and B. below A , points are assigned the value 0 , above B they are assigned 256. in between they are mapped linearly. the position of point A is determined by the argument **bsatur** a value of 0 puts A on the background prototype , a value of 0.5 puts it half way between the 2

prototypes. Point B is controlled similarly: 0 puts it at the foreground prototype and 0.5 puts it halfway. `colors` is a 2x3 matrix of rgb color clusters for the background and foreground respectively.

(**rgbaim-projectcolors** `rgbaim rgbaim2 proto polarity bsatur fsatur`)
(lsh/libimage/rgbaim-greypage.lsh) Last Modified: 2002-10-14 03:30:51

Basically does the same thing as `rgbaim-greys`, but builds a color image instead of a grey image.

See: (`rgbaim-greys` `rgbaim ubim proto polarity bsatur fsatur`)

Color Segmentation on RGBA float images

(lsh/libimage/rgbafim-greypage.lsh) Last Modified: 2002-10-09 22:44:26

a set of function for detecting the background and foreground colors in a document image and for transforming them into grey-level images suitable for recognition, segmentation, etc... essentially sets the background to 0, the foreground to 256, the rest in the middle.

(**rgbafim-cluster-colors** `rgbafim ubim polarity bsatur fsatur proto weight variance`) (lsh/libimage/rgbafim-greypage.lsh) Last Modified: 2002-10-09 22:44:26

See arguments of `grey-page`. This function performs the actual color clustering, using the K-means algorithm. Results in

`proto` : `n_colors` x 3 matrix = colors of the clusters `weight` : `n_colors` vector = weights of the Gaussian mixture for each cluster `variance` : `n_colors` vector = variance parameter of each cluster Gaussian distribution

See: (`rgbafim-greypage` `rgbafim ubim polarity bsatur fsatur`)

(**rgbafim-project-greys** `rgbafim ubim polarity bsatur fsatur colors`)
(lsh/libimage/rgbafim-greypage.lsh) Last Modified: 2002-10-09 22:44:26

transforms an rgb image into a grey image where the foreground is 256, the background is 0, and the greys in between are between 256 and 0 the resulting image can be used for segmentation and recognition. `rgbafim` an `idx3` of flts containing the input RGBA image. `ubim` an `idx2` of flts that will contain the result it should be the same size as `rgbafim` `polarity` determines the polarity of the image: 0 means light background and dark foreground, 1 means dark background and light foreground, -1 should be passed if the polarity is unknown. In that case, the color cluster with the most numerous pixels is assigned to the background. pixel colors in `rgbafim` are projected on a straight segment between two color prototypes computed with K-means. The output pixels are given a grey value that depends on their position along that segment. The function that maps positions on the line to grey-levels is piece-wise linear saturation-type function. It is controlled by two points A and B. below A, points are assigned the value 0, above B they are assigned 256. in between they are mapped linearly. the position of point A is determined by the argument `bsatur` a value of 0 puts A on the background prototype, a value of 0.5 puts it half way between the 2 prototypes. Point B is controlled similarly: 0 puts it at the foreground prototype and 0.5 puts it halfway. `colors` is a 2x3 matrix of rgb color clusters for the background and foreground respectively.

(`rgbafim-project-colors` `rgbafim` `rgbafim2` `polarity` `bsatur` `fsatur` `colors`) (lsh/libimage/rgbafim-greypage.lsh) Last Modified: 2002-10-09 22:44:26
 project the colors of an image onto a line joining two center clusters. `rgbafim` an `idx3` of `flts` containing the input RGBA image. `rgbafim2` contains the result image on output. it should be the same size as `<rgbafim polarity` determines the polarity of the image: 0 means light background and dark foreground , 1 means dark background and light foreground , -1 should be passed if the polarity is unknown. In that case , the color cluster with the most numerous pixels is assigned to the background. pixel colors in `rgbafim` are projected on a straight segment between two color prototypes computed with K-means. The output pixels are given a grey value that depends on their position along that segment. The function that maps positions on the line to grey-levels is piece-wise linear saturation-type function. It is controlled by two points A and B. below A , points are assigned the value 0 , above B they are assigned 256. in between they are mapped linearly. the position of point A is determined by the argument `bsatur` a value of 0 puts A on the background prototype , a value of 0.5 puts it half way between the 2 prototypes. Point B is controlled similarly: 0 puts it at the foreground prototype and 0.5 puts it halfway. `colors` is a 2x3 matrix of `rgb` color clusters for the background and foreground respectively.

(`rgbafim-quickquant` `rgbafim` `ubim` `ppal` `assign`) (lsh/libimage/rgbafim-greypage.lsh) Last Modified: 2002-10-09 22:44:26

quick color quantization using result of `cqu-kmeans` `rgbafim` is the RGBA image to be quantized `ubim` is the output image (of floats). It must be the same size as `rgbafim` . `ppal` is the color histogram as computed by `rgbafim-histo32` `assign` is the vector of color labels for each color in the histogram this is "quick" only for fairly large images.

(`rgbafim-greysquant` `rgbafim` `ubim` `proto` `polarity` `bsatur` `fsatur`) (lsh/libimage/rgbafim-greypage.lsh) Last Modified: 2002-10-09 22:44:26

transforms and RGB image into a grey image with uniform foreground and background. `proto` is a 2x3 matrix of floats containing two color prototypes for background and foreground colors. transforms an `rgb` image into a grey image where the foreground is 256 , the background is 0 , and the greys in between are between 256 and 0 the resulting image can be used for segmentation and recognition. `rgbafim` an `idx3` of `flts` containing the input RGBA image. `ubim` an `idx2` of `flts` that will contain the result it should be the same size as `rgbafim` `polarity` determines the polarity of the image: 0 means light background and dark foreground , 1 means dark background and light foreground , -1 should be passed if the polarity is unknown. In that case , the color cluster with the most numerous pixels is assigned to the background. pixel colors in `rgbafim` are projected on a straight segment between two color prototypes computed with K-means. The output pixels are given a grey value that depends on their position along that segment. The function that maps positions on the line to grey-levels is piece-wise linear saturation-type function. It is controlled by two points A and B. below A , points are assigned the value 0 , above B they are assigned 256. in between they are mapped linearly. the position of point A is determined by the argument `bsatur` a value of 0 puts A on the background

prototype , a value of 0.5 puts it half way between the 2 prototypes. Point B is controlled similarly: 0 puts it at the foreground prototype and 0.5 puts it halfway. `colors` is a 2x3 matrix of rgb color clusters for the background and foreground respectively.

(`rgbafim-projectcolors` `rgbafim` `rgbafim2` `proto` `polarity` `bsatur` `fsatur`)
(lsh/libimage/rgbafim-greypage.lsh) Last Modified: 2002-10-09 22:44:26

Basically does the same thing as `rgbafim-greyquant` , but builds a color image instead of a grey image.

See: (`rgbafim-greyquant` `rgbafim` `ubim` `proto` `polarity` `bsatur` `fsatur`
)

Chapter 9

Packages

Author(s): Yann LeCun

Packages are libraries that are not considered part of the standard library because they may not be available on all the platforms, and they may not be useful to everybody. Packages are generally built around popular pre-existing C libraries that have been interfaced to Lush with a layer to facilitate their use in Lush applications. Current packages include

- An extensive library of classes for gradient-based machine learning algorithms, including neural nets, convolutional nets, and many others.
- an environment for writing simple video games (SDL)
- an interface to the popular 3D graphic API OpenGL/GLUT
- an interface to the OpenInventor 3D scene rendering engine.
- a full interface to the LAPACK/BLAS Linear algebra libraries
- an interface to the Intel Computer Vision Library
- an interface to the Video4Linux video-grabbing API
- an interface to the Advanced Linux Sound Architecture API (Alsa)

9.1 GBLearn2: Machine Learning Library

Author(s): Yann LeCun

The GBLearn2 Library is an object-oriented framework for building , training and testing machine learning systems and algorithms. The "GB" stands for "gradient-based" , but the library can handle learning algorithms that are not gradient-based.

The basic concepts of GBLearn2 are: modules , states , params , trainers , data sources , and meters.

9.1.1 GBLearn2: Basic Concepts

Modules

Modules (subclasses of `gb-module`) are components of learning machines. Modules generally accept at least the method `fprop`, which computes the output of the module from its inputs. Input and outputs are generally subclasses of `idx-state`. Many commonly used modules are pre-defined, including various types of neural nets layers, convolutional layers, RBFs, Softmax and many others. Many of these modules use inputs and outputs that are of class `idx3-ddstate` (more on this below).

Some modules provide two other methods: `bprop` and `bbprop`. Those methods are defined for modules that are used in gradient-based machine learning algorithms. Method `bprop` back-propagates gradients through the module, while `bbprop` backpropagate diagonal second derivatives.

An large collection of high-level "macro-modules" is provided which include various types of fully-connected neural nets, convolutional networks, time-delay neural nets, etc.

Machine learning systems are usually built by using one of these predefined classes, or by subclassing them. Specific machine are generally created by defining classes whose slots are `gb-modules` and `gb-states`, with appropriate `fprop` and `bprop` methods.

States

Modules exchange information through subclasses of `gb-state`.

Most predefined modules use `idx3-state` and its subclasses as input and output. The class `idx3-ddstate` stores numbers arranged in a tensor with 3 indices. Most of those modules (e.g. as convolutional layers) interpret the first index as a feature index, and the last two indexes as spatial dimensions. This allows to easily replicate the modules spatially (e.g. for applications that require scanning a classifier over all locations of an image, more on this later).

Machines and Cost Modules

Basic modules generally do not assume much about the kind of learning algorithm with which they will be trained. The most common form of training is gradient-based training. gradient-based training consists in finding the set of parameters that minimize a particular energy function (generally computed by averaging over a set of training examples).

Classes are provided to conveniently assemble trainable modules and cost modules (e.g. energy functions) into a complete learning machine. An example of such class is `idx3-supervised-module`. This class can be used for supervised training of classifiers. The `idx3-supervised-module` contains:

- a learning machine with one input and one output (both `idx3-ddstate`).

- a cost module whose first inputs is the output of the machine , whose second input is a desired category (stored in an `idx0` of `int`) , and whose output is an `idx0-ddstate` which contains the value of the energy.
- a classifier module which takes the machine's output and computes a `class-state` which stores the label of the recognized category together with confidence scores and runner-up categories.

Parameters

Modules may have trainable parameters that are computed or adjusted by the various learning algorithms. Learning algorithms are designed to operate on `gb-param` objects , so as to make the algorithm implementations independent of the details of the learning machine's architecture. Whenever a instance of a module is created , a `gb-param` object is passed to the constructor. The modules trainable parameters are added to the `param` object.

Trainers

Learning machines created as above are then inserted in a `gb-trainer` object. The `gb-trainer` class (or rather , its subclasses) define the learning algorithm being used to train the machine. The most commonly used `gb-trainer` is `supervised-gradient` , which implements learning algorithms based on stochastic gradient descent.

The `supervised-gradient` trainer class expects a machine whose `fprop` and `bprop` methods take input , an output , a desired output , and an energy. `supervised-gradient` is an generic interpreted class which makes no hard assumptions on the type of the objects it manipulates as long as those objects understand the required set of methods.

Training a machine in such a trainer is performed by calling one of the training methods (e.g. `train` in the case of `supervised-gradient`) with a data source and a `gb-meter` (and possibly other specific parameters , depending on the trainer and method considered).

Data Sources

A data source is a source of training and testing examples. Data sources are a bit like regular modules , but they generally don't have inputs. Most `gb-trainer` subclasses expect data sources to undertand the method `size` , `seek` , `tell` , and `next` , which they use to sweep through the set of examples.

Since many pre-defined machine expect an `idx3-state` as input , and an `idx0` of `int` as desired output , the data source class `dsource-idx31` exists for that purpose. Subclasses of `dsource-idx31` are defined for databases of examples stored as vectors of floats or ubytes , or for databases of images of variable size.

Performance Meters

Meters are used by `gb-trainer` to measure the performance of learning machines. A commonly used one is the `classifier-meter` which can be used to measure the performance of supervised classifiers.

9.1.2 Examples and Demos

Learning the proverbial XOR with backprop

(`packages/gblearn2/demos/xor.lsh`) Last Modified: 2002-12-18 23:27:31

the content of this file is a complete example of how to build a single hidden layer backprop net to solve the XOR problem. Yeah , it's a lame example. Just have a look at the source code.

Lenet5 Convolutional Neural Network Demo

(`packages/gblearn2/demos/lenet5.lsh`)

Author(s): Yann LeCun , December 2002 Last Modified: 2002-12-18 23:27:31

This is an example of how to train a convolutional network on the MNIST database (handwritten digits). The database can be obtained at <http://yann.lecun.com> A paper describing an experiment similar to the one demonstrated here is described in LeCun , Bottou , Bengio , Haffner: "gradient-based learning applied to document recognition" , Proceedings of the IEEE , Nov 1998. This paper is also available at the above URL. This demo assumes that the MNIST data files are in `LUSHDIR/local/mnist`. If you installed them someplace else , simply do (`setq *mnist-dir* "your-mnist-directory"`) before loading the present demo. The MNIST datafiles are:

- `train-images-idx3-ubyte` : training set , images , 60000 samples.
- `train-labels-idx1-ubyte` : training set , labels , 60000 samples.
- `t10k-images-idx3-ubyte` : testing set , images , 10000 samples.
- `t10k-labels-idx1-ubyte` : testing set , labels , 10000 samples.

9.1.3 Learning Algorithms

(`packages/gblearn2/gb-trainers.lsh`) Author(s): Yann LeCun Last Modified: 2002-12-20 02:32:37

Various learning algorithm classes are defined to train learning machines. Learning machines are generally subclasses of `gb-module` . Learning algorithm classes include gradient descent for supervised learning , and others.

eb-trainer

(packages/gblearn2/gb-trainers.lsh) Last Modified: 2002-12-20 02:32:37

Abstract class for energy-based learning algorithms. The class contains an `input` , a (trainable) parameter `parameter` , and an energy. this is an abstract class from which actual trainers can be derived.

supervised

(packages/gblearn2/gb-trainers.lsh) Last Modified: 2002-12-20 02:32:37

A an abstract trainer class for supervised training with of a feed-forward classifier with discrete class labels. Actual supervised trainers can be derived from this. The machine's `fprop` method must have four arguments: `input` , `output` , `energy` , and `desired output`. A call to the machine's `fprop` must look like this:

```
(==> machine fprop input output desired energy)
```

By default , `output` must be a `class-state` , `desired` an `idx0` of `int` (integer scalar) , and `energy` and `idx0-ddstate` (or subclasses thereof). The meter passed to the training and testing methods should be a `classifier-meter` , or any meter whose update method looks like this:

```
(==> meter update output desired energy)
```

where `output` must be a `class-state` , `desired` an `idx0` of `int` , and `energy` and `idx0-ddstate` .

(**new supervised m p [e in out des]**) Type: CLASS(packages/gblearn2/gb-trainers.lsh) Last Modified: 2002-12-20 02:32:37

create a new **supervised** trainer. Arguments are as follow:

- `m` : machine to be trained.
- `p` : trainable parameter object of the machine.
- `e` : energy object (by default an `idx0-ddstate`).
- `in` : input object (by default an `idx3-ddstate`).
- `out` : output object (by default a `class-state`).
- `des` : desired output (by default an `idx0` of `int`).

(**==> supervised train dsource mtr**) Type: MSG(packages/gblearn2/gb-trainers.lsh) Last Modified: 2002-12-20 02:32:37

train the machine with on the data source `dsource` and measure the performance with `mtr` . This is a dummy method that should be defined by subclasses.

(**==> supervised test dsource mtr**) Type: MSG(packages/gblearn2/gb-trainers.lsh) Last Modified: 2002-12-20 02:32:37

measures the performance over all the samples of data source `dsource` . `mtr` must be an appropriate meter.

(`==> supervised test-sample dsource mtr i`) Type: MSG(`packages/gblearn2/gb-trainers.lsh`) Last Modified: 2002-12-20 02:32:37

measures the performance over a single sample of data source `dsource` . This leaves the internal state of the meter unchanged , and can be used for a quick test of a whether a particular pattern is correctly recognized or not.

supervised-gradient

(`packages/gblearn2/gb-trainers.lsh`) Last Modified: 2002-12-20 02:32:37

A basic trainer object for supervised stochastic gradient training of a classifier with discrete class labels. This is a subclass of `supervised` . The machine's `fprop` method must have four arguments: `input` , `output` , `energy` , and `desired output` . A call to the machine's `fprop` must look like this:

```
(==> machine fprop input output desired energy)
```

where `output` must be a `class-state` , `desired` an `idx0` of `int` (integer scalar) , and `energy` and `idx0-ddstate` . The meter passed to the training and testing methods should be a `classifier-meter` , or any meter whose update method looks like this:

```
(==> meter update output desired energy)
```

where `output` must be a `class-state` , `desired` an `idx0` of `int` , and `energy` and `idx0-ddstate` . The trainable parameter object must understand the following methods:

- (`==> param clear-dx`) : clear the gradients.
- (`==> param update eta inertia`) : update the parameters with learning rate `eta` , and momentum term `inertia` .

If the diagonal hessian estimation is to be used , the param object must also understand:

- (`==> param clear-ddx`) : clear the second derivatives.
- (`==> param update-ddeltas knew kold`) : update average second derivatives.
- (`==> param compute-epsilons mu`) : set the per-parameter learning rates to the inverse of the sum of the second derivative estimates and `mu` .

(`new supervised-gradient m p [e in out des]`) Type: CLASS(`packages/gblearn2/gb-trainers.lsh`) Last Modified: 2002-12-20 02:32:37

create a new `supervised-gradient` trainer. Arguments are as follow:

- **m** : machine to be trained.
- **p** : trainable parameter object of the machine.
- **e** : energy object (by default an `idx0-ddstate`).
- **in** : input object (by default an `idx3-ddstate`).
- **out** : output object (by default a `class-state`).
- **des** : desired output (by default an `idx0` of `int`).

(==> supervised-gradient train-online dsource mtr n eta inertia)
 Type: MSG(packages/gblearn2/gb-trainers.lsh) Last Modified: 2002-12-20 02:32:37
 train with stochastic (online) gradient on the next `n` samples of data source `dsource` with global learning rate `eta` . and "momentum term" `inertia` . Record performance in `mtr` . `mtr` must understand the following methods:

```
(==> mtr update age output desired energy)
(==> mtr info)
```

where `age` is the number of calls to parameter updates so far , `output` is the machine's output (most likely a `class-state`) , `desired` is the desired output (most likely an `idx0` of `int`) , and `energy` is an `idx0-state` . The `info` should return a list of relevant measurements.

(==> supervised-gradient train dsource mtr eta inertia) Type: MSG(packages/gblearn2/gb-trainers.lsh) Last Modified: 2002-12-20 02:32:37
 train the machine on all the samples in data source `dsource` and measure the performance with `mtr` .

(==> supervised-gradient compute-diaghessian dsource n mu) Type: MSG(packages/gblearn2/gb-trainers.lsh) Last Modified: 2002-12-20 02:32:37
 Compute per-parameter learning rates (epsilons) using the stochastic diagonal levenberg marquardt method (as described in LeCun et al. "efficient back-prop" , available at [<http://yann.lecun.com>]). This method computes positive estimates the second derivative of the objective function with respect to each parameter using the Gauss-Newton approximation. `dsource` is a data source , `n` is the number of patterns (starting at the current point in the data source) on which the estimate is to be performed. Each parameter-specific learning rate `epsiloni` is computed as $1/(H_{ii} + \mu)$, where H_{ii} are the diagonal Gauss-Newton estimates and `mu` is the blowup prevention fudge factor.

(==> supervised-gradient saliencies ds n) Type: MSG(packages/gblearn2/gb-trainers.lsh) Last Modified: 2002-12-20 02:32:37

Compute the parameters saliencies as defined in the Optimal Brain Damage algorithm of (LeCun , Denker , Solla , NIPS 1989) , available at <http://yann.lecun.com>. This computes the first and second derivatives of the energy with respect to each parameter averaged over the next `n` patterns of data source `ds` . A vector of saliencies is returned. Component `i` of the vector contains $S_i = -G_i * W_i + 1/2 H_{ii} * W_i^2$, this is an estimate of how much the energy would increase

if the parameter was eliminated (set to zero). Parameters with small saliencies can be eliminated by setting their value and epsilon to zero.

9.1.4 Data Sources

(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

data sources are used in machine learning experiments as sources of training and testing examples. Data sources are pretty much like regular modules. Most data sources should implement the following methods to be compatible with most gradient-based learning algorithms

- (`==> ds fprop arg1 ... argn`) : write the current item into `arg1 ... argn` .
- (`==> ds next`) : go to the next item in the data source.
- (`==> ds seek i`) : go to the *i*-th item in the data source.
- (`==> ds tell`) : return the index of the current item.
- (`==> ds size`) : return the total number of item.

Optionally , some data sources may implement the following methods:

- (`==> ds bprop arg1 ... argn`) : backpropagate gradients.
- (`==> ds bbprop arg1 ... argn`) : backpropagate diagonal second derivatives.

These methods can be used by algorithms that updates samples or preprocessings internal to the data source , based on gradients coming from the learning machine. However , since those methods are used by a very small number of very special learning algorithms , implementing them is generally not necessary.

dsource

(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

Semi-abstract class for a data source. Various subclasses are defined for special cases. This root class defines the methods `seek` , `tell` , `next` , and `size`. Subclasses should define (or redefine) at least `fprop` and `size`.

(`==> dsource seek i`) Type: MSG(packages/gblearn2/data-sources.lsh)

Last Modified: 2002-12-10 10:26:58

set pointer to *i* -th item

(`==> dsource tell`) Type: MSG(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

return index of current item.

(`==> dsource next`) Type: MSG(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

move pointer to next item.

(**==> dsource size**) Type: MSG(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

reutrn the size of the data source.

dsource-idx3l

(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

An abstract data source class appropriate for most supervised learning algorithms that take input data in the form of an `idx3`-state and a label in the form of an `idx0` of int. Most learning machines implemented in the library are designed to be compatible with this data source and its subclasses. This class just defines the `fprop` method prototypes. Subclasses actually implement the functionalities.

(**==> dsource-idx3l fprop out lb1**) Type: MSG(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

get the current item and copy the sample into `out` (an `idx3`-state) and the corresponding label into `lb1` (and `idx0` of int).

dsource-idx3fl

(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

a data source that stores samples as `idx3` of floats. As a subclass of `dsource-idx3l`, this class is compatible with most learning machines defined in the `gblearn2` library.

(**new dsource-idx3fl inp lb1**) Type: CLASS(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

create a `dsource-idx3fl` where the input samples are slices of an `PxDxYxX` `idx4` of floats passed as argument `inp` and labels are slices of an `idx1` of int passed as argument `lb1`. `P` is the number of samples, and `D`, `Y`, `X` the dimensions of each sample. If your data needs less than three dimensions simply set `Y` and/or `X` to 1 (i.e. `inp` would be a `PxDx1x1` matrix).

(**==> dsource-idx3fl fprop out lb1**) Type: MSG(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

get the current item and copy the sample into `out` (an `idx3`-state) and the corresponding label into `lb1` (and `idx0` of int).

dsource-idx3ul

(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

a data source that stores samples as `idx3` of ubytes. As a subclass of `dsource-idx3l`, this class is compatible with most learning machines defined in the `gblearn2` library. The ubyte values (between 0 255) can be shifted and scaled by before being written to the output `idx3`-state. This data source is considerably more economical in term of memory than `dsource-idx3fl` (1 byte per sample per variable, versus 4).

(**new dsource-idx3ul inp lb1 bias coeff**) Type: CLASS(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

create a `dsource-idx3ul` where the input samples are `idx3` slices of an `PxDxYxX` `idx4` of ubytes passed as argument `inp` and labels are slices of an `idx1` of int passed as argument `lb1` . `P` is the number of samples , and `D` , `Y` , `X` the dimensions of each sample. If your data needs less than three dimensions simply set `q` `Y` and/or `X` to 1 (i.e. `inp` would be a `PxDx1x1` matrix). values in the item are shifted by `bias` and multiplied by `coeff` before being copied into the destination by `fprop` .

(`==>` `dsource-idx3ul fprop out lb1`) Type: MSG(`packages/gblearn2/data-sources.lsh`) Last Modified: 2002-12-10 10:26:58

get the current item and copy the sample into `out` (an `idx3`-state) and the corresponding label into `lb1` (and `idx0` of int). Raw values are shifted by the `bias` parameter and multiplied by the `coeff` parameter before being copied into the `x` slot of `out` .

dsource-image

(`packages/gblearn2/data-sources.lsh`) Last Modified: 2002-12-10 10:26:58

A data source that stores images of variable sizes , with any number of components per pixel , and one byte per component. As a subclass of `dsource-idx3l` , this class is compatible with most learning machines defined in the `gblearn2` library. The `fprop` method resizes the output argument to the size of the current image. Pixel values may be shifted and scaled before being written the the output `idx3-state` .

(`new dsource-image bias coeff`) Type: CLASS(`packages/gblearn2/data-sources.lsh`) Last Modified: 2002-12-10 10:26:58

create an empty `dsource-image` . When accessing an item , the values are shifted by `bias` and multiplied by `coeff` .

(`==>` `dsource-image size`) Type: MSG(`packages/gblearn2/data-sources.lsh`) Last Modified: 2002-12-10 10:26:58

reutnrs the size of the database

(`==>` `dsource-image fprop dest lb1`) Type: MSG(`packages/gblearn2/data-sources.lsh`) Last Modified: 2002-12-10 10:26:58

writes the current item into the `x` slot of `dest` (an `idx3`-state) and the corresponding label to `lb1` (and `idx0` of int). `dest` is automatically resized to the size of the current item.

(`==>` `dsource-image load-ppms flist clist`) Type: MSG(`packages/gblearn2/data-sources.lsh`) Last Modified: 2002-12-10 10:26:58

fills up the database with images from a bunch of PPM files `clist` is a list of labels (one integer for each image) or a single category , in which case all the loaded images will be in that category.

(`==>` `dsource-image load-pgms flist clist`) Type: MSG(`packages/gblearn2/data-sources.lsh`) Last Modified: 2002-12-10 10:26:58

fills up the database with images from a bunch of PGM files `clist` is a list of labels (one integer for each image) or a single category , in which case all the loaded images will be in that category.

(**==> dsource-image save basename**) Type: MSG(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

save database into pre-cooked IDX files. These files can be subsequently loaded quickly using the `load` or `map` methods. The database is saved in four files named `basename images.idx` , `basename starts.idx` , `basename sizes.idx` , and `basename labels.idx`

(**==> dsource-image load basename**) Type: MSG(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

load database from pre-cooked IDX files produced through the `save` method. The database will be loaded from four files named `basename images.idx` , `basename starts.idx` , `basename sizes.idx` , and `basename labels.idx`

(**==> dsource-image map basename**) Type: MSG(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

memory-map database from pre-cooked IDX files produced through the `save` method. This is MUCH faster than `load` , and consumes fewer memory/disk bandwidth. The database will be mapped from four files named `basename images.idx` , `basename starts.idx` , `basename sizes.idx` , and `basename labels.idx`

dsource-idx3l-narrow

(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

a data source constructed by taking patterns in an existing data source whose indices are within a given range.

(**new dsource-idx3l-narrow base size offset**) Type: CLASS(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

make a new data source by taking `size` items from the data source passed as argument , starting at item `offset` .

(**==> dsource-idx3l-permute fprop out lb1**) Type: MSG(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

copy current item and label into `out` and `lb1` .

dsource-idx3l-permute

(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

a data source constructed by shuffling the items of a base data source.

(**new dsource-idx3l-permute base**) Type: CLASS(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

make a new data source by taking the `base` data source and building a permutation map of its items. Initially , the permutation map is equal to the identity. It can be shuffled into a random order by calling the `shuffle` method.

(**==> dsource-idx3l-permute fprop out lb1**) Type: MSG(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

copy current item and label into `out` and `lb1` .

(**==> dsource-idx3l-permute shuffle**) Type: MSG(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

randomly shuffles the samples of the db.

dsource-idx3l-concat

(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

a data source constructed by concatenating two base data sources.

(new dsource-idx3l-concat base1 base2) Type: CLASS(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

make a new data source by taking concatenating two base data sources **base1** and **base2** . the Two data sources must be **dsource-idx3l** or subclasses thereof.

(=> dsource-idx3l-concat fprop out lb1) Type: MSG(packages/gblearn2/data-sources.lsh) Last Modified: 2002-12-10 10:26:58

copy item and label into **out** and **lb1** .

9.1.5 Meters

(packages/gblearn2/gb-meters.lsh) Author(s): Yann LeCun Last Modified: 2002-11-26 18:33:14

Meters are classes used to measure the performance of learning machines. There are several types of meters for each specific situation. meters are generally assumed to have at least the following methods:

- **update**: updates the meter with the objects and values passed as argument.
- **clear**: resets the meter , so it can be used for a new series of measurements.
- **test**: simply prints performance information for the data passed as argument. This does not update any internal state.

Methods are provided to compute and display the information measured by a meter.

- **display**: display performance information on the terminal
- **info**: returns a list of the informations printed by display

(same-class? actual desired [dummies])

(packages/gblearn2/gb-meters.lsh) Last Modified: 2002-11-26 18:33:14

return 0 if **actual** equals -1 , otherwise , return 1 if **actual** and **desired** are equal , -1 otherwise.

classifier-meter

(packages/gblearn2/gb-meters.lsh) Last Modified: 2002-11-26 18:33:14

a class that can be used to measure the performance of classifiers. This is a simple version that does not record anything but simply computes performance measures.

(**new classifier-meter** [comparison-function]) Type: CLASS(packages/gblearn2/gb-meters.lsh) Last Modified: 2002-11-26 18:33:14

Create a new **classifier-meter** using **comparison-function** to compare actual and desired answers. By default the **same-class?** function is used for that purpose. It takes two integer arguments , and returns 1 if they are equal , -1 if they are different , and 0 if the first argument is -1 (which means reject).

(**==> classifier-meter clear**) Type: MSG(packages/gblearn2/gb-meters.lsh) Last Modified: 2002-11-26 18:33:14

reset the meter. This must be called before a measurement campaign is started.

(**==> classifier-meter update age actual desired energy**) Type: MSG(packages/gblearn2/gb-meters.lsh) Last Modified: 2002-11-26 18:33:14

update the meter with results from a new sample. **age** is the number of training iterations so far , **actual** (a **class-state**) the actual output of the machine , **desired** (an idx0 of int) the desired category , and **energy** (an idx0-state) the energy.

(**==> classifier-meter info**) Type: MSG(packages/gblearn2/gb-meters.lsh) Last Modified: 2002-11-26 18:33:14

return a list with the age , the number of samples (number of calls to update since the last clear) , the average energy , the percentage of correctly recognize samples , the percentage of erroneously recognized samples , and the percentage of rejected samples.

(**==> classifier-meter display**) Type: MSG(packages/gblearn2/gb-meters.lsh) Last Modified: 2002-11-26 18:33:14

Display the meter's information on the terminal. namely , the age , the number of samples (number of calls to update since the last clear) , the average energy , the percentage of correctly recognize samples , the percentage of erroneously recognized samples , and the percentage of rejected samples.

9.1.6 gb-states

(packages/gblearn2/gb-states.lsh) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

states are the main objects used to store state variables in or between modules.

(**each-idx-slots** o s f1 [f2...[fn]])

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

evaluates lists in **f** for all active idx slots of object **o** , (as returned by (**==>** o idx-slots)) with **s** successively taking the values of the slots in question.

idx-state

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

an idx-state is a state that contains vector or matrix variables.

(==> idx-state idx-slots) Type: MSG(packages/gblearn2/gb-states.lsh)
Last Modified: 2002-12-18 23:27:25

return the complete list of idx slots.

(new idx-state [params]) Type: CLASS(packages/gblearn2/gb-states.lsh)
Last Modified: 2002-12-18 23:27:25

create a new simple-state. if **params** is absent , the simple-state is left unsized. Otherwise **params** must be a list of integer (possibly empty) , which will be used to determine the size of the object. **params** has the same meaning as the 2nd parameter of the functions new-index.

(==> idx-state resize [list-of-dims]) Type: MSG(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

resize all the idx slots (as returned by idx-slots) **list-of-dims** is a list with p numbers specifying the size of the last p dimensions of each slot. If **list-of-dim** is absent, all the active slots are undimmed.

See: **(==> idx-state idx-slots)**

(==> idx-state undim) Type: MSG(packages/gblearn2/gb-states.lsh)
Last Modified: 2002-12-18 23:27:25

undim all the active idx slots of **idx-state** .

(==> idx-state clear) Type: MSG(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

fill all the active idx slots with zeroes

See: **(==> idx-state idx-slots)**

(==> idx-state load s) Type: MSG(packages/gblearn2/gb-states.lsh)
Last Modified: 2002-12-18 23:27:25

fill slot x with content of file **s**

(==> idx-state save s (packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

save content of slot x into file **s** .

(==> idx-state dump s (packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

save the entire object into file **s**

idx0-state

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx0-dstate

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx0-ddstate

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx1-state

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx1-dstate

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx1-ddstate

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx2-state

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx2-dstate

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx2-ddstate

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx3-state

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx3-dstate

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx3-ddstate

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx4-state

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx4-dstate

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

idx4-ddstate

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

class-state

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

a special kind of state used to store the output of a classifier. class-state are generated by modules such as class-max , and used by meters such as classifier-meter. No backprop is possible through a class-state.

lclass-state

(packages/gblearn2/gb-states.lsh) Last Modified: 2002-12-18 23:27:25

a special kind of state used to store the output of a spatial classifier such as ledist-classer.

9.1.7 gb-param

(packages/gblearn2/gb-params.lsh) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

a hierarchy of classes that implements gradient-based learning algorithms. Various subclasses of gb-param can be used for various learning algorithms (stochastic gradient , with separate epsilons , with second derivatives.....) The standard gb-param type contains idx1 slots.

(new-index-offset s dlist o)

(packages/gblearn2/gb-params.lsh) Last Modified: 2002-12-18 23:27:25

like new index , but puts the index at a specific offset in the storage **s** is a storage , **dlist** a list of dimensions , **o** an offset. Compilable macro.

idx1-param

(packages/gblearn2/gb-params.lsh) Last Modified: 2002-12-18 23:27:25

a gb-param whose slots are idx1. no learning algorithm is defined. Only the x slot is present. This class is useful for fixed (non adaptive) parameters.

idx1-dparam

(packages/gblearn2/gb-params.lsh) Last Modified: 2002-12-18 23:27:25

gb-param class for regular stochastic gradient descent algorithm

(new idx1-dparam s sts) Type: CLASS(packages/gblearn2/gb-params.lsh)
Last Modified: 2002-12-18 23:27:25

s is a size (possibly zero) , and **sts** is the initial storage size which must be larger than zero , and (can be larger than **s** to avoid unnecessary reallocs when the size of the param is increased.

(==> idx1-dparam update-deltas knew kold) Type: MSG(packages/gblearn2/gb-params.lsh) Last Modified: 2002-12-18 23:27:25

update the average derivatives **deltas** as follow: **deltas = knew*dx + kold*deltas** . The **update** method calls this whenever it is called with a non-zero **inertia** parameter.

(==> idx1-dparam update eta inertia) Type: MSG(packages/gblearn2/gb-params.lsh) Last Modified: 2002-12-18 23:27:25

simple gradient descent update. This will use momentum if the **inertia** parameter is non zero. CAUTION: the **deltas** slot is not updated if the **inertia** is zero. When **inertia** is non zero , the **deltas** are updated as follows: **deltas**

= $(1 - \text{inertia}) * \text{dx} + \text{inertia} * \text{deltas}$ (where dx is the gradient) and the parameters are subsequently updated as follows: $\mathbf{x} = \text{eta} * \text{deltas}$.

idx1-dparam-eps

(packages/gblearn2/gb-params.lsh) Last Modified: 2002-12-18 23:27:25

gb-param class for gradient descent with separate epsilons for each parameter
(new idx1-dparam-eps s sts) Type: CLASS(packages/gblearn2/gb-params.lsh)
 Last Modified: 2002-12-18 23:27:25

\mathbf{s} is a size (possibly zero) , and \mathbf{sts} is the initial storage size which must be larger than zero , and (can be larger than \mathbf{s} to avoid unnecessary reallocs when the size of the param is increased.

