

FreeMat v4.0 Documentation

Samit Basu

October 4, 2009

Contents

1	Introduction and Getting Started	39
1.1	INSTALL Installing FreeMat	39
1.1.1	General Instructions	39
1.1.2	Linux	39
1.1.3	Windows	40
1.1.4	Mac OS X	40
1.1.5	Source Code	40
2	Variables and Arrays	41
2.1	CELL Cell Array Definitions	41
2.1.1	Usage	41
2.1.2	Examples	41
2.2	Function Handles	42
2.2.1	Usage	42
2.3	GLOBAL Global Variables	42
2.3.1	Usage	42
2.3.2	Example	42
2.4	INDEXING Indexing Expressions	43
2.4.1	Usage	43
2.4.2	Array Indexing	43
2.4.3	Cell Indexing	46
2.4.4	Structure Indexing	47
2.4.5	Complex Indexing	49
2.5	MATRIX Matrix Definitions	49
2.5.1	Usage	49
2.5.2	Examples	49
2.6	PERSISTENT Persistent Variables	51
2.6.1	Usage	51
2.6.2	Example	51
2.7	STRUCT Structure Array Constructor	51
2.7.1	Usage	51
2.7.2	Example	52

3	Functions and Scripts	55
3.1	ANONYMOUS Anonymous Functions	55
3.1.1	Usage	55
3.1.2	Examples	55
3.2	FUNCTION Function Declarations	56
3.2.1	Usage	56
3.2.2	Examples	57
3.3	KEYWORDS Function Keywords	60
3.3.1	Usage	60
3.3.2	Example	60
3.4	NARGIN Number of Input Arguments	61
3.4.1	Usage	61
3.4.2	Example	61
3.5	NARGOUT Number of Output Arguments	62
3.5.1	Usage	62
3.5.2	Example	62
3.6	SCRIPT Script Files	63
3.6.1	Usage	63
3.6.2	Example	63
3.7	SPECIAL Special Calling Syntax	64
3.7.1	Usage	64
3.7.2	Example	64
3.8	VARARGIN Variable Input Arguments	65
3.8.1	Usage	65
3.8.2	Example	65
3.9	VARARGOUT Variable Output Arguments	66
3.9.1	Usage	66
3.9.2	Example	66
4	Mathematical Operators	67
4.1	COLON Index Generation Operator	67
4.1.1	Usage	67
4.1.2	Function Internals	67
4.1.3	Examples	69
4.2	COMPARISONOPS Array Comparison Operators	69
4.2.1	Usage	69
4.2.2	Examples	70
4.3	DOTLEFTDIVIDE Element-wise Left-Division Operator	71
4.3.1	Usage	71
4.3.2	Function Internals	71
4.3.3	Examples	71
4.4	DOTPOWER Element-wise Power Operator	73
4.4.1	Usage	73
4.4.2	Function Internals	73
4.4.3	Examples	73
4.5	DOTRIGHTDIVIDE Element-wise Right-Division Operator	74

4.5.1	Usage	74
4.5.2	Function Internals	75
4.5.3	Examples	75
4.6	DOTTIMES Element-wise Multiplication Operator	76
4.6.1	Usage	76
4.6.2	Function Internals	77
4.6.3	Examples	77
4.7	HERMITIAN Matrix Hermitian (Conjugate Transpose) Operator	78
4.7.1	Usage	78
4.7.2	Function Internals	78
4.7.3	Examples	78
4.8	LEFTDIVIDE Matrix Equation Solver/Divide Operator	79
4.8.1	Usage	79
4.8.2	Function Internals	79
4.8.3	Examples	80
4.9	LOGICALOPS Logical Array Operators	81
4.9.1	Usage	81
4.9.2	Examples	82
4.10	MATRIXPOWER Matrix Power Operator	83
4.10.1	Usage	83
4.10.2	Function Internals	84
4.10.3	Examples	84
4.11	MINUS Subtraction Operator	85
4.11.1	Usage	85
4.11.2	Function Internals	85
4.11.3	Examples	86
4.12	PLUS Addition Operator	87
4.12.1	Usage	87
4.12.2	Function Internals	87
4.12.3	Examples	87
4.13	RIGHTDIVIDE Matrix Equation Solver/Divide Operator	88
4.13.1	Usage	88
4.13.2	Function Internals	89
4.13.3	Examples	89
4.14	TIMES Matrix Multiply Operator	89
4.14.1	Usage	89
4.14.2	Function Internals	90
4.14.3	Examples	90
4.15	TRANSPOSE Matrix Transpose Operator	91
4.15.1	Usage	91
4.15.2	Function Internals	91
4.15.3	Examples	91
4.16	TYPERULES Type Rules for Operators	92
4.16.1	Usage	92

5	Flow Control	93
5.1	BREAK Exit Execution In Loop	93
5.1.1	Usage	93
5.1.2	Example	93
5.2	CONTINUE Continue Execution In Loop	94
5.2.1	Usage	94
5.2.2	Example	94
5.3	ERROR Causes an Error Condition Raised	95
5.3.1	Usage	95
5.3.2	Example	95
5.4	FOR For Loop	96
5.4.1	Usage	96
5.4.2	Examples	96
5.5	IF-ELSEIF-ELSE Conditional Statements	97
5.5.1	Usage	97
5.5.2	Examples	97
5.6	KEYBOARD Initiate Interactive Debug Session	98
5.6.1	Usage	98
5.6.2	Example	98
5.7	LASTERR Retrieve Last Error Message	100
5.7.1	Usage	100
5.7.2	Example	100
5.8	RECALL Return From All Keyboard Sessions	100
5.8.1	Usage	100
5.8.2	Example	101
5.9	RETURN Return From Function	101
5.9.1	Usage	101
5.9.2	Example	101
5.10	SWITCH Switch statement	103
5.10.1	Usage	103
5.10.2	Examples	103
5.11	TRY-CATCH Try and Catch Statement	104
5.11.1	Usage	104
5.11.2	Examples	104
5.12	WARNING Emits a Warning Message	105
5.12.1	Usage	105
5.13	WHILE While Loop	105
5.13.1	Usage	105
5.13.2	Examples	105
6	FreeMat Functions	107
6.1	ADDPATH Add	107
6.1.1	Usage	107
6.2	ASSIGNIN Assign Variable in Workspace	107
6.2.1	Usage	107
6.3	BLASLIB Select BLAS library	108

6.3.1	Usage	108
6.4	BUILTIN Evaluate Builtin Function	108
6.4.1	Usage	108
6.5	CLC Clear Display	108
6.5.1	Usage	108
6.6	CLOCK Get Current Time	108
6.6.1	Usage	108
6.6.2	Example	109
6.7	CLOCKTOTIME Convert Clock Vector to Epoch Time	109
6.7.1	Usage	109
6.7.2	Example	109
6.8	COMPUTER Computer System FreeMat is Running On	110
6.8.1	Usage	110
6.9	DIARY Create a Log File of Console	110
6.9.1	Usage	110
6.10	DOCLI Start a Command Line Interface	111
6.10.1	Usage	111
6.11	EDIT Open Editor Window	111
6.11.1	Usage	111
6.12	EDITOR Open Editor Window	111
6.12.1	Usage	111
6.13	ERRORCOUNT Retrieve the Error Counter for the Interpreter	111
6.13.1	Usage	111
6.14	ETIME Elapsed Time Function	111
6.14.1	Usage	111
6.14.2	Example	112
6.15	EVAL Evaluate a String	112
6.15.1	Usage	112
6.15.2	Example	112
6.16	EVALIN Evaluate a String in Workspace	113
6.16.1	Usage	113
6.17	EXIT Exit Program	114
6.17.1	Usage	114
6.18	FEVAL Evaluate a Function	114
6.18.1	Usage	114
6.18.2	Example	114
6.19	FILESEP Directory Separation Character	115
6.19.1	Usage	115
6.20	HELP Help	115
6.20.1	Usage	115
6.21	HELPWIN Online Help Window	116
6.21.1	Usage	116
6.22	JITCONTROL Control the Just In Time Compiler	116
6.22.1	Usage	116
6.23	MFILENAME Name of Current Function	121
6.23.1	Usage	121

6.24	PATH Get or Set FreeMat Path	121
6.24.1	Usage	121
6.25	PATHSEP Path Directories Separation Character	122
6.25.1	Usage	122
6.26	PATHTOOL Open Path Setting Tool	122
6.26.1	Usage	122
6.27	PCODE Convert a Script or Function to P-Code	122
6.27.1	Usage	122
6.28	PROFILER Control Profiling	123
6.28.1	Usage	123
6.29	QUIET Control the Verbosity of the Interpreter	123
6.29.1	Usage	123
6.30	QUIT Quit Program	123
6.30.1	Usage	123
6.31	REHASH Rehash Directory Caches	124
6.31.1	Usage	124
6.32	RESCAN Rescan M Files for Changes	124
6.32.1	Usage	124
6.33	ROOTPATH Set FreeMat Root Path	124
6.33.1	Usage	124
6.34	SIMKEYS Simulate Keypresses from the User	125
6.34.1	Usage	125
6.35	SLEEP Sleep For Specified Number of Seconds	125
6.35.1	Usage	125
6.36	SOURCE Execute an Arbitrary File	125
6.36.1	Usage	125
6.36.2	Example	125
6.37	STARTUP Startup Script	126
6.37.1	Usage	126
6.38	TIC Start Stopwatch Timer	126
6.38.1	Usage	126
6.38.2	Example	126
6.39	TOC Stop Stopwatch Timer	126
6.39.1	Usage	126
6.39.2	Example	126
6.40	TYPE Type	127
6.40.1	Usage	127
6.41	VERSION The Current Version Number	127
6.41.1	Usage	127
6.41.2	Example	127
6.42	VERSTRING The Current Version String	127
6.42.1	Usage	127
6.42.2	Example	127

7	Debugging FreeMat Code	129
7.1	DBAUTO Control Dbauto Functionality	129
7.1.1	Usage	129
7.2	DBDELETE Delete a Breakpoint	130
7.2.1	Usage	130
7.3	DBDown Move Down One Debug Level	130
7.3.1	Usage	130
7.4	DBLIST List Breakpoints	130
7.4.1	Usage	130
7.5	DBSTEP Step N Statements	130
7.5.1	Usage	130
7.6	DBSTOP	130
7.6.1	Usage	130
7.7	DBUP Move Up One Debug Level	131
7.7.1	Usage	131
7.8	FDUMP Dump Information on Function	131
7.8.1	Usage	131
8	Sparse Matrix Support	133
8.1	EIGS Sparse Matrix Eigendecomposition	133
8.1.1	Usage	133
8.1.2	Example	134
8.2	FULL Convert Sparse Matrix to Full Matrix	135
8.2.1	Usage	135
8.2.2	Example	135
8.3	SPARSE Construct a Sparse Matrix	136
8.3.1	Usage	136
8.4	SPEYE Sparse Identity Matrix	136
8.4.1	Usage	136
8.4.2	Example	136
8.5	SPONES Sparse Ones Function	137
8.5.1	Usage	137
8.5.2	Examples	137
8.6	SPRAND Sparse Uniform Random Matrix	138
8.6.1	Usage	138
8.6.2	Examples	138
8.7	SPRANDN Sparse Normal Random Matrix	139
8.7.1	Usage	139
8.7.2	Examples	139
8.8	SPY Visualize Sparsity Pattern of a Sparse Matrix	140
8.8.1	Usage	140
8.8.2	Example	140

9	Mathematical Functions	143
9.1	ACOS Inverse Trigonometric Arccosine Function	143
9.1.1	Usage	143
9.1.2	Function Internals	143
9.1.3	Example	143
9.2	ACOSD Inverse Cosine Degrees Function	144
9.2.1	Usage	144
9.2.2	Examples	144
9.3	ACOSH Inverse Hyperbolic Cosine Function	144
9.3.1	Usage	144
9.3.2	Function Internals	145
9.3.3	Examples	145
9.4	ACOT Inverse Cotangent Function	145
9.4.1	Usage	145
9.4.2	Function Internals	145
9.4.3	Examples	146
9.5	ACOTD Inverse Cotangent Degrees Function	146
9.5.1	Usage	146
9.6	ACOTH Inverse Hyperbolic Cotangent Function	146
9.6.1	Usage	146
9.6.2	Function Internals	146
9.6.3	Examples	147
9.7	ACSC Inverse Cosecant Function	147
9.7.1	Usage	147
9.7.2	Function Internals	147
9.7.3	Examples	147
9.8	ACSCD Inverse Cosecant Degrees Function	148
9.8.1	Usage	148
9.8.2	Examples	148
9.9	ACSCH Inverse Hyperbolic Cosecant Function	148
9.9.1	Usage	148
9.9.2	Function Internals	149
9.9.3	Examples	149
9.10	ANGLE Phase Angle Function	149
9.10.1	Usage	149
9.10.2	Function Internals	149
9.10.3	Example	150
9.11	ASEC Inverse Secant Function	150
9.11.1	Usage	150
9.11.2	Function Internals	150
9.11.3	Examples	150
9.12	ASECD Inverse Secant Degrees Function	151
9.12.1	Usage	151
9.12.2	Examples	151
9.13	ASECH Inverse Hyperbolic Secant Function	151
9.13.1	Usage	151

9.13.2	Function Internals	152
9.13.3	Examples	152
9.14	ASIN Inverse Trigonometric Arcsine Function	152
9.14.1	Usage	152
9.14.2	Function Internals	152
9.14.3	Example	153
9.15	ASIND Inverse Sine Degrees Function	153
9.15.1	Usage	153
9.15.2	Examples	153
9.16	ASINH Inverse Hyperbolic Sine Function	153
9.16.1	Usage	153
9.16.2	Function Internals	154
9.16.3	Examples	154
9.17	ATAN Inverse Trigonometric Arctangent Function	154
9.17.1	Usage	154
9.17.2	Function Internals	154
9.17.3	Example	155
9.18	ATAN2 Inverse Trigonometric 4-Quadrant Arctangent Function	155
9.18.1	Usage	155
9.18.2	Function Internals	155
9.18.3	Example	156
9.19	ATAND Inverse Tangent Degrees Function	156
9.19.1	Usage	156
9.19.2	Examples	156
9.20	ATANH Inverse Hyperbolic Tangent Function	157
9.20.1	Usage	157
9.20.2	Function Internals	157
9.20.3	Examples	157
9.21	COS Trigonometric Cosine Function	157
9.21.1	Usage	157
9.21.2	Function Internals	158
9.21.3	Example	158
9.22	COSD Cosine Degrees Function	158
9.22.1	Usage	158
9.22.2	Examples	158
9.23	COSH Hyperbolic Cosine Function	159
9.23.1	Usage	159
9.23.2	Function Internals	159
9.23.3	Examples	159
9.24	COT Trigonometric Cotangent Function	160
9.24.1	Usage	160
9.24.2	Function Internals	160
9.24.3	Example	160
9.25	COTD Cotangent Degrees Function	160
9.25.1	Usage	160
9.25.2	Examples	161

9.26	COTD Cotangent Degrees Function	161
9.26.1	Usage	161
9.26.2	Examples	161
9.27	COTH Hyperbolic Cotangent Function	161
9.27.1	Usage	161
9.27.2	Function Internals	161
9.27.3	Examples	161
9.28	CROSS Cross Product of Two Vectors	162
9.28.1	Usage	162
9.29	CSC Trigonometric Cosecant Function	162
9.29.1	Usage	162
9.29.2	Function Internals	162
9.29.3	Example	162
9.30	CSCD Cosecant Degrees Function	163
9.30.1	Usage	163
9.31	CSCH Hyperbolic Cosecant Function	163
9.31.1	Usage	163
9.31.2	Function Internals	163
9.31.3	Examples	163
9.32	DEG2RAD Convert From Degrees To Radians	164
9.32.1	Usage	164
9.32.2	Example	164
9.33	ERF Error Function	164
9.33.1	Usage	164
9.33.2	Function Internals	164
9.33.3	Example	165
9.34	ERFC Complimentary Error Function	165
9.34.1	Usage	165
9.34.2	Function Internals	165
9.34.3	Example	165
9.35	EXP Exponential Function	166
9.35.1	Usage	166
9.35.2	Function Internals	166
9.35.3	Example	166
9.36	EXPM1 Exponential Minus One Function	167
9.36.1	Usage	167
9.37	FIX Round Towards Zero	167
9.37.1	Usage	167
9.37.2	Example	168
9.38	GAMMA Gamma Function	168
9.38.1	Usage	168
9.38.2	Function Internals	168
9.38.3	Example	169
9.39	GAMMALN Log Gamma Function	169
9.39.1	Usage	169
9.39.2	Example	169

9.40	IDIV Integer Division Operation	170
9.40.1	Usage	170
9.40.2	Example	170
9.41	LOG Natural Logarithm Function	170
9.41.1	Usage	170
9.41.2	Function Internals	171
9.41.3	Example	171
9.42	LOG10 Base-10 Logarithm Function	171
9.42.1	Usage	171
9.42.2	Example	172
9.43	LOG1P Natural Logarithm of 1+P Function	172
9.43.1	Usage	172
9.44	LOG2 Base-2 Logarithm Function	172
9.44.1	Usage	172
9.44.2	Example	173
9.45	MOD Modulus Operation	173
9.45.1	Usage	173
9.45.2	Example	173
9.46	MPOWER Matrix Power Function	174
9.46.1	Usage	174
9.47	POWER Element-wise Power Function	174
9.47.1	Usage	174
9.48	RAD2DEG Radians To Degrees Conversion Function	175
9.48.1	Usage	175
9.48.2	Examples	175
9.49	REM Remainder After Division	175
9.49.1	Usage	175
9.49.2	Example	176
9.50	SEC Trigonometric Secant Function	176
9.50.1	Usage	176
9.50.2	Function Internals	176
9.50.3	Example	177
9.51	SECD Secant Degrees Function	177
9.51.1	Usage	177
9.52	SECH Hyperbolic Secant Function	177
9.52.1	Usage	177
9.52.2	Function Internals	177
9.52.3	Examples	178
9.53	SIN Trigonometric Sine Function	178
9.53.1	Usage	178
9.53.2	Function Internals	178
9.53.3	Example	178
9.54	SIND Sine Degrees Function	179
9.54.1	Usage	179
9.54.2	Examples	179
9.55	SINH Hyperbolic Sine Function	179

9.55.1	Usage	179
9.55.2	Function Internals	180
9.55.3	Examples	180
9.56	SQRT Square Root of an Array	180
9.56.1	Usage	180
9.56.2	Example	180
9.57	TAN Trigonometric Tangent Function	181
9.57.1	Usage	181
9.57.2	Function Internals	181
9.57.3	Example	182
9.58	TAND Tangent Degrees Function	182
9.58.1	Usage	182
9.58.2	Examples	182
9.59	TANH Hyperbolic Tangent Function	182
9.59.1	Usage	182
9.59.2	Function Internals	183
9.59.3	Examples	183
10	Base Constants	185
10.1	E Euler Constant (Base of Natural Logarithm)	185
10.1.1	Usage	185
10.1.2	Example	185
10.2	EPS Double Precision Floating Point Relative Machine Precision Epsilon	185
10.2.1	Usage	185
10.2.2	Example	186
10.3	FALSE Logical False	186
10.3.1	Usage	186
10.4	FEPS Single Precision Floating Point Relative Machine Precision Epsilon	186
10.4.1	Usage	186
10.4.2	Example	187
10.5	I-J Square Root of Negative One	187
10.5.1	Usage	187
10.5.2	Example	187
10.6	INF Infinity Constant	188
10.6.1	Usage	188
10.6.2	Function Internals	189
10.6.3	Example	189
10.7	PI Constant Pi	190
10.7.1	Usage	190
10.7.2	Example	190
10.8	TEPS Type-based Epsilon Calculation	190
10.8.1	Usage	190
10.8.2	Example	191
10.9	TRUE Logical TRUE	191
10.9.1	Usage	191

11 Elementary Functions	193
11.1 ABS Absolute Value Function	193
11.1.1 Usage	193
11.1.2 Example	193
11.2 ALL All True Function	194
11.2.1 Usage	194
11.2.2 Function Internals	194
11.2.3 Example	194
11.3 ANY Any True Function	195
11.3.1 Usage	195
11.3.2 Function Internals	195
11.3.3 Example	195
11.4 CEIL Ceiling Function	196
11.4.1 Usage	196
11.4.2 Example	196
11.5 CONJ Conjugate Function	197
11.5.1 Usage	197
11.5.2 Example	197
11.6 CUMPROD Cumulative Product Function	198
11.6.1 Usage	198
11.6.2 Function Internals	198
11.6.3 Example	198
11.7 CUMSUM Cumulative Summation Function	199
11.7.1 Usage	199
11.7.2 Function Internals	199
11.7.3 Example	200
11.8 DEAL Multiple Simultaneous Assignments	201
11.8.1 Usage	201
11.9 DEC2HEX Convert Decimal Number to Hexadecimal	201
11.9.1 Usage	201
11.9.2 Example	201
11.10DIFF Difference Function	202
11.10.1 Usage	202
11.11DOT Dot Product Function	202
11.11.1 Usage	202
11.12FLOOR Floor Function	203
11.12.1 Usage	203
11.12.2 Example	203
11.13GETFIELD Get Field Contents	203
11.13.1 Usage	203
11.14HEX2DEC Convert Hexadecimal Numbers To Decimal	204
11.14.1 Usage	204
11.14.2 Examples	204
11.15IMAG Imaginary Function	204
11.15.1 Usage	204
11.15.2 Example	205

11.16	MAX Maximum Function	205
	11.16.1 Usage	205
	11.16.2 Function Internals	206
	11.16.3 Example	206
11.17	MEAN Mean Function	207
	11.17.1 Usage	207
	11.17.2 Function Internals	208
	11.17.3 Example	208
11.18	MIN Minimum Function	208
	11.18.1 Usage	208
	11.18.2 Function Internals	209
	11.18.3 Example	209
11.19	NUM2HEX Convert Numbers to IEEE Hex Strings	211
	11.19.1 Usage	211
	11.19.2 Example	211
11.20	PROD Product Function	212
	11.20.1 Usage	212
	11.20.2 Example	212
11.21	REAL Real Function	213
	11.21.1 Usage	213
	11.21.2 Example	213
11.22	ROUND Round Function	213
	11.22.1 Usage	213
	11.22.2 Example	214
11.23	STD Standard Deviation Function	214
	11.23.1 Usage	214
	11.23.2 Example	215
11.24	SUB2IND Convert Multiple Indexing To Linear Indexing	215
	11.24.1 Usage	215
	11.24.2 Example	216
11.25	SUM Sum Function	216
	11.25.1 Usage	216
	11.25.2 Function Internals	217
	11.25.3 Example	217
11.26	TEST Test Function	217
	11.26.1 Usage	217
11.27	VAR Variance Function	218
	11.27.1 Usage	218
	11.27.2 Function Internals	218
	11.27.3 Example	218
11.28	VEC Reshape to a Vector	219
	11.28.1 Usage	219
	11.28.2 Example	219

12 Inspection Functions	221
12.1 CLEAR Clear or Delete a Variable	221
12.1.1 Usage	221
12.1.2 Example	222
12.2 END End Function	222
12.2.1 Usage	222
12.3 EXIST Test for Existence	222
12.3.1 Usage	222
12.3.2 Example	223
12.4 FIELDNAMES Fieldnames of a Structure	224
12.4.1 Usage	224
12.4.2 Example	224
12.5 ISA Test Type of Variable	224
12.5.1 Usage	224
12.5.2 Examples	225
12.6 ISCELL Test For Cell Array	226
12.6.1 Usage	226
12.6.2 Example	226
12.7 ISCELLSTR Test For Cell Array of Strings	226
12.7.1 Usage	226
12.7.2 Example	227
12.8 ISCHAR Test For Character Array (string)	227
12.8.1 Usage	227
12.9 ISEMPY Test For Variable Empty	227
12.9.1 Usage	227
12.9.2 Examples	227
12.10 ISEQUAL Test For Matrix Equality	228
12.10.1 Usage	228
12.11 ISFIELD Test for Existence of a Structure Field	228
12.11.1 Usage	228
12.11.2 Example	228
12.12 ISHANDLE Test for Graphics Handle	229
12.12.1 Usage	229
12.13 ISINF Test for infinities	229
12.13.1 Usage	229
12.13.2 Example	230
12.14 ISINTTYPE Test For Integer-type Array	230
12.14.1 Usage	230
12.15 ISLOGICAL Test for Logical Array	230
12.15.1 Usage	230
12.16 ISMATRIX Test For a 2D Matrix	230
12.16.1 Usage	230
12.17 ISNAN Test for Not-a-Numbers	231
12.17.1 Usage	231
12.17.2 Example	231
12.18 ISNUMERIC Test for Numeric Array	231

12.18.1 Usage	231
12.19ISREAL Test For Real Array	231
12.19.1 Usage	231
12.20ISSAME Test If Two Arrays Are Identical	232
12.20.1 Usage	232
12.21ISSCALAR Test For Scalar	232
12.21.1 Usage	232
12.22ISSET Test If Variable Set	232
12.22.1 Usage	232
12.22.2 Example	232
12.23ISSPARSE Test for Sparse Matrix	233
12.23.1 Usage	233
12.23.2 Example	233
12.24ISSQUARE Test For a Square matrix	234
12.24.1 Usage	234
12.25ISSTR Test For Character Array (string)	234
12.25.1 Usage	234
12.26ISSTRUCT Test For Structure Array	234
12.26.1 Usage	234
12.27ISVECTOR Test For a Vector	234
12.27.1 Usage	234
12.28LENGTH Length of an Array	234
12.28.1 Usage	234
12.28.2 Example	235
12.29MAXDIM Maximum Dimension in Array	235
12.29.1 Usage	235
12.30NDIMS Number of Dimensions in Array	235
12.30.1 Usage	235
12.31NNZ Number of Nonzeros	235
12.31.1 Usage	235
12.31.2 Example	236
12.32NUMEL Number of Elements in an Array	236
12.32.1 Usage	236
12.32.2 Example	236
12.33SIZE Size of a Variable	237
12.33.1 Usage	237
12.33.2 Example	237
12.34TYPEOF Determine the Type of an Argument	238
12.34.1 Usage	238
12.34.2 Example	238
12.35WHERE Get Information on Program Stack	240
12.35.1 Usage	240
12.35.2 Example	240
12.36WHICH Get Information on Function	241
12.36.1 Usage	241
12.36.2 Example	241

12.37WHO Describe Currently Defined Variables	241
12.37.1 Usage	241
12.37.2 Example	242
12.38WHOS Describe Currently Defined Variables With Memory Usage	242
12.38.1 Usage	242
13 Type Conversion Functions	245
13.1 BIN2DEC Convert Binary String to Decimal	245
13.1.1 USAGE	245
13.1.2 Example	245
13.2 BIN2INT Convert Binary Arrays to Integer	245
13.2.1 Usage	245
13.2.2 Example	246
13.2.3 Tets	247
13.3 CAST Typecast Variable to Specified Type	247
13.3.1 Usage	247
13.3.2 Example	248
13.4 CHAR Convert to character array or string	248
13.4.1 Usage	248
13.4.2 Example	249
13.5 COMPLEX Create a Complex Number	249
13.5.1 Usage	249
13.6 DCOMPLEX Convert to Double Precision (deprecated)	249
13.6.1 Usage	249
13.7 DEC2BIN Convert Decimal to Binary String	250
13.7.1 USAGE	250
13.7.2 Example	250
13.8 DOUBLE Convert to 64-bit Floating Point	250
13.8.1 Usage	250
13.8.2 Example	251
13.9 FLOAT Convert to 32-bit Floating Point	251
13.9.1 Usage	251
13.9.2 Example	252
13.10INT16 Convert to Signed 16-bit Integer	252
13.10.1 Usage	252
13.10.2 Example	253
13.11INT2BIN Convert Integer Arrays to Binary	254
13.11.1 Usage	254
13.11.2 Example	254
13.12INT32 Convert to Signed 32-bit Integer	255
13.12.1 Usage	255
13.12.2 Example	255
13.13INT64 Convert to Signed 64-bit Integer	256
13.13.1 Usage	256
13.13.2 Example	256
13.14INT8 Convert to Signed 8-bit Integer	257

13.14.1 Usage	257
13.14.2 Example	258
13.15 LOGICAL Convert to Logical	259
13.15.1 Usage	259
13.15.2 Example	259
13.16 SINGLE Convert to 32-bit Floating Point	259
13.16.1 Usage	259
13.17 STRING Convert Array to String	259
13.17.1 Usage	259
13.17.2 Example	260
13.18 UINT16 Convert to Unsigned 16-bit Integer	260
13.18.1 Usage	260
13.18.2 Example	260
13.19 UINT32 Convert to Unsigned 32-bit Integer	261
13.19.1 Usage	261
13.19.2 Example	261
13.20 UINT64 Convert to Unsigned 64-bit Integer	262
13.20.1 Usage	262
13.20.2 Example	263
13.21 UINT8 Convert to Unsigned 8-bit Integer	264
13.21.1 Usage	264
13.21.2 Example	264
14 Array Generation and Manipulations	267
14.1 ASSIGN Making assignments	267
14.1.1 Usage	267
14.2 CELL Cell Array of Empty Matrices	268
14.2.1 Usage	268
14.2.2 Example	268
14.3 CIRCSHIFT Circularly Shift an Array	269
14.3.1 USAGE	269
14.3.2 Example	269
14.4 COND Condition Number of a Matrix	272
14.4.1 Usage	272
14.4.2 Function Internals	273
14.4.3 Example	273
14.5 DET Determinant of a Matrix	273
14.5.1 Usage	273
14.5.2 Function Internals	274
14.5.3 Example	274
14.6 DIAG Diagonal Matrix Construction/Extraction	274
14.6.1 Usage	274
14.6.2 Examples	275
14.7 EXPM Matrix Exponential	276
14.7.1 Usage	276
14.7.2 Example	276

14.8 EYE Identity Matrix	276
14.8.1 USAGE	276
14.8.2 Example	276
14.9 FIND Find Non-zero Elements of An Array	277
14.9.1 Usage	277
14.9.2 Example	277
14.10FLIPDIM Reverse a Matrix Along a Given Dimension	280
14.10.1 USAGE	280
14.10.2 Example	280
14.11FLIPLR Reverse the Columns of a Matrix	282
14.11.1 USAGE	282
14.11.2 Example	282
14.12FLIPUD Reverse the Columns of a Matrix	283
14.12.1 USAGE	283
14.12.2 Example	284
14.13IPERMUTE Array Inverse Permutation Function	285
14.13.1 Usage	285
14.13.2 Example	285
14.14ISFLOAT Test for Floating Point Array	287
14.14.1 Usage	287
14.15ISINTEGER Test for Integer Array	287
14.15.1 Usage	287
14.16Linspace Linearly Spaced Vector	287
14.16.1 Usage	287
14.16.2 Examples	288
14.17LOGSPACE Logarithmically Spaced Vector	288
14.17.1 Usage	288
14.17.2 Example	288
14.18MESHGRID Generate Grid Mesh For Plots	288
14.18.1 Usage	288
14.18.2 Example	289
14.19NAN Not-a-Number Constant	290
14.19.1 Usage	290
14.19.2 Example	291
14.20NDGRID Generate N-Dimensional Grid	292
14.20.1 Usage	292
14.20.2 Example	292
14.21NONZEROS Retrieve Nonzero Matrix Entries	294
14.21.1 USAGE	294
14.21.2 Example	294
14.22NORM Norm Calculation	295
14.22.1 Usage	295
14.22.2 Examples	296
14.23NUM2STR Convert Numbers To Strings	297
14.23.1 Usage	297
14.24ONES Array of Ones	298

14.24.1 Usage	298
14.24.2 Example	298
14.25 PERMUTE Array Permutation Function	299
14.25.1 Usage	299
14.25.2 Example	299
14.26 PINV Moore-Penrose Pseudoinverse	300
14.26.1 Usage	300
14.26.2 Function Internals	301
14.26.3 Examples	301
14.27 RANK Calculate the Rank of a Matrix	302
14.27.1 Usage	302
14.27.2 Examples	303
14.28 RCOND Reciprocal Condition Number Estimate	303
14.28.1 Usage	303
14.28.2 Example	304
14.29 REPMAT Array Replication Function	304
14.29.1 Usage	304
14.29.2 Example	305
14.30 RESHAPE Reshape An Array	306
14.30.1 Usage	306
14.30.2 Example	306
14.31 RREF Reduced Row Echelon Form of a Matrix	307
14.31.1 Usage	307
14.32 SHIFTDIM Shift Array Dimensions Function	307
14.32.1 Usage	307
14.32.2 Example	308
14.33 SORT Sort	309
14.33.1 Usage	309
14.33.2 Example	310
14.34 SQUEEZE Remove Singleton Dimensions of an Array	311
14.34.1 Usage	311
14.34.2 Example	311
14.35 SUBSREF Array Dereferencing	312
14.35.1 Usage	312
14.36 TRANSPOSE Matrix Transpose	312
14.36.1 Usage	312
14.36.2 Example	312
14.37 UNIQUE Unique	313
14.37.1 Usage	313
14.37.2 Example	313
14.38 XNRM2 BLAS Norm Calculation	316
14.38.1 Usage	316
14.39 ZEROS Array of Zeros	317
14.39.1 Usage	317
14.39.2 Example	317

15 Random Number Generation	319
15.1 RAND Uniform Random Number Generator	319
15.1.1 Usage	319
15.1.2 Example	320
15.2 RANDBETA Beta Deviate Random Number Generator	321
15.2.1 Usage	321
15.2.2 Function Internals	321
15.2.3 Example	321
15.3 RANDBIN Generate Binomial Random Variables	322
15.3.1 Usage	322
15.3.2 Function Internals	322
15.3.3 Example	322
15.4 RANDCHI Generate Chi-Square Random Variable	323
15.4.1 Usage	323
15.4.2 Function Internals	323
15.4.3 Example	323
15.5 RANDEXP Generate Exponential Random Variable	324
15.5.1 Usage	324
15.5.2 Function Internals	324
15.5.3 Example	324
15.6 RANDF Generate F-Distributed Random Variable	325
15.6.1 Usage	325
15.6.2 Function Internals	325
15.6.3 Example	325
15.7 RANDGAMMA Generate Gamma-Distributed Random Variable	326
15.7.1 Usage	326
15.7.2 Function Internals	326
15.7.3 Example	326
15.8 RANDI Uniformly Distributed Integer	327
15.8.1 Usage	327
15.8.2 Example	327
15.9 RANDMULTI Generate Multinomial-distributed Random Variables	327
15.9.1 Usage	327
15.9.2 Function Internals	328
15.9.3 Example	328
15.10 RANDN Gaussian (Normal) Random Number Generator	328
15.10.1 Usage	328
15.10.2 Function Internals	329
15.10.3 Example	329
15.11 RANDNBIN Generate Negative Binomial Random Variables	330
15.11.1 Usage	330
15.11.2 Function Internals	331
15.11.3 Example	331
15.12 RANDNCHI Generate Noncentral Chi-Square Random Variable	331
15.12.1 Usage	331
15.12.2 Function Internals	331

15.12.3 Examples	332
15.13RANDNF Generate Noncentral F-Distribution Random Variable	332
15.13.1 Usage	332
15.13.2 Function Internals	332
15.13.3 Example	332
15.14RANDP Generate Poisson Random Variable	333
15.14.1 Usage	333
15.14.2 Function Internals	333
15.14.3 Example	333
15.15RANDPERM Random permutation	333
15.15.1 USAGE	333
15.15.2 Example	334
15.16SEED Seed the Random Number Generator	334
15.16.1 Usage	334
15.16.2 Example	334
16 Input/Ouput Functions	335
16.1 CSVREAD Read Comma Separated Value (CSV) File	335
16.1.1 Usage	335
16.1.2 Example	335
16.2 CSVWRITE Write Comma Separated Value (CSV) File	336
16.2.1 Usage	336
16.2.2 Example	337
16.3 DISP Display a Variable or Expression	337
16.3.1 Usage	337
16.3.2 Example	338
16.4 DLMREAD Read ASCII-delimited File	338
16.4.1 Usage	338
16.5 FCLOSE File Close Function	339
16.5.1 Usage	339
16.5.2 Example	339
16.6 FEOF End Of File Function	339
16.6.1 Usage	339
16.6.2 Example	339
16.7 FFLUSH Force File Flush	340
16.7.1 Usage	340
16.8 FGETLINE Read a String from a File	340
16.8.1 Usage	340
16.8.2 Example	340
16.9 FOPEN File Open Function	341
16.9.1 Usage	341
16.9.2 Examples	342
16.10FORMAT Control the Format of Matrix Display	343
16.10.1 Usage	343
16.10.2 Example	344
16.11FPRINTF Formated File Output Function (C-Style)	347

16.11.1 Usage	347
16.11.2 Examples	347
16.12FREAD File Read Function	347
16.12.1 Usage	347
16.12.2 Example	348
16.13FSCANF Formatted File Input Function (C-Style)	349
16.13.1 Usage	349
16.14FSEEK Seek File To A Given Position	349
16.14.1 Usage	349
16.14.2 Example	350
16.15FTELL File Position Function	350
16.15.1 Usage	350
16.15.2 Example	351
16.16FWRITE File Write Function	351
16.16.1 Usage	351
16.16.2 Example	352
16.17GETLINE Get a Line of Input from User	352
16.17.1 Usage	352
16.18GETPRINTLIMIT Get Limit For Printing Of Arrays	352
16.18.1 Usage	352
16.18.2 Example	352
16.19HTMLREAD Read an HTML Document into FreeMat	353
16.19.1 Usage	353
16.20IMREAD Read Image File To Matrix	353
16.20.1 Usage	353
16.21IMWRITE Write Matrix to Image File	354
16.21.1 Usage	354
16.21.2 Example	354
16.22INPUT Get Input From User	354
16.22.1 Usage	354
16.23LOAD Load Variables From A File	355
16.23.1 Usage	355
16.23.2 Example	356
16.24PAUSE Pause Script Execution	357
16.24.1 Usage	357
16.25PRINTF Formated Output Function (C-Style)	357
16.25.1 Usage	357
16.25.2 Format of the format string	358
16.25.3 The flag characters	358
16.25.4 The field width	358
16.25.5 The precision	359
16.25.6 The conversion specifier	359
16.25.7 Example	360
16.26RAWREAD Read N-dimensional Array From File	360
16.26.1 Usage	360
16.27RAWWRITE Write N-dimensional Array From File	361

16.27.1 Usage	361
16.28SAVE Save Variables To A File	361
16.28.1 Usage	361
16.28.2 Example	362
16.29SETPRINTLIMIT Set Limit For Printing Of Arrays	363
16.29.1 Usage	363
16.29.2 Example	363
16.30PRINTF Formated String Output Function (C-Style)	364
16.30.1 Usage	364
16.30.2 Examples	364
16.31SSCANF Formated String Input Function (C-Style)	364
16.31.1 Usage	364
16.32STR2NUM Convert a String to a Number	364
16.32.1 Usage	364
16.33URLWRITE Retrieve a URL into a File	365
16.33.1 Usage	365
16.34WAVPLAY	365
16.34.1 Usage	365
16.35WAVREAD Read a WAV Audio File	365
16.35.1 Usage	365
16.36WAVRECORD	366
16.36.1 Usage	366
16.37WAVWRITE Write a WAV Audio File	367
16.37.1 Usage	367
16.38XMLREAD Read an XML Document into FreeMat	367
16.38.1 Usage	367
17 String Functions	369
17.1 BLANKS Create a blank string	369
17.1.1 Usage	369
17.1.2 Example	369
17.2 CELLSTR Convert character array to cell array of strings	369
17.2.1 Usage	369
17.2.2 Example	370
17.3 DEBLANK Remove trailing blanks from a string	370
17.3.1 Usage	370
17.3.2 Example	370
17.4 ISALPHA Test for Alpha Characters in a String	371
17.4.1 Usage	371
17.4.2 Example	371
17.5 ISDIGIT Test for Digit Characters in a String	371
17.5.1 Usage	371
17.5.2 Example	371
17.6 ISSPACE Test for Space Characters in a String	371
17.6.1 Usage	371
17.6.2 Example	372