(==> idx1-dparam-eps update eta inertia) Type: MSG(packages/gblearn2/gb-params.lsh) Last Modified: 2002-12-18 23:27:25

simple gradient descent update with one individual learning rate per parameter. eta is a global learning rate by which each individual parameter learning rate will be multiplied. This will perform an update "with momentum" if the inertia parameter is non zero. CAUTION: the deltas slot is not updated if inertia is zero. When inertia is non zero , the deltas are updated as follows: $\text{deltas} = (1 - \text{inertia}) * \text{dx} + \text{inertia} * \text{deltas}$ (where dx is the gradient) and the parameters are subsequently updated as follows: $\mathbf{x} = \text{eta} * \text{deltas}$.

(==> idx1-dparam-eps set-epsilons m) Type: MSG(packages/gblearn2/gb-params.lsh) Last Modified: 2002-12-18 23:27:25

copy the values in vector \mathbf{m} to epsilons

(==> idx1-dparam-eps set-epsilons m) Type: MSG(packages/gblearn2/gb-params.lsh) Last Modified: 2002-12-18 23:27:25

copy the values in vector \mathbf{m} to epsilons

idx1-ddparam

(packages/gblearn2/gb-params.lsh) Last Modified: 2002-12-18 23:27:25

a gb-param class for the stochastic diagonal levenberg-marquardt algorithm. In addition to the usual update method , it has an update-bbprop method for computing the second derivatives , and a set-epsilons method to set the epsilons using the second derivatives.

(new idx1-ddparam s alloc) Type: CLASS(packages/gblearn2/gb-params.lsh)
 Last Modified: 2002-12-18 23:27:25

\mathbf{s} is the size (can be 0) \mathbf{alloc} is the size of storages to be preallocated. This will prevent memory fragmentation when the size of the gb-param is subsequently increased.

(==> idx1-ddparam update-ddeltas knew kold) Type: MSG(packages/gblearn2/gb-params.lsh) Last Modified: 2002-12-18 23:27:25

update average second derivative ddeltas as follows: $\text{ddeltas} = \text{knew} * \text{ddx} + \text{kold} * \text{ddeltas}$. where ddx is the instantaneous second derivative.

(`==> idx1-ddparam saliencies`) Type: MSG(packages/gblearn2/gb-params.lsh)
 Last Modified: 2002-12-18 23:27:25

Compute the parameters saliencies as defined by the Optimal Brain Damage algorithm of (LeCun , Denker , Solla , NIPS 1989). This uses the average first and second derivatives of the energy with respect to each parameter to compute a saliency for each weight. The saliency is an estimate of how much the energy would increase by if the parameter was set to zero. It is computed as: $S_i = -G_i * W_i + 1/2 H_{ii} * W_i^2$. A vector of saliencies is returned. The `deltas` and `ddeltas` field must have relevant values before this function is called.

(`==> idx1-ddparam compute-epsilons mu`) Type: MSG(packages/gblearn2/gb-params.lsh) Last Modified: 2002-12-18 23:27:25

compute and set the epsilons using the second derivative. this method should be called after a few iteration of update-bbprop

9.1.8 Module Libraries

modules

(packages/gblearn2/gb-modules.lsh) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

In Lush, building and training a complex system is done by assembling basic blocks, called modules. Modules are subclasses of the class `gb-module`. Though there are several predefined module classes, you can define your own pretty easily. modules must understand 2 basic methods, `fprop` and `bprop`, whose arguments are the inputs and outputs of the module. Optionally most module should understand a `bbprop` method for computing diagonal second derivatives. Modules can have as many input/output "ports" as desired. these "ports" are passed as arguments to the methods that require them, such as `fprop` and `bprop`. In most cases, these arguments belong to the class `idx-state`, or one of its subclasses. Some modules may have internal trainable parameters. When this is the case, an `idx3-ddparam` object must be passed to the constructor. Internal parameters will then be allocated in that param. the `bprop` and `bbprop` methods ACCUMULATE gradients in these parameters so multiple modules can share a single parameter and automatically compute the correct gradient. Gradients on input ports are NOT accumulated.

A special class called `trainer` provides a convenient way to train and test a module combined with pre- and post-processors. Once a module has been created, inserting it in an instance of the `trainer` class is the easiest and fastest way to train it on a database and to measure its performance. the `trainer` class understands methods such as `train`, `test` etc... Most of these methods take instances of database as argument. They also take another argument called a meter. A meter is an object whose role is to keep track of the performance of the machine during a training or test session. trainers, meters, and databases can be put in an instance of `workbench` that handles standard learning sequences (estimate second derivatives, train, test....)

A certain number of predefined basic modules are provided in `modules.sn`.

This includes `idx3-module`, a "root" class of modules with one input and one output, both of type `idx3-state`. Many predefined modules are subclasses of `idx3-module`. Also included are `idx3-squasher` (sigmoid layer), `logadd-layer` (transforms a `idx3-state` into an `idx1-state` by log-adding over the 2 spatial dimensions), `mle-cost` (cost module for minimizing the cost of the desired output).

gb-module (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

The class `gb-module` is the basic class for objects that can be used with the library of training routines. Specific trainable modules and cost functions etc... are subclasses of `gb-module` and can be combined to build complex adaptive machines. `gb-module` are expected to accept at least the methods `fprop`, `bprop`, and optionally the following methods: `bbprop`, `load`, and `save`. The external "plugs" of a `gb-module` are passed as argument to the methods. For example, the `fprop` method of a module with one input vector and one output vector, and one parameter vector can be called with

```
(==> <gb-module> fprop <input> <parameter> <output>)
```

where `input`, `parameter` and `output` are instances of the `gb-state` or one of its subclasses. As a convention, the methods `fprop`, `bprop`, and `bbprop` take the same arguments in the same order. Results of these methods are accumulated in the appropriate slot of the objects passed as parameters. This allows modules to share inputs and outputs while preserving the correctness of forward and backward propagations

a few convenient subclasses of `gb-module` are predefined in the `gblearn2` library. This includes cost functions, classifiers, and others.

(==> `gb-module fprop [args]`) Type: MSG(packages/gblearn2/gb-modules.lsh)
Last Modified: 2002-12-18 23:27:25

performs a forward propagation on the `gb-module`. `args` are optional arguments which represent the external "plugs" of the module. When possible, modules with variable size outputs resize their output ports automatically.

See: (==> `gb-module bprop [args]`)

See: (==> `gb-module bbprop [args]`)

(==> `gb-module bprop [args]`) Type: MSG(packages/gblearn2/gb-modules.lsh)
Last Modified: 2002-12-18 23:27:25

performs a backward propagation on the `gb-module` (propagates derivatives). `args` are optional arguments which represent the external "plugs" of the module. By convention, the list of `args` is the same as for the `fprop` method. `bprop` assumes `fprop` has been called beforehand. If the module has internal parameters, the `bprop` method will ACCUMULATE the gradients in it, so that multiple modules can share the same parameters

See: (==> `gb-module fprop [args]`)

See: (==> `gb-module bbprop [args]`)

(==> `gb-module bbprop [args]`) Type: MSG(packages/gblearn2/gb-modules.lsh)
Last Modified: 2002-12-18 23:27:25

performs a backward propagation of second derivatives on the `gb-module` `args` are optional arguments which represent the external "plugs" of the module. By convention, the list of `args` is the same as for the `fprop` method. `bbprop` assumes `fprop` and `bprop` have been run beforehand. If the module has internal parameters, the `bbprop` method will ACCUMULATE second derivatives in it, so that multiple modules can share the same parameters

See: (`==>` `gb-module` `fprop` [`args`])

See: (`==>` `gb-module` `bprop` [`args`])

noop-module (`packages/gblearn2/gb-modules.lsh`) Last Modified: 2002-12-18 23:27:25

a module that does not do anything (a place-holder). This is NOT an identity-function module not compilable

See: `id-module`

See: `gb-module`

id-module (`packages/gblearn2/gb-modules.lsh`) Last Modified: 2002-12-18 23:27:25

identity function module. It's a straight pass-through forward and backward. arguments must be `idx-ddstates` non compilable.

idx4-module (`packages/gblearn2/gb-modules.lsh`) Last Modified: 2002-12-18 23:27:25

a basic "root" class for modules that have one single `idx-state` input and one single `idx4-state` output. the `fprop` , `bprop` and `bbprop` methods of this root class merely act as identity functions

idx3-module (`packages/gblearn2/gb-modules.lsh`) Last Modified: 2002-12-18 23:27:25

a basic "root" class for modules that have one single `idx-state` input and one single `idx3-state` output. the `fprop` , `bprop` and `bbprop` methods of this root class merely act as identity functions

idx2-module (`packages/gblearn2/gb-modules.lsh`) Last Modified: 2002-12-18 23:27:25

a basic "root" class for modules that have one single `idx-state` input and one single `idx2-state` output. the `fprop` , `bprop` and `bbprop` methods of this root class merely act as identity functions

idx1-module (`packages/gblearn2/gb-modules.lsh`) Last Modified: 2002-12-18 23:27:25

a basic "root" class for modules that have one single `idx-state` input and one single `idx1-state` output. the `fprop` , `bprop` and `bbprop` methods of this root class merely act as identity functions

idx4-squasher (`packages/gblearn2/gb-modules.lsh`) Last Modified: 2002-12-18 23:27:25

a basic squashing function layer for `idx4-state` you can udefine subclasses of this to change the squashing function

idx3-squasher (`packages/gblearn2/gb-modules.lsh`) Last Modified: 2002-12-18 23:27:25

a basic squashing function layer for `idx3-state` you can udefine subclasses of this to change the squashing function

idx4-halfsquare (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

takes half square of each component.

idx3-halfsquare (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

takes half square of each component.

logadd-layer (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

performs a log-add over spatial dimensions of an idx3-state output is an idx1-state

cost (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25
costs are a special type of modules (although there is no definite subclass for them) with two inputs and one output. the output is an idx0-ddstate which stores a cost or energy. one of the inputs is meant to be the output of another module (e.g. a network) , and the other input a desired output (or any kind of supervisor signal like a reinforcement). the gradient slot (dx) of the output state is generally filled with +1. That way , the bprop method of the cost module automatically computes the gradient.

idx3-cost (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

abstract class for a cost function that takes an idx3-state as input , an int as desired output , and an idx0-state as energy.

mle-cost (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

a cost module that propagates the output corresponding to the desired label. If the output is interpreted as a negative log likelihood , minimizing this output is equivalent to maximizing the likelihood. outputs are log-added over spatial dimensions in case of spatial replication.

(**new mle-cost classes si sj**) Type: CLASS(packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

make a new mle-cost. **classes** is an integer vector which contains the labels associated with each output. From that vector , the reverse table is constructed to map labels to class indices. Elements in **classes** must be positive or 0 , and not be too large , as a table as large as the maximum value is allocated. **si** and **sj** are the expected maximum sizes in the spatial dimensions (used for preallocation to prevent memory fragmentation).

mmi-cost (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

a cost function that maximizes the mutual information between the actual output and the desired output. This assumes that the outputs are costs , or negative log likelihoods. this modules accepts spatially replicated inputs.

(**new mmi-cost classes priors si sj prm**) Type: CLASS(packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

make a new mmi-cost. **classes** is an integer vector which contains the labels associated with each output. From that vector , the reverse table is constructed to map labels to class indices. Elements in **classes** must be positive or 0 ,

and not be too large , as a table as large as the maximum value is allocated. **priors** : an `idx1` of `gbtypes` , whose size must be the size of `classes` +1. It specifies the prior probability for each class , and for the junk class. The prior for the junk class must be in the last element. In absence of a better guess , the prior vector should be filled with $1/n$, where n is its size. **si** and **sj** are the expected maximum sizes in the spatial dimensions (used for preallocation to prevent memory fragmentation). **prm** is an `idx1-ddparam` in which the value that determines the constant cost of the junk class will be stored. If the system is to be trained without junk examples , this parameter can be set to a very large value , and not be trained. The effect of setting this parameter to a fixed value is to softly saturate the costs of all the class to the half-square of that value (the overall energy will never be significantly larger then the half-square of the set value) , and to softly clip the gradients , i.e. the units whose cost is higher than the half-square of the set value will receive negligible gradients. The parameter can be learned ONLY IF junk examples (with label -1) are present in the training set. There is a method , called `set-junk-cost` that allows to directly set the value of the junk without having to figure out the half-square business.

(`==> mmi-cost set-junk-cost c`) Type: MSG(`packages/gblearn2/gb-modules.lsh`)

Last Modified: 2002-12-18 23:27:25

set the constant cost of the junk class to `c` . the underlying parameter is given the value ($\sqrt{* 2 c}$) , so `c` must be positive.

edist-cost (`packages/gblearn2/gb-modules.lsh`) Last Modified: 2002-12-18 23:27:25

a replicable Euclidean distance cost function. computes the log-sum over the 2D spatial output of the half squared error between the output and the prototype with the desired label. this does not generate gradients on the prototypes

(`new edist-cost classes si sj p`) Type: CLASS(`packages/gblearn2/gb-modules.lsh`) Last Modified: 2002-12-18 23:27:25

make a new `edist-cost`. `classes` is an integer vector which contains the labels associated with each output. From that vector , the reverse table is constructed to map labels to class indices. Elements in `classes` must be positive or 0 , and not be too large , as a table as large as the maximum value is allocated. **si** and **sj** are the expected maximum sizes in the spatial dimensions (used for preallocation to prevent memory fragmentation). **p** is an `idx2` containing the prototype for each class label. The first dimension of **p** should be equal to the dimension of `classes` . the second dimension of **p** should be equal to the number of outputs of the previous module. The costs are "log-summed" over spatial dimensions

wedist-cost (`packages/gblearn2/gb-modules.lsh`) Last Modified: 2002-12-18 23:27:25

a replicable weighted Euclidean distance cost function. computes the log-sum over the 2D spatial output of the weighted half squared error between the output and the prototype with the desired label. this does not generate gradients on the prototypes.

(`new wedist-cost classes si sj p w`) Type: CLASS(`packages/gblearn2/gb-modules.lsh`) Last Modified: 2002-12-18 23:27:25

make a new `wedist-cost`. `classes` is an integer vector which contains the labels associated with each output. From that vector, the reverse table is constructed to map labels to class indices. Elements in `classes` must be positive or 0, and not be too large, as a table as large as the maximum value is allocated. `si` and `sj` are the expected maximum sizes in the spatial dimensions (used for preallocation to prevent memory fragmentation). `p` is an `idx2` containing the prototype for each class label, and `w` is an `idx2` with a single weight for each of these prototype and each of its elements. The first dimension of `p` (and `w`) should be equal to the dimension of `classes`. the second dimension of `p` (and `w`) should be equal to the number of outputs of the previous module. The costs are "log-summed" over spatial dimensions

weighted-mse-cost (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

This is similar to `wedist-cost` but the weights matrix may run over patterns. The desired output vector has size two: the first element gives the class label, and the second element gives the position (row) in the weights matrix to use for the weighted euclidean distance. It is a replicable weighted Euclidean distance cost function. computes the log-sum over the 2D spatial output of the weighted half squared error between the output and the prototype with the desired label. this does not generate gradients on the prototypes

(**new weighted-mse-cost** `classes si sj p w`) Type: CLASS(packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

make a new `weighted-mse-cost`. `classes` is an integer vector which contains the labels associated with each output. From that vector, the reverse table is constructed to map labels to class indices. Elements in `classes` must be positive or 0, and not be too large, as a table as large as the maximum value is allocated. `si` and `sj` are the expected maximum sizes in the spatial dimensions (used for preallocation to prevent memory fragmentation). `p` is an `idx2` containing the prototype for each class label, and `w` is an `idx2` with a single weight for each of these prototype and each of its elements. The first dimension of `p` (and `w`) should be equal to the dimension of `classes`. the second dimension of `p` (and `w`) should be equal to the number of outputs of the previous module. The costs are "log-summed" over spatial dimensions

ledist-cost (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

a replicable Euclidean distance cost function with LOCAL TARGETS at each position. Target prototypes are associated to classes. The cost is the sum over the 2D output of the half squared error between the local output and the prototype with the desired label at that position. This does not generate gradients on the prototypes

(**new ledist-cost** `classes p`) Type: CLASS(packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

make a new `ledist-cost`. `classes` is an integer vector which contains the labels associated with each output. From that vector, the reverse table is constructed to map labels to class indices. Elements in `classes` must be positive or 0, and not be too large, as a table as large as the maximum value is allocated.

`p` is an `idx2` containing the prototype for each class label. The first dimension of `p` should be equal to the dimension of `classes`. The second dimension of `p` should be equal to the number of outputs of the previous module. The costs are summed over spatial dimensions.

Classifiers

(packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

idx3-classifier (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

The `idx3-classifier` module take an `idx3-state` as input and produce a `class-state` on output. A `class-state` is used to represent the output of classifiers with a discrete set of class labels.

min-classer (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

a module that takes an `idx3-state`, finds the lowest value and output the label associated with the index (in the first dimension of the state) of this lowest value. It actually sorts the labels according to their score (or costs) and outputs the sorted list.

(new min-classer classes) Type: CLASS(packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

makes a new min-classer. `classes` is an integer vector which contains the labels associated with each output.

max-classer (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

a module that takes an `idx3-state`, finds the lowest value and output the label associated with the index (in the first dimension of the state) of this lowest value. It actually sorts the labels according to their score (or costs) and outputs the sorted list.

(new max-classer classes) Type: CLASS(packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

makes a new max-classer. `classes` is an integer vector which contains the labels associated with each output.

edist-classer (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

a replicable Euclidean distance pattern matcher, which finds the class prototype "closest" to the output, where "close" is based on the log-added euclidean distances between the prototype and the output at various positions. This corresponds to finding the class whose a-posteriori probability is largest, when $P(c|data) = \sum_{[position=x]} P(c \text{ at } x | data \text{ at } x) / n_{\text{positions}}$ and the priors over classes are uniform, and the local class likelihoods $P(data \text{ at } x | c \text{ at } x)$ are Gaussian with unit variance and mean = `prototype(c)`.

ledist-classer (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

a replicable Euclidean distance pattern matcher, which finds the class prototype closest to the output for the vectors at each position in the output image.

(**new ledist-classer classes p**) Type: CLASS(packages/gblearn2/gb-modules.lsh)
Last Modified: 2002-12-18 23:27:25

make a new ledist-classer. **classes** is an integer vector which contains the labels associated with each prototype. **p** is an idx2 containing the prototype for each class label. The first dimension of **p** should be equal to the dimension of **classes** . the second dimension of **p** should be equal to the number of outputs of the previous module.

mmi-classer (packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

a classifier that computes class scores based on an MMI type criterion (a kind of softmax in log) It gives scores (cost) for all classes including junk. It should be used in conjunction with mmi-cost. This assumes that the output of the previous module are costs , or negative log likelihoods. this modules accepts spatially replicated inputs.

(**new mmi-classer classes priors si sj prm**) Type: CLASS(packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

makes a new mmi-classer. The arguments are identical to that of mmi-cost. In fact if an mmi-classer is to used in conjunction with an mmi-cost , they should share the prior vector and the parameter. sharing the parameter can be done by first building the classer , then reducing the size of the parameter by one , then creating the cost.

(**==> mmi-classer set-junk-cost c**) Type: MSG(packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

set the constant cost of the junk class to **c** . the underlying parameter is given the value ($\sqrt{2c}$) , so **c** must be positive. BE CAREFUL that the junk parameter of an mmi-classer is usually shared by an mmi-cost , changing one will change the other.

(**build-ascii-proto targets charset**) (packages/gblearn2/gb-modules.lsh)
Last Modified: 2002-12-18 23:27:25

idx3-supervised-module

(packages/gblearn2/gb-modules.lsh) Last Modified: 2002-12-18 23:27:25

a module that takes an idx3 as input , runs it through a machine , and runs the output of the machine through a cost function whose second output is the desired label stored in an idx0 of int.

Basic Neural Net Modules

(packages/gblearn2/gb-modules-nn.lsh) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

a bunch of standard compiled modules for building neural net architectures: sigmoids , RBF , fully connected , convolutions , subsampling layers. Each of these classes has methods for fprop , bprop , bbprop , load , save , forget , and setting various parameters.

f-layer (packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

full connection between replicable 3D layers the full connection is only across the first dimension of the input and output layer. the other two dimensions are spatial dimensions across which the weight matrix is shared. This is much more efficient than using a c-layer with a 1x1 convolution and a full-table.

(**new f-layer tin tout si sj sqsh**) Type: CLASS(packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

makes a new replicable fully-connected layer. **tin** and **tout** are the number of units in the input and output. **si** and **sj** are the vertical and horizontal sizes used for pre-allocating the internal state variables (they , as well as the output. are automatically resized , but the resizing is more efficient of it is a downsizing).

(**==> f-layer save file**) Type: MSG(packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

save weights and biases

(**==> f-layer load file**) Type: MSG(packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

load weights and biases

(**==> f-layer set-squash squash**) Type: MSG(packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

sets the squashing function to **squash** . **squash** must be an instance of the the class **idx3-module** .

(**==> f-layer forget v p**) Type: MSG(packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

initialize weights to random values uniformly distributed between -c and c , where c is $v/(fanin^{1/p})$. reasonable values for v and p are 2.5 and 2.

(**==> f-layer fprop in out**) Type: MSG(packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

forward prop with **in** as input and **out** as output. **in** and **out** must both be **idx3-state** .

(**==> f-layer bprop in out**) Type: MSG(packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

backward prop with **in** as input and **out** as output. **in** and **out** must both be **idx3-state** .

(**==> f-layer bbprop in out**) Type: MSG(packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

backward prop of second derivatives with **in** as input and **out** as output. **in** and **out** must both be **idx3-state** .

e-layer (packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

full connection with Euclidean distance between replicable 3D layers. This is like a layer of RBF WITHOUT NON-LINEAR FUNCTIONS. the output is the square Euclidean distance between the input and the weight the full connection is only across the first dimension of the input and output layer. the other two dimensions are spatial dimensions across which the weight matrix is shared.

(**new e-layer tin tout prm**) Type: CLASS(packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

new e-layer (Eucliden distance RBF). **tin** is the thickness of the input layer , **tout** is the thickness of the output layer , **prm** is the parameter.

e-layer-sparse (packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

like e-layer , but each RBF is only connected to a subset of the inputs this should be used for "multiple bitmap" outputs.

(**new e-layer-sparse tin off prm**) Type: CLASS(packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

new e-layer (Eucliden distance RBF). **tin** is the thickness of the input layer (dimension of prototypes) , **off** is an idx1 with as many elements as there are output. each element contains the offset in the input for the input vector of each output. **prm** is the parameter.

c-layer (packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

convolutional layer module. Performs multiple convolutions between an idx3-state input and an idx3-state output. includes bias and sigmoid.

(**new c-layer ki kj ri rj tbl thick si sj sqsh**) Type: CLASS(packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

Creates a new convolution layer. **ki** (int) vertical kernel size. **kj** (int) horizontal kernel size. **ri** (int) vertical stride (number of pixels by which the kernels are stepped) <**rj** (int) horizontal stride <**tbl** (idx2) N by 2 matrix containing source and destination feature maps for coresponding kernel **thick** (int) thickness of output layer (number of feature maps) **si** (int) vertical size for preallocation of internal state **sj** (int) horizontal size for preallocation of internal state **sqsh** (idx3-module) a squashing function module that operates on idx3-state. **prm** and idx1-ddparam from which the parameters will be allocated

s-layer (packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

subsampling layer class

(**new s-layer ki kj thick si sj sqsh prm**) Type: CLASS(packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

Creates a new subsampling layer for convolutional nets **ki** (int) vertical subsampling ratio. **kj** (int) horizontal subsampling ratio. **thick** (int) thickness of output layer (number of feature maps) **si** (int) vertical size for preallocation of internal state **sj** (int) horizontal size for preallocation of internal state **sqsh** (idx3-module) a squashing function module that operates on idx3-state. **prm** and idx1-ddparam from which the parameters will be allocated

poly2-layer (packages/gblearn2/gb-modules-nn.lsh) Last Modified: 2002-12-18 23:27:25

output all input variables and all cross-products of two variables

net-ff

(packages/gblearn2/net-ff.lsh) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

standard one-hidden layer, fully-connected neural net. this network is replicable.

(**new net-ff** thickin thick0 thick1 ini inj prm) Type: CLASS(packages/gblearn2/net-ff.lsh) Last Modified: 2002-12-18 23:27:25

create a new net-ff network.

<thickin>: number of inputs.

<thick0>: size of hidden layer

<thick1>: number of outputs

<ini> <inj>: expected max size of input over which the network should be replicated.

<prm> an idx1-ddparam in which the parameters will be allocated.

net-c

(packages/gblearn2/net-c.lsh) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

a neural net class with one convolutional layer and one fully connected layer. The main purpose of this class is to make replicable fully connected networks. Unlike with net-ff, the input spatial dimensions are treated correctly here. So if you want to recognize characters, use this class rather than net-ff.

(**new net-c** ini inj ki0 kj0 inthick outthick prm) Type: CLASS(packages/gblearn2/net-c.lsh) Last Modified: 2002-12-18 23:27:25

makes a new net-cscscf module. **ini inj** : expected max size of input for preallocation of internal states **ki0 kj0** : kernel size for first convolutional layer a standard fully-connected network can be obtained when **ini = ki0** and **inj = kj0** . **inthick** : number of input slices **outthick** : number of outputs. **prm** an idx1-ddparam in which the parameters will be allocated.

net-cf

(packages/gblearn2/net-cf.lsh) Last Modified: 2002-12-18 23:27:25

a neural net class with one convolutional layer and one fully connected layer. The main purpose of this class is to make replicable fully connected networks. Unlike with net-ff , the input spatial dimensions are treated correctly here. So if you want to recognize characters , use this class rather than net-ff.

(**new net-cf** ini inj ki0 kj0 tb10 outthick prm) Type: CLASS(packages/gblearn2/net-cf.lsh) Last Modified: 2002-12-18 23:27:25

makes a new net-cscscf module. **ini inj** : expected max size of input for preallocation of internal states **ki0 kj0** : kernel size for first convolutional layer a standard fully-connected network can be obtained when **ini = ki0** and **inj = kj0** . **tb10** : table of connections between input anf feature maps for

first layer `outthick` : number of outputs. `prm` an `idx1-ddparam` in which the parameters will be allocated.

net-cfe

(`packages/gblearn2/net-cfe.lsh`) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

neural net with one convolutional layer, one (replicable) fully-connected layer, and one RBF (euclidean distance) layer.

(**new net-cfe** `n e thick si sj`) Type: CLASS(`packages/gblearn2/net-cfe.lsh`) Last Modified: 2002-12-18 23:27:25

create a new net-cfe. `n` is a net-cf, `e` is a e-layer, `thick` is the number of outputs, and `si`, `sj` are the number of output replications.

net-cff

(`packages/gblearn2/net-cff.lsh`) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

a neural net class with one convolutional layer and two fully connected layers. The main purpose of this class is to make replicable 2-hidden layer fully connected networks.

(**new net-cff** `ini inj ki0 kj0 tb10 f1thick outthick prm`) Type: CLASS(`packages/gblearn2/net-cff.lsh`) Last Modified: 2002-12-18 23:27:25

makes a new net-cff module. `ini inj` : expected max size of input for preallocation of internal states `ki0 kj0` : kernel size for first convolutional layer a standard fully-connected network can be obtained when `ini = ki0` and `inj = kj0`. `tb10` : table of connections between input and feature maps for first layer `f1thick` : number of hidden units in second hidden layer. `outthick` : number of outputs. `prm` an `idx1-ddparam` in which the parameters will be allocated.

net-ccc

(`packages/gblearn2/net-ccc.lsh`) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

convolutional net with 3 convolutional layers.

(**new net-ccc** `ini inj ki0 kj0 tb10 ki1 kj1 tb11 ki2 kj2 tb12 prm`) Type: CLASS(`packages/gblearn2/net-ccc.lsh`) Last Modified: 2002-12-18 23:27:25

makes a new net-ccc module. `ini inj` : expected max size of input for preallocation of internal states `ki0 kj0` : kernel size for first convolutional layer `tb10` : table of connections between input and feature maps for first layer `ki1 kj1 tb11` : kernel and table for next layer `ki2 kj2 tb12` : same for last convolution layer `prm` an `idx1-ddparam` in which the parameters will be allocated.

net-cscf

(packages/gblearn2/net-cscf.lsh) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

convolutional net where the layers are convolution, subsampling, convolution, full connect.

(**new net-cscf ini inj ki0 kj0 tb10 si0 sj0 ki2 kj2 tb12 outthick prm**) Type: CLASS(packages/gblearn2/net-cscf.lsh) Last Modified: 2002-12-18 23:27:25

makes a new net-cscf module. **ini inj** : expected max size of input for preallocation of internal states **ki0 kj0** : kernel size for first convolutional layer **tb10** : table of connections between input and feature maps for first layer **si0 sj0** : subsampling for first layer **ki2 kj2 tb12** : same for last convolution layer **outthick** : number of outputs. **prm** an idx1-ddparam in which the parameters will be allocated.

net-cscfe

(packages/gblearn2/net-cscfe.lsh) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

convolutional net where the layers are convolution, subsampling, convolution, full connect, and euclidean distance RBF.

(**new net-cscfe n e thick si sj**) Type: CLASS(packages/gblearn2/net-cscfe.lsh) Last Modified: 2002-12-18 23:27:25

n is a net-cscf, **e** is an e-dist layer, **thick** is the number of outputs. **si**, **sj** are the initial replication factors.

net-csccf

(packages/gblearn2/net-csccf.lsh) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

a standard LeNet2/LeNet5 architecture without the last RBF layer (look into net-cscscfe for that).

(**new net-csccf ini inj ki0 kj0 tb10 si0 sj0 ki1 kj1 tb11 ki2 kj2 tb12 outthick prm**) Type: CLASS(packages/gblearn2/net-csccf.lsh) Last Modified: 2002-12-18 23:27:25

makes a new net-csccf module. **ini inj** : expected max size of input for preallocation of internal states **ki0 kj0** : kernel size for first convolutional layer **tb10** : table of connections between input and feature maps for first layer **si0 sj0** : subsampling for first layer **ki1 kj1 tb11** : kernel and table for next layer **ki2 kj2 tb12** : same for last convolution layer **outthick** : number of outputs. **prm** an idx1-ddparam in which the parameters will be allocated.

net-cscscf

(packages/gblearn2/net-cscscf.lsh) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

Standard LeNet5-type architecture without the final e-dist RBF layer.

(**new net-cscscf** ini inj ki0 kj0 tbl0 si0 sj0 ki1 kj1 tbl1 si1 sj1 ki2 kj2 tbl2 outthick prm) Type: CLASS(packages/gblearn2/net-cscscf.lsh) Last Modified: 2002-12-18 23:27:25

makes a new net-cscscf module. **ini inj** : expected max size of input for preallocation of internal states **ki0 kj0** : kernel size for first convolutional layer **tbl0** : table of connections between input and feature maps for first layer **si0 sj0** : subsampling for first layer **ki1 kj1 tbl1 si1 sj1** : same for next 2 layers **ki2 kj2 tbl2** : same for last convolution layer **outthick** : number of outputs. **prm** an idx1-ddparam in which the parameters will be allocated.

net-cscscfe

(packages/gblearn2/net-cscscfe.lsh) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

Standard LeNet5-type architecture.

(**new net-cscscfe** n e thick si sj) Type: CLASS(packages/gblearn2/net-cscscfe.lsh) Last Modified: 2002-12-18 23:27:25

n is a net-cscscf , **e** is an e-dist layer , **thick** is the number of outputs. **si** , **sj** are the initial replication factors.

net-cscscscf

(packages/gblearn2/net-cscscscf.lsh) Author(s): Yann LeCun Last Modified: 2002-11-26 15:22:59

A deep convolutional net, with one convo/subsamp set than LeNet5.

(**new net-cscscscf** ini inj ki0 kj0 tbl0 si0 sj0 ki1 kj1 tbl1 si1 sj1 ki2 kj2 tbl2 si2 sj2 ki3 kj3 tbl3 outthick prm) Type: CLASS(packages/gblearn2/net-cscscscf.lsh) Last Modified: 2002-11-26 15:22:59

makes a new net-cscscf module. **ini inj** : expected max size of input for preallocation of internal states **ki0 kj0** : kernel size for first convolutional layer **tbl0** : table of connections between input and feature maps for first layer **si0 sj0** : subsampling for first layer **ki1 kj1 tbl1 si1 sj1** : same for next 2 layers... **ki2 kj2 tbl2 si2 sj2** : ...and the next 2 **ki3 kj3 tbl3** : same yet again for last convolution layer **outthick** : number of outputs. **prm** an idx1-ddparam in which the parameters will be allocated.

lenet5

(packages/gblearn2/net-lenet5.lsh) Last Modified: 2002-11-26 12:24:37

LeNet5 is a convolutional network architecture described in several publications , notably in [LeCun , Bottou , Bengio and Haffner 1998]: "gradient-based learning applied to document recognition" , Proc IEEE , Nov 1998. The paper is also available at [<http://yann.lecun.com>].

(**new-lenet5** image-height image-width ki0 kj0 si0 sj0 ki1 kj1 si1 sj1 hid output-size net-param) (packages/gblearn2/net-lenet5.lsh) Last Modified: 2002-11-26 12:24:37

create a new instance of `net-cscscf` implementing a LeNet-5 type convolutional neural net. This network has regular sigmoid units on the output, not an extra RBF layer as described in the Proc. IEEE paper. The network has 6 feature maps at the first layer and 16 feature maps at the second layer with a connection matrix between feature maps as described in the paper. Arguments:

```
<image-height> <image-width>: height and width of input image
<ki0> <kj0>: height and with of convolutional kernel, first layer.
<si0> <sj0>: subsampling ratio of subsampling layer, second layer.
<ki1> <kj1>: height and with of convolutional kernel, third layer.
<si1> <sj1>: subsampling ratio of subsampling layer, fourth layer.
<hid>: number of hidden units, fifth layer
<output-size>: number of output units
<net-param>: idx1-ddparam that will hold the trainable parameters
              of the network
```

example

```
(setq p (new idx1-ddparam 0 0.1 0.02 0.02 80000))
(setq z (new-lenet5 32 32 5 5 2 2 5 5 2 2 120 10 p))
```

ccc-tdnn

(`packages/gblearn2/ccc-tdnn.lsh`) Author(s): Yann LeCun Last Modified: 2002-12-17 21:36:10

Class for a Time Delay Neural Network with 3 convolutional layers (with local connectivity). This type of network is appropriate for classifying sequences of time/frequency representations such as cepstrum, MEL Scale spectrum, and other.

(`build-ccc-tdnn` `params` `freqs` `max-seq-len` `fk1` `fs1` `nh1` `tk1` `fk2` `fs2` `nh2` `tk2` `fk3` `fs3` `nh3` `tk3`) (`packages/gblearn2/ccc-tdnn.lsh`) Last Modified: 2002-12-17 21:36:10

Build a Time-Delay Neural Network for data such as spectral sequences in which the features have a topology, choosing the connection tables to obtain local connections in feature space (as specified by kernel size and stride for the frequency axis, for each of these layers). The input to the net is assumed to be $F \times T$ where F is for example the number of frequency channels (spectral representation) and T is the length of the sequence. Here, the network has 3 convolutional layers. The arguments are the following:

- `<params>` is a `idx1-ddparam` on which to allocate parameters for the layers.
- `<freqs>` is the number of input frequency channels.
- `<max-seq-len>` = maximum sequence length.
- `<fk1>`,`<fk2>`,`<fk3>` = sizes of frequency kernels (= width of local freq. windows)
- `<fs1>`,`<fs2>`,`<fs3>` = step sizes which separate the successive frequency windows

- <nh1>, <nh2>, <nh3> = number of hidden units per frequency channel, for each layer
- <tk1>, <tk2>, <tk3> = sizes of the temporal kernels, for each layer.

(**tdnn-present-pattern** from to mean idev temporal-window-size) (packages/gblearn2/ccc-tdnn.lsh) Last Modified: 2002-12-17 21:36:10

from is a $T \times f$ source matrix, **to** is a $f \times 1 \times T'$ destination state ($T' = T + \text{temporal-window-size} - 1$), **mean** is a f -vector to subtract from source, **idev** is a f -vector to multiply by source. **temporal-window-size** is the length of the input of the network yielding an output of length 1.

norm-ftdnn (packages/gblearn2/ccc-tdnn.lsh) Last Modified: 2002-12-17 21:36:10

Wrapper around `ccc-tdnn` which does input normalization and pads the input according to the network architecture.

(**new norm-ftdnn** n-inputs nh1 nh2 n-outputs weight-file norm-file [(-tk1 5)(-tk2 8)(-tk3 12) (-fk1 6)(-fk2 3)(-fk3 3) (-fs1 3)(-fs2 2)(-fs3 1) (max-seq-len 2000)]) Type: CLASS(packages/gblearn2/ccc-tdnn.lsh) Last Modified: 2002-12-17 21:36:10

net-xx

(packages/gblearn2/net-xx.lsh) Last Modified: 2002-12-18 23:27:25

generic `idx3-module` with two sub-modules in series.

(**new net-xx** mod1 thick1 si1 sj1 mod2) Type: CLASS(packages/gblearn2/net-xx.lsh) Last Modified: 2002-12-18 23:27:25

create new `net-xx` with `mod1` as first module and `mod2` as second module. `thick1`, `si1`, `sj1` are the initial sizes of the intermediate `idx3-ddstate` which is the output of `mod1` and the input of `mod2`.

net-xxx

(packages/gblearn2/net-xxx.lsh) Last Modified: 2002-12-18 23:27:25

generic `idx3-module` with three sub-modules in series.

(**new net-xx** mod1 thick1 si1 sj1 mod2 thick2 si2 sj2 mod3) Type: CLASS(packages/gblearn2/net-xxx.lsh) Last Modified: 2002-12-18 23:27:25

create new `net-xx` with `mod1` as first module `mod2` as second module, and `mod3` as the third module. `thick1`, `si1`, `sj1` are the initial sizes of the intermediate `idx3-ddstate` which is the output of `mod1` and the input of `mod2`. `thick2`, `si2`, `sj2` are the initial sizes of the intermediate `idx3-ddstate` which is the output of `mod2` and the input of `mod3`.

gb-conx

(packages/gblearn2/gb-conx.lsh) Author(s): Yann LeCun Last Modified: 2002-12-18 23:27:25

A module architecture that connects inputs to outputs (with exactly one input per output but maybe more than one output per input, i.e. a one to many map). The inputs and outputs are bidimensional (i.e. sequences of vectors). It

has no adjustable parameters. Instead it has internal non-adjustable parameters specifying the input/output connectivity.

(**new gb-conx n-outputs n-inputs indices**) Type: CLASS(packages/gblearn2/gb-conx.lsh) Last Modified: 2002-12-18 23:27:25

Construct a gb-conx architecture. The parameters are given when constructing the corresponding machine.

See: machine

See: gb-module

See: gb-conx

(**==> gb-conx fprop parameters input output**) Type: MSG(packages/gblearn2/gb-conx.lsh) Last Modified: 2002-12-18 23:27:25

Fprop method for a gb-conx architecture. Just gather the inputs into the output vectors. Both the input and output are assumed to be sequences of vectors [i,t] (2D matrices).

See: gb-conx

(**==> gb-conx bprop input output**) Type: MSG(packages/gblearn2/gb-conx.lsh) Last Modified: 2002-12-18 23:27:25

Bprop method for a gb-conx architecture. Just scatter with accumulate the derivatives of the outputs into the derivative of the inputs. Both the input and output are assumed to be sequences of vectors [i,t] (2D matrices).

See: gb-conx

(**==> gb-conx generate input output**) Type: MSG(packages/gblearn2/gb-conx.lsh) Last Modified: 2002-12-18 23:27:25

:output:x is a matrix whose columns have zeros and a single one. According to the permutation matrix **parameters**, a similar matrix :input:x is produced.

sum-arch (packages/gblearn2/gb-conx.lsh) Last Modified: 2002-12-18 23:27:25

An architecture that takes the sum of its inputs. The input is a 2D matrix whose first index is **n-inputs** and second index is variable (e.g. time). The output is zero-dimensional (an idx0). There are no parameters.

See: gb-module

See: (new sum-arch)

See: (**==> sum-arch fprop parameters input output**)

See: (**==> sum-arch bprop parameters input output**)

See: (**==> sum-arch output-dim inputx**)

See: (**==> sum-arch generate parameters input output**)

(**new sum-arch**) Type: CLASS(packages/gblearn2/gb-conx.lsh) Last Modified: 2002-12-18 23:27:25

Construct a sum-arch gb-module.

See: machine

See: gb-module

See: sum-arch

(**==> sum-arch output-dim inputx**) Type: MSG(packages/gblearn2/gb-conx.lsh) Last Modified: 2002-12-18 23:27:25

(**==> sum-arch fprop parameters input output**) Type: MSG(packages/gblearn2/gb-conx.lsh) Last Modified: 2002-12-18 23:27:25

Fprop method for a sum-arch architecture. Just add all the inputs. The input is assumed to be any matrix (except an idx0). The output is an idx0.

See: sum-arch

(==> **sum-arch bprop parameters input output**) Type: MSG(packages/gblearn2/gb-conx.lsh) Last Modified: 2002-12-18 23:27:25

Bprop method for a sum-arch architecture.

See: prod

(==> **sum-arch generate parameters input output**) Type: MSG(packages/gblearn2/gb-conx.lsh) Last Modified: 2002-12-18 23:27:25

Generate a 2-d sequence of inputs. A prior distribution on sequence length is chosen here to be uniform in [5 , 6 , ..10]. Ignore the output. For each time step , the input vector is a vector of 0's and a single 1 , chosen uniformly.

9.2 GBLearning: Gradient-Based Learning Library

Author(s): Yann LeCun

NOTE: this library is unsupported and obsolete. Please use the much improved `gblearn2` library instead.

The GBLearning Library is an object-oriented framework for building, training and testing gradient-based machine learning systems and algorithms.

It provides an extensive library of modules, such as neural net layers, convolutional layers, RBFs, SoftMax and many others. Each module is implemented as a Lush class with an "fprop" (forward propagation), a "bprop" (backward propagation of gradients), and a "bbprop" (backward propagation of diagonal Hessian) method. Modules can be assembled into large learning machines. An extensive collection of "macro-modules" are provided, including various types of convolutional networks (SDNN, TDNN, etc).

A "trainer" class is provided to handle such things as computing a cost function, computing gradients, and running learning procedures (parameter updates).

A "db" class is provided to handle training and test databases.

A "workbench" class is provided to handle sequential access to training and test databases as well a to handle performance measurements and reporting.

9.2.1 Examples and Demos

`gblearning/demos/xor-example.lsh`

9.2.2 modules

(packages/gblearning/modules.lsh) Author(s): Yann LeCun Last Modified: 2002-10-14 03:30:51

In Lush, building and training a complex system is done by assembling basic blocks, called modules. Modules are subclasses of the class `sn-module`. Though

there are several predefined module classes, you can define your own pretty easily. modules must understand 2 basic methods, fprop and bprop, whose arguments are the inputs and outputs of the module. Optionally most module should understand a bbprop method for computing diagonal second derivatives. Modules can have as many input/output "ports" as desired. these "ports" are passed as arguments to the methods that require them, such as fprop and bprop. In most cases, these arguments belong to the class idx-state, or one of its subclasses. Some modules may have internal trainable parameters. When this is the case, an idx3-ddparam object must be passed to the constructor. Internal parameters will then be allocated in that param. the bprop and bbprop methods ACCUMULATE gradients in these parameters so multiple modules can share a single parameter and automatically compute the correct gradient. Gradients on input ports are NOT accumulated.

A special class called trainer provides a convenient way to train and test a module combined with pre- and post-processors. Once a module has been created, inserting it in an instance of the trainer class is the easiest and fastest way to train it on a database and to measure its performance. the trainer class understands methods such as train, test etc... Most of these methods take instances of database as argument. They also take another argument called a meter. A meter is an object whose role is to keep track of the performance of the machine during a training or test session. trainers, meters, and databases can be put in an instance of workbench that handles standard learning sequences (estimate second derivatives, train, test....)

A certain number of predefined basic modules are provided in modules.sn. This includes idx3-module, a "root" class of modules with one input and one output, both of type idx3-state. Many predefined modules are subclasses of idx3-module. Also included are idx3-squasher (sigmoid layer), logadd-layer (transforms a idx3-state into an idx1-state by log-adding over the 2 spatial dimensions), mle-cost (cost module for minimizing the cost of the desired output).

sn-module

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

The class sn-module is the basic class for objects that can be used with the library of training routines. Classifiers, neural networks, HMM, cost functions etc... are subclasses of sn-module and can be combined to build complex adaptive machines. sn-module are required to accept at least the following methods: fprop, bprop and optionally the following methods: bbprop, load, save, recognize. the external "plugs" of an sn-module are passed as argument to the methods. For example, the fprop method of a module with one input vector and one output vector, and one parameter vector can be called with

```
(==> <sn-module> fprop <input> <parameter> <output>)
```

where `input` , `parameter` and `output` are instances of the class -state. As

a convention, the methods `fprop`, `bprop`, and `bbprop` take the same arguments in the same order.

a few convenient subclasses of `sn-module` are predefined in the library. This includes cost functions, classifiers, trainers, and others. A convenient way of training an `sn-module` is to insert it in an instance of the class `trainer`.

(`==> sn-module fprop [args]`) Type: MSG(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

performs a forward propagation on the `sn-module`. `args` are optional arguments which represent the external "plugs" of the module. When possible, modules with variable size outputs resize their output ports automatically.

See: (`==> sn-module bprop [args]`)

See: (`==> sn-module bbprop [args]`)

(`==> sn-module bprop [args]`) Type: MSG(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

performs a backward propagation on the `sn-module` (propagates derivatives). `args` are optional arguments which represent the external "plugs" of the module. By convention, the list of `args` is the same as for the `fprop` method. `bprop` assumes `fprop` has been called beforehand. If the module has internal parameters, the `bprop` method will ACCUMULATE the gradients in it, so that multiple modules can share the same parameters

See: (`==> sn-module fprop [args]`)

See: (`==> sn-module bbprop [args]`)

(`==> sn-module bbprop [args]`) Type: MSG(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

performs a backward propagation of second derivatives on the `sn-module`. `args` are optional arguments which represent the external "plugs" of the module. By convention, the list of `args` is the same as for the `fprop` method. `bbprop` assumes `fprop` and `bprop` have been run beforehand. If the module has internal parameters, the `bbprop` method will ACCUMULATE second derivatives in it, so that multiple modules can share the same parameters

See: (`==> sn-module fprop [args]`)

See: (`==> sn-module bprop [args]`)

noop-module

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a module that does not do anything (a place-holder). This is NOT an identity-function module not compilable

See: `id-module`

See: `sn-module`

id-module

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

identity function module. It's a straight pass-through forward and backward. arguments must be `idx-states` non compilable.

idx3-module

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a basic "root" class for modules that have one single idx-state input and one single idx3-state output. the fprop , bprop and bbprop methods of this root class merely act as identity functions

idx2-module

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a basic "root" class for modules that have one single idx-state input and one single idx2-state output. the fprop , bprop and bbprop methods of this root class merely act as identity functions

idx1-module

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a basic "root" class for modules that have one single idx-state input and one single idx1-state output. the fprop , bprop and bbprop methods of this root class merely act as identity functions

idx3-squasher

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a basic squashing function layer for idx3-state you can udefine subclasses of this to change the squashing function

idx3-halfsquare

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

takes half square of each component.

logadd-layer

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

performs a log-add over spatial dimensions of an idx3-state output is an idx1-state

cost (packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

costs are a special type of modules (although there is no definite subclass for them) with two inputs and one output. the output is an idx0-ddstate which stores a cost or energy. one of the inputs is meant to be the output of another module (e.g. a network) , and the other input a desired output (or any kind of supervisor signal like a reinforcement). the gradient slot (dx) of the output state is generally filled with +1. That way , the bprop method of the cost module automatically computes the gradient.

mle-cost

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a cost module that propagates the output corresponding to the desired label. If the output is interpreted as a negative log likelihood, minimizing this output is equivalent to maximizing the likelihood. outputs are log-added over spatial dimensions in case of spatial replication.

(**new mle-cost classes si sj**) Type: CLASS(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

make a new mle-cost. **classes** is an integer vector which contains the labels associated with each output. From that vector, the reverse table is constructed to map labels to class indices. Elements in **classes** must be positive or 0, and not be too large, as a table as large as the maximum value is allocated. **si** and **sj** are the expected maximum sizes in the spatial dimensions (used for preallocation to prevent memory fragmentation).

mmi-cost

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a cost function that maximizes the mutual information between the actual output and the desired output. This assumes that the outputs are costs, or negative log likelihoods. this modules accepts spatially replicated inputs.

(**new mmi-cost classes priors si sj prm**) Type: CLASS(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

make a new mmi-cost. **classes** is an integer vector which contains the labels associated with each output. From that vector, the reverse table is constructed to map labels to class indices. Elements in **classes** must be positive or 0, and not be too large, as a table as large as the maximum value is allocated. **priors**: an idx1 of floats, whose size must be the size of **classes** + 1. It specifies the prior probability for each class, and for the junk class. The prior for the junk class must be in the last element. In absence of a better guess, the prior vector should be filled with $1/n$, where n is its size. **si** and **sj** are the expected maximum sizes in the spatial dimensions (used for preallocation to prevent memory fragmentation). **prm** is an idx1-ddparam in which the value that determines the constant cost of the junk class will be stored. If the system is to be trained without junk examples, this parameter can be set to a very large value, and not be trained. The effect of setting this parameter to a fixed value is to softly saturate the costs of all the class to the half-square of that value (the overall energy will never be significantly larger than the half-square of the set value), and to softly clip the gradients, i.e. the units whose cost is higher than the half-square of the set value will receive negligible gradients. The parameter can be learned ONLY IF junk examples (with label -1) are present in the training set. There is a method, called set-junk-cost that allows to directly set the value of the junk without having to figure out the half-square business.

(**==> mmi-cost set-junk-cost c**) Type: MSG(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

set the constant cost of the junk class to c . the underlying parameter is given the value $(\sqrt{2c})$, so c must be positive.

edist-cost

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a replicable Euclidean distance cost function. computes the log-sum over the 2D spatial output of the half squared error between the output and the prototype with the desired label. this does not generate gradients on the prototypes

(new edist-cost classes si sj p) Type: CLASS(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

make a new edist-cost. **classes** is an integer vector which contains the labels associated with each output. From that vector , the reverse table is constructed to map labels to class indices. Elements in **classes** must be positive or 0 , and not be too large , as a table as large as the maximum value is allocated. **si** and **sj** are the expected maximum sizes in the spatial dimensions (used for preallocation to prevent memory fragmentation). **p** is an idx2 containing the prototype for each class label. The first dimension of **p** should be equal to the dimension of **classes** . the second dimension of **p** should be equal to the number of outputs of the previous module. The costs are "log-summed" over spatial dimensions

wedist-cost

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a replicable weighted Euclidean distance cost function. computes the log-sum over the 2D spatial output of the weighted half squared error between the output and the prototype with the desired label. this does not generate gradients on the prototypes.

(new wedist-cost classes si sj p w) Type: CLASS(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

make a new wedist-cost. **classes** is an integer vector which contains the labels associated with each output. From that vector , the reverse table is constructed to map labels to class indices. Elements in **classes** must be positive or 0 , and not be too large , as a table as large as the maximum value is allocated. **si** and **sj** are the expected maximum sizes in the spatial dimensions (used for preallocation to prevent memory fragmentation). **p** is an idx2 containing the prototype for each class label , and **w** is an idx2 with a single weight for each of these prototype and each of its elements. The first dimension of **p** (and **w**) should be equal to the dimension of **classes** . the second dimension of **p** (and **w**) should be equal to the number of outputs of the previous module. The costs are "log-summed" over spatial dimensions

weighted-mse-cost

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

This is similar to `wedist-cost` but the weights matrix may run over patterns. The desired output vector has size two: the first element gives the class label, and the second element gives the position (row) in the weights matrix to use for the weighted euclidean distance. It is a replicable weighted Euclidean distance cost function. computes the log-sum over the 2D spatial output of the weighted half squared error between the output and the prototype with the desired label. this does not generate gradients on the prototypes

(new weighted-mse-cost classes si sj p w) Type: CLASS(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

make a new weighted-mse-cost. `classes` is an integer vector which contains the labels associated with each output. From that vector, the reverse table is constructed to map labels to class indices. Elements in `classes` must be positive or 0, and not be too large, as a table as large as the maximum value is allocated. `si` and `sj` are the expected maximum sizes in the spatial dimensions (used for preallocation to prevent memory fragmentation). `p` is an `idx2` containing the prototype for each class label, and `w` is an `idx2` with a single weight for each of these prototype and each of its elements. The first dimension of `p` (and `w`) should be equal to the dimension of `classes`. the second dimension of `p` (and `w`) should be equal to the number of outputs of the previous module. The costs are "log-summed" over spatial dimensions

ledist-cost

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a replicable Euclidean distance cost function with LOCAL TARGETS at each position. Target prototypes are associated to classes. The cost is the sum over the 2D output of the half squared error between the local output and the prototype with the desired label at that position. This does not generate gradients on the prototypes

(new ledist-cost classes p) Type: CLASS(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

make a new ledist-cost. `classes` is an integer vector which contains the labels associated with each output. From that vector, the reverse table is constructed to map labels to class indices. Elements in `classes` must be positive or 0, and not be too large, as a table as large as the maximum value is allocated. `p` is an `idx2` containing the prototype for each class label. The first dimension of `p` should be equal to the dimension of `classes`. the second dimension of `p` should be equal to the number of outputs of the previous module. The costs are summed over spatial dimensions.

postproc (packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

classifier (packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

Postprocs and classifiers are modules through which no backprop are performed. postprocs have an `fprop` method that takes the output of a learning machine as argument and produces a "postprocessed" output. For example a particular subclass of `postproc-module` might take the output vectors from a network and produce a sorted list class labels with confidence factors. Another

one might take the output of an SDNN , run multiple-constrained-viterbi with a dictionary and return the best candidates. The output of a postproc is meant to be used directly by a "meter" in conjunction with a desired output for performance measures.

classifiers are special postprocessors that produce a class-state object on their output (which contain sorted class labels and confidence factors). Their outputs can be used by classifier-meter to measure the performance of a machine.

min-classer

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a module that takes an idx3-state , finds the lowest value and output the label associated with the index (in the first dimension of the state) of this lowest value. It actually sorts the labels according to their score (or costs) and outputs the sorted list.

(new min-classer classes) Type: CLASS(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

makes a new min-classer. **classes** is an integer vector which contains the labels associated with each output.

max-classer

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a module that takes an idx3-state , finds the lowest value and output the label associated with the index (in the first dimension of the state) of this lowest value. It actually sorts the labels according to their score (or costs) and outputs the sorted list.

(new max-classer classes) Type: CLASS(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

makes a new max-classer. **classes** is an integer vector which contains the labels associated with each output.

edist-classer

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a replicable Euclidean distance pattern matcher , which finds the class prototype "closest" to the output , where "close" is based on the log-added euclidean distances between the prototype and the output at various positions. This corresponds to finding the class whose a-posteriori probability is largest , when $P(c\text{---}data) = \sum_{[position=x]} P(c \text{ at } x \text{ ---}data \text{ at } x) / n_positions$ and the priors over classes are uniform , and the local class likelihoods $P(data \text{ at } x \text{ ---} c \text{ at } x)$ are Gaussian with unit variance and mean = prototype(c).

ledist-classer

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a replicable Euclidean distance pattern matcher , which finds the class prototype closest to the output for the vectors at each position in the output image.

(**new ledist-classer classes p**) Type: CLASS(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

make a new ledist-classer. **classes** is an integer vector which contains the labels associated with each prototype. **p** is an idx2 containing the prototype for each class label. The first dimension of **p** should be equal to the dimension of **classes** . the second dimension of **p** should be equal to the number of outputs of the previous module.

mmi-classer

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a classifier that computes class scores based on an MMI type criterion (a kind of softmax in log) It gives scores (cost) for all classes including junk. It should be used in conjunction with mmi-cost. This assumes that the output of the previous module are costs , or negative log likelihoods. this modules accepts spatially replicated inputs.

(**new mmi-classer classes priors si sj prm**) Type: CLASS(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

makes a new mmi-classer. The arguments are identical to that of mmi-cost. In fact if an mmi-classer is to used in conjunction with an mmi-cost , they should share the prior vector and the parameter. sharing the parameter can be done by first building the classer , then reducing the size of the parameter by one , then creating the cost.

(**==> mmi-classer set-junk-cost c**) Type: MSG(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

set the constant cost of the junk class to **c** . the underlying parameter is given the value ($\sqrt{2c}$) , so **c** must be positive. BE CAREFUL that the junk parameter of an mmi-classer is usually shared by an mmi-cost , changing one will change the other.

nearest-neighbor

(packages/gblearning/modules.lsh) Last Modified: 2002-10-14 03:30:51

a 1 nearest-neighbor classifier. (old and obsolete , still here for compatibility).

(**new nearest-neighbor pt proto labels [dist]**) Type: CLASS(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

create a new nearest-neighbor classifier using the prototypes in **proto** , and their associated labels in **labels** . **labels** must be a 1d vector of numbers (class indices). **dist** is an optional function for computing the distance between incoming patterns and prototypes. **dist** is an optional distance function to be used for the comparison. It must take two input arguments and one output argument that must be an idx0 (the distance between the first two arguments). The default for **dist** is idx-sqrdist (Euclidian distance between matrices of any

dimension). `proto` must be a matrix with one more dimension than the parameter of the distance function. `pt` is a punting threshold (nil means never punt).

(==> nearest-neighbor set-labels l) Type: MSG(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51
set the label matrix to `l`

(==> nearest-neighbor fprop m cs) Type: MSG(packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

`m` is an sn-state object , and `cs` a class-state object. compare slot `x` of `m` with the stored list of prototypes , and fill the slots of `cs` with the result. return a list with the class index , the confidence , the sorted list of classes and sorted list of scores. if the top score is below the punting threshold , -1 is returned as the class index.

(build-ascii-proto targets charset) (packages/gblearning/modules.lsh)
Last Modified: 2002-10-14 03:30:51

9.2.3 workbench

(packages/gblearning/workbench.lsh) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

the main class for training a machine. a workbench contains a trainer, two databases, and two meters. It provides a convenient user interface for many standard training. it keeps a logbook of all the operations performed on it.

(new workbench [logbook trainer traindb testdb trainmeter testmeter])

Type: CLASS(packages/gblearning/workbench.lsh) Last Modified: 2002-08-19 03:04:06

makes a new workbench for training and testing `trainer` on databases `traindb` and `testdb` . Measurements are performed with `trainmeter` and `testmeter` . All operations are recorded in file `logbook` .

(==> workbench init)

Type: MSG(packages/gblearning/workbench.lsh) Last Modified: 2002-08-19 03:04:06

reset pointers on training and test db. reset meters. reinitialize parameters (weights) of machine in trainer (using the forget method).

(==> workbench compute-diaghessian [n gamma mu])

Type: MSG(packages/gblearning/workbench.lsh) Last Modified: 2002-08-19 03:04:06

compute diagonal hessian and set the learning rates using the result. this should be called every few epochs to readjust the learning rates. `n` is the number of patterns in the training set (starting from pattern 0) on which the second derivatives will be calculated. `gamma` is the running average constant for the

second derivatives (it should be larger than $1/n$). and `mu` is the failsafe constant to prevent to learning rates from blowing up if the Hessians are small. default values for the 3 arguments are 200 , 0.01 , and 0.02.

```
(==> workbench train n [eta])
```

Type: MSG(packages/gblearning/workbench.lsh) Last Modified: 2002-08-19 03:04:06

perform `n` training iterations using learning rate `eta` . the meter is cleared before starting.

```
(==> workbench test [n])
```

Type: MSG(packages/gblearning/workbench.lsh) Last Modified: 2002-08-19 03:04:06

measure performance on first `n` patterns of the test set. If `n` is not given , the performance is measured on the entire set.

```
(==> workbench test-pattern n [db])
```

Type: MSG(packages/gblearning/workbench.lsh) Last Modified: 2002-08-19 03:04:06

runs pattern `n` of database `db` through the machine. If `db` is not given , the test database is used.

9.2.4 trainer

(packages/gblearning/trainer.lsh) Author(s): Yann LeCun Last Modified: 2002-10-14 03:30:51

a subclass of `sn-module` that can be used to conveniently train and test a learning machine with redefinable preprocessor, cost function and postprocessor. most trainer methods take a database as argument, or at least an object that can understand the following three messages:

```
(==> <db> get-input-current)
```

```
(==> <db> get-label-current)
```

```
(==> <db> seek <n>)
```

```
(==> <db> next)
```

```
(==> <db> tell)
```

slots are provided to store conversion function from what comes out of the database to what goes into the machine. Methods are provided to test a machine and measure its performance, as well as train it. Three different main methods are provided: `test`, `train`, and `train-bbprop`. in the first method only an `fprop` is performed on the preprocessor, the machine the post-processor and the cost function. the second does an `fprop`, a `bprop`, and a parameter update, the third

one does an fprop, a bprop, a bbprop (for 2nd derivative) and an update-bbprop (to update the parameter's second derivatives).

(new trainer pre mchn prm post cst ppat pdes raw in out pout des)

Type: CLASS(packages/gblearning/trainer.lsh) Last Modified: 2002-10-14 03:30:51

create a new trainer where `mchn` is the learning machine. It must be an object that understands the methods `fprop`, `bprop` and optionally `bbprop` each of which must take one input and one output as arguments. `prm` is a param (that can understand `update` and `update-bbprop` messages). `cst` is a cost module: a module with two inputs and one output. `pre` is a preprocessing module, and `post` is a postprocessing module (including a classifier, or something similar not subject to updates or backprops). `ppat` is a present-pattern function (see help), `pdes` : a present-desired function (see help), `raw` , `in` , `out` and `des` are objects that will be used to store the raw input, the machine input (or preprocessor output), the output and the desired states of the learning machine. `pout` is an object used to store the output of the postprocessor (probably, this is a class-state)

See: `trainer`

See: `(present-pattern from to)`

See: `(present-idx from to [dummy])`

See: `(present-desired from to actual)`

See: `(present-classindex class-index to [extra])`

(same-class? actual desired [dummies])

(packages/gblearning/trainer.lsh) Last Modified: 2002-10-14 03:30:51

return 0 if `actual` equals -1 , otherwise , return 1 if `actual` and `desired` are equal , -1 otherwise.

(present-pattern from to)

(packages/gblearning/trainer.lsh) Last Modified: 2002-10-14 03:30:51

`present pattern` is a slot in the `trainer` class which is supposed to contain a user-defined function with two arguments. It provides a way for the network designer to interface a database to a particular learning machine , i.e. transform an input from whatever comes from the database to whatever the learning machine (or rather , the preprocessing module) understands. The first argument is an item from the database , and the second argument is the input state of the preprocessing module (usually an `sn-state`)

(present-desired from to)

(packages/gblearning/trainer.lsh) Last Modified: 2002-10-14 03:30:51

`present-desired` is a slot in the `trainer` class which is supposed to contain a user-defined function with two arguments. It provides a way for the network designer to interface a database to a particular learning machine , i.e. transform

a desired output from whatever comes from the database to whatever the learning machine understands. The first argument is an item from the "desired" slot of the database , and the second argument is the desired state of the machine to be trained (usually an sn-state).

sn-meter

(packages/gblearning/trainer.lsh) Last Modified: 2002-10-14 03:30:51

a class that is used to measure the performance of a learning machine. sn-meters have at least three methods: clear, fprop and update. the update method must take six arguments (see the doc of the fprop method for more details). most subclasses of sn-meter can be used as "recorders" as well as meters. If they are put in record mode, they keep a log of everything that happens.

See: (==> **sn-meter fprop posto postd machd nrgy ptr**)

(==> **sn-meter fprop trainer**) Type: MSG(packages/gblearning/trainer.lsh)

Last Modified: 2002-10-14 03:30:51

updates **sn-meter** with data extracted from the trainer **trainer**

(==> **sn-meter clear**) Type: MSG(packages/gblearning/trainer.lsh) Last Modified: 2002-10-14 03:30:51

clears the content of **sn-meter** , and restart the measurement process.

(==> **sn-meter sprint-info**) Type: MSG(packages/gblearning/trainer.lsh)

Last Modified: 2002-10-14 03:30:51

returns a string with most interesting info in **sn-meter**

classifier-meter

(packages/gblearning/trainer.lsh) Last Modified: 2002-10-14 03:30:51

a special sn-meter that can be used to measure the performance of a classifier.

small-classifier-meter

(packages/gblearning/trainer.lsh) Last Modified: 2002-10-14 03:30:51

a special sn-meter that can be used to measure the performance of a classifier. Like classifier-meter but wastes less memory.

multi-classifier-meter

(packages/gblearning/trainer.lsh) Last Modified: 2002-10-14 03:30:51

a special sn-meter that can be used to measure the performance of a classifier along multiple axes.

See: (new multi-classifier-meter **list-of-correctp-functions**)

(**new multi-classifier-meter list-of-correctp-functions**) Type: CLASS(packages/gblearning/trainer.lsh)

Last Modified: 2002-10-14 03:30:51

Given a **list-of-correctp-functions** , return a new multi-classifier-meter. Each of these functions takes two arguments (correct and desired class) and returns 1 if the answer is correct , 0 if punt , -1 if wrong.

9.2.5 Databases

NOTE: This package is not a traditional "database" interface package. It is a library simply designed as a data handling mechanism for the (now obsolete) `gblearning` machine learning package. That packages has been replaces by `gblearn2` which does not use the present library but uses its own mechanism called data sources.

In the `gblearning` library , the workbench and trainer classes access training and testing samples through a "database". Several base classes are provided here to make it easy to build databases of training or test samples. The workbench/trainer classes accept databases as long as they understand the following methods:

- `(==> db size)` : number of element in database (-1 if infinite)
- `(==> db seek i)` : go to element `i`
- `(==> db tell)` : return index of current element
- `(==> db next)` : go to next element
- `(==> db get-input-current [dest])` : write input component of current element into `dest` , or return result if no `dest` .
- `(==> db get-label-current [dest])` : write label component of current element into `dest` , or return result if no `dest` .

`base-db` is the base class from which actual db classes can be derived (though some db classes are created from scratch so as to be compilable). Several pre-defined types of base-db compatible classes are provided for combining the content of existing base-db compatible objects.

base-db

(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

This is the base class used by the trainer classes as training and test databases. real databases can be derived from this class. and override just the get-input, get-label, get-info and size methods. The other methods rarely need to be overrided.

`(new base-db inp lab)` Type: CLASS(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

make a new base-db. `inp` must be a 4D tensor of floats where each 3D slice is a sample. `lab` must be a 2D int-matrix where each slice is a target vector (or something).

`(==> base-db get-input n [dest])` Type: MSG(lsh/libdb/base-db.lsh)
Last Modified: 2002-11-25 19:05:53

returns `n` -th item

`(==> base-db get-label n [dest])` Type: MSG(lsh/libdb/base-db.lsh)
Last Modified: 2002-11-25 19:05:53

returns the `n` -th label.
 (`==>` **base-db get-info n**) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53
 returns extra information about `n` -th item
 (`==>` **base-db size**) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53
 returns the size of the database
 (`==>` **base-db tell i**) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53
 set pointer to `i` -th item
 (`==>` **base-db next**) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53
 move pointer to next item
 (`==>` **base-db seek i**) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53
 set pointer to `i` -th item
 (`==>` **base-db get-input-current [dest]**) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53
 returns current item as a rawchar
 (`==>` **base-db get-label-current [dest]**) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53
 returns labels (ASCII codes) of `n` -th item in an `idx1` of ints.
 (`==>` **base-db get-info-current**) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53
 returns extra information about current item
 (`==>` **db-filter db func**) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53
 return a matrix with the indices of all the patterns for which (`func label`) returns non-nil
 (`==>` **base-db get-current what**) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

This is an obsolete method provided for backward compatibility with ancient code. `get-current` item. if `what` is equal to `'input` , the input item is returned (most likely a rawchar) , if it is equal to `'desired` , the label is returned (most likely an ASCII code) , if it is `'info` extra information on the item is returned (format is not standardized).

idx-db

(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53
 a database where each pattern is an `idx`, and each label is also an `idx` (of any type). All patterns must have the same size.
 (**new idx-db inp lab**) Type: CLASS(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

make a new `idx-db`. `inp` must be a vector/matrix/tensor where each slice is a sample. `lab` must be a vector/matrix/tensor where each slice is a target vector (or something).

(`==>` `idx-db size`) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

returns the size of the database

(`==>` `idx-db get-input n [dest]`) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

returns `n` -th item , or copies it into the `x` slot of `dest` .

(`==>` `idx-db get-label n [dest]`) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

returns labels (ASCII codes) of `n` -th item in an `idx1` of ints.

multi-db

(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

A db formed by the concatenation of multiple dbs. The dbs must understand the methods `get-input` , `get-label` , `get-info` , and `size`.

(`new multi-db base-list`) Type: CLASS(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

make a new base-db compatible object by concatenating all the base-db compatible objects in list `base-list` .

narrow-db

(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

a db constructed by taking patterns in an existing db whose indices are within a given range that what is passed is a db

(`new narrow-db base size offset`) Type: CLASS(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

make a new base-db compatible object by taking a range of items from the db passed as argument

permute-db

(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

a db built by permuting the samples of an existing db no actual movement of data is performed , the samples are simply accessed through a permutation table.

(`new permute-db base-db`) Type: CLASS(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

create a new permute-db. The permutation map is initially set to the identity. the `shuffle` method performs a random shuffling.

(`==>` `permute-db shuffle`) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

randomly shuffles the samples of the db.

(==> **permute-db set-map (m)**) Type: MSG(lsh/libdb/base-db.lsh) Last Modified: 2002-11-25 19:05:53

sets the permutation map to `m`. `m` must be a vector of ints. The size of the db will be the dimension of `m`, and the `i`-th element of the db will be the (`m i`)-th element of the base db.

image-db

(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

This db-like class represents a data-base of examples which are pairs of images (of any size and depth) and classes. The images may have different dimensions, but are returned as `idx3` of `ubytes` (the last dimension being the depth).

This is not a subclass of `base-db`, but it implements the interface and can be used as a base db for a `narrow-db` or `perturb-db` as well as fed directly to a trainer object.

(**new image-db**) Type: CLASS(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

creates an empty image-db. The images can be subsequently loaded through the `load-ppms`, `load-pgms`, `load` or `map` methods.

(==> **<image-db load-ppms flist clist**) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

fills up the database with images from a bunch of PPM files `clist` is a list of labels (one integer for each image) or a single category, in which case all the loaded images will be in that category.

(==> **<image-db load-pgms flist clist**) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

fills up the database with images from a bunch of PGM files `clist` is a list of labels (one integer for each image) or a single category, in which case all the loaded images will be in that category.

(==> **image-db save basename**) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

save database into pre-cooked `IDX` files. These files can be subsequently loaded quickly using the `load` or `map` methods. The database is saved in four files named `basename images.idx`, `basename starts.idx`, `basename sizes.idx`, and `basename classes.idx`

(==> **image-db load basename**) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

load database from pre-cooked `IDX` files produced through the `save` method. The database will be loaded from four files named `basename images.idx`, `basename starts.idx`, `basename sizes.idx`, and `basename classes.idx`

(==> **image-db map basename**) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

memory-map database from pre-cooked `IDX` files produced through the `save` method. This is MUCH faster than `load`, and consumes fewer memory/disk

bandwidth. The database will be mapped from four files named `basename images.idx` , `basename starts.idx` , `basename sizes.idx` , and `basename classes.idx`

(==> `image-db size`) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

returns the size of the database

(==> `image-db get-input n dest`) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

copies `n` -th item as a [`n-rows` , `n-columns` , 3] RGB ubyte image into the `x` slot of `dest` . `dest` must be an `idx3-state` or a subclass thereof. It is resized if necessary.

(==> `image-db get-label n dest`) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

returns labels of `n` -th example as an int.