17.7 LOWER Convert strings to lower case	372
17.7.1 Usage	372
17.7.2 Example	372
17.8 REGEXP Regular Expression Matching Function	373
17.8.1 Usage	373
17.8.2 Example	374
17.9 REGEXPREP Regular Expression Replacement Function	375
17.9.1 Usage	375
17.10 STRCMP String Compare Function	375
17.10.1 USAGE	375
17.10.2 Example	375
17.11 STRCMPI String Compare Case Insensitive Function	376
17.11.1 Usage	376
17.12 STRFIND Find Substring in a String	376
17.12.1 Usage	376
17.12.2 Example	377
17.13 STRNCMP String Compare Function To Length N	377
17.13.1 USAGE	377
17.13.2 Example	377
17.14 STRREP String Replace Function	378
17.14.1 Usage	378
17.14.2 Example	379
17.15 STRSTR String Search Function	379
17.15.1 Usage	379
17.15.2 Example	379
17.16 STRTRIM Trim Spaces from a String	380
17.16.1 Usage	380
17.16.2 Example	380
17.17 UPPER Convert strings to upper case	380
17.17.1 Usage	380
17.17.2 Example	380
18 Transforms/Decompositions	383
18.1 EIG Eigendecomposition of a Matrix	383
18.1.1 Usage	383
18.1.2 Function Internals	384
18.1.3 Example	384
18.2 FFT (Inverse) Fast Fourier Transform Function	387
18.2.1 Usage	387
18.2.2 Function Internals	387
18.2.3 Example	388
18.3 FFTN N-Dimensional Forward FFT	389
18.3.1 Usage	389
18.4 FFTSHIFT Shift FFT Output	390
18.4.1 Usage	390
18.5 HILBERT Hilbert Transform	390

18.5.1 Usage	390
18.6 IFFTN N-Dimensional Inverse FFT	390
18.6.1 Usage	390
18.7 IFFTSHIFT Inverse Shift FFT Output	391
18.7.1 Usage	391
18.8 INV Invert Matrix	391
18.8.1 Usage	391
18.8.2 Example	391
18.9 LU LU Decomposition for Matrices	392
18.9.1 Usage	392
18.9.2 Example	392
18.10QR QR Decomposition of a Matrix	395
18.10.1 Usage	395
18.11SVD Singular Value Decomposition of a Matrix	396
18.11.1 Usage	396
18.11.2 Function Internals	396
18.11.3 Examples	397
19 Signal Processing Functions	399
19.1 CONV Convolution Function	399
19.1.1 Usage	399
19.2 CONV2 Matrix Convolution	399
19.2.1 Usage	399
19.2.2 Function Internals	400
20 Numerical Methods	401
20.1 ODE45 Numerical Solution of ODEs	401
20.1.1 Usage	401
21 Operating System Functions	415
21.1 CD Change Working Directory Function	415
21.1.1 Usage	415
21.1.2 Example	415
21.2 COPYFILE Copy Files	416
21.2.1 Usage	416
21.3 DELETE Delete a File	417
21.3.1 Usage	417
21.4 DIR List Files Function	417
21.4.1 Usage	417
21.5 DIRSEP Director Seperator	418
21.5.1 Usage	418
21.6 FILEPARTS Extract Filename Parts	418
21.6.1 Usage	418
21.7 FULLFILE Build a Full Filename From Pieces	418
21.7.1 Usage	418
21.7.2 Example	418

21.8	GETPATH Get Current Search Path	418
21.8.1	Usage	418
21.8.2	Example	419
21.9	LS List Files Function	419
21.9.1	Usage	419
21.9.2	Example	419
21.10	MKDIR Make Directory	420
21.10.1	Usage	420
21.11	PWD Print Working Directory Function	420
21.11.1	Usage	420
21.11.2	Example	420
21.12	RMDIR Remove Directory	420
21.12.1	Usage	420
21.13	SETPATH Set Current Search Path	421
21.13.1	Usage	421
21.14	SYSTEM Call an External Program	421
21.14.1	Usage	421
21.14.2	Example	421
22	Optimization and Curve Fitting	427
22.1	FITFUN Fit a Function	427
22.1.1	Usage	427
22.2	GAUSFIT Gaussian Curve Fit	428
22.2.1	Usage	428
22.2.2	Example	429
22.3	INTERPLIN1 Linear 1-D Interpolation	429
22.3.1	Usage	429
22.3.2	Example	430
22.4	POLY Convert Roots To Polynomial Coefficients	431
22.4.1	Usage	431
22.4.2	Example	431
22.5	POLYDER Polynomial Coefficient Differentiation	432
22.5.1	Usage	432
22.5.2	Example	432
22.6	POLYFIT Fit Polynomial To Data	433
22.6.1	Usage	433
22.6.2	Function Internals	433
22.6.3	Example	433
22.7	POLYINT Polynomial Coefficient Integration	435
22.7.1	Usage	435
22.7.2	Example	435
22.8	POLYVAL Evaluate Polynomial Fit at Selected Points	435
22.8.1	Usage	435
22.8.2	Function Internals	436
22.8.3	Example	436
22.9	ROOTS Find Roots of Polynomial	436

22.9.1 Usage	436
22.9.2 Function Internals	437
22.9.3 Example	437
23 Handle-Based Graphics	439
23.1 AXES Create Handle Axes	439
23.1.1 Usage	439
23.2 AXIS Setup Axis Behavior	439
23.2.1 Usage	439
23.3 AXISPROPERTIES Axis Object Properties	441
23.3.1 Usage	441
23.4 CLA Clear Current Axis	447
23.4.1 Usage	447
23.5 CLABEL Add Labels To Contour Plot	447
23.5.1 Usage	447
23.5.2 Example	447
23.6 CLF Clear Figure	448
23.6.1 Usage	448
23.7 CLIM Adjust Color limits of plot	448
23.7.1 Usage	448
23.7.2 Example	449
23.8 CLOSE Close Figure Window	450
23.8.1 Usage	450
23.9 COLORBAR Add Colorbar to Current Plot	450
23.9.1 Usage	450
23.10COLORMAP Image Colormap Function	451
23.10.1 Usage	451
23.10.2 Function Internals	451
23.10.3 Examples	451
23.11COLORSPEC Color Property Description	454
23.11.1 Usage	454
23.12CONTOUR Contour Plot Function	455
23.12.1 Usage	455
23.12.2 Example	456
23.13CONTOUR3 3D Contour Plot Function	457
23.13.1 Usage	457
23.13.2 Example	457
23.14COPPER Copper Colormap	458
23.14.1 Usage	458
23.14.2 Example	458
23.15COPY Copy Figure Window	459
23.15.1 Usage	459
23.16COUNTOUR Contour Object Properties	459
23.16.1 Usage	459
23.17DRAWNOW Flush the Event Queue	460
23.17.1 Usage	460

23.18FIGLOWER Lower a Figure Window	460
23.18.1 Usage	460
23.19FIGRAISE Raise a Figure Window	461
23.19.1 Usage	461
23.20FIGURE Figure Window Select and Create Function	461
23.20.1 Usage	461
23.21FIGUREPROPERTIES Figure Object Properties	461
23.21.1 Usage	461
23.22GCA Get Current Axis	462
23.22.1 Usage	462
23.23GCF Get Current Figure	462
23.23.1 Usage	462
23.24GET Get Object Property	463
23.24.1 Usage	463
23.25GLSHOW Show a GL Assembly in a GL Window	463
23.25.1 Usage	463
23.26GRAY Gray Colormap	463
23.26.1 Usage	463
23.26.2 Example	463
23.27GRID Plot Grid Toggle Function	464
23.27.1 Usage	464
23.27.2 Example	464
23.28HCONTOUR Create a contour object	465
23.28.1 Usage	465
23.29HIMAGE Create a image object	465
23.29.1 Usage	465
23.30HIST Histogram Function	466
23.30.1 Usage	466
23.31HLINE Create a line object	466
23.31.1 Usage	466
23.32HOLD Plot Hold Toggle Function	466
23.32.1 Usage	466
23.32.2 Function Internals	467
23.32.3 Example	467
23.33HPATCH Create a patch object	468
23.33.1 Usage	468
23.34HPOINT Get Point From Window	468
23.34.1 Usage	468
23.35HRAWPLOT Generate a Raw Plot File	468
23.35.1 Usage	468
23.36HSURFACE Create a surface object	469
23.36.1 Usage	469
23.37HTEXT Create a text object	469
23.37.1 Usage	469
23.38HTEXTBITMAP Get Text Rendered as a Bitmap	469
23.38.1 Usage	469

23.39IMAGE Image Display Function	469
23.39.1 Usage	469
23.39.2 Example	470
23.40IMAGEPROPERTIES Image Object Properties	471
23.40.1 Usage	471
23.41IMAGESC Image Display Function	472
23.41.1 Usage	472
23.41.2 Example	472
23.42IS2DVIEW Test Axes For 2D View	473
23.42.1 Usage	473
23.43ISHOLD Test Hold Status	473
23.43.1 Usage	473
23.44LEGEND Add Legent to Plot	473
23.44.1 Usage	473
23.45LINE Line Display Function	474
23.45.1 Usage	474
23.46LINEPROPERTIES Line Series Object Properties	475
23.46.1 Usage	475
23.47LOGLOG Log-Log Plot Function	476
23.47.1 Usage	476
23.47.2 Example	476
23.48NEWPLOT Get Handle For Next Plot	476
23.48.1 Usage	476
23.49PATCH Patch Graphics Function	477
23.49.1 Usage	477
23.49.2 Example	477
23.50PCOLOR Pseudocolor Plot	478
23.50.1 Usage	478
23.51PLOT Plot Function	478
23.51.1 Usage	478
23.51.2 Example	480
23.52PLOT3 Plot 3D Function	482
23.52.1 Usage	482
23.52.2 Example	482
23.53POINT Get Axis Position From Mouse Click	483
23.53.1 Usage	483
23.54PRINT Print a Figure To A File	483
23.54.1 Usage	483
23.54.2 Example	484
23.55PVALID Validate Property Name	484
23.55.1 Usage	484
23.55.2 Example	485
23.56SEMILOGX Semilog X Axis Plot Function	485
23.56.1 Usage	485
23.56.2 Example	485
23.57SEMILOGY Semilog Y Axis Plot Function	486

23.57.1 Usage	486
23.57.2 Example	486
23.58SET Set Object Property	487
23.58.1 Usage	487
23.59SIZEFIG Set Size of Figure	487
23.59.1 Usage	487
23.60SUBPLOT Subplot Function	488
23.60.1 Usage	488
23.60.2 Example	488
23.61SURF Surface Plot Function	490
23.61.1 Usage	490
23.61.2 Example	490
23.62SURFACEPROPERTIES Surface Object Properties	491
23.62.1 Usage	491
23.63SURFACEPROPERTIES Surface Object Properties	494
23.63.1 Usage	494
23.64TEXT Add Text Label to Plot	496
23.64.1 Usage	496
23.64.2 Example	496
23.65TEXTPROPERTIES Text Object Properties	497
23.65.1 Usage	497
23.66TITLE Plot Title Function	498
23.66.1 Usage	498
23.66.2 Example	499
23.67TUBEPLT Creates a Tubeplot	500
23.67.1 Usage	500
23.67.2 Example	500
23.68UICONTROL Create a UI Control object	501
23.68.1 Usage	501
23.69UICONTROLPROPERTIES UI Control Properties	501
23.69.1 Usage	501
23.70VIEW Set Graphical View	504
23.70.1 Usage	504
23.70.2 Example	504
23.71WINLEV Image Window-Level Function	505
23.71.1 Usage	505
23.71.2 Function Internals	506
23.71.3 Examples	506
23.72XLABEL Plot X-axis Label Function	507
23.72.1 Usage	507
23.72.2 Example	507
23.73XLIM Adjust X Axis limits of plot	508
23.73.1 Usage	508
23.73.2 Example	508
23.74YLABEL Plot Y-axis Label Function	509
23.74.1 Usage	509

23.74.2 Example	509
23.75YLIM Adjust Y Axis limits of plot	510
23.75.1 Usage	510
23.75.2 Example	510
23.76ZLABEL Plot Z-axis Label Function	511
23.76.1 Usage	511
23.76.2 Example	512
23.77ZLIM Adjust Z Axis limits of plot	512
23.77.1 Usage	512
23.77.2 Example	513
23.78ZOOM Image Zoom Function	514
23.78.1 Usage	514
23.78.2 Example	514
23.79ZPLANE Zero-pole plot	517
23.79.1 Usage	517
24 OpenGL Models	519
24.1 GLASSEMBLY Create a GL Assembly	519
24.1.1 Usage	519
24.2 GLCLUMP Create a GL Clump	519
24.2.1 Usage	519
24.3 GLDEFMATERIAL Defines a GL Material	520
24.3.1 Usage	520
24.4 GLLINES Create a GL Lineset	520
24.4.1 Usage	520
24.5 GLNODE Create a GL Node	520
24.5.1 Usage	520
25 Object Oriented Programming	521
25.1 AND Overloaded Logical And Operator	521
25.1.1 Usage	521
25.2 CAT Concatenation of Arrays	521
25.2.1 Usage	521
25.3 CLASS Class Support Function	521
25.3.1 Usage	521
25.4 COLON Overloaded Colon Operator	522
25.4.1 Usage	522
25.5 CONSTRUCTORS Class Constructors	522
25.5.1 Usage	522
25.6 CTRANSPOSE Overloaded Conjugate Transpose Operator	523
25.6.1 Usage	523
25.7 EQ Overloaded Equals Comparison Operator	523
25.7.1 Usage	523
25.8 GE Overloaded Greater-Than-Equals Comparison Operator	523
25.8.1 Usage	523
25.9 GT Overloaded Greater Than Comparison Operator	524

25.9.1 Usage	524
25.10HORIZCAT Horizontal Array Concatenation	524
25.10.1 Usage	524
25.11HORIZCAT Horizontal Array Concatenation	524
25.11.1 Usage	524
25.12LDIVIDE Overloaded Left Divide Operator	524
25.12.1 Usage	524
25.13LE Overloaded Less-Than-Equals Comparison Operator	525
25.13.1 Usage	525
25.14LT Overloaded Less Than Comparison Operator	525
25.14.1 Usage	525
25.15MINUS Overloaded Addition Operator	525
25.15.1 Usage	525
25.16MLDIVIDE Overloaded Matrix Left Divide Operator	525
25.16.1 Usage	525
25.17MPOWER Overloaded Matrix Power Operator	526
25.17.1 Usage	526
25.18MRDIVIDE Overloaded Matrix Right Divide Operator	526
25.18.1 Usage	526
25.19MTIMES Overloaded Matrix Multiplication Operator	526
25.19.1 Usage	526
25.20NE Overloaded Not-Equals Comparison Operator	526
25.20.1 Usage	526
25.21NOT Overloaded Logical Not Operator	527
25.21.1 Usage	527
25.22OR Overloaded Logical Or Operator	527
25.22.1 Usage	527
25.23PLUS Overloaded Addition Operator	527
25.23.1 Usage	527
25.24POWER Overloaded Power Operator	527
25.24.1 Usage	527
25.25RDIVIDE Overloaded Right Divide Operator	528
25.25.1 Usage	528
25.26SUBSASGN Overloaded Class Assignment	528
25.26.1 Usage	528
25.27SUBSINDEX Overloaded Class Indexing	528
25.27.1 Usage	528
25.28SUBSREF Overloaded Class Indexing	529
25.28.1 Usage	529
25.29TIMES Overloaded Multiplication Operator	529
25.29.1 Usage	529
25.30TRANPOSE Overloaded Transpose Operator	529
25.30.1 Usage	529
25.31UMINUS Overloaded Unary Minus Operator	530
25.31.1 Usage	530
25.32UPLUS Overloaded Unary Plus Operator	530

25.32.1 Usage	530
25.33VERTCAT Horizontal Array Concatenation	530
25.33.1 Usage	530
25.34VERTCAT Horizontal Array Concatenation	530
25.34.1 Usage	530
26 Bitwise Operations	531
26.1 BITAND Bitwise Boolean And Operation	531
26.1.1 Usage	531
26.1.2 Example	531
26.2 BITCMP Bitwise Boolean Complement Operation	531
26.2.1 Usage	531
26.2.2 Example	532
26.3 BITOR Bitwise Boolean Or Operation	532
26.3.1 Usage	532
26.3.2 Example	532
26.4 BITXOR Bitwise Boolean Exclusive-Or (XOR) Operation	533
26.4.1 Usage	533
26.4.2 Example	533
27 FreeMat Threads	535
27.1 THREADCALL Call Function In A Thread	535
27.1.1 Usage	535
27.1.2 Example	535
27.2 THREADFREE Free thread resources	536
27.2.1 Usage	536
27.3 THREADID Get Current Thread Handle	536
27.3.1 Usage	536
27.3.2 Example	536
27.4 THREADKILL Halt execution of a thread	537
27.4.1 Usage	537
27.4.2 Example	537
27.5 THREADNEW Create a New Thread	538
27.5.1 Usage	538
27.6 THREADSTART Start a New Thread Computation	538
27.6.1 Usage	538
27.6.2 Example	539
27.7 THREADVALUE Retrieve the return values from a thread	541
27.7.1 Usage	541
27.7.2 Example	541
27.8 THREADWAIT Wait on a thread to complete execution	542
27.8.1 Usage	542
27.8.2 Example	542

28 Function Related Functions	545
28.1 INLINE Construct Inline Function	545
28.1.1 Usage	545
28.1.2 Example	545
28.2 SYMVAR Find Symbolic Variables in an Expression	546
28.2.1 Usage	546
28.2.2 Example	547
29 FreeMat External Interface	549
29.1 CENUM Lookup Enumerated C Type	549
29.1.1 Usage	549
29.2 CTYPECAST Cast FreeMat Structure to C Structure	549
29.2.1 Usage	549
29.3 CTYPEDEFINE Define C Type	550
29.3.1 Usage	550
29.4 CTYPEFREEZE Convert FreeMat Structure to C Type	550
29.4.1 Usage	550
29.5 CTYPENEW Create New Instance of C Structure	551
29.5.1 Usage	551
29.6 CTYPEPRINT Print C Type	551
29.6.1 Usage	551
29.7 CTYPEPREAD Read a C Structure From File	551
29.7.1 Usage	551
29.8 CYPESIZE Compute Size of C Struct	552
29.8.1 Usage	552
29.9 CYPEHAW Convert C Struct to FreeMat Structure	552
29.9.1 Usage	552
29.10 CTYPEWRITE Write a C Typedef To File	553
29.10.1 Usage	553
29.11 IMPORT Foreign Function Import	553
29.11.1 Usage	553
29.11.2 Example	555
29.12 LOADLIB Load Library Function	556
29.12.1 Usage	556

Chapter 1

Introduction and Getting Started

1.1 INSTALL Installing FreeMat

1.1.1 General Instructions

Here are the general instructions for installing FreeMat. First, follow the instructions listed below for the platform of interest. Then, run the

```
-->pathtool
```

which brings up the path setup tool. More documentation on the GUI elements (and how to use them) will be forthcoming.

1.1.2 Linux

For Linux, FreeMat is now provided as a binary installation. To install it simply download the binary using your web browser, and then unpack it

```
tar xvfz FreeMat-<VERSION_NUMBER>-Linux-Binary.tar.gz
```

You can then run FreeMat directly without any additional effort

```
FreeMat-<VERSION_NUMBER>-Linux-Binary/Contents/bin/FreeMat
```

will start up FreeMat as an X application. If you want to run it as a command line application (to run from within an xterm), use the `nogui` flag

```
FreeMat-<VERSION_NUMBER>-Linux-Binary/Contents/bin/FreeMat -nogui
```

If you do not want FreeMat to use X at all (no graphics at all), use the `noX` flag

```
FreeMat-<VERSION_NUMBER>-Linux-Binary/Contents/bin/FreeMat -noX
```

For convenience, you may want to add FreeMat to your path. The exact mechanism for doing this depends on your shell. Assume that you have unpacked `FreeMat-<VERSION_NUMBER>-Linux-Binary.tar.gz` into the directory `/home/myname`. Then if you use `csh` or its derivatives (like `tcsh`) you should add the following line to your `.cshrc` file:

```
set path=($path /home/myname/FreeMat-<VERSION_NUMBER>-Linux/Binary/Contents/bin)
```

If you use `bash`, then add the following line to your `.bash_profile`

```
PATH=$PATH:/home/myname/FreeMat-<VERSION_NUMBER>-Linux/Binary/Contents/bin
```

If the prebuilt binary package does not work for your Linux distribution, you will need to build FreeMat from source (see the source section below). When you have FreeMat running, you can setup your path using the `pathtool`. Note that the `FREEMAT_PATH` is no longer used by FreeMat. You must use the `pathtool` to adjust the path.

1.1.3 Windows

For Windows, FreeMat is installed via a binary installer program. To use it, simply download the setup program `FreeMat-<VERSION_NUMBER>-Setup.exe`, and double click it. Follow the instructions to do the installation, then setup your path using `pathtool`.

1.1.4 Mac OS X

For Mac OS X, FreeMat is distributed as an application bundle. To install it, simply download the compressed disk image file `FreeMat-<VERSION_NUMBER>.dmg`, double click to mount the disk image, and then copy the application `FreeMat-<VERSION_NUMBER>` to some convenient place. To run FreeMat, simply double click on the application. Run `pathtool` to setup your FreeMat path.

1.1.5 Source Code

The source code build is a little more complicated than previous versions of FreeMat. Here are the current build instructions for all platforms.

1. Build and install Qt 4.3 or later - <http://trolltech.com/developer/downloads/opensource>
2. Install g77 or gfortran (use fink for Mac OS X, use gcc-g77 package for MinGW)
3. Download the source code `FreeMat-<VERSION_NUMBER>-src.tar.gz`.
4. Unpack the source code: `tar xvfz FreeMat-<VERSION_NUMBER>-src.tar.gz`.
5. For Windows, you will need to install MSYS as well as MINGW to build FreeMat. You will also need unzip to unpack the enclosed `matio.zip` archive. Alternately, you can cross-build the Windows version of FreeMat under Linux (this is how I build it now).
6. If you are extraordinarily lucky (or prepared), you can issue the usual `./configure`, then the `make` and `make install`. This is not likely to work because of the somewhat esoteric dependencies of FreeMat. The `configure` step will probably fail and indicate what external dependencies are still needed.
7. I assume that you are familiar with the process of installing dependencies if you are trying to build FreeMat from source.

To build a binary distributable (app bundle on the Mac, setup installer on win32, and a binary distribution on Linux), you will need to run `make package` instead of `make install`.

Chapter 2

Variables and Arrays

2.1 CELL Cell Array Definitions

2.1.1 Usage

The cell array is a fairly powerful array type that is available in FreeMat. Generally speaking, a cell array is a heterogenous array type, meaning that different elements in the array can contain variables of different type (including other cell arrays). For those of you familiar with C, it is the equivalent to the `void *` array. The general syntax for their construction is

```
A = {row_def1;row_def2;...;row_defN}
```

where each row consists of one or more elements, seperated by commas

```
row_defi = element_i1,element_i2,...,element_iM
```

Each element can be any type of FreeMat variable, including matrices, arrays, cell-arrays, structures, strings, etc. The restriction on the definition is that each row must have the same number of elements in it.

2.1.2 Examples

Here is an example of a cell-array that contains a number, a string, and an array

```
--> A = {14,'hello',[1:10]}
```

```
A =  
[14] [hello] [1 10 double array]
```

Note that in the output, the number and string are explicitly printed, but the array is summarized. We can create a 2-dimensional cell-array by adding another row definition

```
--> B = {pi,i;e,-1}
```

```
B =
```

```
[3.14159] [0+1i]
[2.71828] [-1]
```

Finally, we create a new cell array by placing A and B together

```
--> C = {A,B}
```

```
C =
[1 3 cell array] [2 2 cell array]
```

2.2 Function Handles

2.2.1 Usage

Starting with version 1.11, FreeMat now supports **function handles**, or **function pointers**. A **function handle** is an alias for a function or script that is stored in a variable. First, the way to assign a function handle is to use the notation

```
handle = @func
```

where **func** is the name to point to. The function **func** must exist at the time we make the call. It can be a local function (i.e., a subfunction). To use the **handle**, we can either pass it to **feval** via

```
[x,y] = feval(handle,arg1,arg2).
```

Alternately, you can the function directly using the notation

```
[x,y] = handle(arg1,arg2)
```

2.3 GLOBAL Global Variables

2.3.1 Usage

Global variables are shared variables that can be seen and modified from any function or script that declares them. The syntax for the **global** statement is

```
global variable_1 variable_2 ...
```

The **global** statement must occur before the variables appear.

2.3.2 Example

Here is an example of two functions that use a global variable to communicate an array between them. The first function sets the global variable.

```
set_global.m
function set_global(x)
    global common_array
    common_array = x;
```

The second function retrieves the value from the global variable

```

    get_global.m
function x = get_global
    global common_array
    x = common_array;

```

Here we exercise the two functions

```

--> set_global('Hello')
--> get_global

```

```

ans =
Hello

```

2.4 INDEXING Indexing Expressions

2.4.1 Usage

There are three classes of indexing expressions available in FreeMat: `()`, `{}`, and `.`. Each is explained below in some detail, and with its own example section.

2.4.2 Array Indexing

We start with array indexing `()`, which is the most general indexing expression, and can be used on any array. There are two general forms for the indexing expression - the N-dimensional form, for which the general syntax is

```
variable(index_1,index_2,...,index_n)
```

and the vector form, for which the general syntax is

```
variable(index)
```

Here each index expression is either a scalar, a range of integer values, or the special token `:`, which is shorthand for `1:end`. The keyword `end`, when included in an indexing expression, is assigned the length of the array in that dimension. The concept is easier to demonstrate than explain. Consider the following examples:

```
--> A = zeros(4)
```

```

A =
0 0 0 0
0 0 0 0
0 0 0 0
0 0 0 0

```

```
--> B = float(randn(2))
```

```
B =
    0.6869   -0.0908
    1.2610   -1.6104
```

```
--> A(2:3,2:3) = B
```

```
A =
    0         0         0         0
    0   0.6869   -0.0908   0
    0   1.2610   -1.6104   0
    0         0         0         0
```

Here the array indexing was used on the left hand side only. It can also be used for right hand side indexing, as in

```
--> C = A(2:3,1:end)
```

```
C =
    0   0.6869   -0.0908   0
    0   1.2610   -1.6104   0
```

Note that we used the `end` keyword to avoid having to know that `A` has 4 columns. Of course, we could also use the `:` token instead:

```
--> C = A(2:3,:)
```

```
C =
    0   0.6869   -0.0908   0
    0   1.2610   -1.6104   0
```

An extremely useful example of `:` with array indexing is for slicing. Suppose we have a 3-D array, that is `2 x 2 x 3`, and we want to set the middle slice:

```
--> D = zeros(2,2,3)
```

```
D =

(:, :, 1) =
    0  0
    0  0

(:, :, 2) =
    0  0
    0  0

(:, :, 3) =
    0  0
```

```

0 0
--> D(:,:,2) = int32(10*rand(2,2))
D =
(:,:,1) =
 0 0
 0 0
(:,:,2) =
 5 4
 2 6
(:,:,3) =
 0 0
 0 0

```

In another level of nuance, the assignment expression will automatically fill in the indexed rectangle on the left using data from the right hand side, as long as the lengths match. So we can take a vector and roll it into a matrix using this approach:

```

--> A = zeros(4)

A =
 0 0 0 0
 0 0 0 0
 0 0 0 0
 0 0 0 0

--> v = [1;2;3;4]

v =
 1
 2
 3
 4

--> A(2:3,2:3) = v

A =
 0 0 0 0
 0 1 3 0
 0 2 4 0
 0 0 0 0

```

The N-dimensional form of the variable index is limited to accessing only (hyper-) rectangular regions of the array. You cannot, for example, use it to access only the diagonal elements of the

array. To do that, you use the second form of the array access (or a loop). The vector form treats an arbitrary N-dimensional array as though it were a column vector. You can then access arbitrary subsets of the arrays elements (for example, through a `find` expression) efficiently. Note that in vector form, the `end` keyword takes the meaning of the total length of the array (defined as the product of its dimensions), as opposed to the size along the first dimension.

2.4.3 Cell Indexing

The second form of indexing operates, to a large extent, in the same manner as the array indexing, but it is by no means interchangeable. As the name implies, `cell`-indexing applies only to `cell` arrays. For those familiar with C, `cell`-indexing is equivalent to pointer dereferencing in C. First, the syntax:

```
variable{index_1,index_2,...,index_n}
```

and the vector form, for which the general syntax is

```
variable{index}
```

The rules and interpretation for N-dimensional and vector indexing are identical to `()`, so we will describe only the differences. In simple terms, applying `()` to a `cell`-array returns another `cell` array that is a subset of the original array. On the other hand, applying `{}` to a `cell`-array returns the contents of that `cell` array. A simple example makes the difference quite clear:

```
--> A = {1, 'hello', [1:4]}
```

```
A =
```

```
[1] [hello] [1 4 double array]
```

```
--> A(1:2)
```

```
ans =
```

```
[1] [hello]
```

```
--> A{1:2}
```

```
ans =
```

```
1 of 2:
```

```
1
```

```
2 of 2:
```

```
hello
```

You may be surprised by the response to the last line. The output is multiple assignments to `ans!`. The output of a `cell`-array dereference can be used anywhere a list of expressions is required. This includes arguments and returns for function calls, matrix construction, etc. Here is an example of using `cell`-arrays to pass parameters to a function:

```
--> A = {[1,3,0],[5,2,7]}

A =
  [1 3 double array] [1 3 double array]

--> max(A{1:end})

ans =
  5 3 7
```

And here, cell-arrays are used to capture the return.

```
--> [K{1:2}] = max(randn(1,4))
K =
  [1.06943] [4]
```

Here, cell-arrays are used in the matrix construction process:

```
--> C = [A{1};A{2}]

C =
  1 3 0
  5 2 7
```

Note that this form of indexing is used to implement variable length arguments to function. See `varargin` and `varargout` for more details.

2.4.4 Structure Indexing

The third form of indexing is structure indexing. It can only be applied to structure arrays, and has the general syntax

```
variable.fieldname
```

where `fieldname` is one of the fields on the structure. Note that in FreeMat, fields are allocated dynamically, so if you reference a field that does not exist in an assignment, it is created automatically for you. If variable is an array, then the result of the `.` reference is an expression list, exactly like the `{}` operator. Hence, we can use structure indexing in a simple fashion:

```
--> clear A
--> A.color = 'blue'

A =
  color: blue
--> B = A.color

B =
  blue
```

Or in more complicated ways using expression lists for function arguments

```

--> clear A
--> A(1).maxargs = [1,6,7,3]

A =
    maxargs: 1 4 double array
--> A(2).maxargs = [5,2,9,0]

A =
    Fields
    maxargs
--> max(A.maxargs)

ans =
    5 6 9 3

```

or to store function outputs

```

--> clear A
--> A(1).maxreturn = [];
--> A(2).maxreturn = [];
--> [A.maxreturn] = max(randn(1,4))
A =
    Fields
    maxreturn

```

FreeMat now also supports the so called dynamic-field indexing expressions. In this mode, the fieldname is supplied through an expression instead of being explicitly provided. For example, suppose we have a set of structure indexed by color,

```

--> x.red = 430;
--> x.green = 240;
--> x.blue = 53;
--> x.yello = 105

x =
    red: 430
    green: 240
    blue: 53
    yello: 105

```

Then we can index into the structure `x` using a dynamic field reference:

```

--> y = 'green'

y =
    green
--> a = x.(y)

```

```
a =
  240
```

Note that the indexing expression has to resolve to a string for dynamic field indexing to work.

2.4.5 Complex Indexing

The indexing expressions described above can be freely combined to affect complicated indexing expressions. Here is an example that exercises all three indexing expressions in one assignment.

```
--> Z{3}.foo(2) = pi
```

```
Z =
 [0] [0] [1 1 struct array]
```

From this statement, FreeMat infers that Z is a cell-array of length 3, that the third element is a structure array (with one element), and that this structure array contains a field named 'foo' with two double elements, the second of which is assigned a value of pi.

2.5 MATRIX Matrix Definitions

2.5.1 Usage

The matrix is the basic datatype of FreeMat. Matrices can be defined using the following syntax

```
A = [row_def1;row_def2;... ,row_defN]
```

where each row consists of one or more elements, seperated by commas

```
row_defi = element_i1,element_i2,...,element_iM
```

Each element can either be a scalar value or another matrix, provided that the resulting matrix definition makes sense. In general this means that all of the elements belonging to a row have the same number of rows themselves, and that all of the row definitions have the same number of columns. Matrices are actually special cases of N-dimensional arrays where $N \leq 2$. Higher dimensional arrays cannot be constructed using the bracket notation described above. The type of a matrix defined in this way (using the bracket notation) is determined by examining the types of the elements. The resulting type is chosen so no information is lost on any of the elements (or equivalently, by choosing the highest order type from those present in the elements).

2.5.2 Examples

Here is an example of a matrix of `int32` elements (note that untyped integer constants default to type `int32`).

```
--> A = [1,2;5,8]
```

```
A =
  1 2
  5 8
```

Now we define a new matrix by adding a column to the right of A, and using float constants.

```
--> B = [A,[3.2f;5.1f]]
```

```
B =
  1.0000    2.0000    3.2000
  5.0000    8.0000    5.1000
```

Next, we add extend B by adding a row at the bottom. Note how the use of an untyped floating point constant forces the result to be of type `double`

```
--> C = [B;5.2,1.0,0.0]
```

```
C =
  1.0000    2.0000    3.2000
  5.0000    8.0000    5.1000
  5.2000    1.0000         0
```

If we instead add a row of `complex` values (recall that `i` is a `complex` constant, not a `dcomplex` constant)

```
--> D = [B;2.0f+3.0f*i,i,0.0f]
```

```
D =
  1.0000 + 0.0000i    2.0000 + 0.0000i    3.2000 + 0.0000i
  5.0000 + 0.0000i    8.0000 + 0.0000i    5.1000 + 0.0000i
  2.0000 + 3.0000i    0.0000 + 1.0000i         0
```

Likewise, but using `dcomplex` constants

```
--> E = [B;2.0+3.0*i,i,0.0]
```

```
E =
  1.0000 + 0.0000i    2.0000 + 0.0000i    3.2000 + 0.0000i
  5.0000 + 0.0000i    8.0000 + 0.0000i    5.1000 + 0.0000i
  2.0000 + 3.0000i    0.0000 + 1.0000i         0
```

Finally, in FreeMat, you can construct matrices with strings as contents, but you have to make sure that if the matrix has more than one row, that all the strings have the same length.

```
--> F = ['hello';'there']
```

```
F =
hello
there
```

2.6 PERSISTENT Persistent Variables

2.6.1 Usage

Persistent variables are variables whose value persists between calls to a function or script. The general syntax for its use is

```
persistent variable1 variable2 ... variableN
```

The `persistent` statement must occur before the variable is tagged as persistent. Per the MATLAB API documentation an empty variable is created when the `persistent` statement is called.

2.6.2 Example

Here is an example of a function that counts how many times it has been called.

```
count_calls.m
function count_calls
persistent ccount
if (~exist('ccount')) ccount = 0; end;
ccount = ccount + 1;
printf('Function has been called %d times\n',ccount);
```

We now call the function several times:

```
--> for i=1:10; count_calls; end
Function has been called 0 times
```

2.7 STRUCT Structure Array Constructor

2.7.1 Usage

Creates an array of structures from a set of field, value pairs. The syntax is

```
y = struct(n_1,v_1,n_2,v_2,...)
```

where `n_i` are the names of the fields in the structure array, and `v_i` are the values. The values `v_i` must either all be scalars, or be cell-arrays of all the same dimensions. In the latter case, the output structure array will have dimensions dictated by this common size. Scalar entries for the `v_i` are

replicated to fill out their dimensions. An error is raised if the inputs are not properly matched (i.e., are not pairs of field names and values), or if the size of any two non-scalar values cell-arrays are different.

Another use of the `struct` function is to convert a class into a structure. This allows you to access the members of the class, directly but removes the class information from the object.

2.7.2 Example

This example creates a 3-element structure array with three fields, `foo`, `bar` and `key`, where the contents of `foo` and `bar` are provided explicitly as cell arrays of the same size, and the contents of `bar` are replicated from a scalar.

```
--> y = struct('foo',{1,3,4},'bar',{'cheese','cola','beer'},'key',508)
```

```
y =
  Fields
    foo
    bar
    key
--> y(1)
```

```
ans =
  foo: 1
  bar: cheese
  key: 508
--> y(2)
```

```
ans =
  foo: 3
  bar: cola
  key: 508
--> y(3)
```

```
ans =
  foo: 4
  bar: beer
  key: 508
```

An alternate way to create a structure array is to initialize the last element of each field of the structure

```
--> Test(2,3).Type = 'Beer';
--> Test(2,3).Ounces = 12;
--> Test(2,3).Container = 'Can';
--> Test(2,3)
```

```
ans =
```

```
Type: Beer
Ounces: 12
Container: Can
--> Test(1,1)

ans =
Type: 0
Ounces: 0
Container: 0
```


Chapter 3

Functions and Scripts

3.1 ANONYMOUS Anonymous Functions

3.1.1 Usage

Anonymous functions are simple, nameless functions that can be defined anywhere (in a script, function, or at the prompt). They are intended to supplant `inline` functions. The syntax for an anonymous function is simple:

```
y = @(arg1,arg2,...,argn) expression
```

where `arg1,arg2,...,argn` is a list of valid identifiers that define the arguments to the function, and `expression` is the expression to compute in the function. The returned value `y` is a function handle for the anonymous function that can then be used to evaluate the expression. Note that `y` will capture the value of variables that are not indicated in the argument list from the current scope or workspace at the time it is defined. So, for example, consider the simple anonymous function definition

```
y = @(x) a*(x+b)
```

In order for this definition to work, the variables `a` and `b` need to be defined in the current workspace. Whatever value they have is captured in the function handle `y`. To change the values of `a` and `b` in the anonymous function, you must recreate the handle using another call. See the examples section for more information. In order to use the anonymous function, you can use it just like any other function handle. For example,

```
p = y(3)
p = y()
p = feval(y,3)
```

are all examples of using the `y` anonymous function to perform a calculation.