(==> `base-db seek i`) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

set pointer to `i` -th item

(==> `image-db tell i`) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

set pointer to `i` -th item

(==> `base-db next`) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

move pointer to next item

(==> `image-db get-input-current dest`) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

get current pattern.

(==> `image-db get-label-current`) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

get current label.

(==> `image-db get-info-current`) Type: MSG(lsh/libdb/image-db.lsh) Last Modified: 2002-11-25 12:49:17

get current info.

9.2.6 param

(packages/gblearning/params.lsh) Author(s): Yann LeCun Last Modified: 2002-10-14 03:30:51

a hierarchy of classes that implements gradient-based learning algorithms. Various subclasses of `param` can be used for various learning algorithms (stochastic gradient, with separate epsilons, with second derivatives.....) The standard `param` type contains `idx1` slots.

(`new-index-offset s dlist o`)

(packages/gblearning/params.lsh) Last Modified: 2002-10-14 03:30:51

like `new index` , but puts the index at a specific offset in the storage `s` is a storage , `dlist` a list of dimensions , `o` an offset. Compilable macro.

idx1-param

(packages/gblearning/params.lsh) Last Modified: 2002-10-14 03:30:51
 a param whose slots are idx1. no learning algorithm is defined. Only the x slot is present. This class is useful for fixed (non adaptive) parameters.

idx1-dparam

(packages/gblearning/params.lsh) Last Modified: 2002-10-14 03:30:51
 param class for regular stochastic gradient descent algorithm
 (**new idx1-dparam s e**) Type: CLASS(packages/gblearning/params.lsh)
 Last Modified: 2002-10-14 03:30:51
 s is a size , e is a learning rate.
 (==> **idx1-dparam update**) Type: MSG(packages/gblearning/params.lsh)
 Last Modified: 2002-10-14 03:30:51
 simple gradient descent update.

idx1-dparam-eps

(packages/gblearning/params.lsh) Last Modified: 2002-10-14 03:30:51
 param class for gradient descent with separate epsilons for each parameter
 (**new idx1-dparam-eps s e**) Type: CLASS(packages/gblearning/params.lsh)
 Last Modified: 2002-10-14 03:30:51
 s is a size , e is a learning rate. the epsilons are set to 1 by default.
 (==> **idx1-dparam-eps update**) Type: MSG(packages/gblearning/params.lsh)
 Last Modified: 2002-10-14 03:30:51
 performs one step of gradient descent with separate learning rates for each parameter.
 (==> **idx1-dparam-eps set-epsilons m**) Type: MSG(packages/gblearning/params.lsh)
 Last Modified: 2002-10-14 03:30:51
 copy the values in vector m to epsilons

idx1-ddparam

(packages/gblearning/params.lsh) Last Modified: 2002-10-14 03:30:51
 a param class for the stochastic diagonal levenberg-marquardt algorithm. In addition to the usual update method , it has an update-bbprop method for computing the second derivatives , and a set-epsilons method to set the epsilons using the second derivatives.
 (**new idx1-ddparam s e g m alloc**) Type: CLASS(packages/gblearning/params.lsh)
 Last Modified: 2002-10-14 03:30:51
 s is the size (can be 0) , e is the learning rate , g is the gamma parameter used for the on-line update of the second derivatives , m is the mu parameters added to the second derivatives to prevent the epsilons from blowing up. alloc is the size of storages to be preallocated. This will prevent memory fragmentation when the size of the param is subsequently increased.

(**==> idx1-ddparam update-bbprop**) Type: MSG(packages/gblearning/params.lsh)
 Last Modified: 2002-10-14 03:30:51

update second derivative estimates on-line , using current value in the instantaneous second derivative field ddx

(**==> idx1-ddparam compute-epsilons**) Type: MSG(packages/gblearning/params.lsh)
 Last Modified: 2002-10-14 03:30:51

compute and set the epsilons using the second derivative. this method should be called after a few iteration of update-bbprop

9.2.7 sn-states

(packages/gblearning/states.lsh) Author(s): Yann LeCun Last Modified: 2002-10-18 22:02:23

states are the main objects used to store state variables in or between modules.

(**each-idx-slots** o s f1 [f2...[fn]])

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

evaluates lists in **f** for all active idx slots of object **o** , (as returned by (**==>** o idx-slots)) with **s** successively taking the values of the slots in question.

idx-state

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

an idx-state is a state that contains vector or matrix variables.

(**==> idx-state idx-slots**) Type: MSG(packages/gblearning/states.lsh)
 Last Modified: 2002-10-18 22:02:23

return the complete list of idx slots.

(**new idx-state** [params]) Type: CLASS(packages/gblearning/states.lsh)
 Last Modified: 2002-10-18 22:02:23

create a new simple-state. if **params** is absent, the simple-state is left unsized. Otherwise **params** must be a list of integer (possibly empty), which will be used to determine the size of the object. **params** has the same meaning as the 2nd parameter of the functions new-index.

See: simple-state

See: Index.

(**==> idx-state resize** [list-of-dims]) Type: MSG(packages/gblearning/states.lsh)
 Last Modified: 2002-10-18 22:02:23

resize all the idx slots (as returned by idx-slots) **list-of-dims** is a list with p numbers specifying the size of the last p dimensions of each slot. If **list-of-dim** is absent, all the active slots are undimmed.

See: (**==> idx-state idx-slots**)

(**==> idx-state undim**) Type: MSG(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

undim all the active idx slots of **idx-state** .

(==> **idx-state clear**) Type: MSG(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

fill all the active idx slots with zeroes

See: (==> **idx-state idx-slots**)

(==> **idx-state load s**) Type: MSG(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

fill slot x with content of file s

(==> **idx-state save s** (packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

save content of slot x into file s .

(==> **idx-state dump s** (packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

save the entire object into file s

idx0-state

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx0-dstate

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx0-ddstate

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx1-state

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx1-dstate

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx1-ddstate

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx2-state

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx2-dstate

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx2-ddstate

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx3-state

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx3-dstate

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx3-ddstate

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx4-state

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx4-dstate

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

idx4-ddstate

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

class-state

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

a special kind of state used to store the output of a classifier. class-state are generated by modules such as class-max , and used by meters such as classifier-meter. No backprop is possible through a class-state.

lclass-state

(packages/gblearning/states.lsh) Last Modified: 2002-10-18 22:02:23

a special kind of state used to store the output of a spatial classifier such as ledist-classer.

9.2.8 libnets

(packages/gblearning/libnets.lsh) Author(s): Yann LeCun Last Modified: 2002-10-29 18:12:07

a bunch of standard compiled modules for building neural net architectures: sigmoids, RBF, fully connected, convolutions, subsampling layers. Each of these classes has methods for fprop, bprop, bbprop, load, save, forget, and setting various parameters.

(read-matrix f x)

(packages/gblearning/libnets.lsh) Last Modified: 2002-10-29 18:12:07

read matrix file from file **f** and copy the content into matrix **x** useful for filling up an existing matrix.

(midx-m2oversample input nlin ncol output)

(packages/gblearning/libnets.lsh) Last Modified: 2002-10-29 18:12:07

macro for 2D oversampling

f-layer

(packages/gblearning/libnets.lsh) Last Modified: 2002-10-29 18:12:07

full connection between replicable 3D layers the full connection is only across the first dimension of the input and output layer. the other two dimensions are spatial dimensions across which the weight matrix is shared. This is much more efficient than using a c-layer with a 1x1 convolution and a full-table.

(new f-layer tin tout si sj sqsh) Type: CLASS(packages/gblearning/libnets.lsh)
Last Modified: 2002-10-29 18:12:07

makes a new replicable fully-connected layer. **tin** and **tout** are the number of units in the input and output. **si** and **sj** are the vertical and horizontal sizes used for pre-allocating the internal state variables (they , as well as the output. are automatically resized , but the resizing is more efficient of it is a downsizing).

e-layer

(packages/gblearning/libnets.lsh) Last Modified: 2002-10-29 18:12:07

full connection with Euclidean distance between replicable 3D layers. This is like a layer of RBF WITHOUT NON-LINEAR FUNCTIONS. the output is the square Euclidean distance between the input and the weight the full connection is only across the first dimension of the input and output layer. the other two dimensions are spatial dimensions across which the weight matrix is shared.

(new e-layer tin tout prm) Type: CLASS(packages/gblearning/libnets.lsh)
Last Modified: 2002-10-29 18:12:07

new e-layer (Euclidean distance RBF). **tin** is the thickness of the input layer , **tout** is the thickness of the output layer , **prm** is the parameter.

e-layer-sparse

(packages/gblearning/libnets.lsh) Last Modified: 2002-10-29 18:12:07

like e-layer , but each RBF is only connected to a subset of the inputs this should be used for "multiple bitmap" outputs.

(new e-layer-sparse tin off prm) Type: CLASS(packages/gblearning/libnets.lsh)
Last Modified: 2002-10-29 18:12:07

new e-layer (Eucliden distance RBF). `tin` is the thickness of the input layer (dimension of prototypes) , `off` is an `idx1` with as many elements as there are output. each element contains the offset in the input for the input vector of each output. `prm` is the parameter.

c-layer

(packages/gblearning/libnets.lsh) Last Modified: 2002-10-29 18:12:07

convolutional layer module. Performs multiple convolutions between an `idx3`-state input and an `idx3`-state output. includes bias and sigmoid.

(**new c-layer** `ki kj ri rj tbl thick si sj sqsh`) Type: CLASS(packages/gblearning/libnets.lsh)
Last Modified: 2002-10-29 18:12:07

Creates a new convolution layer. `ki` (int) vertical kernel size. `kj` (int) horizontal kernel size. `ri` (int) vertical stride (number of pixels by which the kernels are stepped) `<rj` (int) horizontal stride `<tbl` (`idx2`) N by 2 matrix containing source and destination feature maps for coresponding kernel `thick` (int) thickness of output layer (number of feature maps) `si` (int) vertical size for preallocation of internal state `sj` (int) horizontal size for preallocation of internal state `sqsh` (`idx3`-module) a squashing function module that operates on `idx3`-state. `prm` and `idx1`-ddparam from which the parameters will be allocated

s-layer

(packages/gblearning/libnets.lsh) Last Modified: 2002-10-29 18:12:07

subsampling layer class

(**new s-layer** `ki kj thick si sj sqsh prm`) Type: CLASS(packages/gblearning/libnets.lsh)
Last Modified: 2002-10-29 18:12:07

Creates a new subsampling layer for convolutional nets `ki` (int) vertical subsampling ratio. `kj` (int) horizontal subsampling ratio. `thick` (int) thickness of output layer (number of feature maps) `si` (int) vertical size for preallocation of internal state `sj` (int) horizontal size for preallocation of internal state `sqsh` (`idx3`-module) a squashing function module that operates on `idx3`-state. `prm` and `idx1`-ddparam from which the parameters will be allocated

9.2.9 net-ccc

(packages/gblearning/net-ccc.lsh) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

convolutional net with 3 convolutional layers.

(**new net-ccc** `ini inj ki0 kj0 tbl0 ki1 kj1 tbl1 ki2 kj2 tbl2 prm`)

Type: CLASS(packages/gblearning/net-ccc.lsh) Last Modified: 2002-08-19 03:04:06

makes a new net-ccc module. `ini inj` : expected max size of input for pre-allocation of internal states `ki0 kj0` : kernel size for first convolutional layer `tbl0` : table of connections between input anf feature maps for first layer `ki1`

`ki1 tbl1` : kernel and table for next layer `ki2 kj2 tbl2` : same for last convolution layer `prm` an `idx1-ddparam` in which the parameters will be allocated.

9.2.10 net-cfe

(`packages/gblearning/net-cfe.lsh`) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

neural net with one convolutional layer, one (replicable) fully-connected layer, and one RBF (euclidean distance) layer.

(`new net-cfe n e thick si sj`)

Type: CLASS(`packages/gblearning/net-cfe.lsh`) Last Modified: 2002-08-19 03:04:06

create a new `net-cfe`. `n` is a `net-cf`, `e` is a `e-layer`, `thick` is the number of outputs, and `si`, `sj` are the number of output replications.

9.2.11 net-cff

(`packages/gblearning/net-cff.lsh`) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

a neural net class with one convolutional layer and two fully connected layers. The main purpose of this class is to make replicable 2-hidden layer fully connected networks.

(`new net-cff ini inj ki0 kj0 tbl0 f1thick outthick prm`)

Type: CLASS(`packages/gblearning/net-cff.lsh`) Last Modified: 2002-08-19 03:04:06

makes a new `net-cff` module. `ini inj` : expected max size of input for preallocation of internal states `ki0 kj0` : kernel size for first convolutional layer a standard fully-connected network can be obtained when `ini = ki0` and `inj = kj0`. `tbl0` : table of connections between input and feature maps for first layer `f1thick` : number of hidden units in second hidden layer. `outthick` : number of outputs. `prm` an `idx1-ddparam` in which the parameters will be allocated.

9.2.12 net-cf

(`packages/gblearning/net-cf.lsh`) Last Modified: 2002-08-19 03:04:06

a neural net class with one convolutional layer and one fully connected layer. The main purpose of this class is to make replicable fully connected networks. Unlike with `net-ff`, the input spatial dimensions are treated correctly here. So if you want to recognize characters, use this class rather than `net-ff`.

(`new net-cf ini inj ki0 kj0 tbl0 outthick prm`)

Type: CLASS(`packages/gblearning/net-cf.lsh`) Last Modified: 2002-08-19 03:04:06

makes a new net-cscscf module. `ini inj` : expected max size of input for preallocation of internal states `ki0 kj0` : kernel size for first convolutional layer a standard fully-connected network can be obtained when `ini = ki0` and `inj = kj0` . `tb10` : table of connections between input anf feature maps for first layer `outthick` : number of outputs. `prm` an `idx1-ddparam` in which the parameters will be allocated.

9.2.13 net-c

(`packages/gblearning/net-c.lsh`) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

a neural net class with one convolutional layer and one fully connected layer. The main purpose of this class is to make replicable fully connected networks. Unlike with `net-ff`, the input spatial dimensions are treated correctly here. So if you want to recognize characters, use this class rather than `net-ff`.

(`new net-c ini inj ki0 kj0 inthick outthick prm`)

Type: CLASS(`packages/gblearning/net-c.lsh`) Last Modified: 2002-08-19 03:04:06

makes a new net-cscscf module. `ini inj` : expected max size of input for preallocation of internal states `ki0 kj0` : kernel size for first convolutional layer a standard fully-connected network can be obtained when `ini = ki0` and `inj = kj0` . `inthick` : number of input slices `outthick` : number of outputs. `prm` an `idx1-ddparam` in which the parameters will be allocated.

9.2.14 net-csccf

(`packages/gblearning/net-csccf.lsh`) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

a standard LeNet2/LeNet5 architecture without the last RBF layer (look into `net-cscscfe` for that).

(`new net-csccf ini inj ki0 kj0 tb10 si0 sj0 ki1 kj1 tb11 ki2 kj2 tb12 outthick prm`)

Type: CLASS(`packages/gblearning/net-csccf.lsh`) Last Modified: 2002-08-19 03:04:06

makes a new net-csccf module. `ini inj` : expected max size of input for preallocation of internal states `ki0 kj0` : kernel size for first convolutional layer `tb10` : table of connections between input anf feature maps for first layer `si0 sj0` : subsampling for first layer `ki1 kj1 tb11` : kernel and table for next layer `ki2 kj2 tb12` : same for last convolution layer `outthick` : number of outputs. `prm` an `idx1-ddparam` in which the parameters will be allocated.

9.2.15 net-cscfe

(packages/gblearning/net-cscfe.lsh) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

convolutional net where the layers are convolution, subsampling, convolution, full connect, and euclidean distance RBF.

(new net-cscfe n e thick si sj)

Type: CLASS(packages/gblearning/net-cscfe.lsh) Last Modified: 2002-08-19 03:04:06

n is a net-cscf, e is an e-dist layer, thick is the number of outputs. si, sj are the initial replication factors.

9.2.16 net-cscf

(packages/gblearning/net-cscf.lsh) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

convolutional net where the layers are convolution, subsampling, convolution, full connect.

(new net-cscf ini inj ki0 kj0 tb10 si0 sj0 ki2 kj2 tb12 outthick prm)

Type: CLASS(packages/gblearning/net-cscf.lsh) Last Modified: 2002-08-19 03:04:06

makes a new net-cscf module. ini inj : expected max size of input for preallocation of internal states ki0 kj0 : kernel size for first convolutional layer tb10 : table of connections between input and feature maps for first layer si0 sj0 : subsampling for first layer ki2 kj2 tb12 : same for last convolution layer outthick : number of outputs. prm an idx1-ddparam in which the parameters will be allocated.

9.2.17 net-cscscf

(packages/gblearning/net-cscscf.lsh) Author(s): Yann LeCun Last Modified: 2002-08-27 03:27:02

Standard LeNet5-type architecture without the final e-dist RBF layer.

(new net-cscscf ini inj ki0 kj0 tb10 si0 sj0 ki1 kj1 tb11 si1 sj1 ki2 kj2 tb12 outthick prm)

Type: CLASS(packages/gblearning/net-cscscf.lsh) Last Modified: 2002-08-27 03:27:02

makes a new net-cscscf module. ini inj : expected max size of input for preallocation of internal states ki0 kj0 : kernel size for first convolutional layer tb10 : table of connections between input and feature maps for first layer si0 sj0 : subsampling for first layer ki1 kj1 tb11 si1 sj1 : same for next 2

layers `ki2 kj2 tbl2` : same for last convolution layer `outthick` : number of outputs. `prm` an `idx1-ddparam` in which the parameters will be allocated.

9.2.18 net-cscscfe

(`packages/gblearning/net-cscscfe.lsh`) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

Standard LeNet5-type architecture.

(`new net-cscscfe n e thick si sj`)

Type: CLASS(`packages/gblearning/net-cscscfe.lsh`) Last Modified: 2002-08-19 03:04:06

`n` is a net-cscscf, `e` is an e-dist layer, `thick` is the number of outputs. `si`, `sj` are the initial replication factors.

9.2.19 net-cscscscf

(`packages/gblearning/net-cscscscf.lsh`) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

A deep convolutional net, with one convo/subsamp set than LeNet5.

(`new net-cscscscf ini inj ki0 kj0 tbl0 si0 sj0 ki1 kj1 tbl1 si1 sj1 ki2 kj2 tbl2 si2 sj2 ki3 kj3 tbl3 outthick prm`)

Type: CLASS(`packages/gblearning/net-cscscscf.lsh`) Last Modified: 2002-08-19 03:04:06

makes a new net-cscscf module. `ini inj` : expected max size of input for preallocation of internal states `ki0 kj0` : kernel size for first convolutional layer `tbl0` : table of connections between input and feature maps for first layer `si0 sj0` : subsampling for first layer `ki1 kj1 tbl1 si1 sj1` : same for next 2 layers... `ki2 kj2 tbl2 si2 sj2` : ...and the next 2 `ki3 kj3 tbl3` : same yet again for last convolution layer `outthick` : number of outputs. `prm` an `idx1-ddparam` in which the parameters will be allocated.

9.2.20 net-fff

(`packages/gblearning/net-fff.lsh`) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

standard two-hidden layer, fully-connected neural net. this network is replicable.

(`new net-fff thickin thick0 thick1 thick2 ini inj prm`)

Type: CLASS(`packages/gblearning/net-fff.lsh`) Last Modified: 2002-08-19 03:04:06
create a new net-ff network.

<thickin>: number of inputs.
 <thick0>: size of hidden layer
 <thick1>: size of second hidden layer
 <thick2>: number of outputs
 <ini> <inj>: expected max size of input over which the network
 should be replicated.
 <prm> an idx1-ddparam in which the parameters will be allocated.

9.2.21 net-ff

(packages/gblearning/net-ff.lsh) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

standard one-hidden layer, fully-connected neural net. this network is replicable.

(new net-ff thickin thick0 thick1 ini inj prm)

Type: CLASS(packages/gblearning/net-ff.lsh) Last Modified: 2002-08-19 03:04:06
 create a new net-ff network.

<thickin>: number of inputs.
 <thick0>: size of hidden layer
 <thick1>: number of outputs
 <ini> <inj>: expected max size of input over which the network
 should be replicated.
 <prm> an idx1-ddparam in which the parameters will be allocated.

9.2.22 lclass-meter

(packages/gblearning/lclass-meter.lsh) Last Modified: 2002-10-14 03:30:51
 a special sn-meter that can be used to measure the performance of a classifier.
 (needs more documentation of the methods).

(new lclass-meter sz [cp])

Type: CLASS(packages/gblearning/lclass-meter.lsh) Last Modified: 2002-10-14 03:30:51

create a new lclass-meter with a pre-allocated number of samples `sz` . `cp` is an optional 2-argument function that returns `t` if the two arguments are considered the same class.

9.2.23 ccc-tdnn

(packages/gblearning/ccc-tdnn.lsh) Author(s): Yann LeCun Last Modified: 2002-08-18 21:42:08

Class for a Time Delay Neural Network with 3 convolutional layers (with local connectivity). This type of network is appropriate for classifying sequences of time/frequency representations such as cepstrum, MEL Scale spectrum, and other.

(build-ccc-tdnn params freqs max-seq-len fk1 fs1 nh1 tk1 fk2 fs2 nh2 tk2 fk3 fs3 nh3 tk3)

(packages/gblearning/ccc-tdnn.lsh) Last Modified: 2002-08-18 21:42:08

Build a Time-Delay Neural Network for data such as spectral sequences in which the features have a topology, choosing the connection tables to obtain local connections in feature space (as specified by kernel size and stride for the frequency axis, for each of these layers). The input to the net is assumed to be $F \times 1 \times T$ where F is for example the number of frequency channels (spectral representation) and T is the length of the sequence. Here, the network has 3 convolutional layers. The arguments are the following:

- <params> is a idx1-ddparam on which to allocate parameters for the layers.
- <freqs> is the number of input frequency channels.
- <max-seq-len> = maximum sequence length.
- <fk1>, <fk2>, <fk3> = sizes of frequency kernels (= width of local freq. windows)
- <fs1>, <fs2>, <fs3> = step sizes which separate the successive frequency windows
- <nh1>, <nh2>, <nh3> = number of hidden units per frequency channel, for each layer
- <tk1>, <tk2>, <tk3> = sizes of the temporal kernels, for each layer.

(tdnn-present-pattern from to mean idev temporal-window-size)

(packages/gblearning/ccc-tdnn.lsh) Last Modified: 2002-08-18 21:42:08

from is a $T \times f$ source matrix, **to** is a $f \times 1 \times T'$ destination state ($T' = T + \text{temporal-window-size} - 1$), **mean** is a f -vector to subtract from source, **idev** is a f -vector to multiply by source. **temporal-window-size** is the length of the input of the network yielding an output of length 1.

norm-ftdnn

(packages/gblearning/ccc-tdnn.lsh) Last Modified: 2002-08-18 21:42:08

Wrapper around ccc-tdnn which does input normalization and pads the input according to the network architecture.

(new norm-ftdnn n-inputs nh1 nh2 n-outputs weight-file norm-file [(-tk1 5)(-tk2 8)(-tk3 12) (-fk1 6)(-fk2 3)(-fk3 3) (-fs1 3)(-fs2 2)(-fs3 1) (max-seq-len 2000)])

Type: CLASS(packages/gblearning/ccc-tdnn.lsh) Last Modified: 2002-08-18 21:42:08

9.2.24 conx

(packages/gblearning/conx.lsh) Author(s): Yann LeCun Last Modified: 2002-10-14 03:30:51

An architecture that connects inputs to outputs (with exactly one input per output but maybe more than one output per input, i.e. a one to many map). The inputs and outputs are bidimensional (i.e. sequences of vectors). It has no adjustable parameters. Instead it has internal non-adjustable parameters specifying the input/output connectivity.

(**new conx n-outputs n-inputs indices**)

Type: CLASS(packages/gblearning/conx.lsh) Last Modified: 2002-10-14 03:30:51

Construct a conx architecture. The parameters are given when constructing the corresponding machine.

See: machine

See: sn-module

See: conx

(**==> conx fprop parameters input output**)

Type: MSG(packages/gblearning/conx.lsh) Last Modified: 2002-10-14 03:30:51

Fprop method for a conx architecture. Just gather the inputs into the output vectors. Both the input and output are assumed to be sequences of vectors [i,t] (2D matrices).

See: conx

(**==> conx bprop input output**)

Type: MSG(packages/gblearning/conx.lsh) Last Modified: 2002-10-14 03:30:51

Bprop method for a conx architecture. Just scatter with accumulate the derivatives of the outputs into the derivative of the inputs. Both the input and output are assumed to be sequences of vectors [i,t] (2D matrices).

See: conx

(**==> conx generate input output**)

Type: MSG(packages/gblearning/conx.lsh) Last Modified: 2002-10-14 03:30:51

:output:x is a matrix whose columns have zeros and a single one. According to the permutation matrix **parameters** , a similar matrix :input:x is produced.

sum-arch

(packages/gblearning/conx.lsh) Last Modified: 2002-10-14 03:30:51

An architecture that takes the sum of its inputs. The input is a 2D matrix whose first index is **n-inputs** and second index is variable (e.g. time). The output is zero-dimensional (an idx0). There are no parameters.

See: sn-module

See: (new sum-arch)
 See: (==> **sum-arch fprop** parameters input output)
 See: (==> **sum-arch bprop** parameters input output)
 See: (==> **sum-arch** output-dim inputx)
 See: (==> **sum-arch** generate parameters input output)

(new sum-arch)

Type: CLASS(packages/gblearning/conx.lsh) Last Modified: 2002-10-14 03:30:51

Construct a sum-arch sn-module.

See: machine
 See: sn-module
 See: sum-arch

(==> sum-arch **output-dim** inputx)

Type: MSG(packages/gblearning/conx.lsh) Last Modified: 2002-10-14 03:30:51

(==> sum-arch **fprop** parameters input output)

Type: MSG(packages/gblearning/conx.lsh) Last Modified: 2002-10-14 03:30:51

Fprop method for a sum-arch architecture. Just add all the inputs. The input is assumed to be any matrix (except an idx0). The output is an idx0.

See: sum-arch

(==> sum-arch **bprop** parameters input output)

Type: MSG(packages/gblearning/conx.lsh) Last Modified: 2002-10-14 03:30:51

Bprop method for a sum-arch architecture.

See: prod

(==> sum-arch **generate** parameters input output)

Type: MSG(packages/gblearning/conx.lsh) Last Modified: 2002-10-14 03:30:51

Generate a 2-d sequence of inputs. A prior distribution on sequence length is chosen here to be uniform in [5 , 6 , ..10]. Ignore the output. For each time step , the input vector is a vector of 0's and a single 1 , chosen uniformly.

9.2.25 ftdnn

(packages/gblearning/ftdnn.lsh) Author(s): Yann LeCun Last Modified: 2002-08-19 03:04:06

Time Delay Neural Net class for such things a speech recognition. Somewhat superseded by ccc-tdnn.

(**build-ftdnn** apool params freqs max-seq-len layers-infos)

(packages/gblearning/ftdnn.lsh) Last Modified: 2002-08-19 03:04:06

Build a Time-Delay Neural Network for data such as spectral sequences in which the features have a topology, choosing the connection tables to obtain local connections in feature space (as specified by kernel size and stride for the frequency axis, for each of these layers). The input to the net is assumed to be $F \times 1 \times T$ where F is for example the number of frequency channels (spectral representation) and T is the length of the sequence. The arguments are the following:

- <apool> is a Pool on which to allocate dynamic memory for the internal states and layers.
- <params> is a idx1-ddparam on which to allocate parameters for the layers.
- <freqs> is the number of input frequency channels.
- <max-seq-len> = maximum sequence length.
- <layers-infos> is a list whose elements are lists with the following structure:
 - <layer-type> is a symbol giving the name of an idx3-module subclass (e.g., c-layer, f-layer, s-layer, m-layer, etc...)
- For a c-layer or a f-layer, the following remaining arguments are expected in the list:
 - <freqs-kernel> = size of frequency "kernels" (= width of local freq. windows)
 - <freqs-stride> = step size which separates the successive frequency windows
 - <n-h-per-freq> is the number of hidden units per frequency channel
 - <t-kernel> = size of the temporal kernel
- For a s-layer, the following remaining arguments are expected in the list:
 - <t-subsampling-rate> = temporal subsampling rate
- For a m-layer, the following remaining arguments are expected in the list:
 - <t-window-size> = size of temporal window used for smoothing

9.2.26 seq3-modules

(packages/gblearning/seq3-modules.lsh) Last Modified: 2002-08-19 03:04:06

A set of idx3-modules executed in sequence (the output of one being the input of the next one).

(**new seq3-modules** pool modules-gptr-vector sizes)

Type: CLASS(packages/gblearning/seq3-modules.lsh) Last Modified: 2002-08-19 03:04:06

Construct a new seq3-modules module (it is itself an idx3-module) which will sequence the idx3-modules whose pointers are given in the **modules-gptr-vector** of length N . The **sizes** of the output idx3's of each of the internal modules is given in $(N-1) \times 3$ int-matrix **sizes** . The intermediate idx3-ddstates are allocated on the given **pool** .

9.3 SDL: Simple DirectMedia Layer

Author(s): Yann LeCun

Lush's interface to the popular SDL library provides access to every function of the SDL API. It also provides a set of higher-level classes to make it easy to write 2D video games and other interactive multimedia applications. The high-level interface provides classes for double-buffered SDL screens, event grabbing, and sprites with pixel-accurate collision detection.

The help text associated with each low-level function is from the corresponding .h header file that comes with SDL. If you installed libSDL1.2-devel with RPM, you may find the API documentation in [file:/usr/share/doc/libSDL1.2-devel-1.2.3/index.html] Otherwise, the documentation is available at [http://www.libsdl.org]

The Tutorial section of the Lush manual contains a tutorial on how to write simple video games with Lush and SDL. [Tutorial: Writing Games with Lush and SDL.

]

9.3.1 Installing SDL

To make the SDL interface work with Lush, you need:

- The SDL libraries (libSDL1.2)
- The SDL development libraries and include files (libSDL1.2-devel)
- The SDL_image library (libSDL_image1.2-devel)
- The SDL_image development libraries and include files (libSDL_image1.2-devel)

Those are generally available as RPMs or APTs for most Linux distributions. Otherwise, you can download the latest version from [http://www.libsdl.org]

By default, Lush looks for the SDL .so libraries in "/usr/lib". It expects to find "libSDL-1.2.so.0" and "libSDL_image-1.2.so.0". If these libraries are installed someplace else on your system, you must edit packages/sdl/sdl-config.lsh.

9.3.2 SDL Demos

(sdl-lander-simple)

(packages/sdl/demos/lander-simple.lsh) Last Modified: 2002-08-19 07:54:28

a simple lunar lander game in SDL that uses the high-level interface library (sdl-screen , sdl-sprite , sdl-event). Use the arrows keys to activate the engines. Hit "q" to quit.

Simple Demos of the SDL Lush interface

(packages/sdl/demos/libSDL-demo.lsh) Last Modified: 2002-10-09 22:47:00

(sdl-showimage <filename>) (packages/sdl/demos/libSDL-demo.lsh) Last Modified: 2002-10-09 22:47:00

load an image and display it in an SDL window example:

```
(sdl-showimage (concat lushdir "/packages/sdl/demos/moon.png"))
```

(sdl-bouncy f b n) (packages/sdl/demos/libSDL-demo.lsh) Last Modified: 2002-10-09 22:47:00

bounce a bunch of objects on a background **f** is an image filename for the object **b** is a background image **n** is the number of objects. This uses the low-level SDL interface. example:

```
(bounce (concat lushdir "/packages/sdl/demos/lem.png") (concat lushdir "/packages/sdl/demos/mo
```

9.3.3 LibSDL: High-level Interface to SDL

(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

libSDL is a library of "high level" functions and classes built on top of the SDL (Simple Directmedia Layer) library to facilitate its use from Lush. It includes objects such as screens and sprites with pixel-accurate collision detection. It also includes access functions to keyboard and mouse events, as well as to several common SDL data structures.

(sdl-initialize)

(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

initializes the SDL engine with all its subsystems (timer, audio, video, cdrom, joystick). This function **MUST BE CALLED** before any other SDL function or the interpreter will probably crash.

this function sets the global variable `*sdl-initialized*` to true, and only executes the initialization code if this variable is nil. Returns true on success. If you want to initialize only certain subsystems, call the lower-level function `SDL_Init` with the appropriate flags and set `*sdl-initialized*` to true before opening an SDL screen.

(sdl-terminate)

(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

This shuts down the SDL subsystem. `sdl-initialize` must be called before any new SDL call can be performed again. A call to `sdl-terminate` will close any open SDL screen. In fact, calling `sdl-terminate` is the only way to close an SDL screen.

sdl-screen

(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

a class that creates and manipulates SDL Screens. This class can handle pixel-accurate collision detection. It maintains a timer that measures the time between screen flips. This allows to do time-accurate real-time physical simulation. An SDL Screens is not closed when the corresponding **sdl-screen** object is freed. Closing a screen can be done with a call to (**sdl-terminate**) .

(**new sdl-screen w h cap**) Type: CLASS(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

open a new SDL Screen of size **w** , **h** with title **cap** . This attempts to open a double-buffered 32 bit-per-pixel surface in the video hardware. It reverts to 24 or 16 bits (with 32 bit emulation) if a 32 bit screen cannot be allocated. The function **sdl-initialize** MUST be called before opening an **sdl-screen** , or the interpreter will crash!

(**==> sdl-screen flip**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

flip the buffers. The content of the visible screen is replaced by the content of the back buffer (in which drawing operations are performed). This returns the time since the last call to flip in seconds (as a floating point number). The collision detection array is cleared by this call.

(**==> sdl-screen toggle-fullscreen**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

Toggle full screen mode. Returns 1 on success , 0 on failure.

(**==> sdl-screen clear**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

fill the screen with back. Nothing actually happens on the screen until the next call to flip.

(**new sdlgl-screen w h cap**) Type: CLASS(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

just like **sdl-screen** , except that OpenGL drawing is enabled. Always opens a 32 bit window.

sdlidx-screen

(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

A subclass of **sdl-screen** that associates an SDL screen with an **IDX** of the same size. Whatever is contained in the **IDX** can drawn on the screen by calling the "flip" method. Unlike **sdl-screen** , this class does not use SDL's double buffering mechanism , but uses its own back buffer in which all drawing operations take place. **sdlidx-screen** can be used with sprites. Sprites drawn into the screen will appear in the **IDX**. No collision detection between sprites and the background directly drawn into the **IDX** can be performed (but it works between sprites).

(**new sdlidx-screen m caption**) Type: CLASS(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

opens a new SDL screen where IDX *m* (and *idx3* of ubytes) is used as a back buffer in which images can be drawn. The last dimension of *m* should be 3 or 4. If it is 3, the RGB image in *m* will be drawn opaque, if it is 4, the RGBA image in *m* will be drawn on top of the current image using the A channel as transparency. The screen is updated by a call to the method `flip`.

(==> **sdliidx-screen flip**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

Update screen with content of the IDX back buffer. return number of milliseconds since the last call to `flip`.

sdl-event

(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

a high-level class to manipulate SDL events.

(new **sdl-event**) Type: CLASS(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

create a new `sdl-event` object.

(==> **sdle type**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return type of `sdl-event` **sdle**.

(==> **sdle key-keysym-sym**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return symbol of keyboard key symbol of an `sdl-event`.

(==> **sdle button-x**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return mouse button 1 value of an `sdl-event`.

(==> **sdle button-y**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return mouse button 2 value of an `sdl-event`.

(==> **sdle motion-x**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return horizontal mouse motion of an `sdl-event`.

(==> **sdle motion-y**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return vertical mouse motion of an `sdl-event`.