3.1.2 Examples

Here are some examples of using an anonymous function

```

--> a = 2; b = 4;    % define a and b (slope and intercept)
--> y = @(x) a*x+b  % create the anonymous function

y =
    @(x)  a*x+b    % create the anonymous function
--> y(2)           % evaluate it for x = 2

ans =
     8

--> a = 5; b = 7;    % change a and b
--> y(2)           % the value did not change!  because a=2,b=4 are captured in y

ans =
     8

--> y = @(x) a*x+b  % recreate the function

y =
    @(x)  a*x+b    % recreate the function
--> y(2)           % now the new values are used

ans =
    17

```

3.2 FUNCTION Function Declarations

3.2.1 Usage

There are several forms for function declarations in FreeMat. The most general syntax for a function declaration is the following:

```
function [out_1,...,out_M,varargout] = fname(in_1,...,in_N,varargin)
```

where `out_i` are the output parameters, `in_i` are the input parameters, and `varargout` and `varargin` are special keywords used for functions that have variable inputs or outputs. For functions with a fixed number of input or output parameters, the syntax is somewhat simpler:

```
function [out_1,...,out_M] = fname(in_1,...,in_N)
```

Note that functions that have no return arguments can omit the return argument list (of `out_i`) and the equals sign:

```
function fname(in_1,...,in_N)
```

Likewise, a function with no arguments can eliminate the list of parameters in the declaration:

```
function [out_1,...,out_M] = fname
```

Functions that return only a single value can omit the brackets

```
function out_1 = fname(in_1,...,in_N)
```

In the body of the function `in_i` are initialized with the values passed when the function is called. Also, the function must assign values for `out_i` to pass values to the caller. Note that by default, FreeMat passes arguments by value, meaning that if we modify the contents of `in_i` inside the function, it has no effect on any variables used by the caller. Arguments can be passed by reference by prepending an ampersand `&` before the name of the input, e.g.

```
function [out1,...,out_M] = fname(in_1,&in_2,in_3,...,in_N)
```

in which case `in_2` is passed by reference and not by value. Also, FreeMat works like C in that the caller does not have to supply the full list of arguments. Also, when `keywords` (see help `keywords`) are used, an arbitrary subset of the parameters may be unspecified. To assist in deciphering the exact parameters that were passed, FreeMat also defines two variables inside the function context: `nargin` and `nargout`, which provide the number of input and output parameters of the caller, respectively. See help for `nargin` and `nargout` for more details. In some circumstances, it is necessary to have functions that take a variable number of arguments, or that return a variable number of results. In these cases, the last argument to the parameter list is the special argument `varargin`. Inside the function, `varargin` is a cell-array that contains all arguments passed to the function that have not already been accounted for. Similarly, the function can create a cell array named `varargout` for variable length output lists. See help `varargin` and `varargout` for more details.

The function name `fname` can be any legal FreeMat identifier. Functions are stored in files with the `.m` extension. Note that the name of the file (and not the function name `fname` used in the declaration) is how the function appears in FreeMat. So, for example, if the file is named `foo.m`, but the declaration uses `bar` for the name of the function, in FreeMat, it will still appear as function `foo`. Note that this is only true for the first function that appears in a `.m` file. Additional functions that appear after the first function are known as `helper functions` or `local functions`. These are functions that can only be called by other functions in the same `.m` file. Furthermore the names of these helper functions are determined by their declaration and not by the name of the `.m` file. An example of using helper functions is included in the examples.

Another important feature of functions, as opposed to, say `scripts`, is that they have their own `scope`. That means that variables defined or modified inside a function do not affect the scope of the caller. That means that a function can freely define and use variables without unintentionally using a variable name reserved elsewhere. The flip side of this fact is that functions are harder to debug than scripts without using the `keyboard` function, because the intermediate calculations used in the function are not available once the function exits.

3.2.2 Examples

Here is an example of a trivial function that adds its first argument to twice its second argument:

```
addtest.m
function c = addtest(a,b)
c = a + 2*b;
```

```
--> addtest(1,3)
```

```
ans =  
7
```

```
--> addtest(3,0)
```

```
ans =  
3
```

Suppose, however, we want to replace the value of the first argument by the computed sum. A first attempt at doing so has no effect:

```
addtest2.m  
function addtest2(a,b)  
    a = a + 2*b;
```

```
--> arg1 = 1
```

```
arg1 =  
1
```

```
--> arg2 = 3
```

```
arg2 =  
3
```

```
--> addtest2(arg1,arg2)
```

```
--> arg1
```

```
ans =  
1
```

```
--> arg2
```

```
ans =  
3
```

The values of `arg1` and `arg2` are unchanged, because they are passed by value, so that any changes to `a` and `b` inside the function do not affect `arg1` and `arg2`. We can change that by passing the first argument by reference:

```
addtest3.m  
function addtest3(&a,b)  
    a = a + 2*b
```

Note that it is now illegal to pass a literal value for `a` when calling `addtest3`:

```
--> addtest3(3,4)

a =
  11

Error: Must have lvalue in argument passed by reference
--> addtest3(arg1,arg2)

a =
  7

--> arg1

ans =
  7

--> arg2

ans =
  3
```

The first example fails because we cannot pass a literal like the number 3 by reference. However, the second call succeeds, and note that `arg1` has now changed. Note: please be careful when passing by reference - this feature is not available in MATLAB and you must be clear that you are using it.

As variable argument and return functions are covered elsewhere, as are keywords, we include one final example that demonstrates the use of helper functions, or local functions, where multiple function declarations occur in the same file.

```
    euclidlength.m
function y = foo(x,y)
    square_me(x);
    square_me(y);
    y = sqrt(x+y);

function square_me(&t)
    t = t^2;

--> euclidlength(3,4)

ans =
  5

--> euclidlength(2,0)

ans =
  2
```

3.3 KEYWORDS Function Keywords

3.3.1 Usage

A feature of IDL that FreeMat has adopted is a modified form of **keywords**. The purpose of **keywords** is to allow you to call a function with the arguments to the function specified in an arbitrary order. To specify the syntax of **keywords**, suppose there is a function with prototype

```
function [out_1,...,out_M] = foo(in_1,...,in_N)
```

Then the general syntax for calling function `foo` using keywords is

```
foo(val_1, val_2, /in_k=3)
```

which is exactly equivalent to

```
foo(val_1, val_2, [], [], ..., [], 3),
```

where the 3 is passed as the k-th argument, or alternately,

```
foo(val_1, val_2, /in_k)
```

which is exactly equivalent to

```
foo(val_1, val_2, [], [], ..., [], logical(1)),
```

Note that you can even pass reference arguments using keywords.

3.3.2 Example

The most common use of keywords is in controlling options for functions. For example, the following function takes a number of binary options that control its behavior. For example, consider the following function with two arguments and two options. The function has been written to properly use and handle keywords. The result is much cleaner than the MATLAB approach involving testing all possible values of `nargin`, and forcing explicit empty brackets for don't care parameters.

```
keyfunc.m
function c = keyfunc(a,b,operation,printit)
    if (~isset('a') | ~isset('b'))
        error('keyfunc requires at least the first two 2 arguments');
    end;
    if (~isset('operation'))
        % user did not define the operation, default to '+'
        operation = '+';
    end
    if (~isset('printit'))
        % user did not specify the printit flag, default is false
        printit = 0;
    end
    % simple operation...
```

```
eval(['c = a ' operation ' b;']);
if (printit)
    printf('%f %s %f = %f\n',a,operation,b,c);
end
```

Now some examples of how this function can be called using keywords.

```
--> keyfunc(1,3) % specify a and b, defaults for the others
```

```
ans =
    4
```

```
--> keyfunc(1,3,/printit) % specify printit is true
1.000000 + 3.000000 = 4.000000
```

```
ans =
    4
```

```
--> keyfunc(/operation='- ',2,3) % assigns a=2, b=3
```

```
ans =
   -1
```

```
--> keyfunc(4,/operation='*',/printit) % error as b is unspecified
In /home/basu/dev/branches/FreeMat4/help/tmp/keyfunc.m(keyfunc) at line 3
  In scratch() at line 1
  In base(base)
  In base()
  In global()
```

Error: keyfunc requires at least the first two 2 arguments

3.4 NARGIN Number of Input Arguments

3.4.1 Usage

The `nargin` function returns the number of arguments passed to a function when it was called. The general syntax for its use is

```
y = nargin
```

FreeMat allows for fewer arguments to be passed to a function than were declared, and `nargin`, along with `isset` can be used to determine exactly what subset of the arguments were defined.

3.4.2 Example

Here is a function that is declared to take five arguments, and that simply prints the value of `nargin` each time it is called.

```

nargintest.m
function nargintest(a1,a2,a3,a4,a5)
    printf('nargin = %d\n',nargin);

--> nargintest(3);
nargin = 1
--> nargintest(3,'h');
nargin = 2
--> nargintest(3,'h',1.34);
nargin = 3
--> nargintest(3,'h',1.34,pi,e);
nargin = 5

```

3.5 NARGOUT Number of Output Arguments

3.5.1 Usage

The `nargout` function computes the number of return values requested from a function when it was called. The general syntax for its use

```
y = nargout
```

FreeMat allows for fewer return values to be requested from a function than were declared, and `nargout` can be used to determine exactly what subset of the functions outputs are required.

3.5.2 Example

Here is a function that is declared to return five values, and that simply prints the value of `nargout` each time it is called.

```

nargouttest.m
function [a1,a2,a3,a4,a5] = nargouttest
    printf('nargout = %d\n',nargout);
    a1 = 1; a2 = 2; a3 = 3; a4 = 4; a5 = 5;

--> a1 = nargouttest
nargout = 1

a1 =
    1

--> [a1,a2] = nargouttest
nargout = 2
a1 =
    1

a2 =

```

```
2

--> [a1,a2,a3] = nargouttest
nargout = 3
a1 =
  1

a2 =
  2

a3 =
  3

--> [a1,a2,a3,a4,a5] = nargouttest
nargout = 5
a1 =
  1

a2 =
  2

a3 =
  3

a4 =
  4

a5 =
  5
```

3.6 SCRIPT Script Files

3.6.1 Usage

A script is a sequence of FreeMat commands contained in a `.m` file. When the script is called (via the name of the file), the effect is the same as if the commands inside the script file were issued one at a time from the keyboard. Unlike `function` files (which have the same extension, but have a `function` declaration), script files share the same environment as their callers. Hence, assignments, etc, made inside a script are visible to the caller (which is not the case for functions).

3.6.2 Example

Here is an example of a script that makes some simple assignments and `printf` statements.

```
    tscript.m
a = 13;
```

```
printf('a is %d\n',a);
b = a + 32
```

If we execute the script and then look at the defined variables

```
--> tscript
a is 13

b =
  45

--> who
Variable Name      Type   Flags      Size
              a    double
              ans  double
              b    double
              [1 1]
              [0 0]
              [1 1]
```

we see that `a` and `b` are defined appropriately.

3.7 SPECIAL Special Calling Syntax

3.7.1 Usage

To reduce the effort to call certain functions, FreeMat supports a special calling syntax for functions that take string arguments. In particular, the three following syntaxes are equivalent, with one caveat:

```
functionname('arg1','arg2',...,'argn')
```

or the parenthesis and commas can be removed

```
functionname 'arg1' 'arg2' ... 'argn'
```

The quotes are also optional (providing, of course, that the argument strings have no spaces in them)

```
functionname arg1 arg2 ... argn
```

This special syntax enables you to type `hold on` instead of the more cumbersome `hold('on')`. The caveat is that FreeMat currently only recognizes the special calling syntax as the first statement on a line of input. Thus, the following construction

```
for i=1:10; plot(vec(i)); hold on; end
```

would not work. This limitation may be removed in a future version.

3.7.2 Example

Here is a function that takes two string arguments and returns the concatenation of them.

```

    strcattest.m
function strcattest(str1,str2)
    str3 = [str1,str2];
    printf('str1 = %s, str2 = %s, str3 = %s\n',str1,str2,str3);

```

We call `strcattest` using all three syntaxes.

```

--> strcattest('hi','ho')
str1 = hi, str2 = ho, str3 = hiho
--> strcattest 'hi' 'ho'
str1 = hi, str2 = ho, str3 = hiho
--> strcattest hi ho
str1 = hi, str2 = ho, str3 = hiho

```

3.8 VARARGIN Variable Input Arguments

3.8.1 Usage

FreeMat functions can take a variable number of input arguments by setting the last argument in the argument list to `varargin`. This special keyword indicates that all arguments to the function (beyond the last non-`varargin` keyword) are assigned to a cell array named `varargin` available to the function. Variable argument functions are usually used when writing driver functions, i.e., functions that need to pass arguments to another function. The general syntax for a function that takes a variable number of arguments is

```
function [out_1,...,out_M] = fname(in_1,...,in_M,varargin)
```

Inside the function body, `varargin` collects the arguments to `fname` that are not assigned to the `in_k`.

3.8.2 Example

Here is a simple wrapper to `feval` that demonstrates the use of variable arguments functions.

```

    wrapcall.m
function wrapcall(fname,varargin)
    feval(fname,varargin{:});

```

Now we show a call of the `wrapcall` function with a number of arguments

```

--> wrapcall('printf','%f...%f\n',pi,e)
3.141593...2.718282

```

A more serious driver routine could, for example, optimize a one dimensional function that takes a number of auxiliary parameters that are passed through `varargin`.

3.9 VARARGOUT Variable Output Arguments

3.9.1 Usage

FreeMat functions can return a variable number of output arguments by setting the last argument in the argument list to `varargout`. This special keyword indicates that the number of return values is variable. The general syntax for a function that returns a variable number of outputs is

```
function [out_1,...,out_M,varargout] = fname(in_1,...,in_M)
```

The function is responsible for ensuring that `varargout` is a cell array that contains the values to assign to the outputs beyond `out_M`. Generally, variable output functions use `nargout` to figure out how many outputs have been requested.

3.9.2 Example

This is a function that returns a varying number of values depending on the value of the argument.

```
varoutfunc.m
function [varargout] = varoutfunc
    switch(nargout)
        case 1
            varargout = {'one of one'};
        case 2
            varargout = {'one of two','two of two'};
        case 3
            varargout = {'one of three','two of three','three of three'};
    end
```

Here are some examples of exercising `varoutfunc`:

```
--> [c1] = varoutfunc
c1 =
one of one
--> [c1,c2] = varoutfunc
c1 =
one of two
c2 =
two of two
--> [c1,c2,c3] = varoutfunc
c1 =
one of three
c2 =
two of three
c3 =
three of three
```

Chapter 4

Mathematical Operators

4.1 COLON Index Generation Operator

4.1.1 Usage

There are two distinct syntaxes for the colon `:` operator - the two argument form

```
y = a : c
```

and the three argument form

```
y = a : b : c
```

The two argument form is exactly equivalent to `a:1:c`. The output `y` is the vector

$$y = [a, a + b, a + 2b, \dots, a + nb]$$

where `a+nb <= c`. There is a third form of the colon operator, the no-argument form used in indexing (see `indexing` for more details).

4.1.2 Function Internals

The colon operator turns out to be trickier to implement than one might believe at first, primarily because the floating point versions should do the right thing, which is not the obvious behavior. For example, suppose the user issues a three point colon command

```
y = a : b : c
```

The first question that one might need to answer is: how many points in this vector? If you answered

$$n = \frac{c - a}{b} + 1$$

then you would be doing the straightforward, but not correct thing. because `a`, `b`, and `c` are all floating point values, there are errors associated with each of the quantities that can lead to `n` not

being an integer. A better way (and the way FreeMat currently does the calculation) is to compute the bounding values (for b positive)

$$n \in \left[\frac{(c-a) \rightarrow 0}{b \rightarrow \infty}, \frac{(c-a) \rightarrow \infty}{b \rightarrow 0} \right] + 1$$

where

$$x \rightarrow y$$

means we replace x by the floating point number that is closest to it in the direction of y . Once we have determined the number of points we have to compute the intermediate values

$$[a, a + b, a + 2 * b, \dots, a + n * b]$$

but one can readily verify for themselves that this may *not* be the same as the vector

$$\text{fliplr}[c, c - b, c - 2 * b, \dots, c - n * b]$$

even for the case where

$$c = a + n * b$$

for some n . The reason is that the roundoff in the calculations may be different depending on the nature of the sum. FreeMat uses the following strategy to compute the double-colon vector:

1. The value n is computed by taking the floor of the larger value in the interval defined above.
2. If n falls inside the interval defined above, then it is assumed that the user intended $c = a + n * b$, and the symmetric algorithm is used. Otherwise, the nonsymmetric algorithm is used.
3. The symmetric algorithm computes the vector via

$$[a, a + b, a + 2b, \dots, c - 2b, c - b, c]$$

working symmetrically from both ends of the vector (hence the nomenclature), while the nonsymmetric algorithm computes

$$[a, a + b, a + 2b, \dots, a + nb]$$

In practice, the entries are computed by repeated accumulation instead of multiplying the step size by an integer.

4. The real interval calculation is modified so that we get the exact same result with $a:b:c$ and $c:-b:a$ (which basically means that instead of moving towards infinity, we move towards the signed infinity where the sign is inherited from b).

If you think this is all very obscure, it is. But without it, you will be confronted by mysterious vectors where the last entry is dropped, or where the values show progressively larger amounts of accumulated roundoff error.

4.1.3 Examples

Some simple examples of index generation.

```
--> y = 1:4
```

```
y =
  1 2 3 4
```

Now by half-steps:

```
--> y = 1:.5:4
```

```
y =
  1.0000  1.5000  2.0000  2.5000  3.0000  3.5000  4.0000
```

Now going backwards (negative steps)

```
--> y = 4:-.5:1
```

```
y =
  4.0000  3.5000  3.0000  2.5000  2.0000  1.5000  1.0000
```

If the endpoints are the same, one point is generated, regardless of the step size (middle argument)

```
--> y = 4:1:4
```

```
y =
  4
```

If the endpoints define an empty interval, the output is an empty matrix:

```
--> y = 5:4
```

```
y =
  Empty array 1 0
```

4.2 COMPARISONOPS Array Comparison Operators

4.2.1 Usage

There are a total of six comparison operators available in FreeMat, all of which are binary operators with the following syntax

```
y = a < b
y = a <= b
y = a > b
y = a >= b
y = a ~= b
y = a == b
```

where **a** and **b** are numerical arrays or scalars, and **y** is a **logical** array of the appropriate size. Each of the operators has three modes of operation, summarized in the following list:

1. **a** is a scalar, **b** is an n-dimensional array - the output is then the same size as **b**, and contains the result of comparing each element in **b** to the scalar **a**.
2. **a** is an n-dimensional array, **b** is a scalar - the output is the same size as **a**, and contains the result of comparing each element in **a** to the scalar **b**.
3. **a** and **b** are both n-dimensional arrays of the same size - the output is then the same size as both **a** and **b**, and contains the result of an element-wise comparison between **a** and **b**.

The operators behave the same way as in C, with unequal types being promoted using the standard type promotion rules prior to comparisons. The only difference is that in FreeMat, the not-equals operator is `~=` instead of `!=`.

4.2.2 Examples

Some simple examples of comparison operations. First a comparison with a scalar:

```
--> a = randn(1,5)

a =
  -0.1760    0.0212    0.0095    2.0556   -1.1627

--> a>0

ans =
  0 1 1 1 0
```

Next, we construct two vectors, and test for equality:

```
--> a = [1,2,5,7,3]

a =
  1 2 5 7 3

--> b = [2,2,5,9,4]

b =
  2 2 5 9 4

--> c = a == b

c =
  0 1 1 0 0
```

4.3 DOTLEFTDIVIDE Element-wise Left-Division Operator

4.3.1 Usage

Divides two numerical arrays (elementwise) - gets its name from the fact that the divisor is on the left. There are two forms for its use, both with the same general syntax:

```
y = a .\ b
```

where **a** and **b** are **n**-dimensional arrays of numerical type. In the first case, the two arguments are the same size, in which case, the output **y** is the same size as the inputs, and is the element-wise division of **b** by **a**. In the second case, either **a** or **b** is a scalar, in which case **y** is the same size as the larger argument, and is the division of the scalar with each element of the other argument.

The rules for manipulating types has changed in FreeMat 4.0. See `typerules` for more details.

4.3.2 Function Internals

There are three formulae for the dot-left-divide operator, depending on the sizes of the three arguments. In the most general case, in which the two arguments are the same size, the output is computed via:

$$y(m_1, \dots, m_d) = \frac{b(m_1, \dots, m_d)}{a(m_1, \dots, m_d)}$$

If **a** is a scalar, then the output is computed via

$$y(m_1, \dots, m_d) = \frac{b(m_1, \dots, m_d)}{a}$$

On the other hand, if **b** is a scalar, then the output is computed via

$$y(m_1, \dots, m_d) = \frac{b}{a(m_1, \dots, m_d)}.$$

4.3.3 Examples

Here are some examples of using the dot-left-divide operator. First, a straight-forward usage of the `.\` operator. The first example is straightforward:

```
--> 3 .\ 8
ans =
    2.6667
```

We can also divide complex arguments:

```
--> a = 3 + 4*i
a =
    3.0000 + 4.0000i
```

```
--> b = 5 + 8*i
b =
  5.0000 + 8.0000i
--> c = b .\ a
c =
  0.5281 - 0.0449i
```

We can also demonstrate the three forms of the dot-left-divide operator. First the element-wise version:

```
--> a = [1,2;3,4]
a =
  1 2
  3 4
--> b = [2,3;6,7]
b =
  2 3
  6 7
--> c = a .\ b
c =
  2.0000    1.5000
  2.0000    1.7500
```

Then the scalar versions

```
--> c = a .\ 3
c =
  3.0000    1.5000
  1.0000    0.7500
--> c = 3 .\ a
c =
  0.3333    0.6667
  1.0000    1.3333
```

4.4 DOTPOWER Element-wise Power Operator

4.4.1 Usage

Raises one numerical array to another array (elementwise). There are three operators all with the same general syntax:

$$y = a .^ b$$

The result y depends on which of the following three situations applies to the arguments a and b :

1. a is a scalar, b is an arbitrary n -dimensional numerical array, in which case the output is a raised to the power of each element of b , and the output is the same size as b .
2. a is an n -dimensional numerical array, and b is a scalar, then the output is the same size as a , and is defined by each element of a raised to the power b .
3. a and b are both n -dimensional numerical arrays of *the same size*. In this case, each element of the output is the corresponding element of a raised to the power defined by the corresponding element of b .

The rules for manipulating types has changed in FreeMat 4.0. See `typerules` for more details.

4.4.2 Function Internals

There are three formulae for this operator. For the first form

$$y(m_1, \dots, m_d) = a^{b(m_1, \dots, m_d)},$$

and the second form

$$y(m_1, \dots, m_d) = a(m_1, \dots, m_d)^b,$$

and in the third form

$$y(m_1, \dots, m_d) = a(m_1, \dots, m_d)^{b(m_1, \dots, m_d)}.$$

4.4.3 Examples

We demonstrate the three forms of the dot-power operator using some simple examples. First, the case of a scalar raised to a series of values.

```
--> a = 2
```

```
a =
  2
```

```
--> b = 1:4
```

```
b =
  1 2 3 4
```

```
--> c = a.^b
```

```
c =
     2     4     8    16
```

The second case shows a vector raised to a scalar.

```
--> c = b.^a
```

```
c =
     1     4     9    16
```

The third case shows the most general use of the dot-power operator.

```
--> A = [1,2;3,2]
```

```
A =
     1     2
     3     2
```

```
--> B = [2,1.5;0.5,0.6]
```

```
B =
     2.0000     1.5000
     0.5000     0.6000
```

```
--> C = A.^B
```

```
C =
     1.0000     2.8284
     1.7321     1.5157
```

4.5 DOTRIGHTDIVIDE Element-wise Right-Division Operator

4.5.1 Usage

Divides two numerical arrays (elementwise). There are two forms for its use, both with the same general syntax:

```
y = a ./ b
```

where **a** and **b** are *n*-dimensional arrays of numerical type. In the first case, the two arguments are the same size, in which case, the output **y** is the same size as the inputs, and is the element-wise division of **b** by **a**. In the second case, either **a** or **b** is a scalar, in which case **y** is the same size as the larger argument, and is the division of the scalar with each element of the other argument.

The rules for manipulating types has changed in FreeMat 4.0. See `typerules` for more details.

4.5.2 Function Internals

There are three formulae for the dot-right-divide operator, depending on the sizes of the three arguments. In the most general case, in which the two arguments are the same size, the output is computed via:

$$y(m_1, \dots, m_d) = \frac{a(m_1, \dots, m_d)}{b(m_1, \dots, m_d)}$$

If a is a scalar, then the output is computed via

$$y(m_1, \dots, m_d) = \frac{a}{b(m_1, \dots, m_d)}$$

On the other hand, if b is a scalar, then the output is computed via

$$y(m_1, \dots, m_d) = \frac{a(m_1, \dots, m_d)}{b}.$$

4.5.3 Examples

Here are some examples of using the dot-right-divide operator. First, a straight-forward usage of the `./` operator. The first example is straightforward:

```
--> 3 ./ 8
```

```
ans =
    0.3750
```

We can also divide complex arguments:

```
--> a = 3 + 4*i
```

```
a =
    3.0000 + 4.0000i
```

```
--> b = 5 + 8*i
```

```
b =
    5.0000 + 8.0000i
```

```
--> c = a ./ b
```

```
c =
    0.5281 - 0.0449i
```

We can also demonstrate the three forms of the dot-right-divide operator. First the element-wise version:

```
--> a = [1,2;3,4]
```

```

a =
  1 2
  3 4

--> b = [2,3;6,7]

b =
  2 3
  6 7

--> c = a ./ b

c =
  0.5000    0.6667
  0.5000    0.5714

```

Then the scalar versions

```

--> c = a ./ 3

c =
  0.3333    0.6667
  1.0000    1.3333

--> c = 3 ./ a

c =
  3.0000    1.5000
  1.0000    0.7500

```

4.6 DOTTIMES Element-wise Multiplication Operator

4.6.1 Usage

Multiplies two numerical arrays (elementwise). There are two forms for its use, both with the same general syntax:

```
y = a .* b
```

where **a** and **b** are **n**-dimensional arrays of numerical type. In the first case, the two arguments are the same size, in which case, the output **y** is the same size as the inputs, and is the element-wise product of **a** and **b**. In the second case, either **a** or **b** is a scalar, in which case **y** is the same size as the larger argument, and is the product of the scalar with each element of the other argument.

The rules for manipulating types has changed in FreeMat 4.0. See `typerules` for more details.

4.6.2 Function Internals

There are three formulae for the dot-times operator, depending on the sizes of the three arguments. In the most general case, in which the two arguments are the same size, the output is computed via:

$$y(m_1, \dots, m_d) = a(m_1, \dots, m_d) \times b(m_1, \dots, m_d)$$

If a is a scalar, then the output is computed via

$$y(m_1, \dots, m_d) = a \times b(m_1, \dots, m_d).$$

On the other hand, if b is a scalar, then the output is computed via

$$y(m_1, \dots, m_d) = a(m_1, \dots, m_d) \times b.$$

4.6.3 Examples

Here are some examples of using the `dottimes` operator. First, a straight-forward usage of the `.*` operator. The first example is straightforward:

```
--> 3 .* 8
```

```
ans =
    24
```

Next, we multiply a scalar by a vector of values:

```
--> 3.1 .* [2,4,5,6,7]
```

```
ans =
    6.2000    12.4000    15.5000    18.6000    21.7000
```

With complex values

```
--> a = 3 + 4*i
```

```
a =
    3.0000 + 4.0000i
```

```
--> b = a .* 2
```

```
b =
    6.0000 + 8.0000i
```

Finally, the element-wise version:

```
--> a = [1,2;3,4]
```

```
a =
    1 2
```

```
3 4
```

```
--> b = [2,3;6,7]
```

```
b =
```

```
2 3
```

```
6 7
```

```
--> c = a .* b
```

```
c =
```

```
2 6
```

```
18 28
```

4.7 HERMITIAN Matrix Hermitian (Conjugate Transpose) Operator

4.7.1 Usage

Computes the Hermitian of the argument (a 2D matrix). The syntax for its use is

```
y = a';
```

where a is a $M \times N$ numerical matrix. The output y is a numerical matrix of the same type of size $N \times M$. This operator is the conjugating transpose, which is different from the transpose operator `.'` (which does not conjugate complex values).

4.7.2 Function Internals

The Hermitian operator is defined simply as

$$y_{i,j} = \overline{a_{j,i}}$$

where $y_{i,j}$ is the element in the i th row and j th column of the output matrix y .

4.7.3 Examples

A simple transpose example:

```
--> A = [1,2,0;4,1,-1]
```

```
A =
```

```
1 2 0
```

```
4 1 -1
```

```
--> A'
```

```
ans =
  1  4
  2  1
  0 -1
```

Here, we use a complex matrix to demonstrate how the Hermitian operator conjugates the entries.

```
--> A = [1+i,2-i]
```

```
A =
  1.0000 + 1.0000i  2.0000 - 1.0000i
```

```
--> A.'
```

```
ans =
  1.0000 + 1.0000i
  2.0000 - 1.0000i
```

4.8 LEFTDIVIDE Matrix Equation Solver/Divide Operator

4.8.1 Usage

The divide operator `\` is really a combination of three operators, all of which have the same general syntax:

$$Y = A \setminus B$$

where A and B are arrays of numerical type. The result Y depends on which of the following three situations applies to the arguments A and B :

1. A is a scalar, B is an arbitrary n -dimensional numerical array, in which case the output is each element of B divided by the scalar A .
2. A, B are matrices with the same number of rows, i.e., A is of size $M \times K$, and B is of size $M \times L$, in which case the output is of size $K \times L$.

The output follows the standard type promotion rules, although in the first two cases, if A and B are integers, the output is an integer also, while in the third case if A and B are integers, the output is of type `double`.

A few additional words about the third version, in which A and B are matrices. Very loosely speaking, Y is the matrix that satisfies $A * Y = B$. In cases where such a matrix exists. If such a matrix does not exist, then a matrix Y is returned that approximates $A * Y \approx B$.

4.8.2 Function Internals

There are three formulae for the times operator. For the first form

$$Y(m_1, \dots, m_d) = \frac{B(m_1, \dots, m_d)}{A}$$

In the second form, the calculation of the output depends on the size of A . Because each column of B is treated independently, we can rewrite the equation $A Y = B$ as

$$A[y_1, y_2, \dots, y_i] = [b_1, b_2, \dots, b_i]$$

where y_i are the columns of Y , and b_i are the columns of the matrix B . If A is a square matrix, then the LAPACK routine `*gesvx` (where the `*` is replaced with `sdz` depending on the type of the arguments) is used, which uses an LU decomposition of A to solve the sequence of equations sequentially. If A is singular, then a warning is emitted.

On the other hand, if A is rectangular, then the LAPACK routine `*gelsy` is used. Note that these routines are designed to work with matrices A that are full rank - either full column rank or full row rank. If A fails to satisfy this assumption, a warning is emitted. If A has full column rank (and thus necessarily has more rows than columns), then theoretically, this operator finds the columns y_i that satisfy:

$$y_i = \arg \min_y \|Ay - b_i\|_2$$

and each column is thus the Least Squares solution of $A y = b_i$. On the other hand, if A has full row rank (and thus necessarily has more columns than rows), then theoretically, this operator finds the columns y_i that satisfy

$$y_i = \arg \min_{Ay=b_i} \|y\|_2$$

and each column is thus the Minimum Norm vector y_i that satisfies $A y_i = b_i$. In the event that the matrix A is neither full row rank nor full column rank, a solution is returned, that is the minimum norm least squares solution. The solution is computed using an orthogonal factorization technique that is documented in the LAPACK User's Guide (see the References section for details).

4.8.3 Examples

Here are some simple examples of the divide operator. We start with a simple example of a full rank, square matrix:

```
--> A = [1,1;0,1]
```

```
A =
 1 1
 0 1
```

Suppose we wish to solve

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 3 \\ 2 \end{bmatrix}$$

(which by inspection has the solution $y_1 = 1$, $y_2 = 2$). Thus we compute:

```
--> B = [3;2]
```

```
B =
 3
 2
```

```
--> Y = A\B
```

```
Y =
  1
  2
```

Suppose we wish to solve a trivial Least Squares (LS) problem. We want to find a simple scaling of the vector $[1;1]$ that is closest to the point $[2,1]$. This is equivalent to solving

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix} y = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$$

in a least squares sense. For fun, we can calculate the solution using calculus by hand. The error we wish to minimize is

$$\varepsilon(y) = (y - 2)^2 + (y - 1)^2.$$

Taking a derivative with respect to y , and setting to zero (which we must have for an extrema when y is unconstrained)

$$2(y - 2) + 2(y - 1) = 0$$

which we can simplify to $4y = 6$ or $y = 3/2$ (we must, technically, check to make sure this is a minimum, and not a maximum or an inflection point). Here is the same calculation performed using FreeMat:

```
--> A = [1;1]
```

```
A =
  1
  1
```

```
--> B = [2;1]
```

```
B =
  2
  1
```

```
--> A\B
```

```
ans =
  1.5000
```

which is the same solution.

4.9 LOGICALOPS Logical Array Operators

4.9.1 Usage

There are three Boolean operators available in FreeMat. The syntax for their use is:

```

y = ~x
y = a & b
y = a | b

```

where `x`, `a` and `b` are logical arrays. The operators are

- NOT (`~`) - output `y` is true if the corresponding element of `x` is false, and output `y` is false if the corresponding element of `x` is true.
- OR (`|`) - output `y` is true if corresponding element of `a` is true or if corresponding element of `b` is true (or if both are true).
- AND (`&`) - output `y` is true only if both the corresponding elements of `a` and `b` are both true.

The binary operators AND and OR can take scalar arguments as well as vector arguments, in which case, the scalar is operated on with each element of the vector. As of version 1.10, FreeMat supports shortcut evaluation. This means that if we have two expressions

```
if (expr1 & expr2)
```

then if `expr1` evaluates to `false`, then `expr2` is not evaluated at all. Similarly, for the expression

```
if (expr1 | expr2)
```

then if `expr1` evaluates to `true`, then `expr2` is not evaluated at all. Shortcut evaluation is useful for doing a sequence of tests, each of which is not valid unless the prior test is successful. For example,

```
if isa(p,'string') & strcmp(p,'fro')
```

is not valid without shortcut evaluation (if `p` is an integer, for example, the first test returns false, and an attempt to evaluate the second expression would lead to an error). Note that shortcut evaluation only works with scalar expressions.

4.9.2 Examples

Some simple examples of logical operators. Suppose we want to calculate the exclusive-or (XOR) of two vectors of logical variables. First, we create a pair of vectors to perform the XOR operation on:

```
--> a = (randn(1,6)>0)
```

```
a =
 0 1 1 1 0 1
```

```
--> b = (randn(1,6)>0)
```

```
b =
 0 1 0 0 1 0
```

Next, we can compute the OR of `a` and `b`:

```
--> c = a | b
```

```
c =
 0 1 1 1 1 1
```

However, the XOR and OR operations differ on the fifth entry - the XOR would be false, since it is true if and only if exactly one of the two inputs is true. To isolate this case, we can AND the two vectors, to find exactly those entries that appear as true in both **a** and **b**:

```
--> d = a & b
```

```
d =
 0 1 0 0 0 0
```

At this point, we can modify the contents of **c** in two ways – the Boolean way is to AND `\sim d` with **c**, like so

```
--> xor = c & (~d)
```

```
xor =
 0 0 1 1 1 1
```

The other way to do this is simply force `c(d) = 0`, which uses the logical indexing mode of FreeMat (see the chapter on indexing for more details). This, however, will cause **c** to become an `int32` type, as opposed to a logical type.

```
--> c(d) = 0
```

```
c =
 0 0 1 1 1 1
```

4.10 MATRIXPOWER Matrix Power Operator

4.10.1 Usage

The power operator for scalars and square matrices. This operator is really a combination of two operators, both of which have the same general syntax:

$$y = a \wedge b$$

The exact action taken by this operator, and the size and type of the output, depends on which of the two configurations of **a** and **b** is present:

1. **a** is a scalar, **b** is a square matrix
2. **a** is a square matrix, **b** is a scalar

4.10.2 Function Internals

In the first case that \mathbf{a} is a scalar, and \mathbf{b} is a square matrix, the matrix power is defined in terms of the eigenvalue decomposition of \mathbf{b} . Let \mathbf{b} have the following eigen-decomposition (problems arise with non-symmetric matrices \mathbf{b} , so let us assume that \mathbf{b} is symmetric):

$$\mathbf{b} = E \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \lambda_n \end{bmatrix} E^{-1}$$

Then \mathbf{a} raised to the power \mathbf{b} is defined as

$$\mathbf{a}^{\mathbf{b}} = E \begin{bmatrix} a^{\lambda_1} & 0 & \cdots & 0 \\ 0 & a^{\lambda_2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & a^{\lambda_n} \end{bmatrix} E^{-1}$$

Similarly, if \mathbf{a} is a square matrix, then \mathbf{a} has the following eigen-decomposition (again, suppose \mathbf{a} is symmetric):

$$\mathbf{a} = E \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \lambda_n \end{bmatrix} E^{-1}$$

Then \mathbf{a} raised to the power \mathbf{b} is defined as

$$\mathbf{a}^{\mathbf{b}} = E \begin{bmatrix} \lambda_1^{\mathbf{b}} & 0 & \cdots & 0 \\ 0 & \lambda_2^{\mathbf{b}} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \lambda_n^{\mathbf{b}} \end{bmatrix} E^{-1}$$

4.10.3 Examples

We first define a simple 2 x 2 symmetric matrix

```
--> A = 1.5
```

```
A =
    1.5000
```

```
--> B = [1, .2; .2, 1]
```

```
B =
    1.0000    0.2000
    0.2000    1.0000
```

First, we raise B to the (scalar power) A:

```
--> C = B^A
```

```
C =
  1.0150    0.2995
  0.2995    1.0150
```

Next, we raise A to the matrix power B:

```
--> C = A^B
```

```
C =
  1.5049    0.1218
  0.1218    1.5049
```

4.11 MINUS Subtraction Operator

4.11.1 Usage

Subtracts two numerical arrays (elementwise). There are two forms for its use, both with the same general syntax:

```
y = a - b
```

where **a** and **b** are **n**-dimensional arrays of numerical type. In the first case, the two arguments are the same size, in which case, the output **y** is the same size as the inputs, and is the element-wise difference of **a** and **b**. In the second case, either **a** or **b** is a scalar, in which case **y** is the same size as the larger argument, and is the difference of the scalar to each element of the other argument.

The rules for manipulating types has changed in FreeMat 4.0. See `typerules` for more details.

4.11.2 Function Internals

There are three formulae for the subtraction operator, depending on the sizes of the three arguments. In the most general case, in which the two arguments are the same size, the output is computed via:

$$y(m_1, \dots, m_d) = a(m_1, \dots, m_d) - b(m_1, \dots, m_d)$$

If **a** is a scalar, then the output is computed via

$$y(m_1, \dots, m_d) = a - b(m_1, \dots, m_d).$$

On the other hand, if **b** is a scalar, then the output is computed via

$$y(m_1, \dots, m_d) = a(m_1, \dots, m_d) - b.$$

4.11.3 Examples

Here are some examples of using the subtraction operator. First, a straight-forward usage of the minus operator. The first example is straightforward:

```
--> 3 - 8
```

```
ans =  
-5
```

Next, we subtract a vector of values from a scalar:

```
--> 3.1 - [2,4,5,6,7]
```

```
ans =  
1.1000 -0.9000 -1.9000 -2.9000 -3.9000
```

With complex values

```
--> a = 3 - 4*i
```

```
a =  
3.0000 - 4.0000i
```

```
--> b = a - 2
```

```
b =  
1.0000 - 4.0000i
```

Finally, the element-wise version:

```
--> a = [1,2;3,4]
```

```
a =  
1 2  
3 4
```

```
--> b = [2,3;6,7]
```

```
b =  
2 3  
6 7
```

```
--> c = a - b
```

```
c =  
-1 -1  
-3 -3
```

4.12 PLUS Addition Operator

4.12.1 Usage

Adds two numerical arrays (elementwise) together. There are two forms for its use, both with the same general syntax:

$$y = a + b$$

where **a** and **b** are *n*-dimensional arrays of numerical type. In the first case, the two arguments are the same size, in which case, the output **y** is the same size as the inputs, and is the element-wise the sum of **a** and **b**. In the second case, either **a** or **b** is a scalar, in which case **y** is the same size as the larger argument, and is the sum of the scalar added to each element of the other argument.