(==> **sdle free**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

`sdl-event` destructor

(==> **sdl-event get-arrows xyk**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

xyk is an *idx1* of ints of size 3. (*xyk* 0) is set to +1 if right-arrow is pressed and -1 if left-arrow is pressed (*xyk* 1) is set to +1 if up-arrow is pressed and -1 if down-arrow is pressed (*xyk* 2) is set to the keysym of any other key that's pressed

(==> **sdl-event get-keys keymap keystate**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

`keymap` is an `idx1` of `-int-` containing a list keys to be polled. The keys are identified by their SDLK code. `keystate` is an `-idx1-` of `-int-` of the same size. Upon exit , each element of `keystate` whose corresponding key (as specified in `keymap`) is pressed down is set to 1. If the corresponding key is not pressed , it is set to 0. This method can be used for games when simultaneous keypresses of all the keys specified in `keymap` must be detected.

```
(libload "sdl/libsdl")
(sdl-initialize)
(setq scr (new sdl-screen 640 480 "Key Test"))
(setq ev (new sdl-event))
(setq kmap [i @@SDLK_a @@SDLK_d @@SDLK_w @@SDLK_s])
(setq ks (int-matrix 4))
(while t
  (=> ev get-keys kmap ks)
  (when (<> 0 ((idx-sum ks)))
    (print ks)
    (sleep 0.25)))
```

sdl-sprite

(packages/sdl/libsdl.lsh) Last Modified: 2002-12-24 14:21:48

a `sdl-sprite` class for drawing movable objects in an SDL screen with pixel-accurate collision detection. Sprites can contain multiple frames that can be flipped through to make animations. or rotations. Each sprite object can be given an ID (which is preferably between 1 and 32 , but can be larger) , which identifies it during collision detection. For example , all background sprites can be given ID 1 , spaceship 1 ID 2 , spaceship 2 ID 3 , and missiles , ID 4. IDs do not need to be unique (several sprites can share the same ID) , but then the collision detection mechanism won't be able to distinguish them.

(=> **sdl-sprite set-hotpoint** `hx hy`) Type: MSG(packages/sdl/libsdl.lsh)
Last Modified: 2002-12-24 14:21:48

sets the "hot point" (the handle) of the sprite to coordinates `hx` , `hy` .

(=> **sdl-sprite get-hotpoint**) Type: MSG(packages/sdl/libsdl.lsh)
Last Modified: 2002-12-24 14:21:48

get the "hot point" (the handle) of the sprite as a list (`hx` , `hy`).

(=> **sdl-sprite get-width**) Type: MSG(packages/sdl/libsdl.lsh) Last
Modified: 2002-12-24 14:21:48

return width of current frame

(=> **sdl-sprite get-height**) Type: MSG(packages/sdl/libsdl.lsh) Last
Modified: 2002-12-24 14:21:48

return height of current frame

(**new sdl-sprite** `screen id`) Type: CLASS(packages/sdl/libsdl.lsh) Last
Modified: 2002-12-24 14:21:48

create a new sprite on screen `screen` with ID `id` .

(==> **SDL-Sprite alloc-frame**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

private method for allocating new frames in a sprite

(==> **SDL-Sprite make-frame** *img i hx hy*) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

private method to set frame *i* to SDL_Surface *img* . the screen must be initialized before calling this.

(==> **SDL-Sprite load-frame** *file i hx hy*) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

load image file *file* into frame *i* , and set the hot point (handle) to *hx hy* . The images file can be in any format that SDL_image can handle (see *img-load*). The screen must be initialized before calling this

(==> **SDL-Sprite rotscale-frame** *src dst angle coeff*) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

take frame *src* , rotate it by *angle* degrees , and scale it by coefficient *coeff* , then sets the *dst* -th frame to the resulting image. This function is convenient for generating views of a sprite at all possible angles automatically. This must be done in advance as this function is too slow to generate frames on-the-fly while the game is running.

(==> **SDL-Sprite make-frame-idx** *m i hx hy*) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

set frame *i* to the image contained in idx *m* . The idx will be interpreted as an heightxwidthx4 idx of bytes containing an RGBA images. The screen must be initialized before calling this

(==> **SDL-Sprite get-frame-idx** *i*) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

Return an idx containing the image of the *i* -th frame.

(==> **SDL-Sprite set-frame** *i*) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

make frame *i* the current frame (i.e. the one that will be drawn next).

(==> **SDL-Sprite get-frame**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return index of current frame.

(==> **SDL-Sprite move** *lx ly*) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

move sprite to *lx* , *ly* . This does not actually draw anything.

(==> **SDL-Sprite moverel** *dx dy*) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

move sprite by *dx* , *dy* relative to the current position. This does not actually draw anything.

(==> **SDL-Sprite draw**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

draw the sprite on the screen using the current frame and position. This does not do any collision detection.

(==> **SDL-Sprite drawc**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

draw the sprite on the screen using the current frame and position. This performs all the operation necessary to do pixel-accurate collision detection. Any pixel with a non-zero alpha value drawn on top of another previously drawn object triggers the collision detection.

(**==> sdl-sprite test-collision**) Type: MSG(packages/sdl/libsdl.lsh)
Last Modified: 2002-12-24 14:21:48

Prepare to test if the sprite would collide with any sprite already drawn on the screen. Unlike `drawc`, this method does not draw the sprite and does not modify the screen's collision map. It provides a way to test for collision and take actions before the collision actually happens. The function `collided` must be subsequently called to actually test if the current sprite would have collided with any other already drawn sprites.

(**collided sp1 sp2**) (packages/sdl/libsdl.lsh) Last Modified: 2002-12-24 14:21:48
returns true any sprite with the same id as `sp1` has collided with any sprite with the same id as `sp2`. For this to return a meaningful result, both sprites must have been drawn with the `drawc`, or the one must have been drawn with `drawc`, and the other one must have been sent the message `test-collision`. This must be called before the screen flip (but after the calls to `drawc` or `test-collision`) since the screen flip clears the collision detection array.

mover

(packages/sdl/libsdl.lsh) Last Modified: 2002-12-24 14:21:48

a class that can move sprites according to 2D Newtonian dynamics. The simplest use consists in creating a mover, then go around a loop that:

- gets input from the keyboard, mouse or joystick.
- calls the mover methods `push`, and `update`
- calls the mover method `move` with a sprite as argument
- draws the sprite
- flips the screen
- sets the `deltat` of the mover to the value returned by the screen flip method.

(**==> mover set-mass m**) Type: MSG(packages/sdl/libsdl.lsh) Last Modified: 2002-12-24 14:21:48

set the mass of the mover.

(**==> mover get-mass**) Type: MSG(packages/sdl/libsdl.lsh) Last Modified: 2002-12-24 14:21:48

return mass of mover

(**==> mover set-deltat dt**) Type: MSG(packages/sdl/libsdl.lsh) Last Modified: 2002-12-24 14:21:48

sets the time increment between updates (in seconds) to `dt`.

(**==> mover get-deltat**) Type: MSG(packages/sdl/libsdl.lsh) Last Modified: 2002-12-24 14:21:48

gets the time increment between updates.

(**==> mover set-state** *lx ly lvx lvy lax lay*) Type: MSG(packages/sdl/libSDL.lsh)
Last Modified: 2002-12-24 14:21:48

(**==> mover get-state**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

(**==> mover set-state** *x y vx vy ax ay mass deltat*) Type: MSG(packages/sdl/libSDL.lsh)
Last Modified: 2002-12-24 14:21:48

create a new sprite mover with initial position *x* , *y* , initial velocity *vx* , *vy* , initial acceleration *ax* , *ay* , mass *mass* , and time increment *deltat* .

(**==> mover push** *fx fy*) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

apply a force vector *fx fy* to the mover. the unit is mass units times pixels per second per second. In other words , a mover of mass 1 , with a force vector of length 1 , will accelerate 1 pixel per second per second (assuming the time increment is correct).

(**==> mover update**) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

updates the positions and velocity according to 2D Newtonian mechanics.

(**==> mover move-sprite** *sp*) Type: MSG(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

move sprite *sp* to the position of the mover.

SDL_Rect function

(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

functions to manipulate SDL_Rect structures Most users will prefer to use the higher level sdl-rect class rather than these functions

low-level functions on SDL-Rect (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

(**new-sdl-rect** *x y w h*) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

allocate a new SDL_Rect structure , fill it up with *x* , *y* , *w* , *h* (position , width , height) , and return a gptr to it. The result must be freed with free-sdl-rect.

(**set-sdl-rect** *r x y w h*) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

sets the position width and height of SDL_Rect pointed to by *r* to the arguments.

(**free-sdl-rect** *r*) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48
deallocate SDL_Rect structure pointed to by *r* .

sdl-rect (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

a high-level class to manipulate SDL_Rect with automatic garbage collection and the like.

(**new sdl-rect** *x y w h*) Type: CLASS(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

create a new sdl-rect.

SDL_Surface

(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

manipulating SDL_Surface structures

low level functions on SDL_Surface (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

(sdl-surface-w s) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return width of SDL_Surface pointed to by **s** .

(sdl-surface-h s) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return height of SDL_Surface pointed to by **s** .

(sdl-surface-pitch s) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return pitch of SDL_Surface pointed to by **s** .

(sdl-surface-pixels s) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return a pointer to the pixel data of the SDL_Surface pointed to by **s** .

(sdl-surface-offset s) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

(sdl-surface-ptr-idx m) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

create an SDL_Surface from an IDX. the pixel area of the SDL_Surface points to the data of the IDX. CAUTION: this function is a bit dangerous because the idx can be deallocated without the sdl_surface knowing about it.

(idx-to-sdl-surface m) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

Return an RGB SDL_Surface whose pixels are filled up with the content of IDX **m** . **m** must be an idx3 of ubyte whose last dimension must be 4. It will be interpreted as an RGBA image.

(sdl-surface-to-idx surface) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

returns an idx3 of ubytes filled up with the RGBA pixel values of SDL_surface **surface** .

SDL Drawing Functions

(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

to draw into SDL_Surfaces.

(sdl-fill-rect dst x y w h r g) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

fills a rectangle defined by **x y w h** of an SDL_Surface **dst** with RGB values **r g b** .

Collision Detection (low level functions)

(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

a set of function to manage pixel-accurate collision detection between object drawn on an SDL_Surface.

(sdl-collide id src srcrect dst dstrect coll) (packages/sdl/libSDL.lsh)
Last Modified: 2002-12-24 14:21:48

this is a low-level function for pixel-accurate collision detection. The arguments **src srcrect dst dstrect** play the same role as the corresponding arguments of `SDL_BlitSurface`. **dst** is an `idx3` of ints. The vector of ints obtained by fixing the first two indexes and varying the last is interpreted as a bitmap. The **id** -th bit of the bitmap at location **i** , **j** is set when a non-transparent pixel of an object whose ID is **id** has been painted at the corresponding location. **src** is an RLE-encoded mask of an object (a list of runs) as returned by `rle-encode-alpha` which indicates which pixels trigger collision. **srcrect** is a subrectangle in **src** that is actually painted (must be `NULL` in current implementation). **dstrect** is the clipping rectangle in **dst** . **coll** is an `idx1` of int interpreted as a bitmap , which on exit contains a set bit for every already painted object that collided with the current object.

(sdl-test-collision id src srcrect dst dstrect coll) (packages/sdl/libSDL.lsh)
Last Modified: 2002-12-24 14:21:48

this is a low-level function for pixel-accurate collision detection. This function is essentially identical to `sdl-collide` , except it doesn't modify the destination bitmap **dst** . This function can be used to get a list of which sprites would be collided if a sprite were drawn. On output , the **i**-th bit of **coll** (`idx1` of `ubyte`) is set if the any pixel of **src** would collide with any non-transparent pixel of **dst** . The arguments **src srcrect dst dstrect** play the same role as the corresponding arguments of `SDL_BlitSurface`. **dst** is an `idx3` of ints. The vector of ints obtained by fixing the first two indexes and varying the last is interpreted as a bitmap. **src** is an RLE-encoded mask of an object (a list of runs) as returned by `rle-encode-alpha` which indicates which pixels trigger collision. **srcrect** is a subrectangle in **src** that is actually painted (must be `NULL` in current implementation). **dstrect** is the clipping rectangle in **dst** . **coll** is an `idx1` of int interpreted as a bitmap , which on exit contains a set bit for every already painted object that collided with the current object.

SDL_Event manipulation

(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

functions to create and manipulate `SDL_Event` structures for keyboard/mouse/joystick event grabbing.

low level access to SDL_event structures (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

(new-sdl-event) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48
allocate and return a pointer to a new `SDL_Event` structure.

(sdl-event-type e) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48
return the type field of an `SDL_Event` pointed to by **e** .

(sdl-event-key-keysym-sym e) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return the field `key.keysym.sym` of an `SDL_Event` pointed to by **e** .

(sdl-event-button-x e) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return the button.x field of an SDL_Event pointed to by *e* .

(sdl-event-button-y e) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return the button.y field of an SDL_Event pointed to by *e* .

(sdl-event-motion-x e) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return the motion.x field of an SDL_Event pointed to by *e* .

(sdl-event-motion-y e) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

return the motion.y field of an SDL_Event pointed to by *e* .

Miscellaneous/internal functions

(packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

these are not meant to be called directly.

(rle-encode-alpha image thres) (packages/sdl/libSDL.lsh) Last Modified: 2002-12-24 14:21:48

private returns a gp_{tr} on an RLE-encoded binary mask (used for collision detection) where any pixel of SDL_Surface *image* whose alpha value is less than *thres* is considered transparent

9.3.4 Low-Level Interface to SDL

sdl/SDL.lsh

(SDL_Init flags) (packages/sdl/SDL.lsh) Last Modified: 2002-10-22 20:11:32

```
/* This function loads the SDL dynamically linked library and initializes
 * the subsystems specified by 'flags' (and those satisfying dependencies)
 * Unless the SDL_INIT_NOPARACHUTE flag is set, it will install cleanup
 * signal handlers for some commonly ignored fatal signals (like SIGSEGV)
 */
```

```
extern DECLSPEC int SDL_Init(Uint32 flags);
```

Flags are:

```
SDL_INIT_TIMER           0x00000001
SDL_INIT_AUDIO           0x00000010
SDL_INIT_VIDEO           0x00000020
SDL_INIT_CDROM           0x00000100
SDL_INIT_JOYSTICK        0x00000200
SDL_INIT_NOPARACHUTE     0x00100000 /* don't catch fatal signals */
SDL_INIT_EVENTTHREAD     0x01000000 /* Not supported on all OS's */
SDL_INIT EVERYTHING     0x0000FFFF
```

(SDL_InitSubSystem flags) (packages/sdl/SDL.lsh) Last Modified: 2002-10-22 20:11:32

```
/* This function initializes specific SDL subsystems */
extern DECLSPEC int SDL_InitSubSystem(Uint32 flags);
```

(SDL_QuitSubSystem flags) (packages/sdl/SDL.lsh) Last Modified: 2002-10-22 20:11:32

```
/* This function cleans up specific SDL subsystems */
extern DECLSPEC void SDL_QuitSubSystem(Uint32 flags);
```

(SDL_WasInit flags) (packages/sdl/SDL.lsh) Last Modified: 2002-10-22 20:11:32

```
/* This function returns mask of the specified subsystems which have
   been initialized.
   If 'flags' is 0, it returns a mask of all initialized subsystems.
*/
extern DECLSPEC Uint32 SDL_WasInit(Uint32 flags);
```

(SDL_Quit) (packages/sdl/SDL.lsh) Last Modified: 2002-10-22 20:11:32

```
/* This function cleans up all initialized subsystems and unloads the
 * dynamically linked library. You should call it upon all exit conditions.
 */
extern DECLSPEC void SDL_Quit(void);
```

sdl/SDL_active.lsh

(SDL_GetAppState) (packages/sdl/SDL_active.lsh) Last Modified: 2002-08-19 03:04:08

```
/* Function prototypes */
/*
 * This function returns the current state of the application, which is a
 * bitwise combination of SDL_APPMOUSEFOCUS, SDL_APPINPUTFOCUS, and
 * SDL_APPACTIVE. If SDL_APPACTIVE is set, then the user is able to
 * see your application, otherwise it has been iconified or disabled.
 */
extern DECLSPEC Uint8 SDL_GetAppState(void);
```

sdl/SDL_audio.lsh

(SDL_AudioInit driver_name) (packages/sdl/SDL_audio.lsh) Last Modified: 2002-10-01 18:33:50

```
/* These functions are used internally, and should not be used unless you
 * have a specific need to specify the audio driver you want to use.
 * You should normally use SDL_Init() or SDL_InitSubSystem().
 */
```

```
extern DECLSPEC int SDL_AudioInit(const char *driver_name);
```

(SDL_AudioQuit) (packages/sdl/SDL_audio.lsh) Last Modified: 2002-10-01 18:33:50

```
extern DECLSPEC void SDL_AudioQuit(void);
```

(SDL_AudioDriverName namebuf maxlen) (packages/sdl/SDL_audio.lsh) Last Modified: 2002-10-01 18:33:50

```
/* This function fills the given character buffer with the name of the
 * current audio driver, and returns a pointer to it if the audio driver has
 * been initialized. It returns NULL if no driver has been initialized.
 */
```

```
extern DECLSPEC char *SDL_AudioDriverName(char *namebuf, int maxlen);
```

(SDL_OpenAudio desired obtained) (packages/sdl/SDL_audio.lsh) Last Modified: 2002-10-01 18:33:50

```
/*
 * This function opens the audio device with the desired parameters, and
 * returns 0 if successful, placing the actual hardware parameters in the
 * structure pointed to by 'obtained'. If 'obtained' is NULL, the audio
 * data passed to the callback function will be guaranteed to be in the
 * requested format, and will be automatically converted to the hardware
 * audio format if necessary. This function returns -1 if it failed
 * to open the audio device, or couldn't set up the audio thread.
 *
 * When filling in the desired audio spec structure,
 * 'desired->freq' should be the desired audio frequency in samples-per-second.
 * 'desired->format' should be the desired audio format.
 * 'desired->samples' is the desired size of the audio buffer, in samples.
 * This number should be a power of two, and may be adjusted by the audio
 * driver to a value more suitable for the hardware. Good values seem to
 * range between 512 and 8096 inclusive, depending on the application and
```

```

* CPU speed. Smaller values yield faster response time, but can lead
* to underflow if the application is doing heavy processing and cannot
* fill the audio buffer in time. A stereo sample consists of both right
* and left channels in LR ordering.
* Note that the number of samples is directly related to time by the
* following formula:  $ms = (samples * 1000) / freq$ 
* 'desired->size' is the size in bytes of the audio buffer, and is
* calculated by SDL_OpenAudio().
* 'desired->silence' is the value used to set the buffer to silence,
* and is calculated by SDL_OpenAudio().
* 'desired->callback' should be set to a function that will be called
* when the audio device is ready for more data. It is passed a pointer
* to the audio buffer, and the length in bytes of the audio buffer.
* This function usually runs in a separate thread, and so you should
* protect data structures that it accesses by calling SDL_LockAudio()
* and SDL_UnlockAudio() in your code.
* 'desired->userdata' is passed as the first parameter to your callback
* function.
*
* The audio device starts out playing silence when it's opened, and should
* be enabled for playing by calling SDL_PauseAudio(0) when you are ready
* for your audio callback function to be called. Since the audio driver
* may modify the requested size of the audio buffer, you should allocate
* any local mixing buffers after you open the audio device.
*/
extern DECLSPEC int SDL_OpenAudio(SDL_AudioSpec *desired, SDL_AudioSpec *obtained);

(SDL_GetAudioStatus) (packages/sdl/SDL.audio.lsh) Last Modified: 2002-
10-01 18:33:50

/*
* Get the current audio state:
*/
typedef enum {
    SDL_AUDIO_STOPPED = 0,
    SDL_AUDIO_PLAYING,
    SDL_AUDIO_PAUSED
} SDL_audiostatus;
extern DECLSPEC SDL_audiostatus SDL_GetAudioStatus(void);

(SDL_PauseAudio pause_on) (packages/sdl/SDL.audio.lsh) Last Modi-
fied: 2002-10-01 18:33:50

/*
* This function pauses and unpauses the audio callback processing.

```

```

* It should be called with a parameter of 0 after opening the audio
* device to start playing sound. This is so you can safely initialize
* data for your callback function after opening the audio device.
* Silence will be written to the audio device during the pause.
*/
extern DECLSPEC void SDL_PauseAudio(int pause_on);

(SDL_LoadWAV_RW src freesrc spec audio_buf audio_len) (pack-
ages/sdl/SDL_audio.lsh) Last Modified: 2002-10-01 18:33:50

/*
* This function loads a WAVE from the data source, automatically freeing
* that source if 'freesrc' is non-zero. For example, to load a WAVE file,
* you could do:
*     SDL_LoadWAV_RW(SDL_RWFromFile("sample.wav", "rb"), 1, ...);
*
* If this function succeeds, it returns the given SDL_AudioSpec,
* filled with the audio data format of the wave data, and sets
* 'audio_buf' to a malloc()'d buffer containing the audio data,
* and sets 'audio_len' to the length of that audio buffer, in bytes.
* You need to free the audio buffer with SDL_FreeWAV() when you are
* done with it.
*
* This function returns NULL and sets the SDL error message if the
* wave file cannot be opened, uses an unknown data format, or is
* corrupt. Currently raw and MS-ADPCM WAVE files are supported.
*/
extern DECLSPEC SDL_AudioSpec *SDL_LoadWAV_RW(SDL_RWops *src, int freesrc,
        SDL_AudioSpec *spec, Uint8 **audio_buf, Uint32 *audio_len);

(SDL_LoadWAV file spec audio_buf audio_len) (packages/sdl/SDL_audio.lsh)
Last Modified: 2002-10-01 18:33:50

/* Compatibility convenience function -- loads a WAV from a file */

(SDL_FreeWAV <audio_buf?>) (packages/sdl/SDL_audio.lsh) Last Mod-
ified: 2002-10-01 18:33:50

/*
* This function frees data previously allocated with SDL_LoadWAV_RW()
*/
extern DECLSPEC void SDL_FreeWAV(Uint8 *audio_buf);

(SDL_BuildAudioCVT cvt src_format src_channels src_rate dst_format
dst_channels dst_rate) (packages/sdl/SDL_audio.lsh) Last Modified: 2002-
10-01 18:33:50

```

```

/*
 * This function takes a source format and rate and a destination format
 * and rate, and initializes the 'cvt' structure with information needed
 * by SDL_ConvertAudio() to convert a buffer of audio data from one format
 * to the other.
 * This function returns 0, or -1 if there was an error.
 */

```

```

extern DECLSPEC int SDL_BuildAudioCVT(SDL_AudioCVT *cvt,
                                     Uint16 src_format, Uint8 src_channels, int src_rate,
                                     Uint16 dst_format, Uint8 dst_channels, int dst_rate);

```

(**SDL_ConvertAudio** cvt) (packages/sdl/SDL.audio.lsh) Last Modified: 2002-10-01 18:33:50

```

/* Once you have initialized the 'cvt' structure using SDL_BuildAudioCVT(),
 * created an audio buffer cvt->buf, and filled it with cvt->len bytes of
 * audio data in the source format, this function will convert it in-place
 * to the desired format.
 * The data conversion may expand the size of the audio data, so the buffer
 * cvt->buf should be allocated after the cvt structure is initialized by
 * SDL_BuildAudioCVT(), and should be cvt->len*cvt->len_mult bytes long.
 */

```

```

extern DECLSPEC int SDL_ConvertAudio(SDL_AudioCVT *cvt);

```

(**SDL_MixAudio** dst src len volume) (packages/sdl/SDL.audio.lsh) Last Modified: 2002-10-01 18:33:50

```

/*
 * This takes two audio buffers of the playing audio format and mixes
 * them, performing addition, volume adjustment, and overflow clipping.
 * The volume ranges from 0 - 128, and should be set to SDL_MIX_MAXVOLUME
 * for full audio volume. Note this does not change hardware volume.
 * This is provided for convenience -- you can mix your own audio data.
 */

```

(**SDL_LockAudio**) (packages/sdl/SDL.audio.lsh) Last Modified: 2002-10-01 18:33:50

```

/*
 * The lock manipulated by these functions protects the callback function.
 * During a LockAudio/UnlockAudio pair, you can be guaranteed that the
 * callback function is not running. Do not call these from the callback
 * function or you will cause deadlock.
 */

```

```

*/
extern DECLSPEC void SDL_LockAudio(void);
extern DECLSPEC void SDL_UnlockAudio(void);

```

(SDL_UnlockAudio) (packages/sdl/SDL_audio.lsh) Last Modified: 2002-10-01 18:33:50

unlock audio.

(SDL_CloseAudio) (packages/sdl/SDL_audio.lsh) Last Modified: 2002-10-01 18:33:50

```

/*
 * This function shuts down audio processing and closes the audio device.
 */
extern DECLSPEC void SDL_CloseAudio(void);

```

sdl/SDL_byteorder.lsh

(SDL_byteorder) (packages/sdl/SDL_byteorder.lsh) Last Modified: 2002-08-19 03:04:08

```

/* Pardon the mess, I'm trying to determine the endianness of this host.
   I'm doing it by preprocessor defines rather than some sort of configure
   script so that application code can use this too. The "right" way would
   be to dynamically generate this file on install, but that's a lot of work.
 */
#if defined(__i386__) || defined(__ia64__) || defined(WIN32) || \
    (defined(__alpha__) || defined(__alpha)) || \
    defined(__arm__) || \
    (defined(__mips__) && defined(__MIPSEL__)) || \
    defined(__LITTLE_ENDIAN__)
#define SDL_BYTEORDER      SDL_LIL_ENDIAN
#else
#define SDL_BYTEORDER      SDL_BIG_ENDIAN
#endif

```

sdl/SDL_cdrom.lsh

(CD_INDRIVE status) (packages/sdl/SDL_cdrom.lsh) Last Modified: 2002-08-19 03:04:08

```

/* The possible states which a CD-ROM drive can be in. */
typedef enum {

```

```

        CD_TRAYEMPTY,
        CD_STOPPED,
        CD_PLAYING,
        CD_PAUSED,
        CD_ERROR = -1
} CDstatus;
/* Given a status, returns true if there's a disk in the drive */
(SDL_CNumDrives) (packages/sdl/SDL_cdrom.lsh) Last Modified: 2002-
08-19 03:04:08

/* Returns the number of CD-ROM drives on the system, or -1 if
SDL_Init() has not been called with the SDL_INIT_CDROM flag.
*/
extern DECLSPEC int SDL_CNumDrives(void);

(SDL_CDName drive) (packages/sdl/SDL_cdrom.lsh) Last Modified: 2002-
08-19 03:04:08

/* Returns a human-readable, system-dependent identifier for the CD-ROM.
Example:
    "/dev/cdrom"
    "E:"
    "/dev/disk/ide/1/master"
*/
extern DECLSPEC const char * SDL_CDName(int drive);

(SDL_CDOpen drive) (packages/sdl/SDL_cdrom.lsh) Last Modified: 2002-
08-19 03:04:08

/* Opens a CD-ROM drive for access. It returns a drive handle on success,
or NULL if the drive was invalid or busy. This newly opened CD-ROM
becomes the default CD used when other CD functions are passed a NULL
CD-ROM handle.
Drives are numbered starting with 0. Drive 0 is the system default CD-ROM.
*/
extern DECLSPEC SDL_CD * SDL_CDOpen(int drive);

(SDL_CDStatus cdrom) (packages/sdl/SDL_cdrom.lsh) Last Modified: 2002-
08-19 03:04:08

/* This function returns the current status of the given drive.
If the drive has a CD in it, the table of contents of the CD and current
play position of the CD will be stored in the SDL_CD structure.
*/
extern DECLSPEC CDstatus SDL_CDStatus(SDL_CD *cdrom);

```

(SDL_CDPlayTracks cdrom start_track start_frame ntracks nframes)
 (packages/sdl/SDL_cdrom.lsh) Last Modified: 2002-08-19 03:04:08

```
/* Play the given CD starting at 'start_track' and 'start_frame' for 'ntracks'
   tracks and 'nframes' frames.  If both 'ntrack' and 'nframe' are 0, play
   until the end of the CD.  This function will skip data tracks.
   This function should only be called after calling SDL_CDStatus() to
   get track information about the CD.
```

```
For example:
```

```
    // Play entire CD:
    if ( CD_INDRIVE(SDL_CDStatus(cdrom)) )
        SDL_CDPlayTracks(cdrom, 0, 0, 0, 0);
    // Play last track:
    if ( CD_INDRIVE(SDL_CDStatus(cdrom)) ) {
        SDL_CDPlayTracks(cdrom, cdrom->numtracks-1, 0, 0, 0);
    }
    // Play first and second track and 10 seconds of third track:
    if ( CD_INDRIVE(SDL_CDStatus(cdrom)) )
        SDL_CDPlayTracks(cdrom, 0, 0, 2, 10);
```

```
This function returns 0, or -1 if there was an error.
```

```
*/
extern DECLSPEC int SDL_CDPlayTracks(SDL_CD *cdrom,
                                     int start_track, int start_frame, int ntracks, int nframes);
```

(SDL_CDPlay cdrom start length) (packages/sdl/SDL_cdrom.lsh) Last
 Modified: 2002-08-19 03:04:08

```
/* Play the given CD starting at 'start' frame for 'length' frames.
   It returns 0, or -1 if there was an error.
```

```
*/
extern DECLSPEC int SDL_CDPlay(SDL_CD *cdrom, int start, int length);
```

(SDL_CDPause cdrom) (packages/sdl/SDL_cdrom.lsh) Last Modified: 2002-
 08-19 03:04:08

```
/* Pause play -- returns 0, or -1 on error */
extern DECLSPEC int SDL_CDPause(SDL_CD *cdrom);
```

(SDL_CDResume cdrom) (packages/sdl/SDL_cdrom.lsh) Last Modified:
 2002-08-19 03:04:08

```
/* Resume play -- returns 0, or -1 on error */
extern DECLSPEC int SDL_CDResume(SDL_CD *cdrom);
```

(**SDL_CDStop** cdrom) (packages/sdl/SDL_cdrom.lsh) Last Modified: 2002-08-19 03:04:08

```
/* Stop play -- returns 0, or -1 on error */
extern DECLSPEC int SDL_CDStop(SDL_CD *cdrom);
```

(**SDL_CDEject** cdrom) (packages/sdl/SDL_cdrom.lsh) Last Modified: 2002-08-19 03:04:08

```
/* Eject CD-ROM -- returns 0, or -1 on error */
extern DECLSPEC int SDL_CDEject(SDL_CD *cdrom);
```

(**SDL_CDClose** cdrom) (packages/sdl/SDL_cdrom.lsh) Last Modified: 2002-08-19 03:04:08

```
/* Closes the handle for the CD-ROM drive */
extern DECLSPEC void SDL_CDClose(SDL_CD *cdrom);
```

sdl/SDL_endian.lsh

(**SDL_ReadLE16** src) (packages/sdl/SDL_endian.lsh) Last Modified: 2002-08-19 03:04:08

```
/* Read an item of the specified endianness and return in native format */
extern DECLSPEC Uint16 SDL_ReadLE16(SDL_RWops *src);
```

(**SDL_ReadBE16** src) (packages/sdl/SDL_endian.lsh) Last Modified: 2002-08-19 03:04:08

```
extern DECLSPEC Uint16 SDL_ReadBE16(SDL_RWops *src);
```

(**SDL_ReadLE32** src) (packages/sdl/SDL_endian.lsh) Last Modified: 2002-08-19 03:04:08

```
extern DECLSPEC Uint32 SDL_ReadLE32(SDL_RWops *src);
```

(**SDL_ReadBE32** src) (packages/sdl/SDL_endian.lsh) Last Modified: 2002-08-19 03:04:08

```
extern DECLSPEC Uint32 SDL_ReadBE32(SDL_RWops *src);
```

(**SDL_WriteLE16** dst value) (packages/sdl/SDL_endian.lsh) Last Modified: 2002-08-19 03:04:08

```

/* Write an item of native format to the specified endianness */
extern DECLSPEC int SDL_WriteLE16(SDL_RWops *dst, Uint16 value);

    (SDL_WriteBE16 dst value) (packages/sdl/SDL_endian.lsh) Last Modified: 2002-08-19 03:04:08

extern DECLSPEC int SDL_WriteBE16(SDL_RWops *dst, Uint16 value);

    (SDL_WriteLE32 dst value) (packages/sdl/SDL_endian.lsh) Last Modified: 2002-08-19 03:04:08

extern DECLSPEC int SDL_WriteLE32(SDL_RWops *dst, Uint32 value);

    (SDL_WriteBE32 dst value) (packages/sdl/SDL_endian.lsh) Last Modified: 2002-08-19 03:04:08

extern DECLSPEC int SDL_WriteBE32(SDL_RWops *dst, Uint32 value);

sdl/SDL_error.lsh
(SDL_GetError) (packages/sdl/SDL_error.lsh) Last Modified: 2002-08-19 03:04:08

extern DECLSPEC char * SDL_GetError(void);

    (SDL_ClearError) (packages/sdl/SDL_error.lsh) Last Modified: 2002-08-19 03:04:08

extern DECLSPEC void SDL_ClearError(void);

    (SDL_Error code) (packages/sdl/SDL_error.lsh) Last Modified: 2002-08-19 03:04:08

/* Private error message function - used internally */
#define SDL_OutOfMemory()      SDL_Error(SDL_ENOMEM)
typedef enum {
    SDL_ENOMEM,
    SDL_EFREAD,
    SDL_EFWRITE,
    SDL_EFSEEK,
    SDL_LASTERROR
} SDL_errorcode;
extern void SDL_Error(SDL_errorcode code);

```

sdl/SDL_events.lsh

(SDL_PumpEvents) (packages/sdl/SDL_events.lsh) Last Modified: 2002-10-01 18:34:22

```
/* Pumps the event loop, gathering events from the input devices.
   This function updates the event queue and internal input device state.
   This should only be run in the thread that sets the video mode.
*/
```

```
extern DECLSPEC void SDL_PumpEvents(void);
```

(SDL_PeepEvents events numevents action mask) (packages/sdl/SDL_events.lsh)
Last Modified: 2002-10-01 18:34:22

```
/* Checks the event queue for messages and optionally returns them.
   If 'action' is SDL_ADDEVENT, up to 'numevents' events will be added to
   the back of the event queue.
   If 'action' is SDL_PEEKEVENT, up to 'numevents' events at the front
   of the event queue, matching 'mask', will be returned and will not
   be removed from the queue.
   If 'action' is SDL_GETEVENT, up to 'numevents' events at the front
   of the event queue, matching 'mask', will be returned and will be
   removed from the queue.
   This function returns the number of events actually stored, or -1
   if there was an error. This function is thread-safe.
*/
```

```
typedef enum {
    SDL_ADDEVENT,
    SDL_PEEKEVENT,
    SDL_GETEVENT
} SDL_eventaction;
```

```
extern DECLSPEC int SDL_PeepEvents(SDL_Event *events, int numevents,
```

```
    SDL_eventaction action, Uint32 mask);
```

(SDL_PollEvent event) (packages/sdl/SDL_events.lsh) Last Modified: 2002-10-01 18:34:22

```
/* Polls for currently pending events, and returns 1 if there are any pending
   events, or 0 if there are none available. If 'event' is not NULL, the next
   event is removed from the queue and stored in that area.
*/
```

```
extern DECLSPEC int SDL_PollEvent(SDL_Event *event);
```

(SDL_WaitEvent event) (packages/sdl/SDL_events.lsh) Last Modified:
2002-10-01 18:34:22

```
/* Waits indefinitely for the next available event, returning 1, or 0 if there
   was an error while waiting for events. If 'event' is not NULL, the next
   event is removed from the queue and stored in that area.
```

```
*/
extern DECLSPEC int SDL_WaitEvent(SDL_Event *event);
```

```
(SDL_PushEvent event) (packages/sdl/SDL_events.lsh) Last Modified:
2002-10-01 18:34:22
```

```
/* Add an event to the event queue.
   This function returns 0, or -1 if the event couldn't be added to
   the event queue. If the event queue is full, this function fails.
```

```
*/
extern DECLSPEC int SDL_PushEvent(SDL_Event *event);
```

```
(SDL_SetEventFilter filter) (packages/sdl/SDL_events.lsh) Last Mod-
ified: 2002-10-01 18:34:22
```

```
/*
   This function sets up a filter to process all events before they
   change internal state and are posted to the internal event queue.
   The filter is protypted as:
```

```
*/
typedef int (*SDL_EventFilter)(const SDL_Event *event);
```

```
/*
   If the filter returns 1, then the event will be added to the internal queue.
   If it returns 0, then the event will be dropped from the queue, but the
   internal state will still be updated. This allows selective filtering of
   dynamically arriving events.
```

```
WARNING: Be very careful of what you do in the event filter function, as
         it may run in a different thread!
```

```
There is one caveat when dealing with the SDL_QUITEVENT event type. The
event filter is only called when the window manager desires to close the
application window. If the event filter returns 1, then the window will
be closed, otherwise the window will remain open if possible.
```

```
If the quit event is generated by an interrupt signal, it will bypass the
internal queue and be delivered to the application at the next event poll.
```

```
*/
(SDL_GetEventFilter) (packages/sdl/SDL_events.lsh) Last Modified: 2002-
10-01 18:34:22
```

```
/*
```

```

Return the current event filter - can be used to "chain" filters.
If there is no event filter set, this function returns NULL.
*/
extern DECLSPEC SDL_EventFilter SDL_GetEventFilter(void);

(SDL_EventState type state) (packages/sdl/SDL_events.lsh) Last Modified: 2002-10-01 18:34:22

/*
This function allows you to set the state of processing certain events.
If 'state' is set to SDL_IGNORE, that event will be automatically dropped
from the event queue and will not event be filtered.
If 'state' is set to SDL_ENABLE, that event will be processed normally.
If 'state' is set to SDL_QUERY, SDL_EventState() will return the
current processing state of the specified event.
*/

```

sdl/SDL_getenv.lsh

```

(SDL_putenv variable) (packages/sdl/SDL_getenv.lsh) Last Modified: 2002-08-17 23:10:33

```

```

/* Put a variable of the form "name=value" into the environment */
extern DECLSPEC int SDL_putenv(const char *variable);

```

```

(SDL_getenv name) (packages/sdl/SDL_getenv.lsh) Last Modified: 2002-08-17 23:10:33

```

```

/* Retrieve a variable named "name" from the environment */
extern DECLSPEC char *SDL_getenv(const char *name);

```

sdl/SDL_image.lsh

```

(IMG_LoadTyped_RW src freesrc type) (packages/sdl/SDL_image.lsh)
Last Modified: 2002-08-17 23:10:33

```

```

/* Load an image from an SDL data source.
The 'type' may be one of: "BMP", "GIF", "PNG", etc.
If the image format supports a transparent pixel, SDL will set the
colorkey for the surface. You can enable RLE acceleration on the
surface afterwards by calling:
    SDL_SetColorKey(image, SDL_RLEACCEL, image->format->colorkey);
*/
extern DECLSPEC SDL_Surface *IMG_LoadTyped_RW(SDL_RWops *src, int freesrc,
char *type);

```

(IMG_Load file) (packages/sdl/SDL_image.lsh) Last Modified: 2002-08-17 23:10:33

```
/* Convenience functions */
extern DECLSPEC SDL_Surface *IMG_Load(const char *file);
```

(IMG_Load_RW src freesrc) (packages/sdl/SDL_image.lsh) Last Modified: 2002-08-17 23:10:33

```
extern DECLSPEC SDL_Surface *IMG_Load_RW(SDL_RWops *src, int freesrc);
```

(IMG_InvertAlpha on) (packages/sdl/SDL_image.lsh) Last Modified: 2002-08-17 23:10:33

```
/* Invert the alpha of a surface for use with OpenGL
   This function is now a no-op, and only provided for backwards compatibility.
*/
extern DECLSPEC int IMG_InvertAlpha(int on);
```

sdl/SDL_joystick.lsh

(SDL_NumJoysticks) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```
/* Function prototypes */
/*
 * Count the number of joysticks attached to the system
 */
extern DECLSPEC int SDL_NumJoysticks(void);
```

(SDL_JoystickName device_index) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Get the implementation dependent name of a joystick.
 * This can be called before any joysticks are opened.
 * If no name can be found, this function returns NULL.
 */
extern DECLSPEC const char *SDL_JoystickName(int device_index);
```

(SDL_JoystickOpen device_index) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Open a joystick for use - the index passed as an argument refers to
 * the N'th joystick on the system. This index is the value which will
 * identify this joystick in future joystick events.
 *
 * This function returns a joystick identifier, or NULL if an error occurred.
 */
extern DECLSPEC SDL_Joystick *SDL_JoystickOpen(int device_index);
```

(SDL_JoystickOpened device_index) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Returns 1 if the joystick has been opened, or 0 if it has not.
 */
extern DECLSPEC int SDL_JoystickOpened(int device_index);
```

(SDL_JoystickIndex joystick) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Get the device index of an opened joystick.
 */
extern DECLSPEC int SDL_JoystickIndex(SDL_Joystick *joystick);
```

(SDL_JoystickNumAxes joystick) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Get the number of general axis controls on a joystick
 */
extern DECLSPEC int SDL_JoystickNumAxes(SDL_Joystick *joystick);
```

(SDL_JoystickNumBalls joystick) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Get the number of trackballs on a joystick
 * Joystick trackballs have only relative motion events associated
 * with them and their state cannot be polled.
 */
extern DECLSPEC int SDL_JoystickNumBalls(SDL_Joystick *joystick);
```

(**SDL_JoystickNumHats** joystick) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Get the number of POV hats on a joystick
 */
extern DECLSPEC int SDL_JoystickNumHats(SDL_Joystick *joystick);
```

(**SDL_JoystickNumButtons** joystick) (packages/sdl/SDL_joystick.lsh)
Last Modified: 2002-08-19 03:04:08

```
/*
 * Get the number of buttons on a joystick
 */
extern DECLSPEC int SDL_JoystickNumButtons(SDL_Joystick *joystick);
```

(**SDL_JoystickUpdate**) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Update the current state of the open joysticks.
 * This is called automatically by the event loop if any joystick
 * events are enabled.
 */
extern DECLSPEC void SDL_JoystickUpdate(void);
```

(**SDL_JoystickEventState** state) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Enable/disable joystick event polling.
 * If joystick events are disabled, you must call SDL_JoystickUpdate()
 * yourself and check the state of the joystick when you want joystick
 * information.
 * The state can be one of SDL_QUERY, SDL_ENABLE or SDL_IGNORE.
 */
extern DECLSPEC int SDL_JoystickEventState(int state);
```

(**SDL_JoystickGetAxis** joystick axis) (packages/sdl/SDL_joystick.lsh)
Last Modified: 2002-08-19 03:04:08

```
/*
 * Get the current state of an axis control on a joystick
```

```

* The state is a value ranging from -32768 to 32767.
* The axis indices start at index 0.
*/
extern DECLSPEC Sint16 SDL_JoystickGetAxis(SDL_Joystick *joystick, int axis);

```

(SDL_JoystickGetHat joystick hat) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```

/*
* Get the current state of a POV hat on a joystick
* The return value is one of the following positions:
*/

```

(SDL_JoystickGetBall joystick ball dx dy) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```

/*
* Get the ball axis change since the last poll
* This returns 0, or -1 if you passed it invalid parameters.
* The ball indices start at index 0.
*/

```

```
extern DECLSPEC int SDL_JoystickGetBall(SDL_Joystick *joystick, int ball, int *dx, int *dy);
```

(SDL_JoystickGetButton joystick button) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```

/*
* Get the current state of a button on a joystick
* The button indices start at index 0.
*/

```

```
extern DECLSPEC Uint8 SDL_JoystickGetButton(SDL_Joystick *joystick, int button);
```

(SDL_JoystickClose joystick) (packages/sdl/SDL_joystick.lsh) Last Modified: 2002-08-19 03:04:08

```

/*
* Close a joystick previously opened with SDL_JoystickOpen()
*/

```

```
extern DECLSPEC void SDL_JoystickClose(SDL_Joystick *joystick);
```

sdl/SDL_keyboard.lsh

(SDL_EnableUNICODE enable) (packages/sdl/SDL_keyboard.lsh) Last Modified: 2002-08-19 03:04:08

```

* Enable/Disable UNICODE translation of keyboard input.
* This translation has some overhead, so translation defaults off.
* If 'enable' is 1, translation is enabled.
* If 'enable' is 0, translation is disabled.
* If 'enable' is -1, the translation state is not changed.
* It returns the previous state of keyboard translation.
*/
extern DECLSPEC int SDL_EnableUNICODE(int enable);

(SDL_EnableKeyRepeat delay interval) (packages/sdl/SDL_keyboard.lsh)
Last Modified: 2002-08-19 03:04:08

/*
* Enable/Disable keyboard repeat. Keyboard repeat defaults to off.
* 'delay' is the initial delay in ms between the time when a key is
* pressed, and keyboard repeat begins.
* 'interval' is the time in ms between keyboard repeat events.
*/
/*
* If 'delay' is set to 0, keyboard repeat is disabled.
*/
extern DECLSPEC int SDL_EnableKeyRepeat(int delay, int interval);

(SDL_GetKeyState numkeys) (packages/sdl/SDL_keyboard.lsh) Last Mod-
ified: 2002-08-19 03:04:08

/*
* Get a snapshot of the current state of the keyboard.
* Returns an array of keystates, indexed by the SDLK_* syms.
* Used:
*     Uint8 *keystate = SDL_GetKeyState(NULL);
*     if ( keystate[SDLK_RETURN] ) ... <RETURN> is pressed.
*/
extern DECLSPEC Uint8 * SDL_GetKeyState(int *numkeys);

(SDL_GetModState) (packages/sdl/SDL_keyboard.lsh) Last Modified: 2002-
08-19 03:04:08

/*
* Get the current key modifier state
*/
extern DECLSPEC SDLMod SDL_GetModState(void);

```

(SDL_SetModState modstate) (packages/sdl/SDL_keyboard.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Set the current key modifier state
 * This does not change the keyboard state, only the key modifier flags.
 */
extern DECLSPEC void SDL_SetModState(SDLMod modstate);
```

(SDL_GetKeyName key) (packages/sdl/SDL_keyboard.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Get the name of an SDL virtual keysym
 */
extern DECLSPEC char * SDL_GetKeyName(SDLKey key);
```

sdl/SDL_keysym.lsh

SDL_Keysym (packages/sdl/SDL_keysym.lsh) Last Modified: 2002-08-19 03:04:08

each keyboard symbol is associated with a constant whose symbolic name represents the key, and value represents the key code (generally the ASCII code of the character). Key symbols are of the form `SDLK_0`, `SDLK_A`, `SDLK_b`, `SDLK_SPACE`, etc.... Look into `sdl/SDL_keysym.lsh` for a complete list.

sdl/SDL_mouse.lsh

(SDL_GetMouseState x y) (packages/sdl/SDL_mouse.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Retrieve the current state of the mouse.
 * The current button state is returned as a button bitmask, which can
 * be tested using the SDL_BUTTON(X) macros, and x and y are set to the
 * current mouse cursor position. You can pass NULL for either x or y.
 */
extern DECLSPEC Uint8 SDL_GetMouseState(int *x, int *y);
```

(SDL_GetRelativeMouseState x y) (packages/sdl/SDL_mouse.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * Retrieve the current state of the mouse.
 * The current button state is returned as a button bitmask, which can
```

```

    * be tested using the SDL_BUTTON(X) macros, and x and y are set to the
    * mouse deltas since the last call to SDL_GetRelativeMouseState().
    */
extern DECLSPEC Uint8 SDL_GetRelativeMouseState(int *x, int *y);

    (SDL_WarpMouse x y) (packages/sdl/SDL_mouse.lsh) Last Modified: 2002-
    08-19 03:04:08

/*
    * Set the position of the mouse cursor (generates a mouse motion event)
    */
extern DECLSPEC void SDL_WarpMouse(Uint16 x, Uint16 y);

    (SDL_CreateCursor data mask w h hot_x hot_y) (packages/sdl/SDL_mouse.lsh)
    Last Modified: 2002-08-19 03:04:08

/*
    * Create a cursor using the specified data and mask (in MSB format).
    * The cursor width must be a multiple of 8 bits.
    *
    * The cursor is created in black and white according to the following:
    * data mask    resulting pixel on screen
    * 0    1      White
    * 1    1      Black
    * 0    0      Transparent
    * 1    0      Inverted color if possible, black if not.
    *
    * Cursors created with this function must be freed with SDL_FreeCursor().
    */
extern DECLSPEC SDL_Cursor *SDL_CreateCursor
    (Uint8 *data, Uint8 *mask, int w, int h, int hot_x, int hot_y);

    (SDL_SetCursor cursor) (packages/sdl/SDL_mouse.lsh) Last Modified:
    2002-08-19 03:04:08

/*
    * Set the currently active cursor to the specified one.
    * If the cursor is currently visible, the change will be immediately
    * represented on the display.
    */
extern DECLSPEC void SDL_SetCursor(SDL_Cursor *cursor);

    (SDL_GetCursor) (packages/sdl/SDL_mouse.lsh) Last Modified: 2002-08-
    19 03:04:08

```

```

/*
 * Returns the currently active cursor.
 */
extern DECLSPEC SDL_Cursor * SDL_GetCursor(void);

(SDL_FreeCursor cursor) (packages/sdl/SDL_mouse.lsh) Last Modified:
2002-08-19 03:04:08

/*
 * Deallocates a cursor created with SDL_CreateCursor().
 */
extern DECLSPEC void SDL_FreeCursor(SDL_Cursor *cursor);

(SDL_ShowCursor toggle) (packages/sdl/SDL_mouse.lsh) Last Modified:
2002-08-19 03:04:08

/*
 * Toggle whether or not the cursor is shown on the screen.
 * The cursor start off displayed, but can be turned off.
 * SDL_ShowCursor() returns 1 if the cursor was being displayed
 * before the call, or 0 if it was not. You can query the current
 * state by passing a 'toggle' value of -1.
 */
extern DECLSPEC int SDL_ShowCursor(int toggle);

sdl/SDL_mutex.lsh

(SDL_CreateMutex) (packages/sdl/SDL_mutex.lsh) Last Modified: 2002-08-
19 03:04:08

/* Create a mutex, initialized unlocked */
extern DECLSPEC SDL_mutex * SDL_CreateMutex(void);

(SDL_mutexP mutex) (packages/sdl/SDL_mutex.lsh) Last Modified: 2002-
08-19 03:04:08

/* Lock the mutex (Returns 0, or -1 on error) */
#define SDL_LockMutex(m)      SDL_mutexP(m)
extern DECLSPEC int SDL_mutexP(SDL_mutex *mutex);

(SDL_mutexV mutex) (packages/sdl/SDL_mutex.lsh) Last Modified: 2002-
08-19 03:04:08

```

```

/* Unlock the mutex (Returns 0, or -1 on error) */
#define SDL_UnlockMutex(m)      SDL_mutexV(m)
extern DECLSPEC int SDL_mutexV(SDL_mutex *mutex);

```

(**SDL_DestroyMutex** mutex) (packages/sdl/SDL_mutex.lsh) Last Modified: 2002-08-19 03:04:08

```

/* Destroy a mutex */
extern DECLSPEC void SDL_DestroyMutex(SDL_mutex *mutex);

```

(**SDL_CreateSemaphore** initial_value) (packages/sdl/SDL_mutex.lsh)
Last Modified: 2002-08-19 03:04:08

```

/* Create a semaphore, initialized with value, returns NULL on failure. */
extern DECLSPEC SDL_sem * SDL_CreateSemaphore(Uint32 initial_value);

```

(**SDL_DestroySemaphore** sem) (packages/sdl/SDL_mutex.lsh) Last Modified: 2002-08-19 03:04:08

```

/* Destroy a semaphore */
extern DECLSPEC void SDL_DestroySemaphore(SDL_sem *sem);

```

(**SDL_SemWait** sem) (packages/sdl/SDL_mutex.lsh) Last Modified: 2002-08-19 03:04:08

```

/* This function suspends the calling thread until the semaphore pointed
 * to by sem has a positive count. It then atomically decreases the semaphore
 * count.
 */
extern DECLSPEC int SDL_SemWait(SDL_sem *sem);

```

(**SDL_SemTryWait** sem) (packages/sdl/SDL_mutex.lsh) Last Modified: 2002-08-19 03:04:08

```

/* Non-blocking variant of SDL_SemWait(), returns 0 if the wait succeeds,
 * SDL_MUTEX_TIMEDOUT if the wait would block, and -1 on error.
 */
extern DECLSPEC int SDL_SemTryWait(SDL_sem *sem);

```

(**SDL_SemWaitTimeout** sem ms) (packages/sdl/SDL_mutex.lsh) Last Modified: 2002-08-19 03:04:08

```
/* Variant of SDL_SemWait() with a timeout in milliseconds, returns 0 if
the wait succeeds, SDL_MUTEX_TIMEDOUT if the wait does not succeed in
the allotted time, and -1 on error.
On some platforms this function is implemented by looping with a delay
of 1 ms, and so should be avoided if possible.
*/
```

```
extern DECLSPEC int SDL_SemWaitTimeout(SDL_sem *sem, Uint32 ms);
```

```
(SDL_SemPost sem) (packages/sdl/SDL_mutex.lsh) Last Modified: 2002-
08-19 03:04:08
```

```
/* Atomically increases the semaphore's count (not blocking), returns 0,
or -1 on error.
*/
```

```
extern DECLSPEC int SDL_SemPost(SDL_sem *sem);
```

```
(SDL_CreateCond) (packages/sdl/SDL_mutex.lsh) Last Modified: 2002-
08-19 03:04:08
```

```
/* Create a condition variable */
```

```
extern DECLSPEC SDL_cond * SDL_CreateCond(void);
```

```
(SDL_DestroyCond cnd) (packages/sdl/SDL_mutex.lsh) Last Modified:
2002-08-19 03:04:08
```

```
/* Destroy a condition variable */
```

```
extern DECLSPEC void SDL_DestroyCond(SDL_cond *cond);
```

```
(SDL_CondSignal cnd) (packages/sdl/SDL_mutex.lsh) Last Modified: 2002-
08-19 03:04:08
```

```
/* Restart one of the threads that are waiting on the condition variable,
returns 0 or -1 on error.
*/
```

```
extern DECLSPEC int SDL_CondSignal(SDL_cond *cond);
```

```
(SDL_CondBroadcast cnd) (packages/sdl/SDL_mutex.lsh) Last Modified:
2002-08-19 03:04:08
```

```
/* Restart all threads that are waiting on the condition variable,
returns 0 or -1 on error.
*/
```

```
extern DECLSPEC int SDL_CondBroadcast(SDL_cond *cond);
```

(SDL_CondWait cnd mut) (packages/sdl/SDL_mutex.lsh) Last Modified: 2002-08-19 03:04:08

```
/* Wait on the condition variable, unlocking the provided mutex.
   The mutex must be locked before entering this function!
   Returns 0 when it is signaled, or -1 on error.
*/
extern DECLSPEC int SDL_CondWait(SDL_cond *cond, SDL_mutex *mut);
```

(SDL_CondWaitTimeout cnd mutex ms) (packages/sdl/SDL_mutex.lsh)
Last Modified: 2002-08-19 03:04:08

```
/* Waits for at most 'ms' milliseconds, and returns 0 if the condition
   variable is signaled, SDL_MUTEX_TIMEDOUT if the condition is not
   signaled in the allotted time, and -1 on error.
   On some platforms this function is implemented by looping with a delay
   of 1 ms, and so should be avoided if possible.
*/
extern DECLSPEC int SDL_CondWaitTimeout(SDL_cond *cond, SDL_mutex *mutex, Uint32 ms);
```

sdl/SDL_rwops.lsh

SDL_RWops (packages/sdl/SDL_rwops.lsh) Last Modified: 2002-08-19 03:04:08

Functions to create SDL_RWops structures from various data sources

(SDL_RWFromFile file mode) (packages/sdl/SDL_rwops.lsh) Last Modified: 2002-08-19 03:04:08

```
extern DECLSPEC SDL_RWops * SDL_RWFromFile(const char *file, const char *mode);
```

(SDL_RWFromFP fp autoclose) (packages/sdl/SDL_rwops.lsh) Last Modified: 2002-08-19 03:04:08

```
extern DECLSPEC SDL_RWops * SDL_RWFromFP(FILE *fp, int autoclose);
```

(SDL_RWFromMem mem size) (packages/sdl/SDL_rwops.lsh) Last Modified: 2002-08-19 03:04:08

```
extern DECLSPEC SDL_RWops * SDL_RWFromMem(void *mem, int size);
```

(SDL_AllocRW) (packages/sdl/SDL_rwops.lsh) Last Modified: 2002-08-19 03:04:08

```
extern DECLSPEC SDL_RWops * SDL_AllocRW(void);
```

(SDL_FreeRW area) (packages/sdl/SDL_rwops.lsh) Last Modified: 2002-08-19 03:04:08

```
extern DECLSPEC void SDL_FreeRW(SDL_RWops *area);
```

sdl/SDL_syswm.lsh

(SDL_GetWMInfo info) (packages/sdl/SDL_syswm.lsh) Last Modified: 2002-08-19 03:04:08

```
/*
 * This function gives you custom hooks into the window manager information.
 * It fills the structure pointed to by 'info' with custom information and
 * returns 1 if the function is implemented. If it's not implemented, or
 * the version member of the 'info' structure is invalid, it returns 0.
 */
extern DECLSPEC int SDL_GetWMInfo(SDL_SysWMInfo *info);
```

sdl/SDL_thread.lsh

(SDL_CreateThread fn data) (packages/sdl/SDL_thread.lsh) Last Modified: 2002-08-19 03:04:08

```
/* Create a thread */
extern DECLSPEC SDL_Thread * SDL_CreateThread(int (*fn)(void *), void *data);
```

(SDL_ThreadID) (packages/sdl/SDL_thread.lsh) Last Modified: 2002-08-19 03:04:08

```
/* Get the 32-bit thread identifier for the current thread */
extern DECLSPEC Uint32 SDL_ThreadID(void);
```

(SDL_GetThreadID thread) (packages/sdl/SDL_thread.lsh) Last Modified: 2002-08-19 03:04:08

```
/* Get the 32-bit thread identifier for the specified thread,
   equivalent to SDL_ThreadID() if the specified thread is NULL.
 */
extern DECLSPEC Uint32 SDL_GetThreadID(SDL_Thread *thread);
```

(SDL_WaitThread thread status) (packages/sdl/SDL_thread.lsh) Last Modified: 2002-08-19 03:04:08

```

/* Wait for a thread to finish.
   The return code for the thread function is placed in the area
   pointed to by 'status', if 'status' is not NULL.
*/
extern DECLSPEC void SDL_WaitThread(SDL_Thread *thread, int *status);

(SDL_KillThread thread) (packages/sdl/SDL_thread.lsh) Last Modified:
2002-08-19 03:04:08

/* Forcefully kill a thread without worrying about its state */
extern DECLSPEC void SDL_KillThread(SDL_Thread *thread);

sdl/SDL_timer.lsh

(SDL_GetTicks) (packages/sdl/SDL_timer.lsh) Last Modified: 2002-08-19 03:04:08

/* Get the number of milliseconds since the SDL library initialization.
 * Note that this value wraps if the program runs for more than ~49 days.
 */
extern DECLSPEC Uint32 SDL_GetTicks(void);

(SDL_Delay ms) (packages/sdl/SDL_timer.lsh) Last Modified: 2002-08-19
03:04:08

/* Wait a specified number of milliseconds before returning */
extern DECLSPEC void SDL_Delay(Uint32 ms);

(SDL_SetTimer interval callback) (packages/sdl/SDL_timer.lsh) Last
Modified: 2002-08-19 03:04:08

/* Function prototype for the timer callback function */
typedef Uint32 (*SDL_TimerCallback)(Uint32 interval);
/* Set a callback to run after the specified number of milliseconds has
 * elapsed. The callback function is passed the current timer interval
 * and returns the next timer interval. If the returned value is the
 * same as the one passed in, the periodic alarm continues, otherwise a
 * new alarm is scheduled. If the callback returns 0, the periodic alarm
 * is cancelled.
 *
 * To cancel a currently running timer, call SDL_SetTimer(0, NULL);
 *
 * The timer callback function may run in a different thread than your

```

```

* main code, and so shouldn't call any functions from within itself.
*
* The maximum resolution of this timer is 10 ms, which means that if
* you request a 16 ms timer, your callback will run approximately 20 ms
* later on an unloaded system. If you wanted to set a flag signaling
* a frame update at 30 frames per second (every 33 ms), you might set a
* timer for 30 ms:
*   SDL_SetTimer((33/10)*10, flag_update);
*
* If you use this function, you need to pass SDL_INIT_TIMER to SDL_Init().
*
* Under UNIX, you should not use raise or use SIGALRM and this function
* in the same program, as it is implemented using setitimer(). You also
* should not use this function in multi-threaded applications as signals
* to multi-threaded apps have undefined behavior in some implementations.
*/
extern DECLSPEC int SDL_SetTimer(Uint32 interval, SDL_TimerCallback callback);

(SDL_AddTimer interval callback param) (packages/sdl/SDL_timer.lsh)
Last Modified: 2002-08-19 03:04:08

```

```

/* Add a new timer to the pool of timers already running.
   Returns a timer ID, or NULL when an error occurs.
*/
extern DECLSPEC SDL_TimerID SDL_AddTimer(Uint32 interval, SDL_NewTimerCallback callback, void

(SDL_RemoveTimer tim) (packages/sdl/SDL_timer.lsh) Last Modified:
2002-08-19 03:04:08

```

```

/* Remove one of the multiple timers knowing its ID.
   * Returns a boolean value indicating success.
   */
extern DECLSPEC SDL_bool SDL_RemoveTimer(SDL_TimerID t);

```

sdl/SDL_version.lsh

```

(SDL_Linked_Version) (packages/sdl/SDL_version.lsh) Last Modified: 2002-
08-19 03:04:08

```

```

/* This function gets the version of the dynamically linked SDL library.
   it should NOT be used to fill a version structure, instead you should
   use the SDL_Version() macro.
*/
extern DECLSPEC const SDL_version * SDL_Linked_Version(void);

```

sdl/SDL_video.lsh

(SDL_MUSTLOCK surface) (packages/sdl/SDL_video.lsh) Last Modified:
2002-10-28 14:55:43

```
/* Evaluates to true if the surface needs to be locked before access */
#define SDL_MUSTLOCK(surface)    \
    (surface->offset ||         \
     ((surface->flags & (SDL_HWSURFACE|SDL_ASYNCBLIT|SDL_RLEACCEL)) != 0))
```

(SDL_VideoInit driver_name flags) (packages/sdl/SDL_video.lsh) Last
Modified: 2002-10-28 14:55:43

```
/* These functions are used internally, and should not be used unless you
 * have a specific need to specify the video driver you want to use.
 * You should normally use SDL_Init() or SDL_InitSubSystem().
 *
 * SDL_VideoInit() initializes the video subsystem -- sets up a connection
 * to the window manager, etc, and determines the current video mode and
 * pixel format, but does not initialize a window or graphics mode.
 * Note that event handling is activated by this routine.
 *
 * If you use both sound and video in your application, you need to call
 * SDL_Init() before opening the sound device, otherwise under Win32 DirectX,
 * you won't be able to set full-screen display modes.
 */
```

```
extern DECLSPEC int SDL_VideoInit(const char *driver_name, Uint32 flags);
```

(SDL_VideoQuit) (packages/sdl/SDL_video.lsh) Last Modified: 2002-10-
28 14:55:43

```
extern DECLSPEC void SDL_VideoQuit(void);
```

(SDL_VideoDriverName namebuf maxlen) (packages/sdl/SDL_video.lsh)
Last Modified: 2002-10-28 14:55:43

```
/* This function fills the given character buffer with the name of the
 * video driver, and returns a pointer to it if the video driver has
 * been initialized. It returns NULL if no driver has been initialized.
 */
```

```
extern DECLSPEC char *SDL_VideoDriverName(char *namebuf, int maxlen);
```

(SDL_GetVideoSurface) (packages/sdl/SDL_video.lsh) Last Modified: 2002-
10-28 14:55:43

```

/*
 * This function returns a pointer to the current display surface.
 * If SDL is doing format conversion on the display surface, this
 * function returns the publicly visible surface, not the real video
 * surface.
 */
extern DECLSPEC SDL_Surface * SDL_GetVideoSurface(void);

(SDL_GetVideoInfo) (packages/sdl/SDL-video.lsh) Last Modified: 2002-
10-28 14:55:43

/*
 * This function returns a read-only pointer to information about the
 * video hardware. If this is called before SDL_SetVideoMode(), the 'vfmt'
 * member of the returned structure will contain the pixel format of the
 * "best" video mode.
 */
extern DECLSPEC const SDL_VideoInfo * SDL_GetVideoInfo(void);

(SDL_VideoModeOK width height bpp flags) (packages/sdl/SDL-video.lsh)
Last Modified: 2002-10-28 14:55:43

/*
 * Check to see if a particular video mode is supported.
 * It returns 0 if the requested mode is not supported under any bit depth,
 * or returns the bits-per-pixel of the closest available mode with the
 * given width and height. If this bits-per-pixel is different from the
 * one used when setting the video mode, SDL_SetVideoMode() will succeed,
 * but will emulate the requested bits-per-pixel with a shadow surface.
 *
 * The arguments to SDL_VideoModeOK() are the same ones you would pass to
 * SDL_SetVideoMode()
 */
extern DECLSPEC int SDL_VideoModeOK(int width, int height, int bpp, Uint32 flags);

(SDL_ListModes format flags) (packages/sdl/SDL-video.lsh) Last Mod-
ified: 2002-10-28 14:55:43

/*
 * Return a pointer to an array of available screen dimensions for the
 * given format and video flags, sorted largest to smallest. Returns
 * NULL if there are no dimensions available for a particular format,
 * or (SDL_Rect **)-1 if any dimension is okay for the given format.

```

```

*
* If 'format' is NULL, the mode list will be for the format given
* by SDL_GetVideoInfo()->vfmt
*/
extern DECLSPEC SDL_Rect ** SDL_ListModes(SDL_PixelFormat *format, Uint32 flags);

(SDL_SetVideoMode width height bpp flags) (packages/sdl/SDL-video.lsh)
Last Modified: 2002-10-28 14:55:43

```

```

/*
* Set up a video mode with the specified width, height and bits-per-pixel.
*
* If 'bpp' is 0, it is treated as the current display bits per pixel.
*
* If SDL_ANYFORMAT is set in 'flags', the SDL library will try to set the
* requested bits-per-pixel, but will return whatever video pixel format is
* available. The default is to emulate the requested pixel format if it
* is not natively available.
*
* If SDL_HWSURFACE is set in 'flags', the video surface will be placed in
* video memory, if possible, and you may have to call SDL_LockSurface()
* in order to access the raw framebuffer. Otherwise, the video surface
* will be created in system memory.
*
* If SDL_ASYNCBLIT is set in 'flags', SDL will try to perform rectangle
* updates asynchronously, but you must always lock before accessing pixels.
* SDL will wait for updates to complete before returning from the lock.
*
* If SDL_HWPALETTE is set in 'flags', the SDL library will guarantee
* that the colors set by SDL_SetColors() will be the colors you get.
* Otherwise, in 8-bit mode, SDL_SetColors() may not be able to set all
* of the colors exactly the way they are requested, and you should look
* at the video surface structure to determine the actual palette.
* If SDL cannot guarantee that the colors you request can be set,
* i.e. if the colormap is shared, then the video surface may be created
* under emulation in system memory, overriding the SDL_HWSURFACE flag.
*
* If SDL_FULLSCREEN is set in 'flags', the SDL library will try to set
* a fullscreen video mode. The default is to create a windowed mode
* if the current graphics system has a window manager.
* If the SDL library is able to set a fullscreen video mode, this flag
* will be set in the surface that is returned.
*
* If SDL_DOUBLEBUF is set in 'flags', the SDL library will try to set up
* two surfaces in video memory and swap between them when you call

```

```

* SDL_Flip(). This is usually slower than the normal single-buffering
* scheme, but prevents "tearing" artifacts caused by modifying video
* memory while the monitor is refreshing. It should only be used by
* applications that redraw the entire screen on every update.
*
* If SDL_RESIZABLE is set in 'flags', the SDL library will allow the
* window manager, if any, to resize the window at runtime. When this
* occurs, SDL will send a SDL_VIDEORESIZE event to you application,
* and you must respond to the event by re-calling SDL_SetVideoMode()
* with the requested size (or another size that suits the application).
*
* If SDL_NOFRAME is set in 'flags', the SDL library will create a window
* without any title bar or frame decoration. Fullscreen video modes have
* this flag set automatically.
*
* This function returns the video framebuffer surface, or NULL if it fails.
*
* If you rely on functionality provided by certain video flags, check the
* flags of the returned surface to make sure that functionality is available.
* SDL will fall back to reduced functionality if the exact flags you wanted
* are not available.
*/
extern DECLSPEC SDL_Surface *SDL_SetVideoMode(int width, int height, int bpp, Uint32 flags);

(SDL_UpdateRects screen numrects rects) (packages/sdl/SDL-video.lsh)
Last Modified: 2002-10-28 14:55:43

/*
* Makes sure the given list of rectangles is updated on the given screen.
* If 'x', 'y', 'w' and 'h' are all 0, SDL_UpdateRect will update the entire
* screen.
* These functions should not be called while 'screen' is locked.
*/
extern DECLSPEC void SDL_UpdateRects(SDL_Surface *screen, int numrects, SDL_Rect *rects);

(SDL_UpdateRect screen x y w h) (packages/sdl/SDL-video.lsh) Last
Modified: 2002-10-28 14:55:43

extern DECLSPEC void SDL_UpdateRect(SDL_Surface *screen, Sint32 x, Sint32 y, Uint32 w, Uint32

(SDL_Flip screen) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-
28 14:55:43

/*

```

```

* On hardware that supports double-buffering, this function sets up a flip
* and returns. The hardware will wait for vertical retrace, and then swap
* video buffers before the next video surface blit or lock will return.
* On hardware that doesn't support double-buffering, this is equivalent
* to calling SDL_UpdateRect(screen, 0, 0, 0, 0);
* The SDL_DOUBLEBUF flag must have been passed to SDL_SetVideoMode() when
* setting the video mode for this function to perform hardware flipping.
* This function returns 0 if successful, or -1 if there was an error.
*/

```

```
extern DECLSPEC int SDL_Flip(SDL_Surface *screen);
```

(SDL_SetGamma red green blue) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

```

/*
* Set the gamma correction for each of the color channels.
* The gamma values range (approximately) between 0.1 and 10.0
*
* If this function isn't supported directly by the hardware, it will
* be emulated using gamma ramps, if available. If successful, this
* function returns 0, otherwise it returns -1.
*/

```

```
extern DECLSPEC int SDL_SetGamma(float red, float green, float blue);
```

(SDL_SetGammaRamp red green blue) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

```

/*
* Set the gamma translation table for the red, green, and blue channels
* of the video hardware. Each table is an array of 256 16-bit quantities,
* representing a mapping between the input and output for that channel.
* The input is the index into the array, and the output is the 16-bit
* gamma value at that index, scaled to the output color precision.
*
* You may pass NULL for any of the channels to leave it unchanged.
* If the call succeeds, it will return 0. If the display driver or
* hardware does not support gamma translation, or otherwise fails,
* this function will return -1.
*/

```

```
extern DECLSPEC int SDL_SetGammaRamp(Uint16 *red, Uint16 *green, Uint16 *blue);
```

(SDL_GetGammaRamp red green blue) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

```

/*
 * Retrieve the current values of the gamma translation tables.
 *
 * You must pass in valid pointers to arrays of 256 16-bit quantities.
 * Any of the pointers may be NULL to ignore that channel.
 * If the call succeeds, it will return 0. If the display driver or
 * hardware does not support gamma translation, or otherwise fails,
 * this function will return -1.
 */
extern DECLSPEC int SDL_GetGammaRamp(Uint16 *red, Uint16 *green, Uint16 *blue);

```

(**SDL_SetColors** surface colors firstcolor ncolors) (packages/sdl/SDL-video.lsh)
 Last Modified: 2002-10-28 14:55:43

```

/*
 * Sets a portion of the colormap for the given 8-bit surface. If 'surface'
 * is not a palettized surface, this function does nothing, returning 0.
 * If all of the colors were set as passed to SDL_SetColors(), it will
 * return 1. If not all the color entries were set exactly as given,
 * it will return 0, and you should look at the surface palette to
 * determine the actual color palette.
 *
 * When 'surface' is the surface associated with the current display, the
 * display colormap will be updated with the requested colors. If
 * SDL_HWPALETTE was set in SDL_SetVideoMode() flags, SDL_SetColors()
 * will always return 1, and the palette is guaranteed to be set the way
 * you desire, even if the window colormap has to be warped or run under
 * emulation.
 */

```