The rules for manipulating types has changed in FreeMat 4.0. See `typerules` for more details.

4.12.2 Function Internals

There are three formulae for the addition operator, depending on the sizes of the three arguments. In the most general case, in which the two arguments are the same size, the output is computed via:

$$y(m_1, \dots, m_d) = a(m_1, \dots, m_d) + b(m_1, \dots, m_d)$$

If **a** is a scalar, then the output is computed via

$$y(m_1, \dots, m_d) = a + b(m_1, \dots, m_d).$$

On the other hand, if **b** is a scalar, then the output is computed via

$$y(m_1, \dots, m_d) = a(m_1, \dots, m_d) + b.$$

4.12.3 Examples

Here are some examples of using the addition operator. First, a straight-forward usage of the plus operator. The first example is straightforward:

```
--> 3 + 8
```

```
ans =
  11
```

Next, we add a scalar to a vector of values:

```
--> 3.1 + [2,4,5,6,7]
```

```
ans =
  5.1000    7.1000    8.1000    9.1000   10.1000
```

With complex values

```
--> a = 3 + 4*i
a =
  3.0000 + 4.0000i
--> b = a + 2
b =
  5.0000 + 4.0000i
```

Finally, the element-wise version:

```
--> a = [1,2;3,4]
a =
  1 2
  3 4
--> b = [2,3;6,7]
b =
  2 3
  6 7
--> c = a + b
c =
  3 5
  9 11
```

4.13 RIGHTDIVIDE Matrix Equation Solver/Divide Operator

4.13.1 Usage

The divide operator `/` is really a combination of three operators, all of which have the same general syntax:

$$Y = A / B$$

where `A` and `B` are arrays of numerical type. The result `Y` depends on which of the following three situations applies to the arguments `A` and `B`:

1. `A` is a scalar, `B` is an arbitrary `n`-dimensional numerical array, in which case the output is the scalar `A` divided into each element of `B`.
2. `B` is a scalar, `A` is an arbitrary `n`-dimensional numerical array, in which case the output is each element of `A` divided by the scalar `B`.

3. A,B are matrices with the same number of columns, i.e., A is of size K x M, and B is of size L x M, in which case the output is of size K x L.

The output follows the standard type promotion rules, although in the first two cases, if A and B are integers, the output is an integer also, while in the third case if A and B are integers, the output is of type `double`.

4.13.2 Function Internals

There are three formulae for the times operator. For the first form

$$Y(m_1, \dots, m_d) = \frac{A}{B(m_1, \dots, m_d)},$$

and the second form

$$Y(m_1, \dots, m_d) = \frac{A(m_1, \dots, m_d)}{B}.$$

In the third form, the output is defined as:

$$Y = (B' \setminus A)'$$

and is used in the equation $Y B = A$.

4.13.3 Examples

The right-divide operator is much less frequently used than the left-divide operator, but the concepts are similar. It can be used to find least-squares and minimum norm solutions. It can also be used to solve systems of equations in much the same way. Here's a simple example:

```
--> B = [1,1;0,1];
--> A = [4,5]
```

```
A =
  4  5
```

```
--> A/B
```

```
ans =
  4  1
```

4.14 TIMES Matrix Multiply Operator

4.14.1 Usage

Multiplies two numerical arrays. This operator is really a combination of three operators, all of which have the same general syntax:

```
y = a * b
```

where **a** and **b** are arrays of numerical type. The result **y** depends on which of the following three situations applies to the arguments **a** and **b**:

1. **a** is a scalar, **b** is an arbitrary **n**-dimensional numerical array, in which case the output is the element-wise product of **b** with the scalar **a**.
2. **b** is a scalar, **a** is an arbitrary **n**-dimensional numerical array, in which case the output is the element-wise product of **a** with the scalar **b**.
3. **a, b** are conformant matrices, i.e., **a** is of size **M x K**, and **b** is of size **K x N**, in which case the output is of size **M x N** and is the matrix product of **a**, and **b**.

Matrix multiplication is only defined for matrices of type `double` and `single`.

4.14.2 Function Internals

There are three formulae for the times operator. For the first form

$$y(m_1, \dots, m_d) = a \times b(m_1, \dots, m_d),$$

and the second form

$$y(m_1, \dots, m_d) = a(m_1, \dots, m_d) \times b.$$

In the third form, the output is the matrix product of the arguments

$$y(m, n) = \sum_{k=1}^K a(m, k)b(k, n)$$

4.14.3 Examples

Here are some examples of using the matrix multiplication operator. First, the scalar examples (types 1 and 2 from the list above):

```
--> a = [1,3,4;0,2,1]
```

```
a =
  1 3 4
  0 2 1
```

```
--> b = a * 2
```

```
b =
  2 6 8
  0 4 2
```

The matrix form, where the first argument is **2 x 3**, and the second argument is **3 x 1**, so that the product is size **2 x 1**.

```
--> a = [1,2,0;4,2,3]
```

```
a =  
 1 2 0  
 4 2 3
```

```
--> b = [5;3;1]
```

```
b =  
 5  
 3  
 1
```

```
--> c = a*b
```

```
c =  
 11  
 29
```

Note that the output is double precision.

4.15 TRANSPOSE Matrix Transpose Operator

4.15.1 Usage

Performs a transpose of the argument (a 2D matrix). The syntax for its use is

```
y = a.';
```

where \mathbf{a} is a $M \times N$ numerical matrix. The output \mathbf{y} is a numerical matrix of the same type of size $N \times M$. This operator is the non-conjugating transpose, which is different from the Hermitian operator $'$ (which conjugates complex values).

4.15.2 Function Internals

The transpose operator is defined simply as

$$y_{i,j} = a_{j,i}$$

where $y_{i,j}$ is the element in the i th row and j th column of the output matrix \mathbf{y} .

4.15.3 Examples

A simple transpose example:

```
--> A = [1,2,0;4,1,-1]
```

```
A =
```

```

1 2 0
4 1 -1

--> A.'

ans =
1 4
2 1
0 -1

```

Here, we use a complex matrix to demonstrate how the transpose does *not* conjugate the entries.

```

--> A = [1+i,2-i]

A =
1.0000 + 1.0000i 2.0000 - 1.0000i

--> A.'

ans =
1.0000 + 1.0000i
2.0000 - 1.0000i

```

4.16 TYPERULES Type Rules for Operators

4.16.1 Usage

Starting with FreeMat 4.0, the type of y is determined according to the same rules as Matlab. These are the rules:

1. Integer types of the same class can be combined. The answer is the same type as the inputs, and the operation is performed using saturating arithmetic. Integer types can also be combined with double precision values (again, the result is of the integer type).
2. Single precision floating point values can be combined with double precision, logical and character array classes. The result is of class single.
3. Double precision floating point values can be combined with all other types. Except as noted above, the output is of double precision.

These rules look strange, and they are. In general, computations are done in double precision in almost all cases. When single precision values are involved, the computations take place in single precision.

Chapter 5

Flow Control

5.1 BREAK Exit Execution In Loop

5.1.1 Usage

The `break` statement is used to exit a loop prematurely. It can be used inside a `for` loop or a `while` loop. The syntax for its use is

```
break
```

inside the body of the loop. The `break` statement forces execution to exit the loop immediately.

5.1.2 Example

Here is a simple example of how `break` exits the loop. We have a loop that sums integers from 1 to 10, but that stops prematurely at 5 using a `break`. We will use a `while` loop.

```
break_ex.m
function accum = break_ex
    accum = 0;
    i = 1;
    while (i<=10)
        accum = accum + i;
        if (i == 5)
            break;
        end
        i = i + 1;
    end
```

The function is exercised here:

```
--> break_ex
```

```
ans =
```

```

15
--> sum(1:5)

ans =
    15

```

5.2 CONTINUE Continue Execution In Loop

5.2.1 Usage

The `continue` statement is used to change the order of execution within a loop. The `continue` statement can be used inside a `for` loop or a `while` loop. The syntax for its use is

```

    continue

```

inside the body of the loop. The `continue` statement forces execution to start at the top of the loop with the next iteration. The examples section shows how the `continue` statement works.

5.2.2 Example

Here is a simple example of using a `continue` statement. We want to sum the integers from 1 to 10, but not the number 5. We will use a `for` loop and a `continue` statement.

```

    continue_ex.m
function accum = continue_ex
    accum = 0;
    for i=1:10
        if (i==5)
            continue;
        end
        accum = accum + 1; %skipped if i == 5!
    end

```

The function is exercised here:

```

--> continue_ex

ans =
     9

--> sum([1:4,6:10])

ans =
    50

```

5.3 ERROR Causes an Error Condition Raised

5.3.1 Usage

The `error` function causes an error condition (exception to be raised). The general syntax for its use is

```
error(s),
```

where `s` is the string message describing the error. The `error` function is usually used in conjunction with `try` and `catch` to provide error handling. If the string `s`, then (to conform to the MATLAB API), `error` does nothing.

5.3.2 Example

Here is a simple example of an `error` being issued by a function `evenoddttest`:

```
evenoddttest.m
function evenoddttest(n)
    if (n==0)
        error('zero is neither even nor odd');
    elseif ( n ~= fix(n) )
        error('expecting integer argument');
    end;
    if (n==int32(n/2)*2)
        printf('%d is even\n',n);
    else
        printf('%d is odd\n',n);
    end
```

The normal command line prompt `-->` simply prints the error that occurred.

```
--> evenoddttest(4)
4 is even
--> evenoddttest(5)
5 is odd
--> evenoddttest(0)
In /home/basu/dev/branches/FreeMat4/help/tmp/evenoddttest.m(evenoddttest) at line 3
    In scratch() at line 1
    In base(base)
    In base()
    In global()
Error: zero is neither even nor odd
--> evenoddttest(pi)
In /home/basu/dev/branches/FreeMat4/help/tmp/evenoddttest.m(evenoddttest) at line 5
    In scratch() at line 1
    In base(base)
    In base()
```

```
In global()
Error: expecting integer argument
```

5.4 FOR For Loop

5.4.1 Usage

The `for` loop executes a set of statements with an index variable looping through each element in a vector. The syntax of a `for` loop is one of the following:

```
for (variable=expression)
  statements
end
```

Alternately, the parenthesis can be eliminated

```
for variable=expression
  statements
end
```

or alternately, the index variable can be pre-initialized with the vector of values it is going to take:

```
for variable
  statements
end
```

The third form is essentially equivalent to `for variable=variable`, where `variable` is both the index variable and the set of values over which the `for` loop executes. See the examples section for an example of this form of the `for` loop.

5.4.2 Examples

Here we write `for` loops to add all the integers from 1 to 100. We will use all three forms of the `for` statement.

```
--> accum = 0;
--> for (i=1:100); accum = accum + i; end
--> accum
```

```
ans =
  5050
```

The second form is functionally the same, without the extra parenthesis

```
--> accum = 0;
--> for i=1:100; accum = accum + i; end
--> accum
```

```
ans =
  5050
```

In the third example, we pre-initialize the loop variable with the values it is to take

5.5 IF-ELSEIF-ELSE Conditional Statements

5.5.1 Usage

The `if` and `else` statements form a control structure for conditional execution. The general syntax involves an `if` test, followed by zero or more `elseif` clauses, and finally an optional `else` clause:

```
if conditional_expression_1
    statements_1
elseif conditional_expression_2
    statements_2
elseif conditional_expression_3
    statements_3
...
else
    statements_N
end
```

Note that a conditional expression is considered true if the real part of the result of the expression contains any non-zero elements (this strange convention is adopted for compatibility with MATLAB).

5.5.2 Examples

Here is an example of a function that uses an `if` statement

```
if_test.m
function c = if_test(a)
    if (a == 1)
        c = 'one';
    elseif (a==2)
        c = 'two';
    elseif (a==3)
        c = 'three';
    else
        c = 'something else';
    end
```

Some examples of `if_test` in action:

```
--> if_test(1)

ans =
one
--> if_test(2)
```

```

ans =
two
--> if_test(3)

ans =
three
--> if_test(pi)

ans =
something else

```

5.6 KEYBOARD Initiate Interactive Debug Session

5.6.1 Usage

The `keyboard` statement is used to initiate an interactive session at a specific point in a function. The general syntax for the `keyboard` statement is

```
keyboard
```

A `keyboard` statement can be issued in a `script`, in a `function`, or from within another `keyboard` session. The result of a `keyboard` statement is that execution of the program is halted, and you are given a prompt of the form:

```
[scope,n] -->
```

where `scope` is the current scope of execution (either the name of the function we are executing, or `base` otherwise). And `n` is the depth of the `keyboard` session. If, for example, we are in a `keyboard` session, and we call a function that issues another `keyboard` session, the depth of that second session will be one higher. Put another way, `n` is the number of `return` statements you have to issue to get back to the base workspace. Incidentally, a `return` is how you exit the `keyboard` session and resume execution of the program from where it left off. A `retall` can be used to shortcut execution and return to the base workspace.

The `keyboard` statement is an excellent tool for debugging FreeMat code, and along with `eval` provide a unique set of capabilities not usually found in compiled environments. Indeed, the `keyboard` statement is equivalent to a debugger breakpoint in more traditional environments, but with significantly more inspection power.

5.6.2 Example

Here we demonstrate a two-level `keyboard` situation. We have a simple function that calls `keyboard` internally:

```

key_one.m
function c = key_one(a,b)
c = a + b;
keyboard

```

Now, we execute the function from the base workspace, and at the `keyboard` prompt, we call it again. This action puts us at depth 2. We can confirm that we are in the second invocation of the function by examining the arguments. We then issue two `return` statements to return to the base workspace.

```
--> key_one(1,2)
[key_one,3]--> key_one(5,7)
[key_one,3]--> a
```

```
ans =
     5
```

```
[key_one,3]--> b
```

```
ans =
     7
```

```
[key_one,3]--> c
```

```
ans =
    12
```

```
[key_one,3]--> return
```

```
ans =
    12
```

```
[key_one,3]--> a
```

```
ans =
     1
```

```
[key_one,3]--> b
```

```
ans =
     2
```

```
[key_one,3]--> c
```

```
ans =
     3
```

```
[key_one,3]--> return
```

```
ans =
     3
```

5.7 LASTERR Retrieve Last Error Message

5.7.1 Usage

Either returns or sets the last error message. The general syntax for its use is either

```
msg = lasterr
```

which returns the last error message that occurred, or

```
lasterr(msg)
```

which sets the contents of the last error message.

5.7.2 Example

Here is an example of using the `error` function to set the last error, and then retrieving it using `lasterr`.

```
--> try; error('Test error message'); catch; end;  
--> lasterr
```

```
ans =  
Test error message
```

Or equivalently, using the second form:

```
--> lasterr('Test message');  
--> lasterr
```

```
ans =  
Test message
```

5.8 RETALL Return From All Keyboard Sessions

5.8.1 Usage

The `retall` statement is used to return to the base workspace from a nested `keyboard` session. It is equivalent to forcing execution to return to the main prompt, regardless of the level of nesting of `keyboard` sessions, or which functions are running. The syntax is simple

```
retall
```

The `retall` is a convenient way to stop debugging. In the process of debugging a complex program or set of functions, you may find yourself 5 function calls down into the program only to discover the problem. After fixing it, issuing a `retall` effectively forces FreeMat to exit your program and return to the interactive prompt.

5.8.2 Example

Here we demonstrate an extreme example of `retall`. We are debugging a recursive function `self` to calculate the sum of the first N integers. When the function is called, a `keyboard` session is initiated after the function has called itself N times. At this `keyboard` prompt, we issue another call to `self` and get another `keyboard` prompt, this time with a depth of 2. A `retall` statement returns us to the top level without executing the remainder of either the first or second call to `self`:

```

self.m
function y = self(n)
  if (n>1)
    y = n + self(n-1);
    printf('y is %d\n',y);
  else
    y = 1;
    printf('y is initialized to one\n');
    keyboard
  end
end

--> self(4)
y is initialized to one
[self,8]--> self(6)
y is initialized to one
[self,8]--> retall

```

5.9 RETURN Return From Function

5.9.1 Usage

The `return` statement is used to immediately return from a function, or to return from a `keyboard` session. The syntax for its use is

```
return
```

Inside a function, a `return` statement causes FreeMat to exit the function immediately. When a `keyboard` session is active, the `return` statement causes execution to resume where the `keyboard` session started.

5.9.2 Example

In the first example, we define a function that uses a `return` to exit the function if a certain test condition is satisfied.

```

return_func.m
function ret = return_func(a,b)
  ret = 'a is greater';
  if (a > b)
    return;
  end
end

```

```

end
ret = 'b is greater';
printf('finishing up...\n');

```

Next we exercise the function with a few simple test cases:

```

--> return_func(1,3)
finishing up...

```

```

ans =
b is greater
--> return_func(5,2)

```

```

ans =
a is greater

```

In the second example, we take the function and rewrite it to use a `keyboard` statement inside the `if` statement.

```

return_func2.m
function ret = return_func2(a,b)
if (a > b)
ret = 'a is greater';
keyboard;
else
ret = 'b is greater';
end
printf('finishing up...\n');

```

Now, we call the function with a larger first argument, which triggers the `keyboard` session. After verifying a few values inside the `keyboard` session, we issue a `return` statement to resume execution.

```

--> return_func2(2,4)
finishing up...

```

```

ans =
b is greater
--> return_func2(5,1)
[return_func2,4]--> ret

```

```

ans =
a is greater
[return_func2,4]--> a

```

```

ans =
5
[return_func2,4]--> b

```

```
ans =  
    1  
  
[return_func2,4]--> return  
finishing up...  
  
ans =  
a is greater
```

5.10 SWITCH Switch statement

5.10.1 Usage

The `switch` statement is used to selective execute code based on the value of either scalar value or a string. The general syntax for a `switch` statement is

```
switch(expression)  
    case test_expression_1  
        statements  
    case test_expression_2  
        statements  
    otherwise  
        statements  
end
```

The `otherwise` clause is optional. Note that each test expression can either be a scalar value, a string to test against (if the switch expression is a string), or a `cell-array` of expressions to test against. Note that unlike C `switch` statements, the FreeMat `switch` does not have fall-through, meaning that the statements associated with the first matching case are executed, and then the `switch` ends. Also, if the `switch` expression matches multiple `case` expressions, only the first one is executed.

5.10.2 Examples

Here is an example of a `switch` expression that tests against a string input:

```
switch_test.m  
function c = switch_test(a)  
    switch(a)  
        case {'lima beans','root beer'}  
            c = 'food';  
        case {'red','green','blue'}  
            c = 'color';  
        otherwise  
            c = 'not sure';  
    end
```

Now we exercise the switch statements

```
--> switch_test('root beer')
```

```
ans =  
food
```

```
--> switch_test('red')
```

```
ans =  
color
```

```
--> switch_test('carpet')
```

```
ans =  
not sure
```

5.11 TRY-CATCH Try and Catch Statement

5.11.1 Usage

The `try` and `catch` statements are used for error handling and control. A concept present in C++, the `try` and `catch` statements are used with two statement blocks as follows

```
try  
    statements_1  
catch  
    statements_2  
end
```

The meaning of this construction is: try to execute `statements_1`, and if any errors occur during the execution, then execute the code in `statements_2`. An error can either be a FreeMat generated error (such as a syntax error in the use of a built in function), or an error raised with the `error` command.

5.11.2 Examples

Here is an example of a function that uses error control via `try` and `catch` to check for failures in `fopen`.

```
read_file.m  
function c = read_file(filename)  
try  
    fp = fopen(filename,'r');  
    c = fgetline(fp);  
    fclose(fp);  
catch  
    c = ['could not open file because of error : ' lasterr]  
end
```

Now we try it on an example file - first one that does not exist, and then on one that we create (so that we know it exists).

```
--> read_file('this_filename_is_invalid')

c =
could not open file because of error :Access mode r requires file to exist

ans =
could not open file because of error :Access mode r requires file to exist
--> fp = fopen('test_text.txt','w');
--> fprintf(fp,'a line of text\n');
--> fclose(fp);
--> read_file('test_text.txt')

ans =
a line of text
```

5.12 WARNING Emits a Warning Message

5.12.1 Usage

The `warning` function causes a warning message to be sent to the user. The general syntax for its use is

```
warning(s)
```

where `s` is the string message containing the warning.

5.13 WHILE While Loop

5.13.1 Usage

The `while` loop executes a set of statements as long as a the test condition remains `true`. The syntax of a `while` loop is

```
while test_expression
    statements
end
```

Note that a conditional expression is considered true if the real part of the result of the expression contains any non-zero elements (this strange convention is adopted for compatibility with MATLAB).

5.13.2 Examples

Here is a `while` loop that adds the integers from 1 to 100:

```
--> accum = 0;
--> k=1;
--> while (k<100), accum = accum + k; k = k + 1; end
--> accum
```

```
ans =
    4950
```

```
testeq.m
function x = testeq(a,b)
    if (size(a,1) ~= size(b,1) || size(a,2) ~= size(b,2))
        x = 0;
        return;
    end
    d = full(a)-full(b);
    if (strcmp(typeof(d),'double') | strcmp(typeof(d),'dcomplex'))
        x = isempty(find(abs(d)>10*eps));
    else
        x = isempty(find(abs(d)>10*feps));
    end
```

Chapter 6

FreeMat Functions

6.1 ADDPATH Add

6.1.1 Usage

The `addpath` routine adds a set of directories to the current path. The first form takes a single directory and adds it to the beginning or top of the path:

```
addpath('directory')
```

The second form add several directories to the top of the path:

```
addpath('dir1','dir2',...,'dirn')
```

Finally, you can provide a flag to control where the directories get added to the path

```
addpath('dir1','dir2',...,'dirn','-flag')
```

where if `flag` is either `'-0'` or `'-begin'`, the directories are added to the top of the path, and if the `flag` is either `'-1'` or `'-end'` the directories are added to the bottom (or end) of the path.

6.2 ASSIGNIN Assign Variable in Workspace

6.2.1 Usage

The `assignin` function allows you to assign a value to a variable in either the callers work space or the base work space. The syntax for `assignin` is

```
assignin(workspace,variablename,value)
```

The argument `workspace` must be either `'caller'` or `'base'`. If it is `'caller'` then the variable is assigned in the caller's work space. That does not mean the caller of `assignin`, but the caller of the current function or script. On the other hand if the argument is `'base'`, then the assignment is done in the base work space. Note that the variable is created if it does not already exist.

6.3 BLASLIB Select BLAS library

6.3.1 Usage

The `blaslib` function allows you to select the FreeMat blas library. It has two modes of operation. The first is to select blas library:

```
blaslib LIB_NAME
```

If you want to see current blas library selected issue a `blaslib` command with no arguments.

```
blaslib
```

the returned list of libraries will include an asterix next to the library currently selected. If no library is selected, FreeMat will use its internal (reference) BLAS implementation, which is slow, but portable.

6.4 BUILTIN Evaluate Builtin Function

6.4.1 Usage

The `builtin` function evaluates a built in function with the given name, bypassing any overloaded functions. The syntax of `builtin` is

```
[y1,y2,...,yn] = builtin(fname,x1,x2,...,xm)
```

where `fname` is the name of the function to call. Apart from the fact that `fname` must be a string, and that `builtin` always calls the non-overloaded method, it operates exactly like `feval`. Note that unlike MATLAB, `builtin` does not force evaluation to an actual compiled function. It simply subverts the activation of overloaded method calls.

6.5 CLC Clear Display

6.5.1 Usage

The `clc` function clears the current display. The syntax for its use is

```
clc
```

6.6 CLOCK Get Current Time

6.6.1 Usage

Returns the current date and time as a vector. The syntax for its use is

```
y = clock
```

where `y` has the following format:

```
y = [year month day hour minute seconds]
```

6.6.2 Example

Here is the time that this manual was last built:

```
--> clock

ans =

    1.0e+03 *
    2.0090    0.0100    0.0040    0.0220    0.0260    0.0321
```

6.7 CLOCKTOTIME Convert Clock Vector to Epoch Time

6.7.1 Usage

Given the output of the `clock` command, this function computes the epoch time, i.e, the time in seconds since January 1,1970 at 00:00:00 UTC. This function is most useful for calculating elapsed times using the `clock`, and should be accurate to less than a millisecond (although the true accuracy depends on accuracy of the argument vector). The usage for `clocktotime` is

```
y = clocktotime(x)
```

where `x` must be in the form of the output of `clock`, that is

```
x = [year month day hour minute seconds]
```

6.7.2 Example

Here is an example of using `clocktotime` to time a delay of 1 second

```
--> x = clock

x =

    1.0e+03 *
    2.0090    0.0100    0.0040    0.0220    0.0260    0.0322

--> sleep(1)
--> y = clock

y =

    1.0e+03 *
    2.0090    0.0100    0.0040    0.0220    0.0260    0.0332

--> clocktotime(y) - clocktotime(x)

ans =
    1.0010
```

6.8 COMPUTER Computer System FreeMat is Running On

6.8.1 Usage

Returns a string describing the name of the system FreeMat is running on. The exact value of this string is subject to change, although the 'MAC' and 'PCWIN' values are probably fixed.

```
str = computer
```

Currently, the following return values are defined

- 'PCWIN' - MS Windows
- 'MAC' - Mac OS X
- 'UNIX' - All others

6.9 DIARY Create a Log File of Console

6.9.1 Usage

The `diary` function controls the creation of a log file that duplicates the text that would normally appear on the console. The simplest syntax for the command is simply:

```
diary
```

which toggles the current state of the diary command. You can also explicitly set the state of the diary command via the syntax

```
diary off
```

or

```
diary on
```

To specify a filename for the log (other than the default of `diary`), you can use the form:

```
diary filename
```

or

```
diary('filename')
```

which activates the diary with an output filename of `filename`. Note that the `diary` command is thread specific, but that the output is appended to a given file. That means that if you call `diary` with the same filename on multiple threads, their outputs will be intermingled in the log file (just as on the console). Because the `diary` state is tied to individual threads, you cannot retrieve the current diary state using the `get(0, 'Diary')` syntax from MATLAB. Instead, you must call the `diary` function with no inputs and one output:

```
state = diary
```

which returns a logical 1 if the output of the current thread is currently going to a diary, and a logical 0 if not.

6.10 DOCLI Start a Command Line Interface

6.10.1 Usage

The `docli` function is the main function that you interact with when you run FreeMat. I am not sure why you would want to use it, but hey - its there if you want to use it.

6.11 EDIT Open Editor Window

6.11.1 Usage

Brings up the editor window. The arguments of `edit` function are names of files for editing:

```
edit file1 file2 file3
```

6.12 EDITOR Open Editor Window

6.12.1 Usage

Brings up the editor window. The `editor` function takes no arguments:

```
editor
```

6.13 ERRORCOUNT Retrieve the Error Counter for the Interpreter

6.13.1 Usage

This routine retrieves the internal counter for the interpreter, and resets it to zero. The general syntax for its use is

```
count = errorcount
```

6.14 ETIME Elapsed Time Function

6.14.1 Usage

The `etime` calculates the elapsed time between two `clock` vectors `x1` and `x2`. The syntax for its use is

```
y = etime(x1,x2)
```

where `x1` and `x2` are in the `clock` output format

```
x = [year month day hour minute seconds]
```

6.14.2 Example

Here we use `etime` as a substitute for `tic` and `toc`

```
--> x1 = clock;
--> sleep(1);
--> x2 = clock;
--> etime(x2,x1);
```

6.15 EVAL Evaluate a String

6.15.1 Usage

The `eval` function evaluates a string. The general syntax for its use is

```
eval(s)
```

where `s` is the string to evaluate. If `s` is an expression (instead of a set of statements), you can assign the output of the `eval` call to one or more variables, via

```
x = eval(s)
[x,y,z] = eval(s)
```

Another form of `eval` allows you to specify an expression or set of statements to execute if an error occurs. In this form, the syntax for `eval` is

```
eval(try_clause,catch_clause),
```

or with return values,

```
x = eval(try_clause,catch_clause)
[x,y,z] = eval(try_clause,catch_clause)
```

These later forms are useful for specifying defaults. Note that both the `try_clause` and `catch_clause` must be expressions, as the equivalent code is

```
try
    [x,y,z] = try_clause
catch
    [x,y,z] = catch_clause
end
```

so that the assignment must make sense in both cases.

6.15.2 Example

Here are some examples of `eval` being used.

```
--> eval('a = 32')
```

```
a =  
32
```

```
--> b = eval('a')
```

```
b =  
32
```

The primary use of the `eval` statement is to enable construction of expressions at run time.

```
--> s = ['b = a' ' + 2']
```

```
s =  
b = a + 2  
--> eval(s)
```

```
b =  
34
```

Here we demonstrate the use of the catch-clause to provide a default value

```
--> a = 32
```

```
a =  
32
```

```
--> b = eval('a', '1')
```

```
b =  
32
```

```
--> b = eval('z', 'a+1')
```

```
b =  
33
```

Note that in the second case, `b` takes the value of 33, indicating that the evaluation of the first expression failed (because `z` is not defined).

6.16 EVALIN Evaluate a String in Workspace

6.16.1 Usage

The `evalin` function is similar to the `eval` function, with an additional argument up front that indicates the workspace that the expressions are to be evaluated in. The various syntaxes for `evalin` are:

```

evalin(workspace,expression)
x = evalin(workspace,expression)
[x,y,z] = evalin(workspace,expression)
evalin(workspace,try_clause,catch_clause)
x = evalin(workspace,try_clause,catch_clause)
[x,y,z] = evalin(workspace,try_clause,catch_clause)

```

The argument `workspace` must be either 'caller' or 'base'. If it is 'caller', then the expression is evaluated in the caller's work space. That does not mean the caller of `evalin`, but the caller of the current function or script. On the other hand if the argument is 'base', then the expression is evaluated in the base work space. See `eval` for details on the use of each variation.

6.17 EXIT Exit Program

6.17.1 Usage

The usage is

```
exit
```

Quits FreeMat. This script is a simple synonym for `quit`.

6.18 FEVAL Evaluate a Function

6.18.1 Usage

The `feval` function executes a function using its name. The syntax of `feval` is

```
[y1,y2,...,yn] = feval(f,x1,x2,...,xm)
```

where `f` is the name of the function to evaluate, and `xi` are the arguments to the function, and `yi` are the return values.

Alternately, `f` can be a function handle to a function (see the section on `function handles` for more information).

Finally, FreeMat also supports `f` being a user defined class in which case it will attempt to invoke the `subsref` method of the class.

6.18.2 Example

Here is an example of using `feval` to call the `cos` function indirectly.

```
--> feval('cos',pi/4)
```

```
ans =
    0.7071
```

Now, we call it through a function handle

```
--> c = @cos

c =
  @cos
--> feval(c,pi/4)

ans =
  0.7071
```

Here we construct an inline object (which is a user-defined class) and use `feval` to call it

```
--> afunc = inline('cos(t)+sin(t)', 't')

afunc =
  inline function object
  f(t) = cos(t)+sin(t)
--> feval(afunc,pi)

ans =
  -1.0000

--> afunc(pi)

ans =
  -1.0000
```

In both cases, (the `feval` call and the direct invocation), FreeMat calls the `subsref` method of the class, which computes the requested function.

6.19 FILESEP Directory Separation Character

6.19.1 Usage

The `filesep` routine returns the character used to separate directory names on the current platform (basically, a forward slash for Windows, and a backward slash for all other OSes). The syntax is simple:

```
x = filesep
```

6.20 HELP Help

6.20.1 Usage

Displays help on a function available in FreeMat. The help function takes one argument:

```
help topic
```

where `topic` is the topic to look for help on. For scripts, the result of running `help` is the contents of the comments at the top of the file. If FreeMat finds no comments, then it simply displays the function declaration.

6.21 HELPWIN Online Help Window

6.21.1 Usage

Brings up the online help window with the FreeMat manual. The `helpwin` function takes no arguments:

```
helpwin
helpwin FunctionName
```

6.22 JITCONTROL Control the Just In Time Compiler

6.22.1 Usage

The `jitcontrol` functionality in FreeMat allows you to control the use of the Just In Time (JIT) compiler. Starting in FreeMat version 4, the JIT compiler is enabled by default on all platforms where it is successfully built. The JIT compiler should significantly improve the performance of loop intensive, scalar code. As development progresses, more and more functionality will be enabled under the JIT. In the mean time (if you use the GUI version of FreeMat) you can use the JIT chat window to get information on why your code was JIT compiled (or not).

```
do_jit_test.m
function test_val = do_jit_test(test_name)
printf('Running %s without JIT...',test_name);
jitcontrol off
tic; A1 = feval(test_name); nojit = toc;
jitcontrol on
% Run it twice to remove the compile time
tic; A2 = feval(test_name); wjit = toc;
tic; A2 = feval(test_name); wjit = toc;
printf('Speedup is %g...',nojit/wjit);
test_val = issame(A1,A2);
if (test_val)
    printf('PASS\n');
else
    printf('FAIL\n');
end
```

```
jit_test001.m
function A = jit_test001
A = zeros(1,100000);
for i=1:100000; A(i) = i; end;
```

```
jit_test002.m
function A = jit_test002
A = zeros(512);
for i=1:512;
    for j=1:512;
        A(i,j) = i-j;
    end
end
end
```

```
jit_test003.m
function A = jit_test003
A = zeros(512);
for i=1:512
    for j=1:512
        k = i-j;
        if ((k < 5) && (k>-5))
            A(i,j) = k;
        end
    end
end
end
```

```
jit_test004.m
function A = jit_test004
A = zeros(512);
for i=1:512
    for j=1:512
        k = i-j;
        if ((k < 5) && (k>-5))
            A(i,j) = k;
        else
            A(i,j) = i+j;
        end
    end
end
end
```

```
jit_test005.m
function A = jit_test005
A = zeros(1,1000);
for i=1:1000
    for m=1:100
        if (i > 50)
            k = 55;
        elseif (i > 40)
            k = 45;
        elseif (i > 30)
            k = 35;
        else
```

```
        k = 15;
    end
end
A(i) = k;
end

    jit_test006.m
function A = jit_test006
A = zeros(1000);
m = 0;
for i=1:1000;
    for j=1:1000;
        A(i,j) = i-j;
        m = i;
    end
end
A = m;

    jit_test007.m
function A = jit_test007
B = ones(1,100000);
j = 0;
for i=1:100000
    j = j + B(1,i);
end
A = j;

    jit_test008.m
function A = jit_test008
A = zeros(1,10000);
for i=1:10010;
    A(i) = i;
end

    jit_test009.m
function A = jit_test009
B = ones(1,10000);
A = 0;
for k=1:10000
    A = A + B(k);
end

    jit_test010.m
function A = jit_test010
B = 1:100000;
A = zeros(1,100000);
for i=1:100000
```

```
A(i) = B(i);
end

jit_test011.m
function A = jit_test011
A = zeros(1,100000);
B = 5;
for i=1:100000
    C = B + i;
    A(i) = C;
end

jit_test012.m
function A = jit_test012
A = zeros(400);
for i=1:500
    for j=1:500;
        A(i,j) = i-j;
    end
end

jit_test013.m
function A = jit_test013
A = zeros(512);
for i=1:512
    for j=1:512
        A(i,j) = abs(i-j);
    end
end

jit_test014.m
function A = jit_test014
N = 500;
A = zeros(N,N);
for i=1:N;
    for j=1:N;
        A(j,i) = abs(i-j);
    end
end

jit_test015.m
function A = jit_test015
A = 0;
for i=0:200000
    A = A + abs(sin(i/200000*pi));
end
```

```
jit_test016.m
function A = jit_test016
A = 0;
for i=0:200000
    A = A + abs(cos(i/200000*pi));
end
```

```
jit_test017.m
function A = jit_test017
A = 0;
for j=1:100
    B = j;
    for i=1:10000
        A = A + B;
    end
end
```

```
jit_test018.m
function A = jit_test018
A = 0;
for i=1.2:4000.2
    A = A + i;
end
```

```
jit_test019.m
function A = jit_test019
A = 150000;
C = 0;
for b=1:A
    C = C + b;
end
A = C;
```

```
jit_test020.m
function a = jit_test020
a = zeros(1,10000);
for i=1:10000;
    a(i) = sec(i/10000);
end
```

```
jit_test021.m
function A = jit_test021
A = zeros(100,100);
for i=1:(100*100)
    A(i) = i;
end
```

```
    jit_test022.m
function A = jit_test022
    A = 1:1000000;
    B = 0;
    for i=1:1000000;
        B = B + A(i);
    end
    A = B;

    jit_test023.m
function A = jit_test023
    A = zeros(1,100000);
    for i=1:100000;
        A(i) = i/5.0;
    end

    jit_test024.m
function A = jit_test024
    A = zeros(1,10000);
    for j=1:10
        for i=1:10000;
            A(i) = tjit_sum(A(i),i);
        end;
    end;

function y = tjit_sum(a,b)
    y = a + tjit_double(b);

function y = tjit_double(a)
    y = a*2;
```

6.23 MFILENAME Name of Current Function

6.23.1 Usage

Returns a string describing the name of the current function. For M-files this string will be the complete filename of the function. This is true even for subfunctions. The syntax for its use is

```
y = mfilename
```

6.24 PATH Get or Set FreeMat Path

6.24.1 Usage

The `path` routine has one of the following syntaxes. In the first form

```
x = path
```

`path` simply returns the current path. In the second, the current path is replaced by the argument string `'thepath'`

```
path('thepath')
```

In the third form, a new path is appended to the current search path

```
path(path, 'newpath')
```

In the fourth form, a new path is prepended to the current search path

```
path('newpath', path)
```

In the final form, the `path` command prints out the current path

```
path
```

6.25 PATHSEP Path Directories Separation Character

6.25.1 Usage

The `pathsep` routine returns the character used to separate multiple directories on a path string for the current platform (basically, a semicolon for Windows, and a regular colon for all other OSes). The syntax is simple:

```
x = pathsep
```

6.26 PATHTOOL Open Path Setting Tool

6.26.1 Usage

Brings up the `pathtool` dialog. The `pathtool` function takes no arguments:

```
pathtool
```

6.27 PCODE Convert a Script or Function to P-Code

6.27.1 Usage

Writes out a script or function as a P-code function. The general syntax for its use is:

```
pcode fun1 fun2 ...
```

The compiled functions are written to the current directory.

6.28 PROFILER Control Profiling

6.28.1 Usage

The `profile` function allows you to control the FreeMat profiler. It has two modes of operation. The first is to enable-disable the profiler. To turn on profiling:

```
profiler on
```

to turn off profiling, use

```
profiler off
```

Note that regardless of the state of the profiler, only functions and scripts are profiled. Commands entered on the command line are not profiled. To see information that has accumulated in a profile, you use the variant of the command:

```
profiler list symbol1 symbol2 ...
```

where `symbol1`, `symbol2` are the functions or scripts on which profiling information is desired. If you want to see current profile status issue a `profile` command with no arguments.

```
profiler
```

6.29 QUIET Control the Verbosity of the Interpreter

6.29.1 Usage

The `quiet` function controls how verbose the interpreter is when executing code. The syntax for the function is

```
quiet flag
```

where `flag` is one of

- 'normal' - normal output from the interpreter
- 'quiet' - only intentional output (e.g. `printf` calls and `disp` calls) is printed. The output of expressions that are not terminated in semicolons are not printed.
- 'silent' - nothing is printed to the output.

The `quiet` command also returns the current quiet flag.

6.30 QUIT Quit Program

6.30.1 Usage

The `quit` statement is used to immediately exit the FreeMat application. The syntax for its use is

```
quit
```

6.31 REHASH Rehash Directory Caches

6.31.1 Usage

Usually, FreeMat will automatically determine when M Files have changed, and pick up changes you have made to M files. Sometimes, you have to force a refresh. Use the `rehash` command for this purpose. The syntax for its use is

```
rehash
```

6.32 RESCAN Rescan M Files for Changes

6.32.1 Usage

Usually, FreeMat will automatically determine when M Files have changed, and pick up changes you have made to M files. Sometimes, you have to force a refresh. Use the `rescan` command for this purpose. The syntax for its use is

```
rescan
```

6.33 ROOTPATH Set FreeMat Root Path

6.33.1 Usage

In order to function properly, FreeMat needs to know where to find the `toolbox` directory as well as the `help` directory. These directories are located on what is known as the `root path`. Normally, FreeMat should know where these directories are located. However under some circumstances (usually when FreeMat is installed into a non-default location), it may be necessary to indicate a different root path location, or to specify a particular one. Note that on the Mac OS platform, FreeMat is installed as a bundle, and will use the toolbox that is installed in the bundle regardless of the setting for `rootpath`. For Linux, FreeMat will typically use `/usr/local/share/FreeMat-<Version>/` for the root path. Installations from source code will generally work, but binary installations (e.g., from an RPM) may need to have the `rootpath` set.

The `rootpath` function has two forms. The first form takes no arguments and allows you to browse to the `rootpath` directory

```
rootpath
```

The second form will set a `rootpath` directly from the command line

```
rootpath(path)
```

where `path` is the full path to where the `toolbox` and `help` directories are located. For example, `rootpath('/usr/share/FreeMat-4.0')`. Changes to `rootpath` are persistent (you do not need to run it every time you start FreeMat).

6.34 SIMKEYS Simulate Keypresses from the User

6.34.1 Usage

This routine simulates keystrokes from the user on FreeMat. The general syntax for its use is

```
otext = simkeys(text)
```

where `text` is a string to simulate as input to the console. The output of the commands are captured and returned in the string `otext`. This is primarily used by the testing infrastructure.

6.35 SLEEP Sleep For Specified Number of Seconds

6.35.1 Usage

Suspends execution of FreeMat for the specified number of seconds. The general syntax for its use is

```
sleep(n),
```

where `n` is the number of seconds to wait.

6.36 SOURCE Execute an Arbitrary File

6.36.1 Usage

The `source` function executes the contents of the given filename one line at a time (as if it had been typed at the `-->` prompt). The `source` function syntax is

```
source(filename)
```

where `filename` is a string containing the name of the file to process.

6.36.2 Example

First, we write some commands to a file (note that it does not end in the usual `.m` extension):

```
source_test
a = 32;
b = a;
printf('a is %d and b is %d\n',a,b);
```

Now we source the resulting file.

```
--> clear a b
--> source source_test
a is 32 and b is 32
```

6.37 STARTUP Startup Script

6.37.1 Usage

Upon starting, FreeMat searches for a script names `startup.m`, and if it finds it, it executes it. This script can be in the current directory, or on the FreeMat path (set using `setpath`). The contents of `startup.m` must be a valid script (not a function).

6.38 TIC Start Stopwatch Timer

6.38.1 Usage

Starts the stopwatch timer, which can be used to time tasks in FreeMat. The `tic` takes no arguments, and returns no outputs. You must use `toc` to get the elapsed time. The usage is

```
tic
```

6.38.2 Example

Here is an example of timing the solution of a large matrix equation.

```
--> A = rand(100);  
--> b = rand(100,1);  
--> tic; c = A\b; toc
```

```
ans =  
    0.0240
```

6.39 TOC Stop Stopwatch Timer

6.39.1 Usage

Stop the stopwatch timer, which can be used to time tasks in FreeMat. The `toc` function takes no arguments, and returns no outputs. You must use `toc` to get the elapsed time. The usage is

```
toc
```

6.39.2 Example

Here is an example of timing the solution of a large matrix equation.

```
--> A = rand(100);  
--> b = rand(100,1);  
--> tic; c = A\b; toc
```

```
ans =  
    0.0220
```

6.40 TYPE Type

6.40.1 Usage

Displays the content of a m-script file or an ascii file. The type function takes one argument:

```
type filename
type ('filename')
```

This command will first look for filename in FreeMat search path, if not found it will look for m-script file by adding '.m' to filename

6.41 VERSION The Current Version Number

6.41.1 Usage

The `version` function returns the current version number for FreeMat (as a string). The general syntax for its use is

```
v = version
```

6.41.2 Example

The current version of FreeMat is

```
--> version
```

```
ans =
4.0
```

6.42 VERSTRING The Current Version String

6.42.1 Usage

The `verstring` function returns the current version string for FreeMat. The general syntax for its use is

```
version = verstring
```

6.42.2 Example

The current version of FreeMat is

```
--> verstring
```

```
ans =
FreeMat v4.0
```


Chapter 7

Debugging FreeMat Code

7.1 DBAUTO Control Dbauto Functionality

7.1.1 Usage

The `dbauto` functionality in FreeMat allows you to debug your FreeMat programs. When `dbauto` is `on`, then any error that occurs while the program is running causes FreeMat to stop execution at that point and return you to the command line (just as if you had placed a `keyboard` command there). You can then examine variables, modify them, and resume execution using `return`. Alternately, you can exit out of all running routines via a `retall` statement. Note that errors that occur inside of `try/catch` blocks do not (by design) cause auto breakpoints. The `dbauto` function toggles the `dbauto` state of FreeMat. The syntax for its use is

```
dbauto(state)
```

where `state` is either

```
dbauto('on')
```

to activate `dbauto`, or

```
dbauto('off')
```

to deactivate `dbauto`. Alternately, you can use FreeMat's string-syntax equivalence and enter

```
dbauto on
```

or

```
dbauto off
```

to turn `dbauto` on or off (respectively). Entering `dbauto` with no arguments returns the current state (either `'on'` or `'off'`).

7.2 DBDELETE Delete a Breakpoint

7.2.1 Usage

The `dbdelete` function deletes a breakpoint. The syntax for the `dbdelete` function is

```
dbdelete(num)
```

where `num` is the number of the breakpoint to delete.

7.3 DBDown Move Down One Debug Level

7.3.1 Usage

The `dbdown` function moves up one level in the debug hierarchy. The syntax for the `dbdown` function is

```
dbdown
```

7.4 DBLIST List Breakpoints

7.4.1 Usage

List the current set of breakpoints. The syntax for the `dblist` is simply

```
dblist
```

7.5 DBSTEP Step N Statements

7.5.1 Usage

Step `N` statements during debug mode. The syntax for this is either

```
dbstep(N)
```

to step `N` statements, or

```
dbstep
```

to step one statement.

7.6 DBSTOP

7.6.1 Usage

Set a breakpoint. The syntax for this is:

```
dbstop(funcname,linenumber)
```

where `funcname` is the name of the function where we want to set the breakpoint, and `linenumber` is the line number.

7.7 DBUP Move Up One Debug Level

7.7.1 Usage

The `dbup` function moves up one level in the debug hierarchy. The syntax for the `dbup` function is

```
dbup
```

7.8 FDUMP Dump Information on Function

7.8.1 Usage

Dumps information about a function (diagnostic information only)

```
fdump fname
```


Chapter 8

Sparse Matrix Support

8.1 EIGS Sparse Matrix Eigendecomposition

8.1.1 Usage

Computes the eigendecomposition of a sparse square matrix. The `eigs` function has several forms. The most general form is

```
[V,D] = eigs(A,k,sigma)
```

where `A` is the matrix to analyze, `k` is the number of eigenvalues to compute and `sigma` determines which eigenvalues to solve for. Valid values for `sigma` are 'lm' - largest magnitude 'sm' - smallest magnitude 'la' - largest algebraic (for real symmetric problems) 'sa' - smallest algebraic (for real symmetric problems) 'be' - both ends (for real symmetric problems) 'lr' - largest real part 'sr' - smallest real part 'li' - largest imaginary part 'si' - smallest imaginary part scalar - find the eigenvalues closest to `sigma`. The returned matrix `V` contains the eigenvectors, and `D` stores the eigenvalues. The related form

```
d = eigs(A,k,sigma)
```

computes only the eigenvalues and not the eigenvectors. If `sigma` is omitted, as in the forms

```
[V,D] = eigs(A,k)
```

and

```
d = eigs(A,k)
```

then `eigs` returns the largest magnitude eigenvalues (and optionally the associated eigenvectors). As an even simpler form, the forms

```
[V,D] = eigs(A)
```

and

```
d = eigs(A)
```

then `eigs` returns the six largest magnitude eigenvalues of `A` and optionally the eigenvectors. The `eigs` function uses ARPACK to compute the eigenvectors and/or eigenvalues. Note that due to a limitation in the interface into ARPACK from FreeMat, the number of eigenvalues that are to be computed must be strictly smaller than the number of columns (or rows) in the matrix.

8.1.2 Example

Here is an example of using `eigs` to calculate eigenvalues of a matrix, and a comparison of the results with `eig`

```
--> a = sparse(rand(9));
--> eigs(a)
```

```
ans =
  4.2673 + 0.0000i
 -0.8372 + 0.0000i
 -0.7445 + 0.3406i
 -0.7445 - 0.3406i
  0.6653 + 0.2082i
  0.6653 - 0.2082i
```

```
--> eig(full(a))
```

```
ans =
  4.2673 + 0.0000i
 -0.8372 + 0.0000i
 -0.7445 + 0.3406i
 -0.7445 - 0.3406i
  0.1634 + 0.4595i
  0.1634 - 0.4595i
  0.6653 + 0.2082i
  0.6653 - 0.2082i
  0.0479 + 0.0000i
```

Next, we exercise some of the variants of `eigs`:

```
--> eigs(a,4,'sm')
```

```
ans =
  0.0479 + 0.0000i
  0.1634 - 0.4595i
  0.1634 + 0.4595i
  0.6653 - 0.2082i
```

```
--> eigs(a,4,'lr')
```

```
ans =
```

```

4.2673 + 0.0000i
0.6653 - 0.2082i
0.6653 + 0.2082i
0.1634 - 0.4595i

--> eigs(a,4,'sr')
```

```

ans =
-0.8372 + 0.0000i
-0.7445 - 0.3406i
-0.7445 + 0.3406i
0.0479 + 0.0000i
```

8.2 FULL Convert Sparse Matrix to Full Matrix

8.2.1 Usage

Converts a sparse matrix to a full matrix. The syntax for its use is

```
y = full(x)
```

The type of `x` is preserved. Be careful with the function. As a general rule of thumb, if you can work with the `full` representation of a function, you probably do not need the sparse representation.

8.2.2 Example

Here we convert a full matrix to a sparse one, and back again.

```
--> a = [1,0,4,2,0;0,0,0,0,0;0,1,0,0,2]
```

```

a =
1 0 4 2 0
0 0 0 0 0
0 1 0 0 2
```

```
--> A = sparse(a)
```

```

A =
1 1 1
3 2 1
1 3 4
1 4 2
3 5 2
--> full(A)
```

```

ans =
1 0 4 2 0
```

```
0 0 0 0 0
0 1 0 0 2
```

8.3 SPARSE Construct a Sparse Matrix

8.3.1 Usage

Creates a sparse matrix using one of several formats. The first creates a sparse matrix from a full matrix

```
y = sparse(x).
```

The second form creates a sparse matrix containing all zeros that is of the specified size (the sparse equivalent of `zeros`).

```
y = sparse(m,n)
```

where `m` and `n` are integers. Just like the `zeros` function, the sparse matrix returned is of type `float`. The third form constructs a sparse matrix from the IJV syntax. It has two forms. The first version autosizes the sparse matrix

```
y = sparse(i,j,v)
```

while the second version uses an explicit size specification

```
y = sparse(i,j,v,m,n)
```

8.4 SPEYE Sparse Identity Matrix

8.4.1 Usage

Creates a sparse identity matrix of the given size. The syntax for its use is

```
y = speye(m,n)
```

which forms an `m x n` sparse matrix with ones on the main diagonal, or

```
y = speye(n)
```

which forms an `n x n` sparse matrix with ones on the main diagonal. The matrix type is a `float` single precision matrix.

8.4.2 Example

The following creates a 5000 by 5000 identity matrix, which would be difficult to do using `sparse(eye(5000))` because of the large amount of intermediate storage required.

```
--> I = speye(5000);
--> which I
Function I is unknown!
--> full(I(1:10,1:10))
```

```
ans =
 1 0 0 0 0 0 0 0 0 0
 0 1 0 0 0 0 0 0 0 0
 0 0 1 0 0 0 0 0 0 0
 0 0 0 1 0 0 0 0 0 0
 0 0 0 0 1 0 0 0 0 0
 0 0 0 0 0 1 0 0 0 0
 0 0 0 0 0 0 1 0 0 0
 0 0 0 0 0 0 0 1 0 0
 0 0 0 0 0 0 0 0 1 0
 0 0 0 0 0 0 0 0 0 1
```

8.5 SPONES Sparse Ones Function

8.5.1 Usage

Returns a sparse float matrix with ones where the argument matrix has nonzero values. The general syntax for it is

```
y = spones(x)
```

where x is a matrix (it may be full or sparse). The output matrix y is the same size as x , has type float, and contains ones in the nonzero positions of x .

8.5.2 Examples

Here are some examples of the `spones` function

```
--> a = [1,0,3,0,5;0,0,2,3,0;1,0,0,0,1]
```

```
a =
 1 0 3 0 5
 0 0 2 3 0
 1 0 0 0 1
```

```
--> b = spones(a)
```

```
b =
 1 1 1
 1 2 1
 1 3 1
 1 4 1
```

```

1 5 1
--> full(b)

```

```

ans =
1 1 1 1 1
0 0 0 0 0
0 0 0 0 0

```

8.6 SPRAND Sparse Uniform Random Matrix

8.6.1 Usage

Creates a sparse matrix with uniformly distributed random entries (on $[0,1]$). The syntax for its use is

```
y = sprand(x)
```

where x is a sparse matrix, where y is a sparse matrix that has random entries where x is nonzero. The second form specifies the size of the matrix and the density

```
y = sprand(m,n,density)
```

where m is the number of rows in the output, n is the number of columns in the output, and $density$ (which is between 0 and 1) is the density of nonzeros in the resulting matrix. Note that for very high densities the actual density of the output matrix may differ from the density you specify. This difference is a result of the way the random entries into the matrix are generated. If you need a very dense random matrix, it is better to generate a full matrix and zero out the entries you do not need.

8.6.2 Examples

Here we seed `sprand` with a full matrix (to demonstrate how the structure of the output is determined by the input matrix when using the first form).

```
--> x = [1,0,0;0,0,1;1,0,0]
```

```

x =
1 0 0
0 0 1
1 0 0

```

```
--> y = sprand(x)
```

```

y =
1 1 0.478912
3 1 0.039372
2 3 0.806381
--> full(y)

```

```
ans =
    0.4789         0         0
         0         0    0.8064
    0.0394         0         0
```

The more generic version with a density of 0.001. On many systems the following is impossible using full matrices

```
--> y = sprand(10000,10000,.001);
--> nnz(y)/10000^2
```

```
ans =
    9.9946e-04
```

8.7 SPRANDN Sparse Normal Random Matrix

8.7.1 Usage

Creates a sparse matrix with normally distributed random entries (mean 0, sigma 1). The syntax for its use is

```
y = sprandn(x)
```

where **x** is a sparse matrix, where **y** is a sparse matrix that has random entries where **x** is nonzero. The second form specifies the size of the matrix and the density

```
y = sprandn(m,n,density)
```

where **m** is the number of rows in the output, **n** is the number of columns in the output, and **density** (which is between 0 and 1) is the density of nonzeros in the resulting matrix. Note that for very high densities the actual density of the output matrix may differ from the density you specify. This difference is a result of the way the random entries into the matrix are generated. If you need a very dense random matrix, it is better to generate a full matrix and zero out the entries you do not need.

8.7.2 Examples

Here we seed `sprandn` with a full matrix (to demonstrate how the structure of the output is determined by the input matrix when using the first form).

```
--> x = [1,0,0;0,0,1;1,0,0]
```

```
x =
    1 0 0
    0 0 1
    1 0 0
```

```
--> y = sprandn(x)
```

```

y =
  1 1 0.770419
  3 1 0.332694
  2 3 -0.589851
--> full(y)

ans =
  0.7704         0         0
         0         0    -0.5899
  0.3327         0         0

```

The more generic version with a density of 0.001. On many systems the following is impossible using full matrices

```

--> y = sprandn(10000,10000,.001);
--> nnz(y)/10000^2

ans =
  9.9953e-04

```

8.8 SPY Visualize Sparsity Pattern of a Sparse Matrix

8.8.1 Usage

Plots the sparsity pattern of a sparse matrix. The syntax for its use is

```
spy(x)
```

which uses a default color and symbol. Alternately, you can use

```
spy(x,colspec)
```

where `colspec` is any valid color and symbol spec accepted by `plot`.

8.8.2 Example


First, an example of a random sparse matrix.

```

--> y = sprand(1000,1000,.001);
--> spy(y,'ro')

```


which is shown here



Here is a sparse matrix with a little more structure. First we build a sparse matrix with block diagonal structure, and then use `spy` to visualize the structure.

```
--> A = sparse(1000,1000);  
--> for i=1:25; A((1:40) + 40*(i-1), (1:40) + 40*(i-1)) = 1; end;  
--> spy(A, 'gx')
```

with the result shown here



Chapter 9

Mathematical Functions

9.1 ACOS Inverse Trigonometric Arccosine Function

9.1.1 Usage

Computes the `acos` function for its argument. The general syntax for its use is

```
y = acos(x)
```

where `x` is an `n`-dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `acos` function. Output `y` is of the same size and type as the input `x`, (unless `x` is an integer, in which case `y` is a `double` type).

9.1.2 Function Internals

Mathematically, the `acos` function is defined for all arguments `x` as

$$\operatorname{acos} x \equiv \frac{\pi i}{2} + i \log \left(ix + \sqrt{1 - x^2} \right).$$


For real valued variables `x` in the range `[-1, 1]`, the function is computed directly using the standard C library's numerical `acos` function. For both real and complex arguments `x`, note that generally

$$\operatorname{acos}(\cos(x)) \neq x,$$

9.1.3 Example

The following code demonstrates the `acos` function over the range `[-1, 1]`.

```
--> t = linspace(-1,1);  
--> plot(t,acos(t))
```



9.2 ACOSD Inverse Cosine Degrees Function

9.2.1 Usage

Computes the inverse cosine of the argument, but returns the argument in degrees instead of radians (as is the case for `acos`). The syntax for its use is

$$y = \text{acosd}(x)$$

9.2.2 Examples

The inverse cosine of $\text{sqrt}(2)/2$ should be 45 degrees:

```
--> acosd(sqrt(2)/2)
```

```
ans =  
45.0000 + 0.0000i
```

and the inverse cosine of 0.5 should be 60 degrees:

```
--> acosd(0.5)
```

```
ans =  
60.0000 + 0.0000i
```

9.3 ACOSH Inverse Hyperbolic Cosine Function

9.3.1 Usage

Computes the inverse hyperbolic cosine of its argument. The general syntax for its use is

$$y = \text{acosh}(x)$$

where x is an n -dimensional array of numerical type.

9.3.2 Function Internals

The `acosh` function is computed from the formula


$$\cosh^{-1}(x) = \log(x + (x^2 - 1)^{0.5})$$

where the `log` (and square root) is taken in its most general sense.

9.3.3 Examples

Here is a simple plot of the inverse hyperbolic cosine function

```
--> x = linspace(1,pi);
--> plot(x,acosh(x)); grid('on');
```



9.4 ACOT Inverse Cotangent Function

9.4.1 Usage

Computes the inverse cotangent of its argument. The general syntax for its use is

```
y = acot(x)
```

where `x` is an `n`-dimensional array of numerical type.

9.4.2 Function Internals


The `acot` function is computed from the formula

$$\cot^{-1}(x) = \tan^{-1}\left(\frac{1}{x}\right)$$

9.4.3 Examples

Here is a simple plot of the inverse cotangent function

```
--> x1 = -2*pi/pi/30:-0.1;
--> x2 = 0.1:pi/30:2*pi;
--> plot(x1,acot(x1),x2,acot(x2)); grid('on');
```



9.5 ACOTD Inverse Cotangent Degrees Function

9.5.1 Usage

Computes the inverse cotangent of its argument in degrees. The general syntax for its use is

```
y = acotd(x)
```

where x is an n -dimensional array of numerical type.

9.6 ACOTH Inverse Hyperbolic Cotangent Function

9.6.1 Usage

Computes the inverse hyperbolic cotangent of its argument. The general syntax for its use is

```
y = acoth(x)
```

where x is an n -dimensional array of numerical type.

9.6.2 Function Internals


The `acoth` function is computed from the formula

$$\coth^{-1}(x) = \tanh^{-1}\left(\frac{1}{x}\right)$$

9.6.3 Examples

Here is a simple plot of the inverse hyperbolic cotangent function

```
--> x = linspace(1,pi);
--> plot(x,acoth(x)); grid('on');
```



9.7 ACSC Inverse Cosecant Function

9.7.1 Usage

Computes the inverse cosecant of its argument. The general syntax for its use is

```
y = acsc(x)
```

where x is an n -dimensional array of numerical type.

9.7.2 Function Internals


The `acosh` function is computed from the formula

$$\csc^{-1}(x) = \sin^{-1}\left(\frac{1}{x}\right)$$

9.7.3 Examples

Here is a simple plot of the inverse cosecant function

```
--> x1 = -10:.01:-1.01;
--> x2 = 1.01:.01:10;
--> plot(x1,acsc(x1),x2,acsc(x2)); grid('on');
```



9.8 ACSCD Inverse Cosecant Degrees Function

9.8.1 Usage

Computes the inverse cosecant of the argument, but returns the argument in degrees instead of radians (as is the case for `acsc`). The syntax for its use is

$$y = \text{acscd}(x)$$

9.8.2 Examples

The inverse cosecant of $2/\sqrt{2}$ should be 45 degrees:

```
--> acscd(2/sqrt(2))
```

```
ans =
  45.0000
```

and the inverse cosecant of 2 should be 30 degrees:

```
--> acscd(0.5)
```

```
ans =
  90.0000 - 75.4561i
```

9.9 ACSCH Inverse Hyperbolic Cosecant Function

9.9.1 Usage

Computes the inverse hyperbolic cosecant of its argument. The general syntax for its use is

$$y = \text{acsch}(x)$$

where x is an n -dimensional array of numerical type.

9.9.2 Function Internals


The `acsch` function is computed from the formula

$$\operatorname{csch}^{-1}(x) = \sinh^{-1}\left(\frac{1}{x}\right)$$

9.9.3 Examples

Here is a simple plot of the inverse hyperbolic cosecant function

```
--> x1 = -20:.01:-1;
--> x2 = 1:.01:20;
--> plot(x1,acsch(x1),x2,acsch(x2)); grid('on');
```



9.10 ANGLE Phase Angle Function

9.10.1 Usage

Compute the phase angle in radians of a complex matrix. The general syntax for its use is

```
p = angle(c)
```

where `c` is an `n`-dimensional array of numerical type.

9.10.2 Function Internals

For a complex number `x`, its polar representation is given by

$$x = |x|e^{j\theta}$$

and we can compute

$$\theta = \operatorname{atan2}(\Im x, \Re x)$$

9.10.3 Example

Here are some examples of the use of `angle` in the polar decomposition of a complex number.

```
--> x = 3+4*i

x =
  3.0000 + 4.0000i

--> a = abs(x)

a =
  5

--> t = angle(x)

t =
  0.9273

--> a*exp(i*t)

ans =
  3.0000 + 4.0000i
```

M version contributor: M.W. Vogel 01-30-06

9.11 ASEC Inverse Secant Function

9.11.1 Usage

Computes the inverse secant of its argument. The general syntax for its use is

```
y = asec(x)
```

where `x` is an `n`-dimensional array of numerical type.

9.11.2 Function Internals


The `acosh` function is computed from the formula

$$\sec^{-1}(x) = \cos^{-1}\left(\frac{1}{x}\right)$$

9.11.3 Examples

Here is a simple plot of the inverse secant function

```
--> x1 = -5:.01:-1;
--> x2 = 1:.01:5;
--> plot(x1,asec(x1),x2,asec(x2)); grid('on');
```



9.12 ASECD Inverse Secant Degrees Function

9.12.1 Usage

Computes the inverse secant of the argument, but returns the argument in degrees instead of radians (as is the case for `asec`. The syntax for its use is

```
y = asecd(x)
```

9.12.2 Examples

The inverse secant of $2/\sqrt{2}$ should be 45 degrees:

```
--> asecd(2/sqrt(2))
```

```
ans =
  45.0000 + 0.0000i
```

and the inverse secant of 2 should be 60 degrees:

```
--> asecd(2)
```

```
ans =
  60.0000 + 0.0000i
```

9.13 ASECH Inverse Hyperbolic Secant Function

9.13.1 Usage

Computes the inverse hyperbolic secant of its argument. The general syntax for its use is

```
y = asech(x)
```

where x is an n -dimensional array of numerical type.

9.13.2 Function Internals


The `asech` function is computed from the formula

$$\operatorname{sech}^{-1}(x) = \cosh^{-1}\left(\frac{1}{x}\right)$$

9.13.3 Examples

Here is a simple plot of the inverse hyperbolic secant function

```
--> x1 = -20:.01:-1;
--> x2 = 1:.01:20;
--> plot(x1,imag(asech(x1)),x2,imag(asech(x2))); grid('on');
```



9.14 ASIN Inverse Trigonometric Arcsine Function

9.14.1 Usage

Computes the `asin` function for its argument. The general syntax for its use is

```
y = asin(x)
```

where `x` is an `n`-dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `asin` function. Output `y` is of the same size and type as the input `x`, (unless `x` is an integer, in which case `y` is a `double` type).

9.14.2 Function Internals

Mathematically, the `asin` function is defined for all arguments `x` as

$$\operatorname{asin}x \equiv -i \log\left(ix + \sqrt{1-x^2}\right).$$

For real valued variables `x` in the range `[-1, 1]`, the function is computed directly using the standard C library's numerical `asin` function. For both real and complex arguments `x`, note that generally


$$\operatorname{asin}(\sin(x)) \neq x,$$

due to the periodicity of `sin(x)`.

9.14.3 Example

The following code demonstrates the `asin` function over the range $[-1,1]$.

```
--> t = linspace(-1,1);  
--> plot(t,asin(t))
```



9.15 ASIND Inverse Sine Degrees Function

9.15.1 Usage

Computes the inverse sine of the argument, but returns the argument in degrees instead of radians (as is the case for `asin`). The syntax for its use is

```
y = asind(x)
```

9.15.2 Examples

The inverse sine of $\sqrt{2}/2$ should be 45 degrees:

```
--> asind(sqrt(2)/2)
```

```
ans =  
    45.0000
```

and the inverse sine of 0.5 should be 30 degrees:

```
--> asind(0.5)
```

```
ans =  
    30.0000
```

9.16 ASINH Inverse Hyperbolic Sine Function

9.16.1 Usage

Computes the inverse hyperbolic sine of its argument. The general syntax for its use is

```
y = asinh(x)
```

where x is an n -dimensional array of numerical type.

9.16.2 Function Internals

The `asinh` function is computed from the formula


$$\sinh^{-1}(x) = \log(x + (x^2 + 1)^{0.5})$$

where the `log` (and square root) is taken in its most general sense.

9.16.3 Examples

Here is a simple plot of the inverse hyperbolic sine function

```
--> x = -5:.01:5;
--> plot(x,asinh(x)); grid('on');
```



9.17 ATAN Inverse Trigonometric Arctangent Function

9.17.1 Usage

Computes the `atan` function for its argument. The general syntax for its use is

```
y = atan(x)
```

where x is an n -dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `atan` function. Output y is of the same size and type as the input x , (unless x is an integer, in which case y is a `double` type).

9.17.2 Function Internals

Mathematically, the `atan` function is defined for all arguments x as

$$\operatorname{atan}x \equiv \frac{i}{2} (\log(1 - ix) - \log(ix + 1)).$$

For real valued variables x , the function is computed directly using the standard C library's numerical `atan` function. For both real and complex arguments x , note that generally


$$\text{atan}(\tan(x)) \neq x,$$

due to the periodicity of `tan(x)`.

9.17.3 Example

The following code demonstrates the `atan` function over the range $[-1, 1]$.

```
--> t = linspace(-1,1);
--> plot(t,atan(t))
```



9.18 ATAN2 Inverse Trigonometric 4-Quadrant Arctangent Function

9.18.1 Usage

Computes the `atan2` function for its argument. The general syntax for its use is

```
z = atan2(y,x)
```

where x and y are n -dimensional arrays of numerical type. Integer types are promoted to the `double` type prior to calculation of the `atan2` function. The size of the output depends on the size of x and y . If x is a scalar, then z is the same size as y , and if y is a scalar, then z is the same size as x . The type of the output is equal to the type of $-y/x$.

9.18.2 Function Internals

The function is defined (for real values) to return an angle between $-\pi$ and π . The signs of x and y are used to find the correct quadrant for the solution. For complex arguments, the two-argument arctangent is computed via

$$\text{atan2}(y, x) \equiv -i \log \left(\frac{x + iy}{\sqrt{x^2 + y^2}} \right)$$

For real valued arguments x, y , the function is computed directly using the standard C library's numerical `atan2` function. For both real and complex arguments x , note that generally


$$\text{atan2}(\sin(x), \cos(x)) \neq x,$$

due to the periodicities of $\cos(x)$ and $\sin(x)$.

9.18.3 Example

The following code demonstrates the difference between the `atan2` function and the `atan` function over the range $[-\pi, \pi]$.

```
--> x = linspace(-pi, pi);
--> sx = sin(x); cx = cos(x);
--> plot(x, atan(sx./cx), x, atan2(sx, cx))
```



Note how the two-argument `atan2` function (green line) correctly “unwraps” the phase of the angle, while the `atan` function (red line) wraps the angle to the interval $[-\pi/2, \pi/2]$.

9.19 ATAND Inverse Tangent Degrees Function

9.19.1 Usage

Computes the inverse tangent of the argument, but returns the argument in degrees instead of radians (as is the case for `atan`). The syntax for its use is

$$y = \text{atand}(x)$$

9.19.2 Examples

The inverse tangent of 1 should be 45 degrees:

```
--> atand(1)
```

```
ans =
    45
```

9.20 ATANH Inverse Hyperbolic Tangent Function

9.20.1 Usage

Computes the inverse hyperbolic tangent of its argument. The general syntax for its use is

```
y = atanh(x)
```

where x is an n -dimensional array of numerical type.

9.20.2 Function Internals

The `atanh` function is computed from the formula


$$\tanh^{-1}(x) = \frac{1}{2} \log \left(\frac{1+x}{1-x} \right)$$

where the `log` (and square root) is taken in its most general sense.

9.20.3 Examples

Here is a simple plot of the inverse hyperbolic tangent function

```
--> x = -0.99:.01:0.99;  
--> plot(x,atanh(x)); grid('on');
```



9.21 COS Trigonometric Cosine Function

9.21.1 Usage

Computes the `cos` function for its argument. The general syntax for its use is

```
y = cos(x)
```

where x is an n -dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `cos` function. Output y is of the same size and type as the input x , (unless x is an integer, in which case y is a `double` type).

9.21.2 Function Internals

Mathematically, the `cos` function is defined for all real valued arguments x by the infinite summation

$$\cos x \equiv \sum_{n=0}^{\infty} \frac{(-1)^n x^{2n}}{(2n)!}.$$


For complex valued arguments z , the cosine is computed via

$$\cos z \equiv \cos \Re z \cosh \Im z - \sin \Re z \sinh \Im z.$$

9.21.3 Example

The following piece of code plots the real-valued `cos(2 pi x)` function over one period of `[0,1]`:

```
--> x = linspace(0,1);
--> plot(x,cos(2*pi*x))
```



9.22 COSD Cosine Degrees Function

9.22.1 Usage

Computes the cosine of the argument, but takes the argument in degrees instead of radians (as is the case for `cos`). The syntax for its use is

```
y = cosd(x)
```

9.22.2 Examples

The cosine of 45 degrees should be `sqrt(2)/2`

```
--> cosd(45)
```

```
ans =
    0.7071
```

and the cosine of 60 degrees should be 0.5:

```
--> cosd(60)
```

```
ans =  
    0.5000
```

9.23 COSH Hyperbolic Cosine Function

9.23.1 Usage

Computes the hyperbolic cosine of the argument. The syntax for its use is

```
y = cosh(x)
```

9.23.2 Function Internals

The `cosh` function is computed from the formula

$$\cosh(x) = \frac{e^x + e^{-x}}{2}$$


For x complex, it follows that

$$\cosh(a + i * b) = \frac{e^a(\cos(b) + i * \sin(b)) + e^{-a}(\cos(-b) + i * \sin(-b))}{2}$$

9.23.3 Examples

Here is a simple plot of the hyperbolic cosine function

```
--> x = linspace(-5,5);  
--> plot(x,cosh(x)); grid('on');
```



9.24 COT Trigonometric Cotangent Function

9.24.1 Usage

Computes the `cot` function for its argument. The general syntax for its use is

$$y = \text{cot}(x)$$

where x is an n -dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `cot` function. Output y is of the same size and type as the input x , (unless x is an integer, in which case y is a `double` type).

9.24.2 Function Internals

Mathematically, the `cot` function is defined for all arguments x as

$$\cot x \equiv \frac{\cos x}{\sin x}$$


For complex valued arguments z , the cotangent is computed via

$$\cot z \equiv \frac{\cos 2\Re z + \cosh 2\Im z}{\sin 2\Re z + i \sinh 2\Im z}.$$

9.24.3 Example

The following piece of code plots the real-valued `cot(x)` function over the interval $[-1, 1]$:

```
--> t = linspace(-1,1);
--> plot(t,cot(t))
```



9.25 COTD Cotangent Degrees Function

9.25.1 Usage

Computes the cotangent of the argument, but takes the argument in degrees instead of radians (as is the case for `cot`). The syntax for its use is

$$y = \text{cotd}(x)$$

9.25.2 Examples

The cotangent of 45 degrees should be 1.

```
--> cotd(45)
```

```
ans =  
1
```

9.26 COTD Cotangent Degrees Function

9.26.1 Usage

Computes the cotangent of the argument, but takes the argument in degrees instead of radians (as is the case for `cot`). The syntax for its use is

```
y = cotd(x)
```

9.26.2 Examples

The cotangent of 45 degrees should be 1.

```
--> cotd(45)
```

```
ans =  
1
```

9.27 COTH Hyperbolic Cotangent Function

9.27.1 Usage

Computes the hyperbolic cotangent of the argument. The syntax for its use is

```
y = coth(x)
```

9.27.2 Function Internals


The `coth` function is computed from the formula

$$\coth(x) = \frac{1}{\tanh(x)}$$

9.27.3 Examples

Here is a simple plot of the hyperbolic cotangent function

```
--> x1 = -pi+.01:.01:-.01;  
--> x2 = .01:.01:pi-.01;  
--> plot(x1,coth(x1),x2,coth(x2)); grid('on');
```



9.28 CROSS Cross Product of Two Vectors

9.28.1 Usage

Computes the cross product of two vectors. The general syntax for its use is

```
c = cross(a,b)
```

where **a** and **b** are 3-element vectors.

9.29 CSC Trigonometric Cosecant Function

9.29.1 Usage

Computes the `csc` function for its argument. The general syntax for its use is

```
y = csc(x)
```

where **x** is an **n**-dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `csc` function. Output **y** is of the same size and type as the input **x**, (unless **x** is an integer, in which case **y** is a `double` type).

9.29.2 Function Internals


Mathematically, the `csc` function is defined for all arguments as

$$\csc x \equiv \frac{1}{\sin x}.$$

9.29.3 Example

The following piece of code plots the real-valued `csc(2 pi x)` function over the interval of `[-1,1]`:

```
--> t = linspace(-1,1,1000);
--> plot(t,csc(2*pi*t))
--> axis([-1,1,-10,10]);
```



9.30 CSCD Cosecant Degrees Function

9.30.1 Usage

Computes the cosecant of the argument, but takes the argument in degrees instead of radians (as is the case for `csc`). The syntax for its use is

```
y = cscd(x)
```

9.31 CSCH Hyperbolic Cosecant Function

9.31.1 Usage

Computes the hyperbolic cosecant of the argument. The syntax for its use is

```
y = csch(x)
```

9.31.2 Function Internals


The `csch` function is computed from the formula

$$\text{csch}(x) = \frac{1}{\sinh(x)}$$

9.31.3 Examples

Here is a simple plot of the hyperbolic cosecant function

```
--> x1 = -pi+.01:.01:-.01;
--> x2 = .01:.01:pi-.01;
--> plot(x1,csch(x1),x2,csch(x2)); grid('on');
```



9.32 DEG2RAD Convert From Degrees To Radians

9.32.1 Usage

Converts the argument from degrees to radians. The syntax for its use is

$$y = \text{deg2rad}(x)$$

where x is a numeric array. Conversion is done by simply multiplying x by $\pi/180$.

9.32.2 Example

How many radians in a circle:

```
--> deg2rad(360) - 2*pi
```

```
ans =  
0
```

9.33 ERF Error Function

9.33.1 Usage

Computes the error function for real arguments. The `erf` function takes only a single argument

$$y = \text{erf}(x)$$

where x is either a float or double array. The output vector y is the same size (and type) as x .

9.33.2 Function Internals

The erf function is defined by the integral:

$$\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt,$$


and is the integral of the normal distribution.

9.33.3 Example

Here is a plot of the erf function over the range [-5,5].

```
--> x = linspace(-5,5);
--> y = erf(x);
--> plot(x,y); xlabel('x'); ylabel('erf(x)');
```

which results in the following plot.



9.34 ERFC Complimentary Error Function

9.34.1 Usage

Computes the complimentary error function for real arguments. The `erfc` function takes only a single argument

```
y = erfc(x)
```

where `x` is either a `float` or `double` array. The output vector `y` is the same size (and type) as `x`.

9.34.2 Function Internals

The `erfc` function is defined by the integral:

$$\operatorname{erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^{\infty} e^{-t^2} dt,$$


and is the integral of the normal distribution.

9.34.3 Example

Here is a plot of the `erfc` function over the range [-5,5].

```
--> x = linspace(-5,5);
--> y = erfc(x);
--> plot(x,y); xlabel('x'); ylabel('erfc(x)');
```

which results in the following plot.



9.35 EXP Exponential Function

9.35.1 Usage

Computes the `exp` function for its argument. The general syntax for its use is

$$y = \text{exp}(x)$$

where `x` is an `n`-dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `exp` function. Output `y` is of the same size and type as the input `x`, (unless `x` is an integer, in which case `y` is a `double` type).

9.35.2 Function Internals

Mathematically, the `exp` function is defined for all real valued arguments `x` as

$$\text{exp } x \equiv e^x,$$

where

$$e = \sum_{k=0}^{\infty} \frac{1}{k!}$$


and is approximately 2.718281828459045 (returned by the function `e`). For complex values `z`, the famous Euler formula is used to calculate the exponential

$$e^z = e^{|z|} [\cos \Re z + i \sin \Re z]$$

9.35.3 Example


The following piece of code plots the real-valued `exp` function over the interval `[-1,1]`:

```
--> x = linspace(-1,1);
--> plot(x,exp(x))
```



In the second example, we plot the unit circle in the complex plane $e^{-i 2 \pi x}$ for x in $[-1, 1]$.

```
--> x = linspace(-1,1);
--> plot(exp(-i*x*2*pi))
```



9.36 EXPM1 Exponential Minus One Function

9.36.1 Usage

Computes $\exp(x)-1$ function accurately for x small. The syntax for its use is

```
y = expm1(x)
```

where x is an n -dimensional array of numerical type.