```
extern DECLSPEC int SDL_SetColors(SDL_Surface *surface, SDL_Color *colors, int firstcolor, int ncolors);
```

(**SDL_SetPalette** surface flags colors firstcolor ncolors) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

```

/*
 * Sets a portion of the colormap for a given 8-bit surface.
 * 'flags' is one or both of:
 * SDL_LOGPAL -- set logical palette, which controls how blits are mapped
 *              to/from the surface,
 * SDL_PHYSPAL -- set physical palette, which controls how pixels look on
 *              the screen
 * Only screens have physical palettes. Separate change of physical/logical
 * palettes is only possible if the screen has SDL_HWPALETTE set.
 *
 * The return value is 1 if all colours could be set as requested, and 0

```

```

* otherwise.
*
* SDL_SetColors() is equivalent to calling this function with
*   flags = (SDL_LOGPAL|SDL_PHYSPAL).
*/
extern DECLSPEC int SDL_SetPalette(SDL_Surface *surface, int flags,
                                   SDL_Color *colors, int firstcolor,
                                   int ncolors);

```

(**SDL_MapRGB** format r g b) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

```

/*
* Maps an RGB triple to an opaque pixel value for a given pixel format
*/
extern DECLSPEC Uint32 SDL_MapRGB(SDL_PixelFormat *format, Uint8 r, Uint8 g, Uint8

```

(**SDL_MapRGBA** format r g b a) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

```

/*
* Maps an RGBA quadruple to a pixel value for a given pixel format
*/
extern DECLSPEC Uint32 SDL_MapRGBA(SDL_PixelFormat *format, Uint8 r, Uint8 g, Uint8

```

(**SDL_GetRGB** format r g b a) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

```

/*
* Maps a pixel value into the RGB components for a given pixel format
*/
extern DECLSPEC void SDL_GetRGB(Uint32 pixel, SDL_PixelFormat *fmt, Uint8 *r, Uint8

```

(**SDL_GetRGBA** format r g b a) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

```

/*
* Maps a pixel value into the RGBA components for a given pixel format
*/
extern DECLSPEC void SDL_GetRGBA(Uint32 pixel, SDL_PixelFormat *fmt, Uint8 *r, Uint8

```

(**SDL_CreateRGBSurface** flags width height depth rmask gmask bmask amask) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

```

/*
 * Allocate and free an RGB surface (must be called after SDL_SetVideoMode)
 * If the depth is 4 or 8 bits, an empty palette is allocated for the surface.
 * If the depth is greater than 8 bits, the pixel format is set using the
 * flags '[RGB]mask'.
 * If the function runs out of memory, it will return NULL.
 *
 * The 'flags' tell what kind of surface to create.
 * SDL_SWSURFACE means that the surface should be created in system memory.
 * SDL_HWSURFACE means that the surface should be created in video memory,
 * with the same format as the display surface. This is useful for surfaces
 * that will not change much, to take advantage of hardware acceleration
 * when being blitted to the display surface.
 * SDL_ASYNCBLIT means that SDL will try to perform asynchronous blits with
 * this surface, but you must always lock it before accessing the pixels.
 * SDL will wait for current blits to finish before returning from the lock.
 * SDL_SRCOLORKEY indicates that the surface will be used for colorkey blits.
 * If the hardware supports acceleration of colorkey blits between
 * two surfaces in video memory, SDL will try to place the surface in
 * video memory. If this isn't possible or if there is no hardware
 * acceleration available, the surface will be placed in system memory.
 * SDL_SRCALPHA means that the surface will be used for alpha blits and
 * if the hardware supports hardware acceleration of alpha blits between
 * two surfaces in video memory, to place the surface in video memory
 * if possible, otherwise it will be placed in system memory.
 * If the surface is created in video memory, blits will be much faster,
 * but the surface format must be identical to the video surface format,
 * and the only way to access the pixels member of the surface is to use
 * the SDL_LockSurface() and SDL_UnlockSurface() calls.
 * If the requested surface actually resides in video memory, SDL_HWSURFACE
 * will be set in the flags member of the returned surface. If for some
 * reason the surface could not be placed in video memory, it will not have
 * the SDL_HWSURFACE flag set, and will be created in system memory instead.
 */
#define SDL_AllocSurface    SDL_CreateRGBSurface
extern DECLSPEC SDL_Surface *SDL_CreateRGBSurface
    (Uint32 flags, int width, int height, int depth,
     Uint32 Rmask, Uint32 Gmask, Uint32 Bmask, Uint32 Amask);
extern DECLSPEC SDL_Surface *SDL_CreateRGBSurface(Uint32 flags, int width, int height, int dep

    (SDL_CreateRGBSurfaceFrom pixels width height depth pitch rmask
gmask bmask amask) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28
14:55:43

```

```
extern DECLSPEC SDL_Surface *SDL_CreateRGBSurfaceFrom
    (void *pixels, int width, int height, int depth, int pitch,
     Uint32 Rmask, Uint32 Gmask, Uint32 Bmask, Uint32 Amask);
```

(**SDL_FreeSurface** surface) (packages/sdl/SDL-video.lsh) Last Modified:
2002-10-28 14:55:43

```
extern DECLSPEC void SDL_FreeSurface(SDL_Surface *surface);
```

(**SDL_LockSurface** surface) (packages/sdl/SDL-video.lsh) Last Modified:
2002-10-28 14:55:43

```
/*
 * SDL_LockSurface() sets up a surface for directly accessing the pixels.
 * Between calls to SDL_LockSurface()/SDL_UnlockSurface(), you can write
 * to and read from 'surface->pixels', using the pixel format stored in
 * 'surface->format'. Once you are done accessing the surface, you should
 * use SDL_UnlockSurface() to release it.
 *
 * Not all surfaces require locking. If SDL_MUSTLOCK(surface) evaluates
 * to 0, then you can read and write to the surface at any time, and the
 * pixel format of the surface will not change. In particular, if the
 * SDL_HWSURFACE flag is not given when calling SDL_SetVideoMode(), you
 * will not need to lock the display surface before accessing it.
 *
 * No operating system or library calls should be made between lock/unlock
 * pairs, as critical system locks may be held during this time.
 *
 * SDL_LockSurface() returns 0, or -1 if the surface couldn't be locked.
 */
```

```
extern DECLSPEC int SDL_LockSurface(SDL_Surface *surface);
```

(**SDL_UnlockSurface** surface) (packages/sdl/SDL-video.lsh) Last Mod-
ified: 2002-10-28 14:55:43

```
extern DECLSPEC void SDL_UnlockSurface(SDL_Surface *surface);
```

(**SDL_LoadBMP_RW** src freesrc) (packages/sdl/SDL-video.lsh) Last
Modified: 2002-10-28 14:55:43

```
/*
 * Load a surface from a seekable SDL data source (memory or file.)
 * If 'freesrc' is non-zero, the source will be closed after being read.
```

```

* Returns the new surface, or NULL if there was an error.
* The new surface should be freed with SDL_FreeSurface().
*/
extern DECLSPEC SDL_Surface * SDL_LoadBMP_RW(SDL_RWops *src, int freesrc);

(SDL_SaveBMP_RW surface dst freedst) (packages/sdl/SDL-video.lsh)
Last Modified: 2002-10-28 14:55:43

/*
* Save a surface to a seekable SDL data source (memory or file.)
* If 'freedst' is non-zero, the source will be closed after being written.
* Returns 0 if successful or -1 if there was an error.
*/
extern DECLSPEC int SDL_SaveBMP_RW
(SDL_Surface *surface, SDL_RWops *dst, int freedst);

(SDL_SetColorKey surface flag key) (packages/sdl/SDL-video.lsh) Last
Modified: 2002-10-28 14:55:43

/*
* Sets the color key (transparent pixel) in a blittable surface.
* If 'flag' is SDL_SRCCOLORKEY (optionally OR'd with SDL_RLEACCEL),
* 'key' will be the transparent pixel in the source image of a blit.
* SDL_RLEACCEL requests RLE acceleration for the surface if present,
* and removes RLE acceleration if absent.
* If 'flag' is 0, this function clears any current color key.
* This function returns 0, or -1 if there was an error.
*/
extern DECLSPEC int SDL_SetColorKey
(SDL_Surface *surface, Uint32 flag, Uint32 key);
extern DECLSPEC int SDL_SetColorKey(SDL_Surface *surface, Uint32 flag, Uint32 key);

(SDL_SetAlpha surface flag alpha) (packages/sdl/SDL-video.lsh) Last
Modified: 2002-10-28 14:55:43

/*
* This function sets the alpha value for the entire surface, as opposed to
* using the alpha component of each pixel. This value measures the range
* of transparency of the surface, 0 being completely transparent to 255
* being completely opaque. An 'alpha' value of 255 causes blits to be
* opaque, the source pixels copied to the destination (the default). Note
* that per-surface alpha can be combined with colorkey transparency.
*
* If 'flag' is 0, alpha blending is disabled for the surface.

```

```

* If 'flag' is SDL_SRCALPHA, alpha blending is enabled for the surface.
* OR:ing the flag with SDL_RLEACCEL requests RLE acceleration for the
* surface; if SDL_RLEACCEL is not specified, the RLE accel will be removed.
*/
extern DECLSPEC int SDL_SetAlpha(SDL_Surface *surface, Uint32 flag, Uint8 alpha);

```

(**SDL_SetClipRect** surface rect) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

```

/*
* Sets the clipping rectangle for the destination surface in a blit.
*
* If the clip rectangle is NULL, clipping will be disabled.
* If the clip rectangle doesn't intersect the surface, the function will
* return SDL_FALSE and blits will be completely clipped. Otherwise the
* function returns SDL_TRUE and blits to the surface will be clipped to
* the intersection of the surface area and the clipping rectangle.
*
* Note that blits are automatically clipped to the edges of the source
* and destination surfaces.
*/
extern DECLSPEC SDL_bool SDL_SetClipRect(SDL_Surface *surface, SDL_Rect *rect);

```

(**SDL_GetClipRect** surface rect) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

```

/*
* Gets the clipping rectangle for the destination surface in a blit.
* 'rect' must be a pointer to a valid rectangle which will be filled
* with the correct values.
*/
extern DECLSPEC void SDL_GetClipRect(SDL_Surface *surface, SDL_Rect *rect);

```

(**SDL_ConvertSurface** src fmt flags) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

```

/*
* Creates a new surface of the specified format, and then copies and maps
* the given surface to it so the blit of the converted surface will be as
* fast as possible. If this function fails, it returns NULL.
*
* The 'flags' parameter is passed to SDL_CreateRGBSurface() and has those
* semantics. You can also pass SDL_RLEACCEL in the flags parameter and
* SDL will try to RLE accelerate colorkey and alpha blits in the resulting

```

```

* surface.
*
* This function is used internally by SDL_DisplayFormat().
*/
extern DECLSPEC SDL_Surface *SDL_ConvertSurface
(SDL_Surface *src, SDL_PixelFormat *fmt, Uint32 flags);

```

```

(SDL_BlitSurface src srcrect dst dstrect) (packages/sdl/SDL-video.lsh)
Last Modified: 2002-10-28 14:55:43

```

```

/*
* This performs a fast blit from the source surface to the destination
* surface. It assumes that the source and destination rectangles are
* the same size. If either 'srcrect' or 'dstrect' are NULL, the entire
* surface (src or dst) is copied. The final blit rectangles are saved
* in 'srcrect' and 'dstrect' after all clipping is performed.
* If the blit is successful, it returns 0, otherwise it returns -1.
*
* The blit function should not be called on a locked surface.
*
* The blit semantics for surfaces with and without alpha and colorkey
* are defined as follows:
*
* RGBA->RGB:
*   SDL_SRCALPHA set:
*     alpha-blend (using alpha-channel).
*     SDL_SRCCOLORKEY ignored.
*   SDL_SRCALPHA not set:
*     copy RGB.
*     if SDL_SRCCOLORKEY set, only copy the pixels matching the
*     RGB values of the source colour key, ignoring alpha in the
*     comparison.
*
* RGB->RGBA:
*   SDL_SRCALPHA set:
*     alpha-blend (using the source per-surface alpha value);
*     set destination alpha to opaque.
*   SDL_SRCALPHA not set:
*     copy RGB, set destination alpha to opaque.
*   both:
*     if SDL_SRCCOLORKEY set, only copy the pixels matching the
*     source colour key.
*
* RGBA->RGBA:
*   SDL_SRCALPHA set:

```

```

*      alpha-blend (using the source alpha channel) the RGB values;
*      leave destination alpha untouched. [Note: is this correct?]
*      SDL_SRCCOLORKEY ignored.
*      SDL_SRCALPHA not set:
*      copy all of RGBA to the destination.
*      if SDL_SRCCOLORKEY set, only copy the pixels matching the
*      RGB values of the source colour key, ignoring alpha in the
*      comparison.
*
* RGB->RGB:
*      SDL_SRCALPHA set:
*      alpha-blend (using the source per-surface alpha value).
*      SDL_SRCALPHA not set:
*      copy RGB.
*      both:
*      if SDL_SRCCOLORKEY set, only copy the pixels matching the
*      source colour key.
*
* If either of the surfaces were in video memory, and the blit returns -2,
* the video memory was lost, so it should be reloaded with artwork and
* re-blitted:
      while ( SDL_BlitSurface(image, imgrect, screen, dstrect) == -2 ) {
          while ( SDL_LockSurface(image) < 0 )
              Sleep(10);
          -- Write image pixels to image->pixels --
          SDL_UnlockSurface(image);
      }
* This happens under DirectX 5.0 when the system switches away from your
* fullscreen application. The lock will also fail until you have access
* to the video memory again.
*/
/* You should call SDL_BlitSurface() unless you know exactly how SDL
   blitting works internally and how to use the other blit functions.
*/
#define SDL_BlitSurface SDL_UpperBlit

      (SDL_UpperBlit src srcrect dst dstrect) (packages/sdl/SDL-video.lsh)
Last Modified: 2002-10-28 14:55:43

/* This is the public blit function, SDL_BlitSurface(), and it performs
   rectangle validation and clipping before passing it to SDL_LowerBlit()
*/
extern DECLSPEC int SDL_UpperBlit
      (SDL_Surface *src, SDL_Rect *srcrect,
       SDL_Surface *dst, SDL_Rect *dstrect);

```

(SDL_LowerBlit src srcrect dst dstrect) (packages/sdl/SDL_video.lsh)
Last Modified: 2002-10-28 14:55:43

```
/* This is a semi-private blit function and it performs low-level surface
   blitting only.
```

```
*/
extern DECLSPEC int SDL_LowerBlit
    (SDL_Surface *src, SDL_Rect *srcrect,
     SDL_Surface *dst, SDL_Rect *dstrect);
```

(SDL_FillRect dst dstrect color) (packages/sdl/SDL_video.lsh) Last
Modified: 2002-10-28 14:55:43

```
/*
 * This function performs a fast fill of the given rectangle with 'color'
 * The given rectangle is clipped to the destination surface clip area
 * and the final fill rectangle is saved in the passed in pointer.
 * If 'dstrect' is NULL, the whole surface will be filled with 'color'
 * The color should be a pixel of the format used by the surface, and
 * can be generated by the SDL_MapRGB() function.
 * This function returns 0 on success, or -1 on error.
```

```
*/
extern DECLSPEC int SDL_FillRect
    (SDL_Surface *dst, SDL_Rect *dstrect, Uint32 color);
```

(SDL_DisplayFormat surface) (packages/sdl/SDL_video.lsh) Last Mod-
ified: 2002-10-28 14:55:43

```
/*
 * This function takes a surface and copies it to a new surface of the
 * pixel format and colors of the video framebuffer, suitable for fast
 * blitting onto the display surface. It calls SDL_ConvertSurface()
 *
 * If you want to take advantage of hardware colorkey or alpha blit
 * acceleration, you should set the colorkey and alpha value before
 * calling this function.
```

```
*/
extern DECLSPEC SDL_Surface * SDL_DisplayFormat(SDL_Surface *surface);
```

(SDL_DisplayFormatAlpha surface) (packages/sdl/SDL_video.lsh) Last
Modified: 2002-10-28 14:55:43

```

/*
 * This function takes a surface and copies it to a new surface of the
 * pixel format and colors of the video framebuffer (if possible),
 * suitable for fast alpha blitting onto the display surface.
 * The new surface will always have an alpha channel.
 *
 * If you want to take advantage of hardware colorkey or alpha blit
 * acceleration, you should set the colorkey and alpha value before
 * calling this function.
 *
 * If the conversion fails or runs out of memory, it returns NULL
 */
extern DECLSPEC SDL_Surface * SDL_DisplayFormatAlpha(SDL_Surface *surface);

(SDL_CreateYUVOverlay width height format display) (packages/sdl/SDL-video.lsh)
Last Modified: 2002-10-28 14:55:43

/* This function creates a video output overlay
   Calling the returned surface an overlay is something of a misnomer because
   the contents of the display surface underneath the area where the overlay
   is shown is undefined - it may be overwritten with the converted YUV data.
 */
extern DECLSPEC SDL_Overlay *SDL_CreateYUVOverlay(int width, int height,
        Uint32 format, SDL_Surface *display);

(SDL_LockYUVOverlay overlay) (packages/sdl/SDL-video.lsh) Last Mod-
ified: 2002-10-28 14:55:43

/* Lock an overlay for direct access, and unlock it when you are done */
extern DECLSPEC int SDL_LockYUVOverlay(SDL_Overlay *overlay);

(SDL_UnlockYUVOverlay overlay) (packages/sdl/SDL-video.lsh) Last
Modified: 2002-10-28 14:55:43

extern DECLSPEC void SDL_UnlockYUVOverlay(SDL_Overlay *overlay);

(SDL_DisplayYUVOverlay overlay dstrect) (packages/sdl/SDL-video.lsh)
Last Modified: 2002-10-28 14:55:43

/* Blit a video overlay to the display surface.
   The contents of the video surface underneath the blit destination are
   not defined.

```

The width and height of the destination rectangle may be different from that of the overlay, but currently only 2x scaling is supported.

```
*/
extern DECLSPEC int SDL_DisplayYUVOverlay(SDL_Overlay *overlay, SDL_Rect *dstrect);

(SDL_FreeYUVOverlay overlay) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

/* Free a video overlay */
extern DECLSPEC void SDL_FreeYUVOverlay(SDL_Overlay *overlay);

(SDL_GL_LoadLibrary path) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

/*
 * Dynamically load a GL driver, if SDL is built with dynamic GL.
 *
 * SDL links normally with the OpenGL library on your system by default,
 * but you can compile it to dynamically load the GL driver at runtime.
 * If you do this, you need to retrieve all of the GL functions used in
 * your program from the dynamic library using SDL_GL_GetProcAddress().
 *
 * This is disabled in default builds of SDL.
 */
extern DECLSPEC int SDL_GL_LoadLibrary(const char *path);

(SDL_GL_GetProcAddress proc) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

/*
 * Get the address of a GL function (for extension functions)
 */
extern DECLSPEC void *SDL_GL_GetProcAddress(const char* proc);

(SDL_GL_SetAttribute attr value) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43

/*
 * Set an attribute of the OpenGL subsystem before initialization.
 */
extern DECLSPEC int SDL_GL_SetAttribute(SDL_GLAttr attr, int value);

(SDL_GL_GetAttribute attr value) (packages/sdl/SDL-video.lsh) Last Modified: 2002-10-28 14:55:43
```

```

/*
 * Get an attribute of the OpenGL subsystem from the windowing
 * interface, such as glX. This is of course different from getting
 * the values from SDL's internal OpenGL subsystem, which only
 * stores the values you request before initialization.
 *
 * Developers should track the values they pass into SDL_GL_SetAttribute
 * themselves if they want to retrieve these values.
 */
extern DECLSPEC int SDL_GL_GetAttribute(SDL_GLattr attr, int* value);

(SDL_GL_SwapBuffers) (packages/sdl/SDL-video.lsh) Last Modified: 2002-
10-28 14:55:43

/*
 * Swap the OpenGL buffers, if double-buffering is supported.
 */
extern DECLSPEC void SDL_GL_SwapBuffers(void);

(SDL_WM_SetCaption title icon) (packages/sdl/SDL-video.lsh) Last
Modified: 2002-10-28 14:55:43

/*
 * Sets/Gets the title and icon text of the display window
 */
extern DECLSPEC void SDL_WM_SetCaption(const char *title, const char *icon);

(SDL_WM_GetCaption title icon) (packages/sdl/SDL-video.lsh) Last
Modified: 2002-10-28 14:55:43

extern DECLSPEC void SDL_WM_GetCaption(char **title, char **icon);

(SDL_WM_SetIcon icon mask) (packages/sdl/SDL-video.lsh) Last Mod-
ified: 2002-10-28 14:55:43

/*
 * Sets the icon for the display window.
 * This function must be called before the first call to SDL_SetVideoMode().
 * It takes an icon surface, and a mask in MSB format.
 * If 'mask' is NULL, the entire icon surface will be used as the icon.
 */
extern DECLSPEC void SDL_WM_SetIcon(SDL_Surface *icon, Uint8 *mask);

```

(**SDL_WM_IconifyWindow**) (packages/sdl/SDL_video.lsh) Last Modified: 2002-10-28 14:55:43

```
/*
 * This function iconifies the window, and returns 1 if it succeeded.
 * If the function succeeds, it generates an SDL_APPACTIVE loss event.
 * This function is a noop and returns 0 in non-windowed environments.
 */
extern DECLSPEC int SDL_WM_IconifyWindow(void);
```

(**SDL_WM_ToggleFullScreen** surface) (packages/sdl/SDL_video.lsh) Last Modified: 2002-10-28 14:55:43

```
/*
 * Toggle fullscreen mode without changing the contents of the screen.
 * If the display surface does not require locking before accessing
 * the pixel information, then the memory pointers will not change.
 *
 * If this function was able to toggle fullscreen mode (change from
 * running in a window to fullscreen, or vice-versa), it will return 1.
 * If it is not implemented, or fails, it returns 0.
 *
 * The next call to SDL_SetVideoMode() will set the mode fullscreen
 * attribute based on the flags parameter - if SDL_FULLSCREEN is not
 * set, then the display will be windowed by default where supported.
 *
 * This is currently only implemented in the X11 video driver.
 */
extern DECLSPEC int SDL_WM_ToggleFullScreen(SDL_Surface *surface);
```

(**SDL_WM_GrabInput** mode) (packages/sdl/SDL_video.lsh) Last Modified: 2002-10-28 14:55:43

```
/*
 * This function allows you to set and query the input grab state of
 * the application. It returns the new input grab state.
 */
typedef enum {
    SDL_GRAB_QUERY = -1,
    SDL_GRAB_OFF = 0,
    SDL_GRAB_ON = 1,
    SDL_GRAB_FULLSCREEN /* Used internally */
} SDL_GrabMode;
```

9.4 OpenGL

Lush has a full interface to OpenGL , OpenGLU , and GLUT. Please refer to the Lush on-line documentation tool (`helptool`) for more details.

9.5 OpenRM

Lush has a full interface to the OpenRM scene graph rendering engine. Please refer to the Lush on-line documentation tool (`helptool`) for more details.

9.6 Video4Linux: video grabbing

Author(s): Yann LeCun

The `video4linux` package allows video frame grabbing under Linux using the `video4linux` API of the Linux kernel. This interface is compatible with most video grabbing hardware supported by `video4linux`. Notable examples include (among many others) TV/Tuner cards based on the BrookTree BT8x8 chips (such as the WinTV cards from Hauppauge) , and v4l-supported webcams such the Logitech Quickam Pro 3000 and other USB webcams based on the Philips chipset (see list at [<http://www.smcc.demon.nl/webcam>]). NOTE: these Philips-based webcams requires a non-opensource driver called "pwcx" to handle resolutions higher than 160x120. This driver can downloaded from [<http://www.smcc.demon.nl/webcam/release.html>]).

9.6.1 Requirements and Installation

`video4linux` and appropriate device drivers are usually provided with most recent Linux distributions (as of september 2002 , RedHat 7.3 and Mandrake 8.2). No special installation is required.

9.6.2 Video4Linux API

The Video4Linux interface to Lush is implemented through the class `v4ldevice` . Creating a new `v4ldevice` will open and initialize one of the Linux video devices. A number of methods are provided to set the parameters (channel , image size , video mode , etc...) and to grab video frames.

Video Mode Constants

(`packages/video4linux/v4l.lsh`) Last Modified: 2003-01-06 23:43:51

<code>VIDEO_MODE_PAL</code>	0
<code>VIDEO_MODE_NTSC</code>	1
<code>VIDEO_MODE_SECAM</code>	2
<code>VIDEO_MODE_AUTO</code>	3

Video Palette Constants

(packages/video4linux/v4l.lsh) Last Modified: 2003-01-06 23:43:51

```

VIDEO_PALETTE_GREY          1
VIDEO_PALETTE_HI240         2
VIDEO_PALETTE_RGB565        3
VIDEO_PALETTE_RGB24         4
VIDEO_PALETTE_RGB32         5
VIDEO_PALETTE_RGB555        6
VIDEO_PALETTE_YUV422        7
VIDEO_PALETTE_YUVV          8
VIDEO_PALETTE_UYVY          9
VIDEO_PALETTE_YUV420        10
VIDEO_PALETTE_YUV411        11
VIDEO_PALETTE_RAW           12
VIDEO_PALETTE_YUV422P       13
VIDEO_PALETTE_YUV411P       14
VIDEO_PALETTE_YUV420P       15
VIDEO_PALETTE_YUV410P       16
VIDEO_PALETTE_PLANAR        13
VIDEO_PALETTE_COMPONENT     7

```

v4ldevice

(packages/video4linux/v4l.lsh) Last Modified: 2003-01-06 23:43:51

the v4l device is the main class through which video frames can be grabbed using the Video4Linux device. This requires a v4l compatible video grabbing card or webcam.

Here is an example of code that repeatedly grabs an image from the composite video input and displays it on the screen:

```

(de v4l-demo (swidth sheight)
  (let* ((bg (ubyte-matrix sheight swidth 3)) ; make image buffer
        ;; open video device NTSC standard, composite input (channel 1).
        (vdev (new v4ldevice "/dev/video0" "NTSC" 1 swidth sheight)))
    ;; open window
    (setq screen (new-window))
    (while t
      ;; grab frame
      (==> vdev grab-into-rgbx bg)
      ;; display frame
      (rgb-draw-matrix 10 10 bg))))

```

(==> v4ldevice **cleanup**) Type: MSG(packages/video4linux/v4l.lsh) Last Modified: 2003-01-06 23:43:51

deallocate all malloced slots and close device. This is a private method called by the destructor and various error handlers, but rarely called directly.

(==> **v4ldevice set-format f w h**) Type: MSG(packages/video4linux/v4l.lsh)
Last Modified: 2003-01-06 23:43:51

set the video format/palette (@@VIDEO_PALETTE_RGB24, @@VIDEO_PALETTE_RGB32, @@VIDEO_PALETTE_YUV422 etc...) , and the width and the height of the image to be grabbed (see [”Video Palette Constants”]for a list of formats). Returns 0 on success and -1 if the format requested is not supported by the device. This method rarely called directly, unless the low-level **grab** method is to be used. The various **grab-into-rgb** and **grab-into-rgba** methods determine the appropriate format automatically, and all the **grab-xxx-into-xxx** methods set the desired format before grabbing, therefore a prior call to **set-format** is unnecessary if these high-level methods are used.

(**new v4ldevice devname n c w h**) Type: CLASS(packages/video4linux/v4l.lsh)
Last Modified: 2003-01-06 23:43:51

opens a new v4ldevice for video grabbing. Arguments are: device name (e.g. ”/dev/video0”), norm, channel, width, and height. norm can be ”NTSC”, ”PAL”, ”SECAM”, ”AUTO”. The ”channel” c determines which video source the card will listen to. On most bttv cards, 0 means TV tuner, 1 means composite video input, and 2 means S-video input (if present). Tuning to a particular broadcast or cable channel is done with the ”tune” method.

example: open device to grab 320x240 NTSC video from S-video on BTTV/avermedia card (channel 2):

```
(setq a (new v4ldevice "/dev/video0" "NTSC" 2 320 240))
```

(==> **v4ldevice tune freq**) Type: MSG(packages/video4linux/v4l.lsh)
Last Modified: 2003-01-06 23:43:51

tune the tuner to frequency n, expressed in 1/16th of a MHz. Tables that map channels to frequencies are available in file **freq-table.lsh**. Tuning to a channel can only be done if the device was opened to listen to the TV tuner (as opposed to the composite video or S-video).

(==> **v4ldevice grab-into-rgb img**) Type: MSG(packages/video4linux/v4l.lsh)
Last Modified: 2003-01-06 23:43:51

grab an image from the v4l device into a preexisting image (a ubyte-matrix of size heightxwidthxdepth, where depth must be at least 3). Consecutive elements in the last dimension are interpreted as RGB values. Only those elements are modified. If the video image format has not been previously set, this method automatically determines which format is supported by the current device (among RGB24, RGB32 and YUV420P). This allows to support videocards and webcams transparently without having to worry about which formats are supported. Some webcams (like the Logitech Quickcam 3000 pro) only support YUV420 and YUV420P.

(==> **v4ldevice grab-into-rgbx img**) Type: MSG(packages/video4linux/v4l.lsh)
Last Modified: 2003-01-06 23:43:51

same as **grab-into-rgb**.

(==> **v4ldevice grab-into-rgba** *img alpha*) Type: MSG(packages/video4linux/v4l.lsh)
 Last Modified: 2003-01-06 23:43:51

grab an image from the v4l device into a preexisting image (a ubyte-matrix of size *heightxwidthxdepth* , where *depth* must be at least 3). Consecutive elements in the last dimension are interpreted as RGBA values. Only those 4 elements are modified. The A components are filled with **alpha** . If the video image format has not been previously set , this method automatically determines which format is supported by the current device (among RGB24 , RGB32 and YUV420P). This allows to support videocards and webcams transparently without having to worry about which formats are supported. Some webcams (like the Logitech Quickcam 3000 pro) only support YUV420 and YUV420P.

(==> **v4ldevice grab-rgb24-into-rgbx** *img*) Type: MSG(packages/video4linux/v4l.lsh)
 Last Modified: 2003-01-06 23:43:51

grab an image from the v4l device in RGB24 format into a preexisting image (a ubyte-matrix of size *heightxwidthxdepth* , where *depth* must be at least 3). Consecutive elements in the last dimension are interpreted as RGB values. Only those elements are modified.

(==> **v4ldevice grab-rgb32-into-rgbx** *img*) Type: MSG(packages/video4linux/v4l.lsh)
 Last Modified: 2003-01-06 23:43:51

grab an image from the v4l device in RGB32 format into a preexisting image (a ubyte-matrix of size *heightxwidthxdepth* , where *depth* must be at least 3). Consecutive elements in the last dimension are interpreted as RGB values. Only those elements are modified.

(==> **v4ldevice grab-rgb24-into-rgba** *img alpha*) Type: MSG(packages/video4linux/v4l.lsh)
 Last Modified: 2003-01-06 23:43:51

grab an image from the v4l device in RGB32 format into a preexisting image (a ubyte-matrix of size *heightxwidthxdepth* with *depth* equal to at least 4) , Consecutive elements in the last dimension are interpreted as RGBA values. The A values are filled with **alpha** .

(==> **v4ldevice grab-rgb32-into-rgba** *img alpha*) Type: MSG(packages/video4linux/v4l.lsh)
 Last Modified: 2003-01-06 23:43:51

grab an image from the v4l device in RGB32 format into a preexisting image (a ubyte-matrix of size *heightxwidthxdepth* with *depth* equal to at least 4) , Consecutive elements in the last dimension are interpreted as RGBA values. The A values are filled with **alpha** .

(==> **v4ldevice grab-yuv422-into-yuv422** *img*) Type: MSG(packages/video4linux/v4l.lsh)
 Last Modified: 2003-01-06 23:43:51

grab an image from the v4l device in YUV422 format into a preexisting image (a ubyte-matrix of size *heightxwidthx2*. Consecutive elements are interpreted as YUYVYUYV..... according to the YUV422 definition. so (*img y x 0*) is the luminance component of pixel (*x* , *y*).

(==> **v4ldevice grab-yuv420p-into-rgbx** *img*) Type: MSG(packages/video4linux/v4l.lsh)
 Last Modified: 2003-01-06 23:43:51

grab an image from the v4l device in YUV420P (planar) format into a preexisting image (a ubyte-matrix of size *heightxwidthxdepth* with *depth* at least equal to 3). Consecutive elements are interpreted as RGB components. This

mode is useful for certain webcam that do not support RGB grabbing (e.g. Logitech Quickam 3000 Pro).

(==> **v4ldevice grab-yuv420p-into-rgba** *img alpha*) Type: MSG(packages/video4linux/v4l) Last Modified: 2003-01-06 23:43:51

grab an image from the v4l device in YUV420P (planar) format into a pre-existing image (a ubyte-matrix of size *heightxwidthxdepth* with depth at least equal to 4). Consecutive elements are interpreted as RGBA components. the A components are filled with *alpha* . This mode is useful for certain webcam that do not support RGB grabbing (e.g. Logitech Quickam 3000 Pro).

(==> **v4ldevice grab-grey-into-y** *img*) Type: MSG(packages/video4linux/v4l.lsh) Last Modified: 2003-01-06 23:43:51

grab an image from the v4l device in GREY (monochrome) format into a preexisting image (a ubyte-matrix of size *heightxwidth*).

(==> **v4ldevice record-into-rgb** *movie s*) Type: MSG(packages/video4linux/v4l.lsh) Last Modified: 2003-01-06 23:43:51

record successive frames from the v4l device into a preexisting movie (a 4-dimensional ubyte-matrix of size *nframesxheightxwidthx3*). Consecutive elements in the last dimension are interpreted as RGBA values. Consecutive slices in the first dimension are individual frames. The *s* parameter controls the temporal subsampling: every *s* video frame is saved in *movie* .

(==> **v4ldevice monitor** *x y z n w h*) Type: MSG(packages/video4linux/v4l.lsh) Last Modified: 2003-01-06 23:43:51

grab video frames and display them in the current Lush window at position *x* , *y* , with zoom factor *z* . The number of frames grabbed is *n* , and their width/height is *w* / *h* . return the last grabbed frame.

(==> **v4ldevice grab-rgb24-acc-rgbx** *img*) Type: MSG(packages/video4linux/v4l.lsh) Last Modified: 2003-01-06 23:43:51

grab an image from the v4l device in RGB24 format and accumulate the result into a preexisting image (an int-matrix of size *height* x *width* x *depth* where *height* and *width* are the dimensions of the *v4ldevice*) and *depth* must be at least 3. Consecutive elements in the last dimension are interpreted as RGB values.

(==> **v4ldevice grab-rgb24-acc-rgba** *img alpha*) Type: MSG(packages/video4linux/v4l.lsh) Last Modified: 2003-01-06 23:43:51

grab an image from the v4l device and accumulate the result into a preexisting image (an int-matrix of size *height* x *width* x *depth* where *height* and *width* are the dimensions of the *v4ldevice* , and with depth equal to at least 4). Consecutive elements in the last dimension are interpreted as RGBA values. The alpha channel is filled with value *alpha* .

(==> **v4ldevice grab**) Type: MSG(packages/video4linux/v4l.lsh) Last Modified: 2003-01-06 23:43:51

low-level grab of an image from the v4l device. This is a low-level function that simply performs a frame grab in the video format and image size (as set with *set-format*) with no format conversion. Most users will prefer to use the other grab methods which perform appropriate conversions and write their result into an *idx*. The present function simply provides a pointer to the frame buffer. The

grabbed data can be accessed through the `framebuffer` slot of the `v4ldevice` (a `gptr`). The number of bytes in the grab frame is given by the product of the `width`, `height`, and `depth` slots. The format depends on what parameter was passed to `set-format` prior to calling `grab`. Please note that the RGB24 and RGB32 formats of video4linux are really BGR formats from Lush's point of view.