9.37 FIX Round Towards Zero

9.37.1 Usage

Rounds the argument array towards zero. The syntax for its use is

```
y = fix(x)
```

where x is a numeric array. For positive elements of x , the output is the largest integer smaller than x . For negative elements of x the output is the smallest integer larger than x . For complex x , the operation is applied separately to the real and imaginary parts.

9.37.2 Example

Here is a simple example of the `fix` operation on some values

```
--> a = [-1.8,pi,8,-pi,-0.001,2.3+0.3i]
```

```
a =
```

```
Columns 1 to 4
```

```
-1.8000 + 0.0000i  3.1416 + 0.0000i  8.0000 + 0.0000i  -3.1416 + 0.0000i
```

```
Columns 5 to 6
```

```
-0.0010 + 0.0000i  2.3000 + 0.3000i
```

```
--> fix(a)
```

```
ans =
```

```
Columns 1 to 4
```

```
-1.0000 + 0.0000i  3.0000 + 0.0000i  8.0000 + 0.0000i  -3.0000 + 0.0000i
```

```
Columns 5 to 6
```

```
0  2.0000 + 0.0000i
```

9.38 GAMMA Gamma Function

9.38.1 Usage

Computes the gamma function for real arguments. The `gamma` function takes only a single argument

```
y = gamma(x)
```

where `x` is either a `float` or `double` array. The output vector `y` is the same size (and type) as `x`.

9.38.2 Function Internals

The gamma function is defined by the integral:

$$\Gamma(x) = \int_0^{\infty} e^{-t} t^{x-1} dt$$

The gamma function obeys the interesting relationship

$$\Gamma(x) = (x-1)\Gamma(x-1),$$


and for integer arguments, is equivalent to the factorial function.

9.38.3 Example

Here is a plot of the gamma function over the range $[-5,5]$.

```
--> x = linspace(-5,5);
--> y = gamma(x);
--> plot(x,y); xlabel('x'); ylabel('gamma(x)');
--> axis([-5,5,-5,5]);
```

which results in the following plot.



9.39 GAMMALN Log Gamma Function

9.39.1 Usage

Computes the natural log of the gamma function for real arguments. The `gammaln` function takes only a single argument

```
y = gammaln(x)
```


where `x` is either a float or double array. The output vector `y` is the same size (and type) as `x`.

9.39.2 Example

Here is a plot of the `gammaln` function over the range $[-5,5]$.

```
--> x = linspace(0,10);
--> y = gammaln(x);
--> plot(x,y); xlabel('x'); ylabel('gammaln(x)');
```

which results in the following plot.



9.40 IDIV Integer Division Operation

9.40.1 Usage

Computes the integer division of two arrays. The syntax for its use is

```
y = idiv(a,b)
```

where *a* and *b* are arrays or scalars. The effect of the `idiv` is to compute the integer division of *b* into *a*.

9.40.2 Example

The following examples show some uses of `idiv` arrays.

```
--> idiv(27,6)
```

```
ans =  
4
```

```
--> idiv(4,-2)
```

```
ans =  
-2
```

```
--> idiv(15,3)
```

```
ans =  
5
```

9.41 LOG Natural Logarithm Function

9.41.1 Usage

Computes the `log` function for its argument. The general syntax for its use is

$$y = \log(x)$$

where x is an n -dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `log` function. Output y is of the same size as the input x . For strictly positive, real inputs, the output type is the same as the input. For negative and complex arguments, the output is complex.

9.41.2 Function Internals

Mathematically, the `log` function is defined for all real valued arguments x by the integral

$$\log x \equiv \int_1^x \frac{dt}{t}.$$

For complex-valued arguments, z , the complex logarithm is defined as


$$\log z \equiv \log |z| + i \arg z,$$

where `arg` is the complex argument of z .

9.41.3 Example

The following piece of code plots the real-valued `log` function over the interval $[1, 100]$:

```
--> x = linspace(1,100);
--> plot(x,log(x))
--> xlabel('x');
--> ylabel('log(x)');
```



9.42 LOG10 Base-10 Logarithm Function

9.42.1 Usage

Computes the `log10` function for its argument. The general syntax for its use is


$$y = \log_{10}(x)$$

where x is an n -dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `log10` function. Output y is of the same size as the input x . For strictly positive, real inputs, the output type is the same as the input. For negative and complex arguments, the output is complex.

9.42.2 Example

The following piece of code plots the real-valued `log10` function over the interval `[1,100]`:

```
--> x = linspace(1,100);  
--> plot(x,log10(x))  
--> xlabel('x');  
--> ylabel('log10(x)');
```



9.43 LOG1P Natural Logarithm of 1+P Function

9.43.1 Usage

Computes the `log` function for one plus its argument. The general syntax for its use is

$$y = \text{log1p}(x)$$

where x is an n -dimensional array of numerical type.

9.44 LOG2 Base-2 Logarithm Function

9.44.1 Usage

Computes the `log2` function for its argument. The general syntax for its use is


$$y = \text{log2}(x)$$

where x is an n -dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `log2` function. Output y is of the same size as the input x . For strictly positive, real inputs, the output type is the same as the input. For negative and complex arguments, the output is complex.

9.44.2 Example

The following piece of code plots the real-valued `log2` function over the interval `[1,100]`:

```
--> x = linspace(1,100);  
--> plot(x,log2(x))  
--> xlabel('x');  
--> ylabel('log2(x)');
```



9.45 MOD Modulus Operation

9.45.1 Usage

Computes the modulus of an array. The syntax for its use is

```
y = mod(x,n)
```

where `x` is matrix, and `n` is the base of the modulus. The effect of the `mod` operator is to add or subtract multiples of `n` to the vector `x` so that each element `x_i` is between 0 and `n` (strictly). Note that `n` does not have to be an integer. Also, `n` can either be a scalar (same base for all elements of `x`), or a vector (different base for each element of `x`).

Note that the following are defined behaviors:

1. `mod(x,0) = x@`
2. `mod(x,x) = 0@`
3. `mod(x,n)@` has the same sign as `n` for all other cases.

9.45.2 Example

The following examples show some uses of `mod` arrays.

```
--> mod(18,12)
```

```
ans =  
6
```

```
--> mod(6,5)
```

```
ans =  
1
```

```
--> mod(2*pi,pi)
```

```
ans =  
0
```

Here is an example of using `mod` to determine if integers are even or odd:

```
--> mod([1,3,5,2],2)
```

```
ans =  
1 1 1 0
```

Here we use the second form of `mod`, with each element using a separate base.

```
--> mod([9 3 2 0],[1 0 2 2])
```

```
ans =  
0 3 0 0
```

9.46 MPOWER Matrix Power Function

9.46.1 Usage

Computes the matrix power operator for two arrays. It is an M-file version of the `^` operator. The syntax for its use is

```
y = mpower(a,b)
```

where $y=a^b$. See the `matrixpower` documentation for more details on what this function actually does.

9.47 POWER Element-wise Power Function

9.47.1 Usage

Computes the element-wise power operator for two arrays. It is an M-file version of the `.^` operator. The syntax for its use is

```
y = power(a,b)
```

where $y=a.^b$. See the `dotpower` documentation for more details on what this function actually does.

9.48 RAD2DEG Radians To Degrees Conversion Function

9.48.1 Usage

Converts the argument array from radians to degrees. The general syntax for its use is

```
y = rad2deg(x)
```

Note that the output type will be the same as the input type, and that complex arguments are allowed. The output is not wrapped to $[0,360)$.

9.48.2 Examples

Some known conversion factors

```
--> rad2deg(1) % one radian is about 57 degrees
```

```
ans =
    57.2958
```

```
--> rad2deg(pi/4) % should be 45 degrees
```

```
ans =
    45
```

```
--> rad2deg(2*pi) % Note that this is 360 not 0 degrees
```

```
ans =
    360
```

9.49 REM Remainder After Division

9.49.1 Usage

Computes the remainder after division of an array. The syntax for its use is

```
y = rem(x,n)
```

where x is matrix, and n is the base of the modulus. The effect of the `rem` operator is to add or subtract multiples of n to the vector x so that each element x_i is between 0 and n (strictly). Note that n does not have to be an integer. Also, n can either be a scalar (same base for all elements of x), or a vector (different base for each element of x).

Note that the following are defined behaviors:

1. `rem(x,0) = nan@`
2. `rem(x,x) = 0@` for nonzero x
3. `rem(x,n)@` has the same sign as x for all other cases.

Note that `rem` and `mod` return the same value if x and n are of the same sign. But differ by n if x and y have different signs.

9.49.2 Example

The following examples show some uses of `rem` arrays.

```
--> rem(18,12)
```

```
ans =
     6
```

```
--> rem(6,5)
```

```
ans =
     1
```

```
--> rem(2*pi,pi)
```

```
ans =
     0
```

Here is an example of using `rem` to determine if integers are even or odd:

```
--> rem([1,3,5,2],2)
```

```
ans =
     1 1 1 0
```

Here we use the second form of `rem`, with each element using a separate base.

```
--> rem([9 3 2 0],[1 0 2 2])
```

```
ans =
     0 NaN     0     0
```

9.50 SEC Trigonometric Secant Function

9.50.1 Usage

Computes the `sec` function for its argument. The general syntax for its use is

```
y = sec(x)
```

where `x` is an `n`-dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `sec` function. Output `y` is of the same size and type as the input `x`, (unless `x` is an integer, in which case `y` is a `double` type).

9.50.2 Function Internals


Mathematically, the `sec` function is defined for all arguments as

$$\sec x \equiv \frac{1}{\cos x}.$$

9.50.3 Example

The following piece of code plots the real-valued `sec(2 pi x)` function over the interval of `[-1,1]`:

```
--> t = linspace(-1,1,1000);  
--> plot(t,sec(2*pi*t))  
--> axis([-1,1,-10,10]);
```



9.51 SECD Secant Degrees Function

9.51.1 Usage

Computes the secant of the argument, but takes the argument in degrees instead of radians (as is the case for `sec`). The syntax for its use is

$$y = \text{secd}(x)$$

9.52 SECH Hyperbolic Secant Function

9.52.1 Usage

Computes the hyperbolic secant of the argument. The syntax for its use is

$$y = \text{sech}(x)$$

9.52.2 Function Internals


The `sech` function is computed from the formula

$$\text{sech}(x) = \frac{1}{\cosh(x)}$$

9.52.3 Examples

Here is a simple plot of the hyperbolic secant function

```
--> x = -2*pi:.01:2*pi;
--> plot(x,sech(x)); grid('on');
```



9.53 SIN Trigonometric Sine Function

9.53.1 Usage

Computes the `sin` function for its argument. The general syntax for its use is

```
y = sin(x)
```

where `x` is an `n`-dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `sin` function. Output `y` is of the same size and type as the input `x`, (unless `x` is an integer, in which case `y` is a `double` type).

9.53.2 Function Internals

Mathematically, the `sin` function is defined for all real valued arguments `x` by the infinite summation

$$\sin x \equiv \sum_{n=1}^{\infty} \frac{(-1)^{n-1} x^{2n-1}}{(2n-1)!}.$$


For complex valued arguments `z`, the sine is computed via

$$\sin z \equiv \sin \Re z \cosh \Im z - i \cos \Re z \sinh \Im z.$$

9.53.3 Example

The following piece of code plots the real-valued `sin(2 pi x)` function over one period of `[0,1]`:

```
--> x = linspace(0,1);
--> plot(x,sin(2*pi*x))
```



9.54 SIND Sine Degrees Function

9.54.1 Usage

Computes the sine of the argument, but takes the argument in degrees instead of radians (as is the case for `cos`). The syntax for its use is

```
y = sind(x)
```

9.54.2 Examples

The sine of 45 degrees should be `sqrt(2)/2`

```
--> sind(45)
```

```
ans =  
    0.7071
```

and the sine of 30 degrees should be 0.5:

```
--> sind(30)
```

```
ans =  
    0.5000
```

9.55 SINH Hyperbolic Sine Function

9.55.1 Usage

Computes the hyperbolic sine of the argument. The syntax for its use is

```
y = sinh(x)
```

9.55.2 Function Internals


The `sinh` function is computed from the formula

$$\sinh(x) = \frac{e^x - e^{-x}}{2}$$

9.55.3 Examples

Here is a simple plot of the hyperbolic sine function

```
--> x = linspace(-5,5);  
--> plot(x, sinh(x)); grid('on');
```



9.56 SQRT Square Root of an Array

9.56.1 Usage

Computes the square root of the argument matrix. The general syntax for its use is

```
y = sqrt(x)
```

where `x` is an N-dimensional numerical array.

9.56.2 Example

Here are some examples of using `sqrt`

```
--> sqrt(9)
```

```
ans =  
3
```

```
--> sqrt(i)
```

```
ans =
```

```

0.7071 + 0.7071i
--> sqrt(-1)
ans =
0.0000 + 1.0000i
--> x = rand(4)
x =
0.6031    0.7537    0.6479    0.3985
0.8219    0.3124    0.3073    0.5023
0.1780    0.9026    0.5598    0.9462
0.6403    0.4911    0.4890    0.7665
--> sqrt(x)
ans =
0.7766    0.8681    0.8049    0.6313
0.9066    0.5589    0.5543    0.7087
0.4219    0.9501    0.7482    0.9727
0.8002    0.7008    0.6993    0.8755

```

9.57 TAN Trigonometric Tangent Function

9.57.1 Usage

Computes the `tan` function for its argument. The general syntax for its use is

```
y = tan(x)
```

where `x` is an `n`-dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the `tan` function. Output `y` is of the same size and type as the input `x`, (unless `x` is an integer, in which case `y` is a `double` type).

9.57.2 Function Internals

Mathematically, the `tan` function is defined for all real valued arguments `x` by the infinite summation

$$\tan x \equiv x + \frac{x^3}{3} + \frac{2x^5}{15} + \cdots,$$

or alternately by the ratio

$$\tan x \equiv \frac{\sin x}{\cos x}$$


For complex valued arguments `z`, the tangent is computed via

$$\tan z \equiv \frac{\sin 2\Re z + i \sinh 2\Im z}{\cos 2\Re z + \cosh 2\Im z}.$$

9.57.3 Example

The following piece of code plots the real-valued `tan(x)` function over the interval `[-1,1]`:

```
--> t = linspace(-1,1);  
--> plot(t,tan(t))
```



9.58 TAND Tangent Degrees Function

9.58.1 Usage

Computes the tangent of the argument, but takes the argument in degrees instead of radians (as is the case for `cos`). The syntax for its use is

```
y = tand(x)
```

9.58.2 Examples

The tangent of 45 degrees should be 1

```
--> tand(45)
```

```
ans =  
    1.0000
```

9.59 TANH Hyperbolic Tangent Function

9.59.1 Usage

Computes the hyperbolic tangent of the argument. The syntax for its use is

```
y = tanh(x)
```

9.59.2 Function Internals


The `tanh` function is computed from the formula

$$\tanh(x) = \frac{\sinh(x)}{\cosh(x)}$$

9.59.3 Examples

Here is a simple plot of the hyperbolic tangent function

```
--> x = linspace(-5,5);  
--> plot(x,tanh(x)); grid('on');
```



Chapter 10

Base Constants

10.1 E Euler Constant (Base of Natural Logarithm)

10.1.1 Usage

Returns a `double` (64-bit floating point number) value that represents Euler's constant, the base of the natural logarithm. Typical usage

```
y = e
```

This value is approximately 2.718281828459045.

10.1.2 Example

The following example demonstrates the use of the `e` function.

```
--> e
```

```
ans =  
  2.7183
```

```
--> log(e)
```

```
ans =  
  1
```

10.2 EPS Double Precision Floating Point Relative Machine Precision Epsilon

10.2.1 Usage

Returns `eps`, which quantifies the relative machine precision of floating point numbers (a machine specific quantity). The syntax for `eps` is:

```

y = eps
y = eps('double')
y = eps(X)

```

First form returns `eps` for `double` precision values. For most typical processors, this value is approximately 2^{-52} , or $2.2204e-16$. Second form return `eps` for class `double` or `single`. Third form returns distance to the next value greater than `X`.

10.2.2 Example

The following example demonstrates the use of the `eps` function, and one of its numerical consequences.

```

--> eps

ans =
  2.2204e-16

--> 1.0+eps

ans =
  1.0000

--> eps(1000.)

ans =
  1.1369e-13

```

10.3 FALSE Logical False

10.3.1 Usage

Returns a logical 0. The syntax for its use is

```
y = false
```

10.4 FEPS Single Precision Floating Point Relative Machine Precision Epsilon

10.4.1 Usage

Returns `feps`, which quantifies the relative machine precision of floating point numbers (a machine specific quantity). The syntax for `feps` is:

```
y = feps
```

which returns `feps` for `single` precision values. For most typical processors, this value is approximately 2^{-24} , or $5.9604e-8$.

10.4.2 Example

The following example demonstrates the use of the `feps` function, and one of its numerical consequences.

```
--> feps  
  
ans =  
  1.1921e-07  
  
--> 1.0f+eps  
  
ans =  
  1
```

10.5 I-J Square Root of Negative One

10.5.1 Usage

Returns a complex value that represents the square root of -1. There are two functions that return the same value:

```
  y = i  
and  
  y = j.
```

This allows either `i` or `j` to be used as loop indices. The returned value is a 32-bit complex value.

10.5.2 Example

The following examples demonstrate a few calculations with `i`.

```
--> i  
  
ans =  
  0.0000 + 1.0000i  
  
--> i^2  
  
ans =  
 -1.0000 + 0.0000i
```

The same calculations with `j`:

```
--> j  
  
ans =
```

```
0.0000 + 1.0000i
```

```
--> j^2
```

```
ans =
-1.0000 + 0.0000i
```

Here is an example of how `i` can be used as a loop index and then recovered as the square root of `-1`.

```
--> accum = 0; for i=1:100; accum = accum + i; end; accum
```

```
ans =
5050
```

```
--> i
```

```
ans =
100
```

```
--> clear i
```

```
--> i
```

```
ans =
0.0000 + 1.0000i
```

10.6 INF Infinity Constant

10.6.1 Usage

Returns a value that represents positive infinity for both 32 and 64-bit floating point values. There are several forms for the `Inf` function. The first form returns a double precision `Inf`.

```
y = inf
```

The next form takes a class name that can be either `'double'`

```
y = inf('double')
```

or `'single'`:

```
y = inf('single')
```

With a single parameter it generates a square matrix of `infs`.

```
y = inf(n)
```

Alternatively, you can specify the dimensions of the array via

```
y = inf(m,n,p,...)
```

or

```
y = inf([m,n,p,...])
```

Finally, you can add a classname of either 'single' or 'double'.

10.6.2 Function Internals

The infinity constant has several interesting properties. In particular:

$$\begin{aligned} \infty \times 0 &= \text{NaN} \\ \infty \times a &= \infty \text{ for all } a > 0 \\ \infty \times a &= -\infty \text{ for all } a < 0 \\ \infty / \infty &= \text{NaN} \\ \infty / 0 &= \infty \end{aligned}$$

Note that infinities are not preserved under type conversion to integer types (see the examples below).

10.6.3 Example

The following examples demonstrate the various properties of the infinity constant.

```
--> inf*0

ans =
    NaN

--> inf*2

ans =
    Inf

--> inf*-2

ans =
    -Inf

--> inf/inf

ans =
    NaN

--> inf/0

ans =
    Inf
```

```
--> inf/nan
```

```
ans =  
NaN
```

Note that infinities are preserved under type conversion to floating point types (i.e., `float`, `double`, `complex` and `dcomplex` types), but not integer types.

```
--> uint32(inf)
```

```
ans =  
4294967295
```

```
--> complex(inf)
```

```
ans =  
Inf
```

10.7 PI Constant Pi

10.7.1 Usage

Returns a `double` (64-bit floating point number) value that represents pi (ratio between the circumference and diameter of a circle...). Typical usage

```
y = pi
```

This value is approximately 3.141592653589793.

10.7.2 Example

The following example demonstrates the use of the `pi` function.

```
--> pi
```

```
ans =  
3.1416
```

```
--> cos(pi)
```

```
ans =  
-1
```

10.8 TEPS Type-based Epsilon Calculation

10.8.1 Usage

Returns `eps` for double precision arguments and `feps` for single precision arguments. The syntax for `teps` is

```
y = teps(x)
```

The `teps` function is most useful if you need to compute epsilon based on the type of the array.

10.8.2 Example

The following example demonstrates the use of the `teps` function, and one of its numerical consequences.

```
--> teps(float(3.4))
```

```
ans =  
1.1921e-07
```

```
--> teps(complex(3.4+i*2))
```

```
ans =  
2.2204e-16
```

10.9 TRUE Logical TRUE

10.9.1 Usage

Returns a logical 1. The syntax for its use is

```
y = true
```


Chapter 11

Elementary Functions

11.1 ABS Absolute Value Function

11.1.1 Usage

Returns the absolute value of the input array for all elements. The general syntax for its use is

```
y = abs(x)
```

where x is an n -dimensional array of numerical type. The output is the same numerical type as the input, unless the input is `complex` or `dcomplex`. For `complex` inputs, the absolute value is a floating point array, so that the return type is `float`. For `dcomplex` inputs, the absolute value is a double precision floating point array, so that the return type is `double`.

11.1.2 Example

The following demonstrates the `abs` applied to a complex scalar.

```
--> abs(3+4*i)
```

```
ans =  
5
```

The `abs` function applied to integer and real values:

```
--> abs([-2,3,-4,5])
```

```
ans =  
2 3 4 5
```

For a double-precision complex array,

```
--> abs([2.0+3.0*i, i])
```

```
ans =  
3.6056    1.0000
```

11.2 ALL All True Function

11.2.1 Usage

Reduces a logical array along a given dimension by testing for all logical 1s. The general syntax for its use is

```
y = all(x,d)
```

where x is an n -dimensions array of `logical` type. The output is of `logical` type. The argument d is optional, and denotes the dimension along which to operate. The output y is the same size as x , except that it is singular along the operated direction. So, for example, if x is a $3 \times 3 \times 4$ array, and we `all` operation along dimension $d=2$, then the output is of size $3 \times 1 \times 4$.

11.2.2 Function Internals

The output is computed via

$$y(m_1, \dots, m_{d-1}, 1, m_{d+1}, \dots, m_p) = \min_k x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p)$$

If d is omitted, then the minimum is taken over all elements of x .

11.2.3 Example

The following piece of code demonstrates various uses of the `all` function

```
--> A = [1,0,0;1,0,0;0,0,1]
```

```
A =
 1 0 0
 1 0 0
 0 0 1
```

We start by calling `all` without a dimension argument, in which case it defaults to testing all values of A

```
--> all(A)
```

```
ans =
 0 0 0
```

The `all` function is useful in expressions also.

```
--> all(A>=0)
```

```
ans =
 1 1 1
```

Next, we apply the `all` operation along the rows.

```
--> all(A,2)
```

```
ans =
  0
  0
  0
```

11.3 ANY Any True Function

11.3.1 Usage

Reduces a logical array along a given dimension by testing for any logical 1s. The general syntax for its use is

```
y = any(x,d)
```

where x is an n -dimensions array of `logical` type. The output is of `logical` type. The argument d is optional, and denotes the dimension along which to operate. The output y is the same size as x , except that it is singular along the operated direction. So, for example, if x is a $3 \times 3 \times 4$ array, and we `any` operation along dimension $d=2$, then the output is of size $3 \times 1 \times 4$.

11.3.2 Function Internals

The output is computed via

$$y(m_1, \dots, m_{d-1}, 1, m_{d+1}, \dots, m_p) = \max_k x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p)$$

If d is omitted, then the summation is taken along the first non-singleton dimension of x .

11.3.3 Example

The following piece of code demonstrates various uses of the summation function

```
--> A = [1,0,0;1,0,0;0,0,1]
```

```
A =
  1 0 0
  1 0 0
  0 0 1
```

We start by calling `any` without a dimension argument, in which case it defaults to the first nonsingular dimension (in this case, along the columns or $d = 1$).

```
--> any(A)
```

```
ans =
  1 0 1
```

Next, we apply the `any` operation along the rows.

```
--> any(A,2)
```

```
ans =  
  1  
  1  
  1
```

11.4 CEIL Ceiling Function

11.4.1 Usage

Computes the ceiling of an n-dimensional array elementwise. The ceiling of a number is defined as the smallest integer that is larger than or equal to that number. The general syntax for its use is

$$y = \text{ceil}(x)$$

where `x` is a multidimensional array of numerical type. The `ceil` function preserves the type of the argument. So integer arguments are not modified, and `float` arrays return `float` arrays as outputs, and similarly for `double` arrays. The `ceil` function is not defined for `complex` or `dcomplex` types.

11.4.2 Example

The following demonstrates the `ceil` function applied to various (numerical) arguments. For integer arguments, the `ceil` function has no effect:

```
--> ceil(3)
```

```
ans =  
  3
```

```
--> ceil(-3)
```

```
ans =  
 -3
```

Next, we take the `ceil` of a floating point value:

```
--> ceil(float(3.023))
```

```
ans =  
  4
```

```
--> ceil(float(-2.341))
```

```
ans =  
 -2
```

Note that the return type is a `float` also. Finally, for a `double` type:

```
--> ceil(4.312)
```

```
ans =  
5
```

```
--> ceil(-5.32)
```

```
ans =  
-5
```

11.5 CONJ Conjugate Function

11.5.1 Usage

Returns the complex conjugate of the input array for all elements. The general syntax for its use is

$$y = \text{conj}(x)$$

where x is an n -dimensional array of numerical type. The output is the same numerical type as the input. The `conj` function does nothing to real and integer types.

11.5.2 Example

The following demonstrates the complex conjugate applied to a complex scalar.

```
--> conj(3+4*i)
```

```
ans =  
3.0000 - 4.0000i
```

The `conj` function has no effect on real arguments:

```
--> conj([2,3,4])
```

```
ans =  
2 3 4
```

For a double-precision complex array,

```
--> conj([2.0+3.0*i,i])
```

```
ans =  
2.0000 - 3.0000i  0.0000 - 1.0000i
```

11.6 CUMPROD Cumulative Product Function

11.6.1 Usage

Computes the cumulative product of an n-dimensional array along a given dimension. The general syntax for its use is

```
y = cumprod(x,d)
```

where x is a multidimensional array of numerical type, and d is the dimension along which to perform the cumulative product. The output y is the same size of x . Integer types are promoted to `int32`. If the dimension d is not specified, then the cumulative sum is applied along the first non-singular dimension.

11.6.2 Function Internals

The output is computed via

$$y(m_1, \dots, m_{d-1}, j, m_{d+1}, \dots, m_p) = \prod_{k=1}^j x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p).$$

11.6.3 Example

The default action is to perform the cumulative product along the first non-singular dimension.

```
--> A = [5,1,3;3,2,1;0,3,1]
```

```
A =
  5  1  3
  3  2  1
  0  3  1
```

```
--> cumprod(A)
```

```
ans =
  5  1  3
 15  2  3
  0  6  3
```

To compute the cumulative product along the columns:

```
--> cumprod(A,2)
```

```
ans =
  5  5 15
  3  6  6
  0  0  0
```

The cumulative product also works along arbitrary dimensions

```
--> B(:, :, 1) = [5,2;8,9];
--> B(:, :, 2) = [1,0;3,0]
```

```
B =
```

```
(: , : , 1) =
  5  2
  8  9
```

```
(: , : , 2) =
  1  0
  3  0
```

```
--> cumprod(B,3)
```

```
ans =
```

```
(: , : , 1) =
  5  2
  8  9
```

```
(: , : , 2) =
  5  0
 24  0
```

11.7 CUMSUM Cumulative Summation Function

11.7.1 Usage

Computes the cumulative sum of an n-dimensional array along a given dimension. The general syntax for its use is

```
y = cumsum(x,d)
```

where **x** is a multidimensional array of numerical type, and **d** is the dimension along which to perform the cumulative sum. The output **y** is the same size of **x**. Integer types are promoted to `int32`. If the dimension **d** is not specified, then the cumulative sum is applied along the first non-singular dimension.

11.7.2 Function Internals

The output is computed via

$$y(m_1, \dots, m_{d-1}, j, m_{d+1}, \dots, m_p) = \sum_{k=1}^j x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p).$$

11.7.3 Example

The default action is to perform the cumulative sum along the first non-singular dimension.

```
--> A = [5,1,3;3,2,1;0,3,1]
```

```
A =
  5 1 3
  3 2 1
  0 3 1
```

```
--> cumsum(A)
```

```
ans =
  5 1 3
  8 3 4
  8 6 5
```

To compute the cumulative sum along the columns:

```
--> cumsum(A,2)
```

```
ans =
  5 6 9
  3 5 6
  0 3 4
```

The cumulative sum also works along arbitrary dimensions

```
--> B(:,:,1) = [5,2;8,9];
--> B(:,:,2) = [1,0;3,0]
```

```
B =
```

```
(:,:,1) =
  5 2
  8 9
```

```
(:,:,2) =
  1 0
  3 0
```

```
--> cumsum(B,3)
```

```
ans =
```

```
(:,:,1) =
  5 2
```

```

8 9
(:, :, 2) =
6 2
11 9

```

11.8 DEAL Multiple Simultaneous Assignments

11.8.1 Usage

When making a function call, it is possible to assign multiple outputs in a single call, (see, e.g., `max` for an example). The `deal` call allows you to do the same thing with a simple assignment. The syntax for its use is

```
[a,b,c,...] = deal(expr)
```

where `expr` is an expression with multiple values. The simplest example is where `expr` is the dereference of a cell array, e.g. `expr <-- A{:}`. In this case, the `deal` call is equivalent to

```
a = A{1}; b = A{2}; c = A{3};
```

Other expressions which are multivalued are structure arrays with multiple entries (non-scalar), where field dereferencing has been applied.

11.9 DEC2HEX Convert Decimal Number to Hexadecimal

11.9.1 Usage

Converts an integer value into its hexadecimal representation. The syntax for its use is

```
y = dec2hex(x)
```

where `x` is an integer (and is promoted to a 64-bit integer if it is not). The returned value `y` is a string containing the hexadecimal representation of that integer. If you require a minimum length for the hexadecimal representation, you can specify an optional second argument

```
y = dec2hex(x,n)
```

where `n` indicates the minimum number of digits in the representation.

11.9.2 Example

Here are some simple examples:

```
--> dec2hex(1023)
```

```
ans =
3FF
```

```
--> dec2hex(58128493)
```

```
ans =
376F86D
```

11.10 DIFF Difference Function

11.10.1 Usage

```
y=diff(x)
y=diff(x,k)
y=diff(x,k,dim)
```

Produce difference of successive vector elements.

If \mathbf{x} is a vector of length n , `diff (x)` is the vector of first differences

$$[x_2 - x_1, \dots, x_n - x_{n-1}].$$

If \mathbf{x} is a matrix, `diff (x)` is the matrix of column differences along the first non-singleton dimension.

The second argument is optional. If supplied, `diff (x,k)`, where k is a nonnegative integer, returns the k -th differences. It is possible that k is larger than the first non-singleton dimension of the matrix. In this case, `diff` continues to take the differences along the next non-singleton dimension.

The dimension along which to take the difference can be explicitly stated with the optional variable `dim`. In this case the k -th order differences are calculated along this dimension. In the case where k exceeds `size (x, dim)` then an empty matrix is returned.

11.11 DOT Dot Product Function

11.11.1 Usage

Computes the scalar dot product of its two arguments. The general syntax for its use is

```
y = dot(x,z)
```

where \mathbf{x} and \mathbf{z} are numerical vectors of the same length. If \mathbf{x} and \mathbf{z} are multi-dimensional arrays of the same size, then the dot product is taken along the first non-singleton dimension. You can also specify the dimension to take the dot product along using the alternate form

```
y = dot(x,z,dim)
```

where `dim` specifies the dimension to take the dot product along.

11.12 FLOOR Floor Function

11.12.1 Usage

Computes the floor of an n-dimensional array elementwise. The floor of a number is defined as the smallest integer that is less than or equal to that number. The general syntax for its use is

```
y = floor(x)
```

where `x` is a multidimensional array of numerical type. The `floor` function preserves the type of the argument. So integer arguments are not modified, and `float` arrays return `float` arrays as outputs, and similarly for `double` arrays. The `floor` function is not defined for complex types.

11.12.2 Example

The following demonstrates the `floor` function applied to various (numerical) arguments. For integer arguments, the floor function has no effect:

```
--> floor(3)
```

```
ans =  
3
```

```
--> floor(-3)
```

```
ans =  
-3
```

Next, we take the `floor` of a floating point value:

```
--> floor(3.023)
```

```
ans =  
3
```

```
--> floor(-2.341)
```

```
ans =  
-3
```

11.13 GETFIELD Get Field Contents

11.13.1 Usage

Given a structure or structure array, returns the contents of the specified field. The first version is for scalar structures, and has the following syntax

```
y = getfield(x, 'fieldname')
```

and is equivalent to `y = x.fieldname` where `x` is a scalar (1 x 1) structure. If `x` is not a scalar structure, then `y` is the first value, i.e., it is equivalent to `y = x(1).fieldname`. The second form allows you to specify a subindex into a structure array, and has the following syntax

```
y = getfield(x, {m,n}, 'fieldname')
```

and is equivalent to `y = x(m,n).fieldname`. You can chain multiple references together using this syntax.

11.14 HEX2DEC Convert Hexadecimal Numbers To Decimal

11.14.1 Usage

Converts a hexadecimal number (encoded as a string matrix) into integers. The syntax for its use is

```
y = hex2dec(x)
```

where `x` is a character matrix where each row represents an integer in hexadecimal form. The output is of type `Double`.

11.14.2 Examples

```
--> hex2dec('3ff')
```

```
ans =
    1023
```

Or for a more complex example

```
--> hex2dec(['0ff'; '2de'; '123'])
```

```
ans =
    255
    734
    291
```

11.15 IMAG Imaginary Function

11.15.1 Usage

Returns the imaginary part of the input array for all elements. The general syntax for its use is

```
y = imag(x)
```

where `x` is an `n`-dimensional array of numerical type. The output is the same numerical type as the input, unless the input is `complex` or `dcomplex`. For `complex` inputs, the imaginary part is a floating point array, so that the return type is `float`. For `dcomplex` inputs, the imaginary part is a double precision floating point array, so that the return type is `double`. The `imag` function returns zeros for real and integer types.

11.15.2 Example

The following demonstrates `imag` applied to a complex scalar.

```
--> imag(3+4*i)
```

```
ans =
     4
```

The imaginary part of real and integer arguments is a vector of zeros, the same type and size of the argument.

```
--> imag([2,4,5,6])
```

```
ans =
     0     0     0     0
```

For a double-precision complex array,

```
--> imag([2.0+3.0*i,i])
```

```
ans =
     3     1
```

11.16 MAX Maximum Function

11.16.1 Usage

Computes the maximum of an array along a given dimension, or alternately, computes two arrays (entry-wise) and keeps the smaller value for each array. As a result, the `max` function has a number of syntaxes. The first one computes the maximum of an array along a given dimension. The first general syntax for its use is either

$$[y,n] = \max(x, [], d)$$

where `x` is a multidimensional array of numerical type, in which case the output `y` is the maximum of `x` along dimension `d`. The second argument `n` is the index that results in the maximum. In the event that multiple maxima are present with the same value, the index of the first maximum is used. The second general syntax for the use of the `max` function is

$$[y,n] = \max(x)$$

In this case, the maximum is taken along the first non-singleton dimension of `x`. For complex data types, the maximum is based on the magnitude of the numbers. NaNs are ignored in the calculations. The third general syntax for the use of the `max` function is as a comparison function for pairs of arrays. Here, the general syntax is

$$y = \max(x,z)$$

where `x` and `z` are either both numerical arrays of the same dimensions, or one of the two is a scalar. In the first case, the output is the same size as both arrays, and is defined elementwise by the smaller of the two arrays. In the second case, the output is defined elementwise by the smaller of the array entries and the scalar.

11.16.2 Function Internals

In the general version of the `max` function which is applied to a single array (using the `max(x, [], d)` or `max(x)` syntaxes), The output is computed via

$$y(m_1, \dots, m_{d-1}, 1, m_{d+1}, \dots, m_p) = \max_k x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p),$$

and the output array `n` of indices is calculated via

$$n(m_1, \dots, m_{d-1}, 1, m_{d+1}, \dots, m_p) = \arg \max_k x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p)$$

In the two-array version (`max(x, z)`), the single output is computed as

$$y(m_1, \dots, m_{d-1}, 1, m_{d+1}, \dots, m_p) = \begin{cases} x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p) & x(\dots) \leq z(\dots) \\ z(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p) & z(\dots) < x(\dots). \end{cases}$$

11.16.3 Example

The following piece of code demonstrates various uses of the maximum function. We start with the one-array version.

```
--> A = [5,1,3;3,2,1;0,3,1]
```

```
A =
 5 1 3
 3 2 1
 0 3 1
```

We first take the maximum along the columns, resulting in a row vector.

```
--> max(A)
```

```
ans =
 5 3 3
```

Next, we take the maximum along the rows, resulting in a column vector.

```
--> max(A, [], 2)
```

```
ans =
 5
 3
 3
```

When the dimension argument is not supplied, `max` acts along the first non-singular dimension. For a row vector, this is the column direction:

```
--> max([5,3,2,9])
```

```
ans =
 9
```

For the two-argument version, we can compute the smaller of two arrays, as in this example:

```
--> a = int8(100*randn(4))
```

```
a =  
    8  127  -51  -57  
   -5  -25  -62  -14  
  -17  -73  127  -73  
 -102   21   36   39
```

```
--> b = int8(100*randn(4))
```

```
b =  
   59  127   46   24  
  -98  -78  -50 -114  
 -128   57  127   -2  
  -64  -95   30  127
```

```
--> max(a,b)
```

```
ans =  
   59  127   46   24  
   -5  -25  -50  -14  
  -17   57  127   -2  
  -64   21   36  127
```

Or alternately, we can compare an array with a scalar

```
--> a = randn(2)
```

```
a =  
    1.4356    0.4894  
   -0.8816    0.6359
```

```
--> max(a,0)
```

```
ans =  
    1.4356    0.4894  
         0    0.6359
```

11.17 MEAN Mean Function

11.17.1 Usage

Computes the mean of an array along a given dimension. The general syntax for its use is

```
y = mean(x,d)
```

where \mathbf{x} is an n -dimensions array of numerical type. The output is of the same numerical type as the input. The argument \mathbf{d} is optional, and denotes the dimension along which to take the mean. The output \mathbf{y} is the same size as \mathbf{x} , except that it is singular along the mean direction. So, for example, if \mathbf{x} is a $3 \times 3 \times 4$ array, and we compute the mean along dimension $\mathbf{d}=2$, then the output is of size $3 \times 1 \times 4$.

11.17.2 Function Internals

The output is computed via

$$y(m_1, \dots, m_{d-1}, 1, m_{d+1}, \dots, m_p) = \frac{1}{N} \sum_{k=1}^N x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p)$$

If \mathbf{d} is omitted, then the mean is taken along the first non-singleton dimension of \mathbf{x} .

11.17.3 Example

The following piece of code demonstrates various uses of the mean function

```
--> A = [5,1,3;3,2,1;0,3,1]
```

```
A =
  5 1 3
  3 2 1
  0 3 1
```

We start by calling `mean` without a dimension argument, in which case it defaults to the first nonsingular dimension (in this case, along the columns or $\mathbf{d} = 1$).

```
--> mean(A)

ans =
  2.6667    2.0000    1.6667
```

Next, we take the mean along the rows.

```
--> mean(A,2)

ans =
  3.0000
  2.0000
  1.3333
```

11.18 MIN Minimum Function

11.18.1 Usage

Computes the minimum of an array along a given dimension, or alternately, computes two arrays (entry-wise) and keeps the smaller value for each array. As a result, the `min` function has a number

of syntaxes. The first one computes the minimum of an array along a given dimension. The first general syntax for its use is either

$$[y,n] = \min(x,[],d)$$

where x is a multidimensional array of numerical type, in which case the output y is the minimum of x along dimension d . The second argument n is the index that results in the minimum. In the event that multiple minima are present with the same value, the index of the first minimum is used. The second general syntax for the use of the `min` function is

$$[y,n] = \min(x)$$

In this case, the minimum is taken along the first non-singleton dimension of x . For complex data types, the minimum is based on the magnitude of the numbers. NaNs are ignored in the calculations. The third general syntax for the use of the `min` function is as a comparison function for pairs of arrays. Here, the general syntax is

$$y = \min(x,z)$$

where x and z are either both numerical arrays of the same dimensions, or one of the two is a scalar. In the first case, the output is the same size as both arrays, and is defined elementwise by the smaller of the two arrays. In the second case, the output is defined elementwise by the smaller of the array entries and the scalar.

11.18.2 Function Internals

In the general version of the `min` function which is applied to a single array (using the `min(x,[],d)` or `min(x)` syntaxes), The output is computed via

$$y(m_1, \dots, m_{d-1}, 1, m_{d+1}, \dots, m_p) = \min_k x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p),$$

and the output array n of indices is calculated via

$$n(m_1, \dots, m_{d-1}, 1, m_{d+1}, \dots, m_p) = \arg \min_k x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p)$$

In the two-array version (`min(x,z)`), the single output is computed as

$$y(m_1, \dots, m_{d-1}, 1, m_{d+1}, \dots, m_p) = \begin{cases} x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p) & x(\dots) \leq z(\dots) \\ z(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p) & z(\dots) < x(\dots). \end{cases}$$

11.18.3 Example

The following piece of code demonstrates various uses of the minimum function. We start with the one-array version.

```
--> A = [5,1,3;3,2,1;0,3,1]
```

```
A =
  5  1  3
  3  2  1
  0  3  1
```

We first take the minimum along the columns, resulting in a row vector.

```
--> min(A)
```

```
ans =
  0 1 1
```

Next, we take the minimum along the rows, resulting in a column vector.

```
--> min(A, [], 2)
```

```
ans =
  1
  1
  0
```

When the dimension argument is not supplied, `min` acts along the first non-singular dimension. For a row vector, this is the column direction:

```
--> min([5,3,2,9])
```

```
ans =
  2
```

For the two-argument version, we can compute the smaller of two arrays, as in this example:

```
--> a = int8(100*randn(4))
```

```
a =
  27   61 -128  -22
 -54  -74  -27  -61
  66 -128  127   34
 -13 -128   42   86
```

```
--> b = int8(100*randn(4))
```

```
b =
  56   33  -15   17
 127  127   29  -97
  99  -45  117 -128
  56 -114  127  104
```

```
--> min(a,b)
```

```
ans =
  27   33 -128  -22
 -54  -74  -27  -97
  66 -128  117 -128
 -13 -128   42   86
```

Or alternately, we can compare an array with a scalar

```
--> a = randn(2)

a =
   -0.2804    0.3739
   -0.8883   -0.4260

--> min(a,0)

ans =
   -0.2804         0
   -0.8883   -0.4260
```

11.19 NUM2HEX Convert Numbers to IEEE Hex Strings

11.19.1 Usage

Converts single and double precision arrays to IEEE hex strings. The syntax for its use is

```
y = num2hex(x)
```

where `x` is either a `float` or `double` array. The output `y` is a `n-by-p` character array, where `n` is the number of elements in `x`, and `p` is 16 for `double` arrays, and 8 for `single` arrays.

11.19.2 Example

Some interesting numbers

```
--> num2hex([1 0 0.1 -pi inf nan])
```

```
ans =
3ff0000000000000
0000000000000000
3fb999999999999a
c00921fb54442d18
7ff0000000000000
fff8000000000000
```

The same in single precision

```
--> num2hex(float([1 0 0.1 -pi inf nan]))
```

```
ans =
3f800000
00000000
3dcccccd
c0490fdb
```

```
7f800000
fff80000
```

11.20 PROD Product Function

11.20.1 Usage

Computes the product of an array along a given dimension. The general syntax for its use is

```
y = prod(x,d)
```

where x is an n -dimensions array of numerical type. The output is of the same numerical type as the input, except for integer types, which are automatically promoted to `int32`. The argument d is optional, and denotes the dimension along which to take the product. The output is computed via

$$y(m_1, \dots, m_{d-1}, 1, m_{d+1}, \dots, m_p) = \prod_k x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p)$$

If d is omitted, then the product is taken along the first non-singleton dimension of x . Note that by definition (starting with FreeMat 2.1) `prod([]) = 1`.

11.20.2 Example

The following piece of code demonstrates various uses of the product function

```
--> A = [5,1,3;3,2,1;0,3,1]
```

```
A =
 5 1 3
 3 2 1
 0 3 1
```

We start by calling `prod` without a dimension argument, in which case it defaults to the first nonsingular dimension (in this case, along the columns or $d = 1$).

```
--> prod(A)
```

```
ans =
 0 6 3
```

Next, we take the product along the rows.

```
--> prod(A,2)
```

```
ans =
 15
 6
 0
```

11.21 REAL Real Function

11.21.1 Usage

Returns the real part of the input array for all elements. The general syntax for its use is

```
y = real(x)
```

where **x** is an **n**-dimensional array of numerical type. The output is the same numerical type as the input, unless the input is **complex** or **dcomplex**. For **complex** inputs, the real part is a floating point array, so that the return type is **float**. For **dcomplex** inputs, the real part is a double precision floating point array, so that the return type is **double**. The **real** function does nothing to real and integer types.

11.21.2 Example

The following demonstrates the **real** applied to a complex scalar.

```
--> real(3+4*i)
```

```
ans =  
3
```

The **real** function has no effect on real arguments:

```
--> real([2,3,4])
```

```
ans =  
2 3 4
```

For a double-precision complex array,

```
--> real([2.0+3.0*i,i])
```

```
ans =  
2 0
```

11.22 ROUND Round Function

11.22.1 Usage

Rounds an **n**-dimensional array to the nearest integer elementwise. The general syntax for its use is

```
y = round(x)
```

where **x** is a multidimensional array of numerical type. The **round** function preserves the type of the argument. So integer arguments are not modified, and **float** arrays return **float** arrays as outputs, and similarly for **double** arrays. The **round** function is not defined for **complex** or **dcomplex** types.

11.22.2 Example

The following demonstrates the `round` function applied to various (numerical) arguments. For integer arguments, the `round` function has no effect:

```
--> round(3)
```

```
ans =  
3
```

```
--> round(-3)
```

```
ans =  
-3
```

Next, we take the `round` of a floating point value:

```
--> round(3.023f)
```

```
ans =  
3
```

```
--> round(-2.341f)
```

```
ans =  
-2
```

Note that the return type is a `float` also. Finally, for a `double` type:

```
--> round(4.312)
```

```
ans =  
4
```

```
--> round(-5.32)
```

```
ans =  
-5
```

11.23 STD Standard Deviation Function

11.23.1 Usage

Computes the standard deviation of an array along a given dimension. The general syntax for its use is

```
y = std(x,d)
```

where x is an n -dimensions array of numerical type. The output is of the same numerical type as the input. The argument d is optional, and denotes the dimension along which to take the variance. The output y is the same size as x , except that it is singular along the mean direction. So, for example, if x is a $3 \times 3 \times 4$ array, and we compute the mean along dimension $d=2$, then the output is of size $3 \times 1 \times 4$.

11.23.2 Example

The following piece of code demonstrates various uses of the `std` function

```
--> A = [5,1,3;3,2,1;0,3,1]
```

```
A =
  5 1 3
  3 2 1
  0 3 1
```

We start by calling `std` without a dimension argument, in which case it defaults to the first nonsingular dimension (in this case, along the columns or $d = 1$).

```
--> std(A)
```

```
ans =
  2.5166    1.0000    1.1547
```

Next, we take the variance along the rows.

```
--> std(A,2)
```

```
ans =
  2.0000
  1.0000
  1.5275
```

11.24 SUB2IND Convert Multiple Indexing To Linear Indexing

11.24.1 Usage

The `sub2ind` function converts a multi-dimensional indexing expression into a linear (or vector) indexing expression. The syntax for its use is

```
y = sub2ind(sizevec,d1,d2,...,dn)
```

where `sizevec` is the size of the array being indexed into, and each d_i is a vector of the same length, containing index values. The basic idea behind `sub2ind` is that it makes

```
[z(d1(1),d2(1),...,dn(1)),...,z(d1(n),d2(n),...,dn(n))]
```

equivalent to

```
z(sub2ind(size(z),d1,d2,...,dn))
```

where the later form is using vector indexing, and the former one is using native, multi-dimensional indexing.

11.24.2 Example

Suppose we have a simple 3 x 4 matrix A containing some random integer elements

```
--> A = randi(ones(3,4),10*ones(3,4))
```

```
A =
     2     1     4     2
     8    10     4     7
    10     7     4    10
```

We can extract the elements (1,3), (2,3), (3,4) of A via `sub2ind`. To calculate which elements of A this corresponds to, we can use `sub2ind` as

```
--> n = sub2ind(size(A),1:3,2:4)
```

```
n =
     4     8    12
```

```
--> A(n)
```

```
ans =
     1     4    10
```

11.25 SUM Sum Function

11.25.1 Usage

Computes the summation of an array along a given dimension. The general syntax for its use is

```
y = sum(x,d)
```

where `x` is an `n`-dimensions array of numerical type. The output is of the same numerical type as the input. The argument `d` is optional, and denotes the dimension along which to take the summation. The output `y` is the same size as `x`, except that it is singular along the summation direction. So, for example, if `x` is a 3 x 3 x 4 array, and we compute the summation along dimension `d=2`, then the output is of size 3 x 1 x 4.

11.25.2 Function Internals

The output is computed via

$$y(m_1, \dots, m_{d-1}, 1, m_{d+1}, \dots, m_p) = \sum_k x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p)$$

If d is omitted, then the summation is taken along the first non-singleton dimension of x .

11.25.3 Example

The following piece of code demonstrates various uses of the summation function

```
--> A = [5,1,3;3,2,1;0,3,1]
```

```
A =
  5  1  3
  3  2  1
  0  3  1
```

We start by calling `sum` without a dimension argument, in which case it defaults to the first nonsingular dimension (in this case, along the columns or $d = 1$).

```
--> sum(A)
```

```
ans =
  8  6  5
```

Next, we take the sum along the rows.

```
--> sum(A,2)
```

```
ans =
  9
  6
  4
```

11.26 TEST Test Function

11.26.1 Usage

Tests for the argument array to be all logical 1s. It is completely equivalent to the `all` function applied to a vectorized form of the input. The syntax for the `test` function is

```
y = test(x)
```

and the result is equivalent to `all(x(:))`.

11.27 VAR Variance Function

11.27.1 Usage

Computes the variance of an array along a given dimension. The general syntax for its use is

```
y = var(x,d)
```

where \mathbf{x} is an n -dimensions array of numerical type. The output is of the same numerical type as the input. The argument \mathbf{d} is optional, and denotes the dimension along which to take the variance. The output \mathbf{y} is the same size as \mathbf{x} , except that it is singular along the mean direction. So, for example, if \mathbf{x} is a $3 \times 3 \times 4$ array, and we compute the mean along dimension $\mathbf{d}=2$, then the output is of size $3 \times 1 \times 4$.

11.27.2 Function Internals

The output is computed via

$$y(m_1, \dots, m_{d-1}, 1, m_{d+1}, \dots, m_p) = \frac{1}{N-1} \sum_{k=1}^N (x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p) - \bar{x})^2,$$

where

$$\bar{x} = \frac{1}{N} \sum_{k=1}^N x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p)$$

If \mathbf{d} is omitted, then the mean is taken along the first non-singleton dimension of \mathbf{x} .

11.27.3 Example

The following piece of code demonstrates various uses of the `var` function

```
--> A = [5,1,3;3,2,1;0,3,1]
```

```
A =
  5 1 3
  3 2 1
  0 3 1
```

We start by calling `var` without a dimension argument, in which case it defaults to the first nonsingular dimension (in this case, along the columns or $\mathbf{d} = 1$).

```
--> var(A)
```

```
ans =
  6.3333    1.0000    1.3333
```

Next, we take the variance along the rows.

```
--> var(A,2)
```

```
ans =  
    4.0000  
    1.0000  
    2.3333
```

11.28 VEC Reshape to a Vector

11.28.1 Usage

Reshapes an n-dimensional array into a column vector. The general syntax for its use is

```
y = vec(x)
```

where x is an n-dimensional array (not necessarily numeric). This function is equivalent to the expression $y = x(:)$.

11.28.2 Example

A simple example of the `vec` operator reshaping a 2D matrix:

```
--> A = [1,2,4,3;2,3,4,5]
```

```
A =  
    1  2  4  3  
    2  3  4  5
```

```
--> vec(A)
```

```
ans =  
    1  
    2  
    2  
    3  
    4  
    4  
    3  
    5
```


Chapter 12

Inspection Functions

12.1 CLEAR Clear or Delete a Variable

12.1.1 Usage

Clears a set of variables from the current context, or alternately, delete all variables defined in the current context. There are several formats for the function call. The first is the explicit form in which a list of variables are provided:

```
clear a1 a2 ...
```

The variables can be persistent or global, and they will be deleted. The second form

```
clear 'all'
```

clears all variables and libraries from the current context. Alternately, you can use the form:

```
clear 'libs'
```

which will unload any libraries or DLLs that have been `imported`. Optionally, you can specify that persistent variables should be cleared via:

```
clear 'persistent'
```

and similarly for global variables:

```
clear 'global'
```

You can use

```
clear 'classes'
```

to clear all definitions of user-defined classes. With no arguments, `clear` defaults to clearing `'all'`.

12.1.2 Example

Here is a simple example of using `clear` to delete a variable. First, we create a variable called `a`:

```
--> a = 53
```

```
a =  
53
```

Next, we clear `a` using the `clear` function, and verify that it is deleted.

```
--> clear a  
--> a  
Error: Undefined function or variable a
```

12.2 END End Function

12.2.1 Usage

Computes the size of a variable along a given dimension. The syntax for its use is

```
y = end(x,dim,subindexes)
```

where `x` is the array to be analyzed, `dim` is the dimension along which to compute the end, and `subindexes` indicates how many dimensions are involved in the `end` calculation.

12.3 EXIST Test for Existence

12.3.1 Usage

Tests for the existence of a variable, function, directory or file. The general syntax for its use is

```
y = exist(item,kind)
```

where `item` is a string containing the name of the item to look for, and `kind` is a string indicating the type of the search. The `kind` must be one of

- `'builtin'` checks for built-in functions
- `'dir'` checks for directories
- `'file'` checks for files
- `'var'` checks for variables
- `'all'` checks all possibilities (same as leaving out `kind`)

You can also leave the `kind` specification out, in which case the calling syntax is

```
y = exist(item)
```

The return code is one of the following:

- 0 - if `item` does not exist
- 1 - if `item` is a variable in the workspace
- 2 - if `item` is an M file on the search path, a full pathname to a file, or an ordinary file on your search path
- 5 - if `item` is a built-in FreeMat function
- 7 - if `item` is a directory

Note: previous to version 1.10, `exist` used a different notion of existence for variables: a variable was said to exist if it was defined and non-empty. This test is now performed by `isset`.

12.3.2 Example

Some examples of the `exist` function. Note that generally `exist` is used in functions to test for keywords. For example,

```
function y = testfunc(a, b, c)
if (~exist('c'))
    % c was not defined, so establish a default
    c = 13;
end
y = a + b + c;
```

An example of `exist` in action.

```
--> a = randn(3,5,2)

a =

(:, :, 1) =
    0.0744   -0.7000   -1.8022   -0.0255   -0.2933
   -0.0152   -0.5401   -0.2030   -0.5437    0.9261
   -0.4359   -1.2334    0.5156    0.0821   -0.5065

(:, :, 2) =
    0.7408    0.6302   -0.2798    0.7279    1.7674
    0.9063   -1.7321   -0.8851   -0.4091   -0.5516
    1.4880    0.5030    0.6747   -0.5492    1.1867

--> b = []

b =
[]
--> who
```

Variable Name	Type	Flags	Size
a	double		[3 5 2]
b	double		[0 0]

```

--> exist('a')

ans =
    1

--> exist('b')

ans =
    1

--> exist('c')

ans =
    0

```

12.4 FIELDNAMES Fieldnames of a Structure

12.4.1 Usage

Returns a cell array containing the names of the fields in a structure array. The syntax for its use is

```
x = fieldnames(y)
```

where *y* is a structure array of object array. The result is a cell array, with one entry per field in *y*.

12.4.2 Example

We define a simple structure array:

```

--> y.foo = 3; y.goo = 'hello';
--> x = fieldnames(y)

```

```

x =
    [foo]
    [goo]

```

12.5 ISA Test Type of Variable

12.5.1 Usage

Tests the type of a variable. The syntax for its use is

```
y = isa(x,type)
```

where `x` is the variable to test, and `type` is the type. Supported built-in types are

- `'cell'` for cell-arrays
- `'struct'` for structure-arrays
- `'logical'` for logical arrays
- `'uint8'` for unsigned 8-bit integers
- `'int8'` for signed 8-bit integers
- `'uint16'` for unsigned 16-bit integers
- `'int16'` for signed 16-bit integers
- `'uint32'` for unsigned 32-bit integers
- `'int32'` for signed 32-bit integers
- `'uint64'` for unsigned 64-bit integers
- `'int64'` for signed 64-bit integers
- `'single'` for 32-bit floating point numbers
- `'double'` for 64-bit floating point numbers
- `'char'` for string arrays

If the argument is a user-defined type (via the `class` function), then the name of that class is returned.

12.5.2 Examples

Here are some examples of the `isa` call.

```
--> a = {1}

a =
    [1]

--> isa(a,'char')

ans =
    0

--> isa(a,'cell')

ans =
    1
```

Here we use `isa` along with shortcut boolean evaluation to safely determine if a variable contains the string `'hello'`

```
--> a = 'hello'

a =
hello
--> isa(a,'char') && strcmp(a,'hello')

ans =
1
```

12.6 ISCELL Test For Cell Array

12.6.1 Usage

The syntax for `iscell` is

```
x = iscell(y)
```

and it returns a logical 1 if the argument is a cell array and a logical 0 otherwise.

12.6.2 Example

Here are some examples of `iscell`

```
--> iscell('foo')

ans =
0

--> iscell(2)

ans =
0

--> iscell({1,2,3})

ans =
1
```

12.7 ISCELLSTR Test For Cell Array of Strings

12.7.1 Usage

The syntax for `iscellstr` is

```
x = iscellstr(y)
```

and it returns a logical 1 if the argument is a cell array in which every cell is a character array (or is empty), and a logical 0 otherwise.

12.7.2 Example

Here is a simple example

```
--> A = {'Hello', 'Yellow'; 'Mellow', 'Othello'}
```

```
A =  
[Hello] [Yellow]  
[Mellow] [Othello]
```

```
--> iscellstr(A)
```

```
ans =  
1
```

12.8 ISCHAR Test For Character Array (string)

12.8.1 Usage

The syntax for `ischar` is

```
x = ischar(y)
```

and it returns a logical 1 if the argument is a string and a logical 0 otherwise.

12.9 ISEMPTY Test For Variable Empty

12.9.1 Usage

The `isempty` function returns a boolean that indicates if the argument variable is empty or not. The general syntax for its use is

```
y = isempty(x).
```

12.9.2 Examples

Here are some examples of the `isempty` function

```
--> a = []
```

```
a =  
[]  
--> isempty(a)
```

```
ans =  
  1  
  
--> b = 1:3  
  
b =  
  1 2 3  
  
--> isempty(b)  
  
ans =  
  0
```

Note that if the variable is not defined, `isempty` does not return true.

```
--> clear x  
--> isempty(x)  
Error: Undefined function or variable x
```

12.10 ISEQUAL Test For Matrix Equality

12.10.1 Usage

Test two arrays for equality. The general format for its use is

```
y = isequal(a,b)
```

This function returns true if the two arrays are equal (compared element-wise). Unlike `issame` the `isequal` function will type convert where possible to do the comparison.

12.11 ISFIELD Test for Existence of a Structure Field

12.11.1 Usage

Given a structure array, tests to see if that structure array contains a field with the given name. The syntax for its use is

```
y = isfield(x,field)
```

and returns a logical 1 if `x` has a field with the name `field` and a logical 0 if not. It also returns a logical 0 if the argument `x` is not a structure array.

12.11.2 Example

Here we define a simple struct, and then test for some fields

```
--> a.foo = 32

a =
  foo: 32
--> a.goo = 64

a =
  foo: 32
  goo: 64
--> isfield(a,'goo')

ans =
  1

--> isfield(a,'got')

ans =
  0

--> isfield(pi,'round')

ans =
  0
```

12.12 ISHANDLE Test for Graphics Handle

12.12.1 Usage

Given a constant, this routine will test to see if the constant is a valid graphics handle or not. The syntax for its use is

```
y = ishandle(h,type)
```

and returns a logical 1 if *x* is a handle of type *type* and a logical 0 if not.

12.13 ISINF Test for infinities

12.13.1 Usage

Returns true for entries of an array that are infs (i.e., infinities). The usage is

```
y = isinf(x)
```

The result is a logical array of the same size as *x*, which is true if *x* is not-a-number, and false otherwise. Note that for `complex` or `dcomplex` data types that the result is true if either the real or imaginary parts are infinite.

12.13.2 Example

Suppose we have an array of floats with one element that is `inf`:

```
--> a = [1.2 3.4 inf 5]

a =
    1.2000    3.4000 Inf    5.0000

--> isinf(a)

ans =
    0 0 1 0

--> b = 3./[2 5 0 3 1]

b =
    1.5000    0.6000 Inf    1.0000    3.0000
```

12.14 ISINTTYPE Test For Integer-type Array

12.14.1 Usage

The syntax for `isinttype` is

```
x = isinttype(y)
```

and it returns a logical 1 if the argument is an integer type and a logical 0 otherwise. Note that this function only tests the type of the variable, not the value. So if, for example, `y` is a `float` array containing all integer values, it will still return a logical 0.

12.15 ISLOGICAL Test for Logical Array

12.15.1 Usage

The syntax for `islogical` is

```
x = islogical(y)
```

and it returns a logical 1 if the argument is a logical array and a logical 0 otherwise.

12.16 ISMATRIX Test For a 2D Matrix

12.16.1 Usage

This function tests to see if the argument is a matrix. The syntax for `ismatrix` is

```
x = ismatrix(y)
```

and it returns a logical 1 if the argument is size `N x M` or `M x N` and a logical 0 otherwise.

12.17 ISNAN Test for Not-a-Numbers

12.17.1 Usage

Returns true for entries of an array that are NaN's (i.e., Not-a-Numbers). The usage is

```
y = isnan(x)
```

The result is a logical array of the same size as `x`, which is true if `x` is not-a-number, and false otherwise. Note that for complex data types that the result is true if either the real or imaginary parts are NaNs.

12.17.2 Example

Suppose we have an array of floats with one element that is `nan`:

```
--> a = [1.2 3.4 nan 5]

a =
    1.2000    3.4000 NaN    5.0000

--> isnan(a)

ans =
    0 0 1 0
```

12.18 ISNUMERIC Test for Numeric Array

12.18.1 Usage

The syntax for `isnumeric` is

```
x = isnumeric(y)
```

and it returns a logical 1 if the argument is a numeric (i.e., not a structure array, cell array, string or user defined class), and a logical 0 otherwise.

12.19 ISREAL Test For Real Array

12.19.1 Usage

The syntax for `isreal` is

```
x = isreal(y)
```

and it returns a logical 1 if the argument is real valued and a logical 0 otherwise.

12.20 ISSAME Test If Two Arrays Are Identical

12.20.1 Usage

Tests for two arrays to be identical. The syntax for its use is

```
y = issame(a,b)
```

where **a** and **b** are two arrays to compare. This comparison succeeds only if **a** and **b** are of the same data type, size, and contents. Unlike numerical equivalence tests, the `issame` function considers NaN to be equal in both arguments.

12.21 ISSCALAR Test For Scalar

12.21.1 Usage

The syntax for `isscalar` is

```
x = isscalar(y)
```

and it returns a logical 1 if the argument is a scalar, and a logical 0 otherwise.

12.22 ISSET Test If Variable Set

12.22.1 Usage

Tests for the existence and non-emptiness of a variable. the general syntax for its use is

```
y = isset('name')
```

where **name** is the name of the variable to test. This is functionally equivalent to

```
y = exist('name','var') & ~isempty(name)
```

It returns a logical 1 if the variable is defined in the current workspace, and is not empty, and returns a 0 otherwise.

12.22.2 Example

Some simple examples of using `isset`

```
--> who
  Variable Name      Type  Flags      Size
--> isset('a')

ans =
  0

--> a = [];
```

```
--> isset('a')
```

```
ans =  
0
```

```
--> a = 2;  
--> isset('a')
```

```
ans =  
1
```

12.23 ISSPARSE Test for Sparse Matrix

12.23.1 Usage

Test a matrix to see if it is sparse or not. The general format for its use is

```
y = issparse(x)
```

This function returns true if x is encoded as a sparse matrix, and false otherwise.

12.23.2 Example

Here is an example of using `issparse`:

```
--> a = [1,0,0,5;0,3,2,0]
```

```
a =  
1 0 0 5  
0 3 2 0
```

```
--> issparse(a)
```

```
ans =  
0
```

```
--> A = sparse(a)
```

```
A =  
1 1 1  
2 2 3  
2 3 2  
1 4 5
```

```
--> issparse(A)
```

```
ans =  
1
```

12.24 ISSQUARE Test For a Square matrix

12.24.1 Usage

This function tests to see if the argument is a square matrix. The syntax for `issquare` is

```
x = issquare(y)
```

and it returns a logical 1 if the argument is size $N \times N$ logical 0 otherwise.

12.25 ISSTR Test For Character Array (string)

12.25.1 Usage

The syntax for `isstr` is

```
x = isstr(y)
```

and it returns a logical 1 if the argument is a string and a logical 0 otherwise.

12.26 ISSTRUCT Test For Structure Array

12.26.1 Usage

The syntax for `isstruct` is

```
x = isstruct(y)
```

and it returns a logical 1 if the argument is a structure array, and a logical 0 otherwise.

12.27 ISVECTOR Test For a Vector

12.27.1 Usage

This function tests to see if the argument is a vector. The syntax for `isvector` is

```
x = isvector(y)
```

and it returns a logical 1 if the argument is size $N \times 1$ or $1 \times N$ and a logical 0 otherwise.

12.28 LENGTH Length of an Array

12.28.1 Usage

Returns the length of an array `x`. The syntax for its use is

```
y = length(x)
```

and is defined as the maximum length of `x` along any of its dimensions, i.e., `max(size(x))`. If you want to determine the number of elements in `x`, use the `numel` function instead.

12.28.2 Example

For a 4 x 4 x 3 matrix, the length is 4, not 48, as you might expect.

```
--> x = rand(4,4,3);  
--> length(x)
```

```
ans =  
4
```

12.29 MAXDIM Maximum Dimension in Array

12.29.1 Usage

The `maxdim` function returns the lowest order dimension along which an array is largest. The general syntax for its use is

```
n = maxdim(x)
```

and is equivalent to `min(find(size(x) == max(size(x))))`.

12.30 NDIMS Number of Dimensions in Array

12.30.1 Usage

The `ndims` function returns the number of dimensions allocated in an array. The general syntax for its use is

```
n = ndims(x)
```

and is equivalent to `length(size(x))`.

12.31 NNZ Number of Nonzeros

12.31.1 Usage

Returns the number of nonzero elements in a matrix. The general format for its use is

```
y = nnz(x)
```

This function returns the number of nonzero elements in a matrix or array. This function works for both sparse and non-sparse arrays. For

12.31.2 Example

```
--> a = [1,0,0,5;0,3,2,0]
```

```
a =  
 1 0 0 5  
 0 3 2 0
```

```
--> nnz(a)
```

```
ans =  
 4
```

```
--> A = sparse(a)
```

```
A =  
 1 1 1  
 2 2 3  
 2 3 2  
 1 4 5  
--> nnz(A)
```

```
ans =  
 4
```

12.32 NUMEL Number of Elements in an Array

12.32.1 Usage

Returns the number of elements in an array `x`, or in a subindex expression. The syntax for its use is either

```
y = numel(x)
```

or

```
y = numel(x,varargin)
```

Generally, `numel` returns `prod(size(x))`, the number of total elements in `x`. However, you can specify a number of indexing expressions for `varargin` such as `index1`, `index2`, ..., `indexm`. In that case, the output of `numel` is `prod(size(x(index1,...,indexm)))`.

12.32.2 Example

For a 4 x 4 x 3 matrix, the length is 4, not 48, as you might expect, but `numel` is 48.

```
--> x = rand(4,4,3);  
--> length(x)
```

```
ans =  
    4
```

```
--> numel(x)
```

```
ans =  
    48
```

Here is an example of using `numel` with indexing expressions.

```
--> numel(x,1:3,1:2,2)
```

```
ans =  
    6
```

12.33 SIZE Size of a Variable

12.33.1 Usage

Returns the size of a variable. There are two syntaxes for its use. The first syntax returns the size of the array as a vector of integers, one integer for each dimension

```
[d1,d2,...,dn] = size(x)
```

The other format returns the size of `x` along a particular dimension:

```
d = size(x,n)
```

where `n` is the dimension along which to return the size.

12.33.2 Example

```
--> a = randn(23,12,5);  
--> size(a)
```

```
ans =  
    23 12  5
```

Here is an example of the second form of `size`.

```
--> size(a,2)
```

```
ans =  
    12
```

12.34 TYPEOF Determine the Type of an Argument

12.34.1 Usage

Returns a string describing the type of an array. The syntax for its use is

```
y = typeof(x),
```

The returned string is one of

- 'cell' for cell-arrays
- 'struct' for structure-arrays
- 'logical' for logical arrays
- 'uint8' for unsigned 8-bit integers
- 'int8' for signed 8-bit integers
- 'uint16' for unsigned 16-bit integers
- 'int16' for signed 16-bit integers
- 'uint32' for unsigned 32-bit integers
- 'int32' for signed 32-bit integers
- 'float' for 32-bit floating point numbers
- 'double' for 64-bit floating point numbers
- 'string' for string arrays

12.34.2 Example

The following piece of code demonstrates the output of the `typeof` command for each possible type. The first example is with a simple cell array.

```
--> typeof({1})
```

```
ans =  
cell
```

The next example uses the `struct` constructor to make a simple scalar struct.

```
--> typeof(struct('foo',3))
```

```
ans =  
struct
```

The next example uses a comparison between two scalar integers to generate a scalar logical type.

```
--> typeof(3>5)
```

```
ans =  
logical
```

For the integers, the typecast operations are used to generate the arguments.

```
--> typeof(uint8(3))
```

```
ans =  
uint8  
--> typeof(int8(8))
```

```
ans =  
int8  
--> typeof(uint16(3))
```

```
ans =  
uint16  
--> typeof(int16(8))
```

```
ans =  
int16  
--> typeof(uint32(3))
```

```
ans =  
uint32  
--> typeof(int32(3))
```

```
ans =  
int32  
--> typeof(uint64(3))
```

```
ans =  
uint64  
--> typeof(int64(3))
```

```
ans =  
int64
```

Float, and double can be created using the suffixes.

```
--> typeof(1.0f)
```

```
ans =  
single  
--> typeof(1.0D)
```

```
ans =  
double  
--> typeof(1.0f+i)  
  
ans =  
single  
--> typeof(1.0D+2.0D*i)  
  
ans =  
double
```

12.35 WHERE Get Information on Program Stack

12.35.1 Usage

Returns information on the current stack. The usage is

```
where
```

The result is a kind of stack trace that indicates the state of the current call stack, and where you are relative to the stack.

12.35.2 Example

Suppose we have the following chain of functions.

```
chain1.m  
function chain1  
a = 32;  
b = a + 5;  
chain2(b)  
  
chain2.m  
function chain2(d)  
d = d + 5;  
chain3  
  
chain3.m  
function chain3  
g = 54;  
f = g + 1;  
keyboard
```

The execution of the `where` command shows the stack trace.

```
--> chain1
[chain3,4]--> where
In /home/basu/dev/branches/FreeMat4/help/tmp/chain3.m(chain3) at line 4
    In /home/basu/dev/branches/FreeMat4/help/tmp/chain2.m(chain2) at line 4
    In /home/basu/dev/branches/FreeMat4/help/tmp/chain1.m(chain1) at line 4
    In scratch() at line 2
    In base(base)
    In base()
    In global()
[chain3,4]
```

12.36 WHICH Get Information on Function

12.36.1 Usage

Returns information on a function (if defined). The usage is

```
which(fname)
```

where `fname` is a **string** argument that contains the name of the function. For functions and scripts defined via `.m` files, the `which` command returns the location of the source file:

```
y = which(fname)
```

will return the filename for the `.m` file corresponding to the given function, and an empty string otherwise.

12.36.2 Example

First, we apply the `which` command to a built in function.

```
--> which fft
Function fft is a built in function
```

Next, we apply it to a function defined via a `.m` file.

```
--> which fliplr
Function fliplr, M-File function in file '/home/basu/dev/branches/FreeMat4/src/toolbox/array/fliplr.m'
```

12.37 WHO Describe Currently Defined Variables

12.37.1 Usage

Reports information on either all variables in the current context or on a specified set of variables. For each variable, the `who` function indicates the size and type of the variable as well as if it is a global or persistent. There are two formats for the function call. The first is the explicit form, in which a list of variables are provided:

```
who a1 a2 ...
```

In the second form

```
who
```

the `who` function lists all variables defined in the current context (as well as global and persistent variables). Note that there are two alternate forms for calling the `who` function:

```
who 'a1' 'a2' ...
```

and

```
who('a1','a2',...)
```

12.37.2 Example

Here is an example of the general use of `who`, which lists all of the variables defined.

```
--> c = [1,2,3];
--> f = 'hello';
--> p = randn(1,256);
--> who
Variable Name      Type   Flags      Size
                c    double
                f    char
                p    double
                [1 3]
                [1 5]
                [1 256]
```

In the second case, we examine only a specific variable:

```
--> who c
Variable Name      Type   Flags      Size
                c    double
                [1 3]
--> who('c')
Variable Name      Type   Flags      Size
                c    double
                [1 3]
```

12.38 WHOS Describe Currently Defined Variables With Memory Usage

12.38.1 Usage

Reports information on either all variables in the current context or on a specified set of variables. For each variable, the `who` function indicates the size and type of the variable as well as if it is a global or persistent. There are two formats for the function call. The first is the explicit form, in which a list of variables are provided:

```
whos a1 a2 ...
```

12.38. WHOS DESCRIBE CURRENTLY DEFINED VARIABLES WITH MEMORY USAGE²⁴³

In the second form

`whos`

the `whos` function lists all variables defined in the current context (as well as global and persistent variables). Note that there are two alternate forms for calling the `whos` function:

`whos 'a1' 'a2' ...`

and

`whos('a1','a2',...)`

Chapter 13

Type Conversion Functions

13.1 BIN2DEC Convert Binary String to Decimal

13.1.1 USAGE

Converts a binary string to an integer. The syntax for its use is

```
y = bin2dec(x)
```

where *x* is a binary string. If *x* is a matrix, then the resulting *y* is a column vector.

13.1.2 Example

Here we convert some numbers to bits

```
--> bin2dec('101110')
```

```
ans =  
46
```

```
--> bin2dec('010')
```

```
ans =  
2
```

13.2 BIN2INT Convert Binary Arrays to Integer

13.2.1 Usage

Converts the binary decomposition of an integer array back to an integer array. The general syntax for its use is

```
y = bin2int(x)
```

where `x` is a multi-dimensional logical array, where the last dimension indexes the bit planes (see `int2bin` for an example). By default, the output of `bin2int` is unsigned `uint32`. To get a signed integer, it must be typecast correctly. A second form for `bin2int` takes a `'signed'` flag

```
y = bin2int(x,'signed')
```

in which case the output is signed.

13.2.2 Example

The following piece of code demonstrates various uses of the `int2bin` function. First the simplest example:

```
--> A = [2;5;6;2]

A =
     2
     5
     6
     2

--> B = int2bin(A,8)

B =
 0 0 0 0 0 0 1 0
 0 0 0 0 0 1 0 1
 0 0 0 0 0 1 1 0
 0 0 0 0 0 0 1 0

--> bin2int(B)

ans =
     2
     5
     6
     2

--> A = [1;2;-5;2]

A =
     1
     2
    -5
     2

--> B = int2bin(A,8)
```

```

B =
 0 0 0 0 0 0 0 1
 0 0 0 0 0 0 1 0
 1 1 1 1 1 0 1 1
 0 0 0 0 0 0 1 0

--> bin2int(B)

ans =
     1
     2
    251
     2

--> int32(bin2int(B))

ans =
     1
     2
    251
     2

```

13.2.3 Tets

13.3 CAST Typecast Variable to Specified Type

13.3.1 Usage

The `cast` function allows you to typecast a variable from one type to another. The syntax for its use is

```
y = cast(x,toclass)
```

where `toclass` is the name of the class to cast `x` to. Note that the typecast must make sense, and that `toclass` must be one of the builtin types. The current list of supported types is

- 'cell' for cell-arrays
- 'struct' for structure-arrays
- 'logical' for logical arrays
- 'uint8' for unsigned 8-bit integers
- 'int8' for signed 8-bit integers
- 'uint16' for unsigned 16-bit integers
- 'int16' for signed 16-bit integers

- 'uint32' for unsigned 32-bit integers
- 'int32' for signed 32-bit integers
- 'uint64' for unsigned 64-bit integers
- 'int64' for signed 64-bit integers
- 'float' for 32-bit floating point numbers
- 'single' is a synonym for 'float'
- 'double' for 64-bit floating point numbers
- 'char' for string arrays

13.3.2 Example

Here is an example of a typecast from a float to an 8-bit integer

```
--> cast(pi,'uint8')
```

```
ans =
     3
```

and here we cast an array of arbitrary integers to a logical array

```
--> cast([1 0 3 0],'logical')
```

```
ans =
     1 0 1 0
```

13.4 CHAR Convert to character array or string

13.4.1 Usage

The `char` function can be used to convert an array into a string. It has several forms. The first form is

```
y = char(x)
```

where `x` is a numeric array containing character codes. FreeMat does not currently support Unicode, so the character codes must be in the range of `[0,255]`. The output is a string of the same size as `x`. A second form is

```
y = char(c)
```

where `c` is a cell array of strings, creates a matrix string where each row contains a string from the corresponding cell array. The third form is

```
y = char(s1, s2, s3, ...)
```

where `si` are a character arrays. The result is a matrix string where each row contains a string from the corresponding argument.

13.4.2 Example

Here is an example of the first technique being used to generate a string containing some ASCII characters

```
--> char([32:64;65:97])
```

```
ans =
 !"#%&'()*+,-./0123456789:;<=>?@
 ABCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`a
```

In the next example, we form a character array from a set of strings in a cell array. Note that the character array is padded with spaces to make the rows all have the same length.

```
--> char({'hello','to','the','world'})
```

```
ans =
hello
to
the
world
```

In the last example, we pass the individual strings as explicit arguments to `char`

```
--> char('hello','to','the','world')
```

```
ans =
hello
to
the
world
```

13.5 COMPLEX Create a Complex Number

13.5.1 Usage

Converts the two real input arguments into the real and imaginary part (respectively) of a complex number. The syntax for its use is

```
y = complex(x,z)
```

where `x` and `z` are `n`-dimensional numerical arrays. The usual rules for binary operators apply (i.e., one of the arguments can be a scalar, if either is of type `single` the output is `single`, etc.).