```
(libload "video4linux/v4l")
(libload "libc/libc")
(new-window)
;; these parameters are for Quickcam 3000 pro webcam
(setq v (new v4ldevice "/dev/video" "AUTO" 0 320 240))
(setq m (ubyte-matrix 240 320))
;; set format to YUV420P (planar).
(==> v set-format @@VIDEO_PALETTE_YUV420P 320 240)
(==> v grab)
;; copy the Y component into the matrix m
(memcpy (idx-ptr m) :v:framebuffer (* 240 320))
(rgb-draw-matrix 0 0 m)
```

TV Frequency tables

(packages/video4linux/freq-tables.lsh) Last Modified: 2002-08-19 20:44:14

these tables contain the frequencies (in 1/16th of a MHz) that can be passed to the "tune" method of a v4l-device to set the channel of the tuner.

freqtable-us-bcast Type: variable(packages/video4linux/freq-tables.lsh) Last Modified: 2002-08-19 20:44:14

frequency table vector for US broadcast TV. The *i*-th element is the frequency of the channel *i*. When the frequency is 0, the channel does not exist.

freqtable-us-cable Type: variable(packages/video4linux/freq-tables.lsh) Last Modified: 2002-08-19 20:44:14

frequency table vector for US cable TV. The *i*-th element is the frequency of the channel *i*. When the frequency is 0, the channel does not exist.

9.6.3 Demos

Simple Video Capture Demos

(packages/video4linux/demos/v4l-demo.lsh) Last Modified: 2002-11-19 23:32:14

(**v4l-demo-rgb in width height**) (packages/video4linux/demos/v4l-demo.lsh) Last Modified: 2002-11-19 23:32:14

monitors NTSC video captured from `/dev/video on` on a Lush graphic screen. `in` is the card's input from which the video will be captured (=1 for composite input, =2 for S-video input on some bttv cards).

```
(v4l-demo-rgb 1 320 240)
```

bug: the function must be interrupted with CTRL-C.

feature: the function is 6 lines of lush.

(v4l-demo2 in1 in2 width height) (packages/video4linux/demos/v4l-demo.lsh) Last Modified: 2002-11-19 23:32:14

play captured videos on a Lush graphic screen from two video source simultaneously. This assumes that two video capture devices are present in the computer. `in1` and `in2` are the card's inputs from which the videos will be captured (=2 for S-video input on some bttv cards).

```
(v4l-demo2 2 2 320 240)
```

bug: the function must be interrupted with CTRL-C.

(effect-sdl-demo eff)

(packages/video4linux/demos/effect-sdl-demo.lsh) Last Modified: 2002-10-10 20:33:03

This starts a real-time video effect demo that uses video4linux for grabbing video and SDL for displaying. The `eff` parameter must be an instance of a subclass of `effect`. Several subclasses are provided including edge detection, laplacian, image difference, etc...

This function uses three global variables to configure the video grabbing. Their names are: `*video-device*`, `*video-standard*` and `*video-input*`, and their default values are respectively `"/dev/video"`, `"NTSC"`, `1` (1 is the composite input on most TV cards). The default values can be changed by `defvar` ing those variables to the desired value before loading the present file. Example:

```
(effect-sdl-demo (new effect-trail 320 240 20))
(effect-sdl-demo (new effect-edge 320 240))
```

Video Effects (packages/video4linux/demos/effect-sdl-demo.lsh) Last Modified: 2002-10-10 20:33:03

Video effect classes are subclasses of the `effect` class. Video effects are used by calling the `fprop` method which takes one input (RGBA image) and one output (RGBA image). The effect is applied to the input image and written into the output image. The input and output must be `height x width x4 idx3` of ubytes. If used with `effect-sdl-demo` The alpha channel of the output image will determine the blending of the output image with the previously displayed image. Set all the alphas to 255 for full opacity. Effect subclasses must have a constructor that takes at least a width and height parameter (and possibly other initialization parameters). They must have width and height methods that return the width and height.

(new effect width height) Type: CLASS(packages/video4linux/demos/effect-sdl-demo.lsh) Last Modified: 2002-10-10 20:33:03

noop effect: simply copies the input to the output unaffected

(**new effect-edge width height**) Type: CLASS(packages/video4linux/demos/effect-sdl-demo.lsh) Last Modified: 2002-10-10 20:33:03

performs a simple edge detection by applying an on-center/off-surround convolution filter and rectifying , amplifying and saturating the result.

(**new effect-trail width height hl**) Type: CLASS(packages/video4linux/demos/effect-sdl-demo.lsh) Last Modified: 2002-10-10 20:33:03

simply change the transparency to create a "trail" effect. **hl** is the "half life" of the trail: 0 leaves no trail, 1 leaves a trail that decays by half at each frame, and 20 looks funny. Example:

```
(new effect-diff 320 240 20)
```

(**new effect-laplace width height gain**) Type: CLASS(packages/video4linux/demos/effect-sdl-demo.lsh) Last Modified: 2002-10-10 20:33:03

this applies a sort-of laplacian filter (really a set-of-your-pant on-center/off-surround filter) to each RGB component of the image. The **gain** paramter determines the gain of the output. Values around 1 are reasonable.

(**new effect-diff width height alpha**) Type: CLASS(packages/video4linux/demos/effect-sdl-demo.lsh) Last Modified: 2002-10-10 20:33:03

computes a simple difference between consecutive images. **alpha** is the transparency with which the difference image is painted. Setting it to 255 will make it opaque , setting it to 128 will leave a short trail.

(**new effect-bgsubst width height imagefile erosion dilation thres**) Type: CLASS(packages/video4linux/demos/effect-sdl-demo.lsh) Last Modified: 2002-10-10 20:33:03

performs background substitution. This compares the current grabbed image with an image computed as the average of the first 10 frames. If the Euclidean distance in RGB space of a pixel to the average is larger than **thres** , the corresponding pixel of the currently grabbed image is drawn , otherwise the corresponding pixel in the image **imagefile** is drawn. **thres** should be between 500 and 5000 for normal lighting conditions. **erosion** and **dilation** are used to erode and dilate the mask so as to eliminate rogue foreground pixels. Suggested values are 3 and 2 respectively.

realtime video effects

(packages/video4linux/demos/effect-demo.lsh) Last Modified: 2002-10-10 20:33:32

a simple GUI application that shows several real-time video effects , including: color manipulation , image difference , background substitution , edge detection , and image warping.

9.7 BLAS: Basic Linear Algebra Subroutine

Lush has a full interface to the Basic Linear Algebra Subroutines library (BLAS). Please refer to the Lush on-line documentation tool tool (**helptool**) for more details.

9.8 LAPACK

Lush has a full interface to the Linear Algebra Package library (LAPACK). Please refer to the Lush on-line documentation tool tool (`helptool`) for more details.

9.9 GSL: GNU Scientific Library

Lush has a full interface to the GNU Scientific Library (GSL). Please refer to the Lush on-line documentation tool tool (`helptool`) for more details.

9.10 OpenCV: Intel Computer Vision Library

Lush has a full interface to the Intel Open Computer Vision Library (OpenCV). Please refer to the Lush on-line documentation tool tool (`helptool`) for more details.

9.11 ALSA: Advanced Linux Sound Architecture

Lush has a full interface to the Advanced Linux Sound Architecture (ALSA). Please refer to the Lush on-line documentation tool tool (`helptool`) for more details.

Chapter 10

Applications

Applications are Lush programs (generally with a graphical user interface) that are useful by themselves without without requiring any programming.

10.1 Graphic Tools

Chapter 11

Tutorials

11.1 Tutorial: Writing Games with Lush and SDL.

Author(s): Yann LeCun

Lush provides a simple way to write 2D real-time games using the SDL library (Simple Directmedia Layer). What follows is a gentle tutorial on how to do that.

This tutorial was written to be understandable by motivated high school students with some basic familiarity with Lush but with little programming experience. Seasoned programmers will probably want to skip certain paragraphs.

The code for this tutorial is available at [.](#)

11.1.1 A Quick Reminder on Basic Lush Programming

In this tutorial, we make use of some simple object oriented concepts. Creating a new instance of a class is performed with a call to `new` :

```
(setq an-instance (new a-class arg1 arg2 ...))
```

For example, creating a new sprite (a movable object on a graphic screen) is done with:

```
(setq a-new-sprite (new sdl-sprite scr 0))
```

Objects have slots and methods. Slots are variables that are attached to the object, while methods are functions that are attached to the class of the object. Calling a method on an object is like telling it to carry out a particular action. This is done with the function `==>`. Here is an example:

```
(==> a-new-sprite move 300 400)
(==> a-new-sprite draw)
```

This can be interpreted as "tell `a-new-sprite` to move to position 300 , 400" , then "tell `a-new-sprite` to draw itself on the screen".

11.1.2 A Simple Lunar Lander

As a case study , we will describe how to write the core of a simple Lunar Lander game. The amazing thing is that the whole game will fit in less than a page of Lush code.

The code described in this section is available in `/home/yann/lush/packages/sdl/demos/tutorial`.

First , make a working directory , say `game` and `cd` to it.

Creating the Ship and Background Images

Before we start writing code , we need to create images for the ship and the background. This can be done with The Gimp. Start The Gimp , and open a new drawing with a transparent background. Paint your object (or modify an existing digital picture) , and save the image as a PNG file. The advantage of using the PNG format is that it is relatively compact , uses lossless compression , and supports an alpha channel for transparency.

For the time being , instead of creating your own art , you can simply copy the files `lem.png` and `moon.png` from the directory `/home/yann/lush/packages/sdl/demos` into your working directory.

First Implementation of Lander

Getting started

Start your favorite text editor (say `emacs`) , and open a new program file , say `lander.lsh`.

At the top of the file , we put the following line:

```
(libload "sdl/libsdl")
```

which will load the SDL library.

Next we need to define a function that will run our game:

```
(de lander1 ()
```

initializing SDL

The first line of our function MUST be this:

```
(sdl-initialize)
```

This will initialize the SDL subsystem. It's OK to call this function multiple times (it only initialize SDL once).

opening the SDL screen

Opening an SDL screen is done as follows:

```
(setq scr (new sdl-screen 640 480 "Lander"))
```

now the variable `scr` contains the screen object.

creating the background and ship sprites

First , we create a new sprite on the screen `scr` and give it the ID 0. Sprites are composed of one of several frames that can be loaded with images. We can load the image file `moon.png` into frame 0 , and set the hot-point at coordinates 0 0 , (which is the upper-left corner). The hot point is the "handle" by which the sprite is "held". When we move the sprite to coordinate (x , y) , we move its handle to (x , y). Here , we will move this sprite to (0 , 0).

```
;; create background sprite
(setq bgd (new sdl-sprite scr 0))
;; load moon image into frame 0 of bgd sprite
(=> bgd load-frame "moon.png" 0 0 0)
;; move sprite to 0 0
(=> bgd move 0 0)
```

The move method does not actually draw anything on the screen , it merely sets the internal coordinate variables of the sprite to (0 , 0).

Now let's create the ship sprite and create two variables to hold its position:

```
;; create ship sprite
(setq ship (new sdl-sprite scr 1))
;; load lem image into frame 0 of ship sprite
(=> ship load-frame "lem.png" 0 40 35)
;; set position of ship
(setq x 10) (setq y 20)
```

The handle of the ship sprite is a coordinate 40 , 35 (i.e. 40 pixel to the right , and 35 pixels down from the upper left corner of the ship image) , which is roughly at the center of the ship.

creating the event handler

We are going to need to grab inputs from the keyboard. To do so , we create an `sdl-event` object , and create an integer vector with 3 elements to hold the result of grabbing the events:

```
(setq event (new sdl-event))
(setq xyk (int-matrix 3))
```

the main loop

The main loop of our game will look something like this:

- clear the screen
- draw the background sprite
- read the keyboard
- move the ship sprite according to the keys pressed
- draw the ship sprite at its new location
- flip the screens (see below)
- repeat the above step until the user quits

The SDL screen opened through the `sdl-screen` object is *double buffered*. What that means is that all the drawing commands do not directly happen on the visible screen, but happen in a "hidden" screen (often called a back buffer). When all the objects have been drawn, we swap the visible screen and the invisible screen: the screen we just drew into is now visible, and the previously visible screen is now available for drawing into without affecting what's shown on the screen. This technique allows us to take our time drawing all the object without the user seeing a mess of partially drawn things. In other words, the "double buffer" technique avoids the "flickering" that happens on the screen when objects are drawn one by one in sequence. Flipping the screens is performed by calling the `flip` method of the `sdl-screen`. Here the code of our main loop:

```
(while (not stop)
  (=> scr clear)                ; fill image with black
  (=> bgd draw)                 ; draw moon ground
  (=> event get-arrows xyk)    ; read keyboard
  (when (= (xyk 2) @SDLK_q) (setq stop t)) ; stop when q is pressed
  (setq x (+ x (* 10 (xyk 0)))) ; compute new X coordinate
  (setq y (+ y (* 10 (xyk 1)))) ; compute new Y coordinate
  (=> ship move x y)           ; move ship sprite to new position
  (=> ship draw)               ; draw ship in back buffer
  (=> scr flip)                 ; flip screen buffers
  )                             ; loop
```

The "get-arrow" method fills the three elements of the "xyk" vector as follows:

- (xyk 0): -1 if left arrow pressed, +1 if right arrow pressed, 0 otherwise
- (xyk 1): -1 if up arrow pressed, +1 if down arrow pressed, 0 otherwise
- (xyk 2): key symbol of any other key that is pressed simultaneously.

So, the expression `(setq x (+ x (* 10 (xyk 0))))` will move the ship by 10 pixels to the left or right when the left or right arrow keys are pressed.

The expression `(when (= (xyk 2) @SDLK_q) (setq stop t))` tests if the "q" key was pressed, and sets the "stop" variable to true if it was pressed. `SDLK_q` is a constant. The value of a constant is accessed by prepending one or two "@" characters to the name. The `while` loop tests the `stop` variable and exits if it is true (i.e. if the "q" key has been pressed).

Putting it all together

Here is the complete code:

```
(de lander01 ()
  ;; initialize the SDL subsystem. DONT FORGET THIS!!!
  (sdl-initialize)
  (setq scr (new sdl-screen 640 480 "Lander")) ; open screen
  ;; create background sprite
  (setq bgd (new sdl-sprite scr 0))
  (==> bgd load-frame "moon.png" 0 0 0)
  (==> bgd move 0 0)
  ;; create lem sprite
  (setq ship (new sdl-sprite scr 1))
  (==> ship load-frame "lem.png" 0 40 35)
  ;; set position of ship
  (setq x 200) (setq y 100)
  ;; create event object
  (setq event (new sdl-event))
  (setq xyk (int-matrix 3))

  (while (not stop)
    (==> scr clear) ; fill image with black
    (==> bgd draw) ; draw moon ground
    (==> event get-arrows xyk) ; read keyboard
    (when (= (xyk 2) @SDLK_q) (setq stop t)) ; stop when q is pressed
    (setq x (+ x (* 10 (xyk 0))))
    (setq y (+ y (* 10 (xyk 1))))
    (==> ship move x y) ; move ship sprite to position
    (==> ship draw) ; draw ship
    (==> scr flip) ; flip screens
  ))
```

This code has three major problems:

- It's not a real Lunar Lander game, since the ship does not obey Newtonian mechanics.
- It's written in a terribly un-clean style: the code uses lots of global variables, which is universally recognized as bad practice.
- nothing prevents the ship from getting off the screen.

Second Lander: Gravity and Newtonian mechanics

The new implementation replaces all the global variables by local variables created using the `let*` construct. Local variables created with `let*` (or `let`) disappear after the `let*` evaluation completes.

The new implementation also obeys Newtonian mechanics with (gravity, inertia and such). This is done very simply with the following sequence of operations inside the main loop:

- 1 - read the keyboard arrows and determine the engines thrusts
- 2 - compute accelerations from thrust and gravity (apply $accel=force/mass$):
 - set X-acceleration = X-thrust / ship's mass
 - set Y-acceleration = Y-thrust / ship's mass + gravity
- 3 - compute new velocity from acceleration (time integration):
 - set new X-velocity = old X-velocity + X-acceleration * `deltat`
 - set new Y-velocity = old Y-velocity + Y-acceleration * `deltat`
- 4 - compute new position from velocity (time integration):
 - set X-position = X-position + X-velocity * `deltat`
 - set Y-position = Y-position + Y-velocity * `deltat`

The `deltat` variable is the expected time it takes to go around the main loop of the game (more on this below). Here is how the above equations work. **Step 2** computes the accelerations using Newton's law *acceleration equals force divided by mass*. We must add the acceleration of gravity (which is independent from the mass, as Galileo showed by dropping things from the tower of Pisa) to the Y component of the acceleration. Our objects will be positioned on the screen, where the unit of distance is the pixel. Therefore, our unit of speed will be the pixel per second, and our unit of acceleration the pixel per second per second (i.e. by how many *pixels per second* does our speed change every second). Rather than divide the forces by the mass, we precompute the inverse of the mass `mass-inv`, and multiply the force by this value. We do this because multiplication is a lot faster than division. The Lush code segment is:

```
(setq ax (* mass-inv side-thrust (xyk 0)))
(setq ay (+ grav (* mass-inv main-thrust (xyk 1))))
```

Step 3 computes the velocities from the accelerations. The idea is the following, if horizontal acceleration is `X-acceleration`, and we maintain that acceleration for `deltat` seconds, our velocity will have increased by `X-acceleration` times `deltat`. If we assume that the acceleration was constant while our program went around the loop, and that it took `deltat` seconds to go around that loop, then we must increase the velocity by `(* ax deltat)`. Here is the appropriate Lush code:

```
(setq vx (+ vx (* ax deltat))) ; update X-velocity
(setq vy (+ vy (* ay deltat))) ; update Y-velocity
```

Step 4 computes the position from the velocities. The idea is similar: During the time `deltat` that it takes our program to go around its main loop, we assume that the velocity is constant. The position of the ship during that time has changed by the velocity times `deltat`. This is true for the horizontal velocity and position as well as the vertical velocity and position:

```
(setq x (+ x (* vx deltat))) ; update X-position
(setq y (+ y (* vy deltat))) ; update Y-position
```

So, if going around the loop takes 0.05 seconds (20 frames per second), `deltat` should be 0.05. For example, if the X-velocity is 40 pixels per second and going around the loop takes 0.05 seconds, then the X-position should be incremented by $40 \times 0.05 = 2$ pixels each time we go around the loop, hence the formula above. How do we know how long it takes to go around the loop? Fortunately, the `flip` method of `sdl-screen` object sets the `deltat` slot of the `sdl-screen` object to the number of seconds since the last call to `flip` (most likely, and hopefully a number much smaller than 1). So as long as we do one screen flip per cycle around the loop, we can simply set our `deltat` to the screen's `deltat` which we can access with `:scr:deltat`. The complete update code for our main loop is now:

```
[...get keyboard input into xyk here...]
(setq ax (* mass-inv side-thrust (xyk 0))) ; update X-acceleration
(setq ay (+ grav (* mass-inv main-thrust (xyk 1)))) ; update Y-accel
(setq vx (+ vx (* ax deltat))) ; update X-velocity
(setq vy (+ vy (* ay deltat))) ; update Y-velocity
(setq x (+ x (* vx deltat))) ; update X-position
(setq y (+ y (* vy deltat))) ; update Y-position
[...draw all the objects here...]
(==> scr flip) ; flip screens
(setq deltat :scr:deltat) ; get time between screen flips
```

Next, we need some code to bounce the ship around or have it wrap around the screen when it goes off the boundaries. Here is the code below. We are assuming that the variable "ground" contains the value 360 or so (near the bottom of the screen):

```
(when (< x -40) (setq x (+ 640 (- x -40)))) ; wrap around left side
(when (> x 680) (setq x (+ -40 (- x 640)))) ; wrap around right side
(when (> y ground) ; bounce on ground
  (setq vy (* -0.5 vy)) ; divide vertical speed by 2
  (setq vx (* 0.25 vx)) ; divide horiz speed by 4
  (setq y ground)) ; set vert position to ground altitude
```

Here is the new complete code:

```

(de lander02 ()
  ;; initialize the SDL subsystem. DONT FORGET THIS!!!
  (sdl-initialize)
  (let* ((scr (new sdl-screen 640 480 "Lander")) ; open screen
         (bgd (new sdl-sprite scr 0))          ; create background sprite
         (ship (new sdl-sprite scr 1))         ; create lem sprite
         ;; set position, velocity, acceleration of ship
         (x 200) (y 100) (vx 4) (vy 0) (ax 0) (ay 0)
         ;; set mass, inverse mass, and deltat of ship
         (mass 1) (mass-inv (/ 1 mass)) (deltat 0.01)
         (side-thrust 200)                    ; set side engine thrust
         (main-thrust 400)                    ; set main engine thrust
         (grav 200)                           ; set gravity coefficient in pixels/s/s
         (stop ())
         (event (new sdl-event))
         (xyk (int-matrix 3))
         (ground 360))
    (==> bgd load-frame "moon.png" 0 0 0)
    (==> bgd move 0 0)
    (==> ship load-frame "lem.png" 0 40 35)
    (while (not stop)
      (==> scr clear)                          ; fill image with black
      (==> bgd draw)                            ; draw moon ground
      (==> event get-arrows xyk)                ; read keyboard
      (when (= (xyk 2) @SDLK_q) (setq stop t))   ; stop when q is pressed
      (setq ax (* mass-inv side-thrust (xyk 0))) ; update acceleration
      (setq ay (+ grav (* mass-inv main-thrust (xyk 1)))) ; update acceleration
      (setq vx (+ vx (* ax deltat)))            ; update velocity
      (setq vy (+ vy (* ay deltat)))            ; update velocity
      (setq x (+ x (* vx deltat)))              ; update position
      (setq y (+ y (* vy deltat)))              ; update position
      (when (< x -40) (setq x (+ 640 (- x -40)))) ; wrap around left side
      (when (> x 680) (setq x (+ -40 (- x 640)))) ; wrap around right side
      (when (> y ground)                       ; bounce on ground
        (setq vy (* -0.5 vy))                  ; divide vertical speed by 2
        (setq vx (* 0.25 vx))                  ; divide horiz speed by 4
        (setq y ground))                       ; set vert position to ground altitude
      (==> ship move x y)                       ; move ship sprite to position
      (==> ship draw)                           ; draw ship
      (==> scr flip)                             ; flip screens
      (setq deltat :scr:deltat)                 ; update deltat to time between screen flips
    )))

```

Third Lander: a flame and a shadow

Let's add a little coolness and realism. We will add a flame when the engine is on, and shadow of the ship projected on the ground. We will create two additional sprites, one for the flame, one for the shadow. The flame will be drawn at the same location as the ship, whenever the motor is on. The shadow will be drawn at the same horizontal position as the ship, but at the same vertical position as the ground. Here are the relevant lines:

```
(let* ([... allocate screen and bg sprite ...]
      (ship (new sdl-sprite scr 1))      ; create lem sprite
      (flame (new sdl-sprite scr 1))    ; create flame sprite
      (shadow (new sdl-sprite scr 3)) ; create shadow sprite
      [... more initializations...])
  )
  [... load background image ...]
  ;; load image and put the handle at the center of the sprite (40,35)
  (==> ship load-frame "lem.png" 0 40 35)
  ;; the flame is designed to have the handle at the same place
  (==> flame load-frame "lem-flame.png" 0 40 35)
  ;; here is the shadow
  (==> shadow load-frame "lem-shadow.png" 0 40 -6)
  (while (not stop) ; main loop
    [... compute all the coordinates...]
    (==> shadow move x 360) ; move shadow sprite to position
    (==> flame move x y) ; move flame sprite to position
    (==> ship move x y) ; move ship sprite to position
    (==> shadow draw) ; draw shadow
    (when (<> 0 (xyk 1)) (==> flame draw)) ; draw flame if engine is on
    (==> ship draw) ; draw ship
    (==> scr flip) ; flip screens
    (setq deltat :scr:deltat) ; update deltat to time between screen flips
  ))
```

Here is the complete code (which can be found in `/home/yann/lush/packages/sdl/demos/tutorial.lsh`).

```
(de lander03 ()
  ;; initialize the SDL subsystem. DONT FORGET THIS!!!
  (sdl-initialize)
  (let* ((scr (new sdl-screen 640 480 "Lander")) ; open screen
        (bgd (new sdl-sprite scr 0)) ; create background sprite
        (ship (new sdl-sprite scr 1)) ; create lem sprite
        (flame (new sdl-sprite scr 1)) ; create flame sprite
        (shadow (new sdl-sprite scr 3)) ; create shadow sprite
```

```

;; set position, velocity, acceleration of ship
(x 200) (y 100) (vx 4) (vy 0) (ax 0) (ay 0)
;; set mass, inverse mass, and deltat of ship
(mass 1) (mass-inv (/ 1 mass)) (deltat 0.01)
(side-thrust 200) ; set side engine thrust
(main-thrust 400) ; set main engine thrust
(grav 200) ; set gravity coefficient in pixels/s/s
(stop ())
(event (new sdl-event))
(xyk (int-matrix 3))
(ground 360)
(==> bgd load-frame "moon.png" 0 0 0)
(==> bgd move 0 0)
;; put the handle at the center of the sprite (40,35)
(==> ship load-frame "lem.png" 0 40 35)
(==> flame load-frame "lem-flame.png" 0 40 35)
(==> shadow load-frame "lem-shadow.png" 0 40 -6)
(while (not stop)
  (==> scr clear) ; fill image with black
  (==> bgd draw) ; draw moon ground
  (==> event get-arrows xyk) ; read keyboard
  (when (= (xyk 2) @SDLK_q) (setq stop t)) ; stop when q is pressed
  (setq ax (* mass-inv side-thrust (xyk 0))) ; update acceleration
  (setq ay (+ grav (* mass-inv main-thrust (xyk 1)))) ; update acceleration
  (setq vx (+ vx (* ax deltat))) ; update velocity
  (setq vy (+ vy (* ay deltat))) ; update velocity
  (setq x (+ x (* vx deltat))) ; update position
  (setq y (+ y (* vy deltat))) ; update position
  (when (< x -40) (setq x (+ 640 (- x -40)))) ; wrap around left side
  (when (> x 680) (setq x (+ -40 (- x 640)))) ; wrap around right side
  (when (> y ground) ; bounce on ground
    (setq vy (* -0.5 vy))
    (setq vx (* 0.25 vx))
    (setq y ground))
  (==> shadow move x 360) ; move ship sprite to position
  (==> flame move x y) ; move ship sprite to position
  (==> ship move x y) ; move ship sprite to position
  (==> shadow draw)
  (when (<> 0 (xyk 1)) (==> flame draw)) ; draw flame if engine is on
  (==> ship draw)
  (==> scr flip) ; flip screens
  (setq deltat :scr:deltat) ; update deltat to time between screen flips
  )))

```

Fourth Lander: the ship rotates

More realistic lunar landers have no side thruster , but rotate around to direct their main thruster. We will now explain how to create rotated images of the ship. To do so , the `sdl-sprite` has a method called `rotscale-frame` . This method is called as follows:

```
(=> ship rotscale-frame src-frame dst-frame angle scale)
```

This takes the image of a source frame , identified by its number `src-frame` , rotates it by `angle` degrees (clockwise) , scale it by `scale` , and write the resulting image in the frame with index `dst-frame` . The handle (or hotpoint) is left unchanged in the process , i.e. , the handle in the transformed frame is at the same location within the object as in the source frame.

In the code segment below , we take the 0-th frame and create rotated version of it every 10 degrees:

```
(=> ship load-frame "lem.png" 0 40 35)
(let ((i 1))
  ;; make an image every 10 degrees from 10 to 350
  (for (angle 10 350 10)
    ;; take frame 0, rotate by angle, scale by 1,
    ;; and copy into frame i
    (=> ship rotscale-frame 0 i angle 1)
    (incr i)))
```

Within the main loop , we must read the keyboard , and use the result from the left and right arrow keys to rotate the ship. We will use a variable `theta` to store the current angle of the ship. This variable is an integer between 0 and 35 which , when multiplied by 10 gives the ship's angle with the vertical. `theta` is used as the index of the frame for the ship sprite:

```
(while (not stop)
  [... stuff deleted...]
  (=> event get-arrows xyk) ; read keyboard
  ;; update the angle from the left-right arrow keys
  (setq theta (+ theta (xyk 0)))
  ;; bring the angle back to the 0-360 interval
  (while (>= theta 36) (setq theta (- theta 36)))
  (while (< theta 0) (setq theta (+ theta 36)))
  ;; set frame of ship and flame to one matching the angle
  (=> ship set-frame (int theta))
  [... stuff deleted...]
  )
```

Next , we must modify the computation of the X and Y accelerations to take into account the fact that the thrust of the main engine is at an angle. When the angle of the ship with the vertical is `theta` times 10 degrees , the horizontal component of the thrust is `main-thrust * sin(-pi/180 * 10 * theta)` , which in lush is written `(* main-thrust (sin (* pi/180 -10 theta)))` . We must multiply the angle in degree by `pi/180` because the `sin` functions takes angles in radians. Rather than computing `pi/180` every time , we pre-compute the value and put the result in a global variable called `pi/180` :

```
(defvar pi/180 (/ 3.1415927 180))
```

The horizontal acceleration is the horizontal thrust divided by the ship's mass (or multiplied by the inverse of the mass)

```
(setq ax (* mass-inv main-thrust (xyk 1) (sin (* pi/180 -10 theta))))
```

The vertical component of the acceleration is derived similarly , except the `sin` is now a `cos` , and the gravity component is added:

```
;; compute x acceleration
(setq ax (* mass-inv main-thrust (xyk 1) (sin (* pi/180 -10 theta))))
;; compute y acceleration
(setq ay (+ grav (* mass-inv main-thrust
                 (xyk 1) (cos (* pi/180 10 theta)))))
```

The rest of the code is identical to the previous version.

Here is the complete code , which can be found in `/home/yann/lush/packages/sdl/demos/tutor` .):

```
(setq pi/180 (/ 3.1415927 180))
(de lander04 ()
  ;; initialize the SDL subsystem. DONT FORGET THIS!!!
  (sdl-initialize)
  (let* ((scr (new sdl-screen 640 480 "Lander")) ; open screen
        (bgd (new sdl-sprite scr 0))           ; create background sprite
        (ship (new sdl-sprite scr 1))          ; create lem sprite
        (flame (new sdl-sprite scr 1))         ; create flame sprite
        (shadow (new sdl-sprite scr 3))        ; create shadow sprite
        ;; set position, velocity, acceleration of ship
        (x 200) (y 100) (vx 4) (vy 0) (ax 0) (ay 0)
        ;; angle of ship
        (theta 0)
        ;; set mass, inverse mass, and deltat of ship
        (mass 1) (mass-inv (/ 1 mass)) (deltat 0.01)
```

```

    (side-thrust 200)                ; set side engine thrust
    (main-thrust 400)               ; set main engine thrust
    (grav 200)                      ; set gravity coefficient in pixels/s/s
    (stop ())
    (event (new sdl-event))
    (xyk (int-matrix 3))
    (ground 360))
(==> bgd load-frame "moon.png" 0 0 0)
(==> bgd move 0 0)
;; put the handle at the center of the sprite (40,35)
(==> ship load-frame "lem.png" 0 40 35)
(==> flame load-frame "lem-flame.png" 0 40 35)
(==> shadow load-frame "lem-shadow.png" 0 40 -6)
;; fill up frames with rotated lems
(let ((i 1))
  ;; make an image every 10 degrees from 10 to 350
  (for (angle 10 350 10)
    ;; take frame 0, rotate by angle, scale by 1,
    ;; and copy into frame i
    (==> ship rotscale-frame 0 i angle 1)
    ;; same for flame
    (==> flame rotscale-frame 0 i angle 1)
    (incr i)))
(while (not stop)
  (==> scr clear)                ; fill image with black
  (==> bgd draw)                 ; draw moon ground
  (==> event get-arrows xyk)     ; read keyboard
  (when (= (xyk 2) @SDLK_q) (setq stop t)) ; stop when q is pressed
  ;; update the angle from the left-right arrow keys
  (setq theta (+ theta (xyk 0)))
  ;; bring the angle back to the 0-360 interval
  (while (>= theta 36) (setq theta (- theta 36)))
  (while (< theta 0) (setq theta (+ theta 36)))
  ;; set frame of ship and flame to one matching the angle
  (==> ship set-frame (int theta))
  (==> flame set-frame (int theta))
  ;; compute x acceleration
  (setq ax (* mass-inv main-thrust (xyk 1) (sin (* pi/180 -10 theta))))
  ;; compute y acceleration
  (setq ay (+ grav (* mass-inv main-thrust
                    (xyk 1) (cos (* pi/180 10 theta)))))
  (setq vx (+ vx (* ax deltat))) ; update velocity
  (setq vy (+ vy (* ay deltat))) ; update velocity
  (setq x (+ x (* vx deltat)))   ; update position
  (setq y (+ y (* vy deltat)))   ; update position
  (when (< x -40) (setq x (+ 640 (- x -40)))) ; wrap around left side

```

```

(when (> x 680) (setq x (+ -40 (- x 640)))) ; wrap around right side
(when (> y ground) ; bounce on ground
  (setq vy (* -0.5 vy))
  (setq vx (* 0.25 vx))
  (setq theta 0)
  (setq y ground))
(==> shadow move x 360) ; move shadow sprite to position
(==> flame move x y) ; move flame sprite to position
(==> ship move x y) ; move ship sprite to position
;; now draw sprites in the right order, bottom sprite first
(==> shadow draw) ; draw shadow
(when (<> 0 (xyk 1)) (==> flame draw)) ; draw flame if engine is on
(==> ship draw) ; draw ship
(==> scr flip) ; flip screens
(setq deltat :scr:deltat) ; update deltat to time between screen flips
))

```

11.1.3 SpaceWar: missiles and collision detection

Two ships

[under construction]

Shooting missiles

[under construction]

Collision Detection

[under construction]

Chapter 12

FAQ (Frequently Asked Questions)

12.1 What to do when I get this error?

12.1.1 Module has undefined references

when you load a partially linked object file , or when a function you compiled calls a C function that is not defined in any previously loaded library or module you get the above error. To find out which functions are not defined , type `(mod-undefined)` .

12.1.2 compiler : Unknown Type in: ()

This message may occur for a variety of reasons , but the most common one is that the compiler can't figure out the type of the return value of an expression. A frequent cause is an hash-brace construct (inline C code segment) that is the last expression of a function (and therefore its return value). Lush cannot know the type of the value returned by a hash-brace unless you cast it.

Example 1:

```
;; this causes an error
(de foo1-broken (x) ((-double-) x) #{ $x*$x #})
;; this is correct
(de foo1-correk (x) ((-double-) x) (to-double #{ $x*$x #}))
```

Example 2:

```
;; this causes an error (dunno the return type).
(de foo2-broken (x) ((-idx1- (-int-)) x) #{ *(IDX_PTR($x,int)) = 34 #})
;; this is correct (return value is nil).
```

```
(de foo2-correk (x) ((-idx1- (-int-)) x) #{ *(IDX_PTR($x,int)) = 34 #} ( )
```

12.2 How Do I ... ?

12.2.1 Read lines from a file into a list

```
(de read-lines(f)
  (reading f
    (let ((ans ()))
      (while (<> (skip-char "\\n\\r\\f") "\\e")
        (setq ans (cons (read-string) ans)))
      (reverse ans))))
```

12.2.2 Apply a function to all elements of a vector

interpreted lisp way:

```
(idx-bloop ((x v)) (func x))
```

efficient compiled C code way:

```
(cidx-bloop ("i" ("v" v)) #{ *v = my_c_fun(*v); #}) ( )
```

12.2.3 Get a pointer to the raw data of an idx

just use the `idx-ptr` function.

```
(idx-ptr m}
```

12.2.4 Get a pointer to a function written in Lisp

Some functions in popular libraries take function pointers as argument (a typical example is the GSL function minimization routines). To obtain a pointer to the compiled version is a function written in Lisp , simply use the function `to-gptr` :

```
(de myfunc (x) ((-double-) x) (- (* x x) 2))
(dhc-make () myfunc)
(some-root-finding-function-in-C (to-gptr myfunc))
```

12.2.5 Know if a function can be used in compiled code

Just use `compilablep` :

```
? (compilablep +)  
= t
```


Chapter 13

Local Packages

This section contains the documentation of packages that are special to your local installation of Lush. Those packages can be found in the `local` in the Lush directory.