13.6 DCOMPLEX Convert to Double Precision (deprecated)

13.6.1 Usage

The `dcomplex` function used to convert variables into 64-bit complex data types in prior versions of FreeMat. Starting with FreeMat 4, the type rules are the same as Matlab, hence, there is no

distinction between a 64-bit complex type and a 64-bit real type. Thus, the `dcomplex` function is just a synonym for `double`.

13.7 DEC2BIN Convert Decimal to Binary String

13.7.1 USAGE

Converts an integer to a binary string. The syntax for its use is

```
y = dec2bin(x,n)
```

where `x` is the positive integer, and `n` is the number of bits to use in the representation. Alternately, if you leave `n` unspecified,

```
y = dec2bin(x)
```

the minimum number of bits needed to represent `x` are used. If `x` is a vector, then the resulting `y` is a character matrix.

13.7.2 Example

Here we convert some numbers to bits

```
--> dec2bin(56)

ans =
111000
--> dec2bin(1039456)

ans =
11111101110001100000
--> dec2bin([63,73,32],5)

ans =
11111
01001
00000
```

13.8 DOUBLE Convert to 64-bit Floating Point

13.8.1 Usage

Converts the argument to a 64-bit floating point number. The syntax for its use is

```
y = double(x)
```

where `x` is an `n`-dimensional numerical array. Conversion follows the saturation rules. Note that both `NaN` and `Inf` are both preserved under type conversion.

13.8.2 Example

The following piece of code demonstrates several uses of `double`. First, we convert from an integer (the argument is an integer because no decimal is present):

```
--> double(200)
```

```
ans =  
    200
```

In the next example, a single precision argument is passed in (the presence of the `f` suffix implies single precision).

```
--> double(400.0f)
```

```
ans =  
    400
```

In the next example, a complex argument is passed in.

```
--> double(3.0+4.0*i)
```

```
ans =  
    3.0000 + 4.0000i
```

In the next example, a string argument is passed in. The string argument is converted into an integer array corresponding to the ASCII values of each character.

```
--> double('hello')
```

```
ans =  
    104 101 108 111
```

In the last example, a cell-array is passed in. For cell-arrays and structure arrays, the result is an error.

```
--> double({4})
```

```
Error: Cannot perform type conversions with this type
```

13.9 FLOAT Convert to 32-bit Floating Point

13.9.1 Usage

Converts the argument to a 32-bit floating point number. The syntax for its use is

```
y = float(x)
```

where `x` is an `n`-dimensional numerical array. Conversion follows the saturation rules. Note that both `NaN` and `Inf` are both preserved under type conversion.

13.9.2 Example

The following piece of code demonstrates several uses of `float`. First, we convert from an integer (the argument is an integer because no decimal is present):

```
--> float(200)
```

```
ans =
  200
```

In the next example, a double precision argument is passed in

```
--> float(400.0)
```

```
ans =
  400
```

In the next example, a complex argument is passed in.

```
--> float(3.0+4.0*i)
```

```
ans =
  3.0000 + 4.0000i
```

In the next example, a string argument is passed in. The string argument is converted into an integer array corresponding to the ASCII values of each character.

```
--> float('helo')
```

```
ans =
  104 101 108 111
```

In the last example, a cell-array is passed in. For cell-arrays and structure arrays, the result is an error.

```
--> float({4})
```

```
Error: Cannot perform type conversions with this type
```

13.10 INT16 Convert to Signed 16-bit Integer

13.10.1 Usage

Converts the argument to an signed 16-bit Integer. The syntax for its use is

```
y = int16(x)
```

where `x` is an `n`-dimensional numerical array. Conversion follows the saturation rules (e.g., if `x` is outside the normal range for a signed 16-bit integer of `[-32767,32767]`, it is truncated to that range). Note that both `NaN` and `Inf` both map to 0.

13.10.2 Example

The following piece of code demonstrates several uses of `int16`. First, the routine uses

```
--> int16(100)
```

```
ans =  
  100
```

```
--> int16(-100)
```

```
ans =  
 -100
```

In the next example, an integer outside the range of the type is passed in. The result is truncated to the range of the data type.

```
--> int16(40000)
```

```
ans =  
 32767
```

In the next example, a positive double precision argument is passed in. The result is the signed integer that is closest to the argument.

```
--> int16(pi)
```

```
ans =  
  3
```

In the next example, a complex argument is passed in. The result is the signed complex integer that is closest to the argument.

```
--> int16(5+2*i)
```

```
ans =  
 5.0000 + 2.0000i
```

In the next example, a string argument is passed in. The string argument is converted into an integer array corresponding to the ASCII values of each character.

```
--> int16('hello')
```

```
ans =  
 104 101 108 111
```

In the last example, a cell-array is passed in. For cell-arrays and structure arrays, the result is an error.

```
--> int16({4})
```

```
Error: Cannot perform type conversions with this type
```

13.11 INT2BIN Convert Integer Arrays to Binary

13.11.1 Usage

Computes the binary decomposition of an integer array to the specified number of bits. The general syntax for its use is

```
y = int2bin(x,n)
```

where *x* is a multi-dimensional integer array, and *n* is the number of bits to expand it to. The output array *y* has one extra dimension to it than the input. The bits are expanded along this extra dimension.

13.11.2 Example

The following piece of code demonstrates various uses of the `int2bin` function. First the simplest example:

```
--> A = [2;5;6;2]
```

```
A =
     2
     5
     6
     2
```

```
--> int2bin(A,8)
```

```
ans =
 0 0 0 0 0 0 1 0
 0 0 0 0 0 1 0 1
 0 0 0 0 0 1 1 0
 0 0 0 0 0 0 1 0
```

```
--> A = [1;2;-5;2]
```

```
A =
     1
     2
    -5
     2
```

```
--> int2bin(A,8)
```

```
ans =
 0 0 0 0 0 0 0 1
 0 0 0 0 0 0 1 0
```

```
1 1 1 1 1 0 1 1
0 0 0 0 0 0 1 0
```

13.12 INT32 Convert to Signed 32-bit Integer

13.12.1 Usage

Converts the argument to an signed 32-bit Integer. The syntax for its use is

```
y = int32(x)
```

where x is an n -dimensional numerical array. Conversion follows the saturation rules (e.g., if x is outside the normal range for a signed 32-bit integer of $[-2147483647, 2147483647]$, it is truncated to that range). Note that both NaN and Inf both map to 0.

13.12.2 Example

The following piece of code demonstrates several uses of `int32`. First, the routine uses

```
--> int32(100)
```

```
ans =
    100
```

```
--> int32(-100)
```

```
ans =
   -100
```

In the next example, an integer outside the range of the type is passed in. The result is truncated to the range of the data type.

```
--> int32(40e9)
```

```
ans =
 2147483647
```

In the next example, a positive double precision argument is passed in. The result is the signed integer that is closest to the argument.

```
--> int32(pi)
```

```
ans =
     3
```

In the next example, a complex argument is passed in. The result is the signed complex integer that is closest to the argument.

```
--> int32(5+2*i)

ans =
    5.0000 + 2.0000i
```

In the next example, a string argument is passed in. The string argument is converted into an integer array corresponding to the ASCII values of each character.

```
--> int32('hello')

ans =
    104 101 108 111
```

In the last example, a cell-array is passed in. For cell-arrays and structure arrays, the result is an error.

```
--> int32({4})
Error: Cannot perform type conversions with this type
```

13.13 INT64 Convert to Signed 64-bit Integer

13.13.1 Usage

Converts the argument to an signed 64-bit Integer. The syntax for its use is

```
y = int64(x)
```

where x is an n -dimensional numerical array. Conversion follows the saturation rules (e.g., if x is outside the normal range for a signed 64-bit integer of $[-2^{63}+1, 2^{63}-1]$, it is truncated to that range). Note that both `NaN` and `Inf` both map to 0.

13.13.2 Example

The following piece of code demonstrates several uses of `int64`. First, the routine uses

```
--> int64(100)

ans =
    100

--> int64(-100)

ans =
   -100
```

In the next example, an integer outside the range of the type is passed in. The result is truncated to the range of the data type.

```
--> int64(40e9)
```

```
ans =  
40000000000
```

In the next example, a positive double precision argument is passed in. The result is the signed integer that is closest to the argument.

```
--> int64(pi)
```

```
ans =  
3
```

In the next example, a complex argument is passed in. The result is the complex signed integer that is closest to the argument.

```
--> int64(5+2*i)
```

```
ans =  
5.0000 + 2.0000i
```

In the next example, a string argument is passed in. The string argument is converted into an integer array corresponding to the ASCII values of each character.

```
--> int64('helo')
```

```
ans =  
104 101 108 111
```

In the last example, a cell-array is passed in. For cell-arrays and structure arrays, the result is an error.

```
--> int64({4})
```

```
Error: Cannot perform type conversions with this type
```

13.14 INT8 Convert to Signed 8-bit Integer

13.14.1 Usage

Converts the argument to an signed 8-bit Integer. The syntax for its use is

```
y = int8(x)
```

where x is an n -dimensional numerical array. Conversion follows the saturation rules (e.g., if x is outside the normal range for a signed 8-bit integer of $[-127, 127]$, it is truncated to that range. Note that both `NaN` and `Inf` both map to 0.

13.14.2 Example

The following piece of code demonstrates several uses of `int8`. First, the routine uses

```
--> int8(100)
```

```
ans =  
    100
```

```
--> int8(-100)
```

```
ans =  
   -100
```

In the next example, an integer outside the range of the type is passed in. The result is truncated to the range of the type.

```
--> int8(400)
```

```
ans =  
    127
```

In the next example, a positive double precision argument is passed in. The result is the signed integer that is closest to the argument.

```
--> int8(pi)
```

```
ans =  
     3
```

In the next example, a complex argument is passed in. The result is the signed complex integer that is closest to the argument.

```
--> int8(5+2*i)
```

```
ans =  
 5.0000 + 2.0000i
```

In the next example, a string argument is passed in. The string argument is converted into an integer array corresponding to the ASCII values of each character.

```
--> int8('helo')
```

```
ans =  
 104 101 108 111
```

In the last example, a cell-array is passed in. For cell-arrays and structure arrays, the result is an error.

```
--> int8({4})
```

```
Error: Cannot perform type conversions with this type
```

13.15 LOGICAL Convert to Logical

13.15.1 Usage

Converts the argument to a logical array. The syntax for its use is

```
y = logical(x)
```

where x is an n -dimensional numerical array. Any nonzero element maps to a logical 1.

13.15.2 Example

Here we convert an integer array to `logical`:

```
--> logical([1,2,3,0,0,0,5,2,2])
```

```
ans =  
 1 1 1 0 0 0 1 1 1
```

The same example with double precision values:

```
--> logical([pi,pi,0,e,0,-1])
```

```
ans =  
 1 1 0 1 0 1
```

13.16 SINGLE Convert to 32-bit Floating Point

13.16.1 Usage

A synonym for the `float` function, converts the argument to a 32-bit floating point number. The syntax for its use is

```
y = single(x)
```

where x is an n -dimensional numerical array. Conversion follows the general C rules. Note that both NaN and Inf are both preserved under type conversion.

13.17 STRING Convert Array to String

13.17.1 Usage

Converts the argument array into a string. The syntax for its use is

```
y = string(x)
```

where x is an n -dimensional numerical array.

13.17.2 Example

Here we take an array containing ASCII codes for a string, and convert it into a string.

```
--> a = [104,101,108,108,111]
```

```
a =  
 104 101 108 108 111
```

```
--> string(a)
```

```
ans =  
hello
```

13.18 UINT16 Convert to Unsigned 16-bit Integer

13.18.1 Usage

Converts the argument to an unsigned 16-bit Integer. The syntax for its use is

```
y = uint16(x)
```

where x is an n -dimensional numerical array. Conversion follows saturation rules (e.g., if x is outside the normal range for an unsigned 16-bit integer of $[0,65535]$, it is truncated to that range. Note that both `NaN` and `Inf` both map to 0.

13.18.2 Example

The following piece of code demonstrates several uses of `uint16`.

```
--> uint16(200)
```

```
ans =  
 200
```

In the next example, an integer outside the range of the type is passed in. The result is truncated to the maximum value of the data type.

```
--> uint16(99400)
```

```
ans =  
65535
```

In the next example, a negative integer is passed in. The result is truncated to zero.

```
--> uint16(-100)
```

```
ans =  
 0
```

In the next example, a positive double precision argument is passed in. The result is the unsigned integer that is closest to the argument.

```
--> uint16(pi)
```

```
ans =  
    3
```

In the next example, a complex argument is passed in. The result is the complex unsigned integer that is closest to the argument.

```
--> uint16(5+2*i)
```

```
ans =  
    5.0000 + 2.0000i
```

In the next example, a string argument is passed in. The string argument is converted into an integer array corresponding to the ASCII values of each character.

```
--> uint16('hello')
```

```
ans =  
    104 101 108 111
```

In the last example, a cell-array is passed in. For cell-arrays and structure arrays, the result is an error.

```
--> uint16({4})
```

```
Error: Cannot perform type conversions with this type
```

13.19 **UINT32 Convert to Unsigned 32-bit Integer**

13.19.1 **Usage**

Converts the argument to an unsigned 32-bit Integer. The syntax for its use is

```
y = uint32(x)
```

where *x* is an *n*-dimensional numerical array. Conversion follows saturation rules (e.g., if *x* is outside the normal range for an unsigned 32-bit integer of [0,4294967295], it is truncated to that range. Note that both NaN and Inf both map to 0.

13.19.2 **Example**

The following piece of code demonstrates several uses of `uint32`.

```
--> uint32(200)
```

```
ans =  
    200
```

In the next example, an integer outside the range of the type is passed in. The result is truncated to the maximum value of the data type.

```
--> uint32(40e9)
```

```
ans =  
4294967295
```

In the next example, a negative integer is passed in. The result is truncated to zero.

```
--> uint32(-100)
```

```
ans =  
0
```

In the next example, a positive double precision argument is passed in. The result is the unsigned integer that is closest to the argument.

```
--> uint32(pi)
```

```
ans =  
3
```

In the next example, a complex argument is passed in. The result is the complex unsigned integer that is closest to the argument.

```
--> uint32(5+2*i)
```

```
ans =  
5.0000 + 2.0000i
```

In the next example, a string argument is passed in. The string argument is converted into an integer array corresponding to the ASCII values of each character.

```
--> uint32('hello')
```

```
ans =  
104 101 108 111
```

In the last example, a cell-array is passed in. For cell-arrays and structure arrays, the result is an error.

```
--> uint32({4})
```

```
Error: Cannot perform type conversions with this type
```

13.20 UINT64 Convert to Unsigned 64-bit Integer

13.20.1 Usage

Converts the argument to an unsigned 64-bit Integer. The syntax for its use is

```
y = uint64(x)
```

where x is an n -dimensional numerical array. Conversion follows saturation rules (e.g., if x is outside the normal range for an unsigned 64-bit integer of $[0, 2^{64}-1]$, it is truncated to that range. Note that both `NaN` and `Inf` both map to 0.

13.20.2 Example

The following piece of code demonstrates several uses of `uint64`.

```
--> uint64(200)
```

```
ans =  
200
```

In the next example, an integer outside the range of the type is passed in. The result is truncated to the maximum value of the data type.

```
--> uint64(40e9)
```

```
ans =  
40000000000
```

In the next example, a negative integer is passed in. The result is zero.

```
--> uint64(-100)
```

```
ans =  
0
```

In the next example, a positive double precision argument is passed in. The result is the unsigned integer that is closest to the argument.

```
--> uint64(pi)
```

```
ans =  
3
```

In the next example, a complex argument is passed in. The result is the complex unsigned integer that is closest to the argument.

```
--> uint64(5+2*i)
```

```
ans =  
5.0000 + 2.0000i
```

In the next example, a string argument is passed in. The string argument is converted into an integer array corresponding to the ASCII values of each character.

```
--> uint64('helo')
```

```
ans =  
    104 101 108 111
```

In the last example, a cell-array is passed in. For cell-arrays and structure arrays, the result is an error.

```
--> uint64({4})  
Error: Cannot perform type conversions with this type
```

13.21 UINT8 Convert to Unsigned 8-bit Integer

13.21.1 Usage

Converts the argument to an unsigned 8-bit Integer. The syntax for its use is

```
y = uint8(x)
```

where x is an n -dimensional numerical array. Conversion follows saturation rules (e.g., if x is outside the normal range for an unsigned 8-bit integer of $[0, 255]$, it is truncated to that range. Note that both `NaN` and `Inf` both map to 0.

13.21.2 Example

The following piece of code demonstrates several uses of `uint8`.

```
--> uint8(200)
```

```
ans =  
    200
```

In the next example, an integer outside the range of the type is passed in. The result is truncated to the maximum value of the data type.

```
--> uint8(400)
```

```
ans =  
    255
```

In the next example, a negative integer is passed in. The result is truncated to zero.

```
--> uint8(-100)
```

```
ans =  
     0
```

In the next example, a positive double precision argument is passed in. The result is the unsigned integer that is closest to the argument.

```
--> uint8(pi)
```

```
ans =  
    3
```

In the next example, a complex argument is passed in. The result is complex unsigned integer that is closest to the argument.

```
--> uint8(5+2*i)
```

```
ans =  
    5.0000 + 2.0000i
```

In the next example, a string argument is passed in. The string argument is converted into an integer array corresponding to the ASCII values of each character.

```
--> uint8('helo')
```

```
ans =  
    104 101 108 111
```

In the last example, a cell-array is passed in. For cell-arrays and structure arrays, the result is an error.

```
--> uint8({4})  
Error: Cannot perform type conversions with this type
```


Chapter 14

Array Generation and Manipulations

14.1 ASSIGN Making assignments

14.1.1 Usage

FreeMat assignments take a number of different forms, depending on the type of the variable you want to make an assignment to. For numerical arrays and strings, the form of an assignment is either

```
a(ndx) = val
```

where `ndx` is a set of vector indexing coordinates. This means that the values `ndx` takes reference the elements of `a` in column order. So, if, for example `a` is an $N \times M$ matrix, the first column has vector indices $1, 2, \dots, N$, and the second column has indices $N+1, N+2, \dots, 2N$, and so on. Alternately, you can use multi-dimensional indexing to make an assignment:

```
a(ndx_1,ndx_2,...,ndx_m) = val
```

where each indexing expression `ndx_i` corresponds to the *i*-th dimension of `a`. In both cases, (vector or multi-dimensional indexing), the right hand side `val` must either be a scalar, an empty matrix, or of the same size as the indices. If `val` is an empty matrix, the assignment acts like a delete. Note that the type of `a` may be modified by the assignment. So, for example, assigning a `double` value to an element of a `float` array `a` will cause the array `a` to become `double`.

For cell arrays, the above forms of assignment will still work, but only if `val` is also a cell array. If you want to assign the contents of a cell in a cell array, you must use one of the two following forms, either

```
a{ndx} = val
```

or

```
a{ndx_1,ndx_2,...,ndx_m} = val
```

which will modify the contents of the cell.

14.2 CELL Cell Array of Empty Matrices

14.2.1 Usage

Creates a cell array of empty matrix entres. Two separate syntaxes are possible. The first syntax specifies the array dimensions as a sequence of scalar dimensions:

```
y = cell(d1,d2,...,dn).
```

The resulting array has the given dimensions, and is filled with all zeros. The type of `y` is `cell`, a cell array.

The second syntax specifies the array dimensions as a vector, where each element in the vector specifies a dimension length:

```
y = cell([d1,d2,...,dn]).
```

This syntax is more convenient for calling `zeros` using a variable for the argument. In both cases, specifying only one dimension results in a square matrix output.

14.2.2 Example

The following examples demonstrate generation of some zero arrays using the first form.

```
--> cell(2,3,2)
```

```
ans =
```

```
(:,:,1) =
 [] [] []
 [] [] []
```

```
(:,:,2) =
 [] [] []
 [] [] []
```

```
--> cell(1,3)
```

```
ans =
```

```
 [] [] []
```

The same expressions, using the second form.

```
--> cell([2,6])
```

```
ans =
```

```
 [] [] [] [] [] []
 [] [] [] [] [] []
```

```
--> cell([1,3])
```

```
ans =
 [] [] []
```

14.3 CIRCSHIFT Circularly Shift an Array

14.3.1 USAGE

Applies a circular shift along each dimension of a given array. The syntax for its use is

```
y = circshift(x,shiftvec)
```

where **x** is an n-dimensional array, and **shiftvec** is a vector of integers, each of which specify how much to shift **x** along the corresponding dimension.

14.3.2 Example

The following examples show some uses of **circshift** on N-dimensional arrays.

```
--> x = int32(rand(4,5)*10)
```

```
x =
 5  1  5  5  9
 5  7  3  0  5
 8  4  9  9 10
 2  9  7  2  8
```

```
--> circshift(x,[1,0])
```

```
ans =
 2  9  7  2  8
 5  1  5  5  9
 5  7  3  0  5
 8  4  9  9 10
```

```
--> circshift(x,[0,-1])
```

```
ans =
 1  5  5  9  5
 7  3  0  5  5
 4  9  9 10  8
 9  7  2  8  2
```

```
--> circshift(x,[2,2])
```

```
ans =
 9 10  8  4  9
```

```
2 8 2 9 7
5 9 5 1 5
0 5 5 7 3
```

```
--> x = int32(rand(4,5,3)*10)
```

```
x =
```

```
(:,:,1) =
 4 4 5 10 8
 5 1 7 4 6
 8 6 3 5 8
 1 4 3 3 4
```

```
(:,:,2) =
 9 8 8 4 7
 7 7 4 5 8
 8 5 9 5 6
 8 2 5 1 3
```

```
(:,:,3) =
 6 9 2 1 10
 1 6 7 9 9
 10 10 2 6 1
 7 4 0 6 4
```

```
--> circshift(x,[1,0,0])
```

```
ans =
```

```
(:,:,1) =
 1 4 3 3 4
 4 4 5 10 8
 5 1 7 4 6
 8 6 3 5 8
```

```
(:,:,2) =
 8 2 5 1 3
 9 8 8 4 7
 7 7 4 5 8
 8 5 9 5 6
```

```
(:,:,3) =
 7 4 0 6 4
 6 9 2 1 10
 1 6 7 9 9
```

```
10 10 2 6 1
--> circshift(x,[0,-1,0])
```

```
ans =
```

```
(:,:,1) =
  4  5 10  8  4
  1  7  4  6  5
  6  3  5  8  8
  4  3  3  4  1
```

```
(:,:,2) =
  8  8  4  7  9
  7  4  5  8  7
  5  9  5  6  8
  2  5  1  3  8
```

```
(:,:,3) =
  9  2  1 10  6
  6  7  9  9  1
 10  2  6  1 10
  4  0  6  4  7
```

```
--> circshift(x,[0,0,-1])
```

```
ans =
```

```
(:,:,1) =
  9  8  8  4  7
  7  7  4  5  8
  8  5  9  5  6
  8  2  5  1  3
```

```
(:,:,2) =
  6  9  2  1 10
  1  6  7  9  9
 10 10  2  6  1
  7  4  0  6  4
```

```
(:,:,3) =
  4  4  5 10  8
  5  1  7  4  6
  8  6  3  5  8
  1  4  3  3  4
```

```
--> circshift(x,[2,-3,1])
```

```
ans =
```

```
(:,:,1) =
    6    1   10   10    2
    6    4    7    4    0
    1   10    6    9    2
    9    9    1    6    7
```

```
(:,:,2) =
    5    8    8    6    3
    3    4    1    4    3
   10    8    4    4    5
    4    6    5    1    7
```

```
(:,:,3) =
    5    6    8    5    9
    1    3    8    2    5
    4    7    9    8    8
    5    8    7    7    4
```

14.4 COND Condition Number of a Matrix

14.4.1 Usage

Calculates the condition number of a matrix. To compute the 2-norm condition number of a matrix (ratio of largest to smallest singular values), use the syntax

```
y = cond(A)
```

where A is a matrix. If you want to compute the condition number in a different norm (e.g., the 1-norm), use the second syntax

```
y = cond(A,p)
```

where p is the norm to use when computing the condition number. The following choices of p are supported

- $p = 1$ returns the 1-norm, or the max column sum of A
- $p = 2$ returns the 2-norm (largest singular value of A)
- $p = \text{inf}$ returns the infinity norm, or the max row sum of A
- $p = \text{'fro'}$ returns the Frobenius-norm (vector Euclidean norm, or RMS value)

14.4.2 Function Internals

The condition number is defined as

$$\frac{\|A\|_p}{\|A^{-1}\|_p}$$

This equation is precisely how the condition number is computed for the case $p = 2$. For the $p=2$ case, the condition number can be computed much more efficiently using the ratio of the largest and smallest singular values.

14.4.3 Example

The condition number of this matrix is large

```
--> A = [1,1;0,1e-15]
```

```
A =
    1.0000    1.0000
         0    0.0000
```

```
--> cond(A)
```

```
ans =
    2000000000000000
```

```
--> cond(A,1)
```

```
ans =
    20000000000000002
```

You can also (for the case $p=1$ use `rcond` to calculate an estimate of the condition number

```
--> 1/rcond(A)
```

```
ans =
    2.0000e+15
```

14.5 DET Determinant of a Matrix

14.5.1 Usage

Calculates the determinant of a matrix. Note that for all but very small problems, the determinant is not particularly useful. The condition number `cond` gives a more reasonable estimate as to the suitability of a matrix for inversion than comparing `det(A)` to zero. In any case, the syntax for its use is

```
y = det(A)
```

where A is a square matrix.

14.5.2 Function Internals

The determinant is calculated via the LU decomposition. Note that the determinant of a product of matrices is the product of the determinants. Then, we have that

$$LU = PA$$

where L is lower triangular with 1s on the main diagonal, U is upper triangular, and P is a row-permutation matrix. Taking the determinant of both sides yields

$$|LU| = |L||U| = |U| = |PA| = |P||A|$$

where we have used the fact that the determinant of L is 1. The determinant of P (which is a row exchange matrix) is either 1 or -1.

14.5.3 Example

Here we assemble a random matrix and compute its determinant

```
--> A = rand(5);
--> det(A)
```

```
ans =
   -0.1160
```

Then, we exchange two rows of A to demonstrate how the determinant changes sign (but the magnitude is the same)

```
--> B = A([2,1,3,4,5],:);
--> det(B)
```

```
ans =
    0.1160
```

14.6 DIAG Diagonal Matrix Construction/Extraction

14.6.1 Usage

The `diag` function is used to either construct a diagonal matrix from a vector, or return the diagonal elements of a matrix as a vector. The general syntax for its use is

```
y = diag(x,n)
```

If x is a matrix, then y returns the n -th diagonal. If n is omitted, it is assumed to be zero. Conversely, if x is a vector, then y is a matrix with x set to the n -th diagonal.

14.6.2 Examples

Here is an example of `diag` being used to extract a diagonal from a matrix.

```
--> A = int32(10*rand(4,5))
```

```
A =  
 2 8 3 2 2  
 6 8 6 8 5  
 8 9 3 9 0  
 8 3 7 8 2
```

```
--> diag(A)
```

```
ans =  
 2  
 8  
 3  
 8
```

```
--> diag(A,1)
```

```
ans =  
 8  
 6  
 9  
 2
```

Here is an example of the second form of `diag`, being used to construct a diagonal matrix.

```
--> x = int32(10*rand(1,3))
```

```
x =  
 1 1 10
```

```
--> diag(x)
```

```
ans =  
 1 0 0  
 0 1 0  
 0 0 10
```

```
--> diag(x,-1)
```

```
ans =  
 0 0 0 0  
 1 0 0 0
```

```
0 1 0 0
0 0 10 0
```

14.7 EXPM Matrix Exponential

14.7.1 Usage

Calculates e^A for a square, full rank matrix A . The syntax for its use is

```
y = expm(A)
```

Internally, `expm` is mapped to a simple e^A expression (which in turn uses the eigenvalue expansion of A to compute the exponential).

14.7.2 Example

An example of `expm`

```
--> A = [1 1 0; 0 0 2; 0 0 -1]
```

```
A =
 1  1  0
 0  0  2
 0  0 -1
```

```
--> expm(A)
```

```
ans =
 2.7183    1.7183    1.0862
         0    1.0000    1.2642
         0         0    0.3679
```

14.8 EYE Identity Matrix

14.8.1 USAGE

Creates an identity matrix of the specified size. The syntax for its use is

```
y = eye(n)
```

where n is the size of the identity matrix. The type of the output matrix is `float`.

14.8.2 Example

The following example demonstrates the identity matrix.

```
--> eye(3)
```

```
ans =
  1 0 0
  0 1 0
  0 0 1
```

14.9 FIND Find Non-zero Elements of An Array

14.9.1 Usage

Returns a vector that contains the indices of all non-zero elements in an array. The usage is

```
y = find(x)
```

The indices returned are generalized column indices, meaning that if the array x is of size $[d_1, d_2, \dots, d_n]$, and the element $x(i_1, i_2, \dots, i_n)$ is nonzero, then y will contain the integer

$$i_1 + (i_2 - 1)d_1 + (i_3 - 1)d_1d_2 + \dots$$

The second syntax for the `find` command is

```
[r,c] = find(x)
```

which returns the row and column index of the nonzero entries of x . The third syntax for the `find` command also returns the values

```
[r,c,v] = find(x).
```

Note that if the argument is a row vector, then the returned vectors are also row vectors. This form is particularly useful for converting sparse matrices into IJV form.

The `find` command also supports some additional arguments. Each of the above forms can be combined with an integer indicating how many results to return:

```
y = find(x,k)
```

where k is the maximum number of results to return. This form will return the first k results. You can also specify an optional flag indicating whether to take the first or last k values:

```
y = find(x,k,'first')
y = find(x,k,'last')
```

in the case of the `'last'` argument, the last k values are returned.

14.9.2 Example

Some simple examples of its usage, and some common uses of `find` in FreeMat programs.

```
--> a = [1,2,5,2,4];  
--> find(a==2)
```

```
ans =  
 2 4
```

Here is an example of using `find` to replace elements of `A` that are 0 with the number 5.

```
--> A = [1,0,3;0,2,1;3,0,0]
```

```
A =  
 1 0 3  
 0 2 1  
 3 0 0
```

```
--> n = find(A==0)
```

```
n =  
 2  
 4  
 6  
 9
```

```
--> A(n) = 5
```

```
A =  
 1 5 3  
 5 2 1  
 3 5 5
```

Incidentally, a better way to achieve the same concept is:

```
--> A = [1,0,3;0,2,1;3,0,0]
```

```
A =  
 1 0 3  
 0 2 1  
 3 0 0
```

```
--> A(A==0) = 5
```

```
A =  
 1 5 3  
 5 2 1  
 3 5 5
```

Now, we can also return the indices as row and column indices using the two argument form of `find`:

```
--> A = [1,0,3;0,2,1;3,0,0]
```

```
A =  
 1 0 3  
 0 2 1  
 3 0 0
```

```
--> [r,c] = find(A)
```

```
r =  
 1  
 3  
 2  
 1  
 2
```

```
c =  
 1  
 1  
 2  
 3  
 3
```

Or the three argument form of `find`, which returns the value also:

```
--> [r,c,v] = find(A)
```

```
r =  
 1  
 3  
 2  
 1  
 2
```

```
c =  
 1  
 1  
 2  
 3  
 3
```

```
v =  
 1  
 3  
 2  
 3  
 1
```

14.10 FLIPDIM Reverse a Matrix Along a Given Dimension

14.10.1 USAGE

Reverses an array along the given dimension. The syntax for its use is

```
y = flipdim(x,n)
```

where **x** is matrix, and **n** is the dimension to reverse.

14.10.2 Example

The following examples show some uses of `flipdim` on N-dimensional arrays.

```
--> x = int32(rand(4,5,3)*10)
```

```
x =
```

```
(:,:,1) =
  5  4  5  3  9
  7  7  5  6  7
  1  6  8  4  1
  5  7  1  8  8
```

```
(:,:,2) =
  4  1  2  2  2
  5  2  3  5  4
  0  3  6  8  7
  9  2  3 10  8
```

```
(:,:,3) =
  6  8  9  5  5
  5  4  4  2  9
  8  0  3  0 10
  0  2  6  9  0
```

```
--> flipdim(x,1)
```

```
ans =
```

```
(:,:,1) =
  5  7  1  8  8
  1  6  8  4  1
  7  7  5  6  7
  5  4  5  3  9
```

```
(:,:,2) =
```

```
9 2 3 10 8
0 3 6 8 7
5 2 3 5 4
4 1 2 2 2

(:, :, 3) =
0 2 6 9 0
8 0 3 0 10
5 4 4 2 9
6 8 9 5 5

--> flipdim(x,2)

ans =

(:, :, 1) =
9 3 5 4 5
7 6 5 7 7
1 4 8 6 1
8 8 1 7 5

(:, :, 2) =
2 2 2 1 4
4 5 3 2 5
7 8 6 3 0
8 10 3 2 9

(:, :, 3) =
5 5 9 8 6
9 2 4 4 5
10 0 3 0 8
0 9 6 2 0

--> flipdim(x,3)

ans =

(:, :, 1) =
6 8 9 5 5
5 4 4 2 9
8 0 3 0 10
0 2 6 9 0

(:, :, 2) =
4 1 2 2 2
5 2 3 5 4
```

```

0 3 6 8 7
9 2 3 10 8

```

```

(:, :, 3) =
5 4 5 3 9
7 7 5 6 7
1 6 8 4 1
5 7 1 8 8

```

14.11 FLIPLR Reverse the Columns of a Matrix

14.11.1 USAGE

Reverses the columns of a matrix. The syntax for its use is

```
y = fliplr(x)
```

where `x` is matrix. If `x` is an N-dimensional array then the second dimension is reversed.

14.11.2 Example

The following example shows `fliplr` applied to a 2D matrix.

```
--> x = int32(rand(4)*10)
```

```

x =
6 2 7 1
1 10 10 4
7 7 6 3
8 0 8 5

```

```
--> fliplr(x)
```

```

ans =
1 7 2 6
4 10 10 1
3 6 7 7
5 8 0 8

```

For a 3D array, note how the columns in each slice are flipped.

```
--> x = int32(rand(4,4,3)*10)
```

```

x =
(:, :, 1) =
2 10 7 5

```

```

0 4 5 6
6 2 3 5
2 0 9 6

(:, :, 2) =
3 3 8 5
1 7 2 5
4 6 8 4
7 1 6 1

(:, :, 3) =
3 3 9 2
5 3 7 6
3 3 3 9
6 4 0 8

--> fliplr(x)

ans =

(:, :, 1) =
5 7 10 2
6 5 4 0
5 3 2 6
6 9 0 2

(:, :, 2) =
5 8 3 3
5 2 7 1
4 8 6 4
1 6 1 7

(:, :, 3) =
2 9 3 3
6 7 3 5
9 3 3 3
8 0 4 6

```

14.12 FLIPUD Reverse the Columns of a Matrix

14.12.1 USAGE

Reverses the rows of a matrix. The syntax for its use is

```
y = flipud(x)
```

where **x** is matrix. If **x** is an N-dimensional array then the first dimension is reversed.

14.12.2 Example

The following example shows `flipud` applied to a 2D matrix.

```
--> x = int32(rand(4)*10)
```

```
x =  
 0  7  2  6  
 5  7 10  1  
 2  6  8  0  
 1  0  9  8
```

```
--> flipud(x)
```

```
ans =  
 1  0  9  8  
 2  6  8  0  
 5  7 10  1  
 0  7  2  6
```

For a 3D array, note how the rows in each slice are flipped.

```
--> x = int32(rand(4,4,3)*10)
```

```
x =  
  
(:,:,1) =  
 8  4  6  7  
 6  0  7  6  
 8  9  9  0  
 1  2 10  5
```

```
(:,:,2) =  
 9  6  2  8  
 3 10  1  7  
 6 10  9  8  
 3  2  4  2
```

```
(:,:,3) =  
 3  1  5  0  
 0  6  5  1  
 9  4  9  2  
 6 10  8 10
```

```
--> flipud(x)
```

```
ans =
```

```
(:,:,1) =
  1  2 10  5
  8  9  9  0
  6  0  7  6
  8  4  6  7

(:,:,2) =
  3  2  4  2
  6 10  9  8
  3 10  1  7
  9  6  2  8

(:,:,3) =
  6 10  8 10
  9  4  9  2
  0  6  5  1
  3  1  5  0
```

14.13 IPERMUTE Array Inverse Permutation Function

14.13.1 Usage

The `ipermute` function rearranges the contents of an array according to the inverse of the specified permutation vector. The syntax for its use is

```
y = ipermute(x,p)
```

where `p` is a permutation vector - i.e., a vector containing the integers `1...ndims(x)` each occurring exactly once. The resulting array `y` contains the same data as the array `x`, but ordered according to the inverse of the given permutation. This function and the `permute` function are inverses of each other.

14.13.2 Example

First we create a large multi-dimensional array, then permute it and then inverse permute it, to retrieve the original array:

```
--> A = randn(13,5,7,2);
--> size(A)

ans =
  13  5  7  2

--> B = permute(A,[3,4,2,1]);
--> size(B)
```

```
ans =  
    7  2  5 13  
  
--> C = ipermute(B,[3,4,2,1]);  
--> size(C)  
  
ans =  
    13  5  7  2  
  
--> any(A~=C)  
  
ans =  
  
(:,:,1,1) =  
    0  0  0  0  0  
  
(:,:,2,1) =  
    0  0  0  0  0  
  
(:,:,3,1) =  
    0  0  0  0  0  
  
(:,:,4,1) =  
    0  0  0  0  0  
  
(:,:,5,1) =  
    0  0  0  0  0  
  
(:,:,6,1) =  
    0  0  0  0  0  
  
(:,:,7,1) =  
    0  0  0  0  0  
  
(:,:,1,2) =  
    0  0  0  0  0  
  
(:,:,2,2) =  
    0  0  0  0  0  
  
(:,:,3,2) =  
    0  0  0  0  0  
  
(:,:,4,2) =  
    0  0  0  0  0
```

```
(:,:,5,2) =  
0 0 0 0 0
```

```
(:,:,6,2) =  
0 0 0 0 0
```

```
(:,:,7,2) =  
0 0 0 0 0
```

14.14 ISFLOAT Test for Floating Point Array

14.14.1 Usage

The syntax for `isfloat` is

```
x = isfloat(y)
```

and it returns a logical 1 if the argument is a floating point array (i.e., a `single` or `double`), and a logical 0 otherwise.

14.15 ISINTEGER Test for Integer Array

14.15.1 Usage

The syntax for `isinteger` is

```
x = isinteger(y)
```

and it returns a logical 1 if the argument is an integer. The decision of whether the argument is an integer or not is made based on the class of `y`, not on its value.

14.16 Linspace Linearly Spaced Vector

14.16.1 Usage

Generates a row vector with the specified number of elements, with entries uniformly spaced between two specified endpoints. The syntax for its use is either

```
y = linspace(a,b,count)
```

or, for a default `count = 100`,

```
y = linspace(a,b);
```

14.16.2 Examples

Here is a simple example of using `linspace`

```
--> x = linspace(0,1,5)
```

```
x =  
    0    0.2500    0.5000    0.7500    1.0000
```

14.17 LOGSPACE Logarithmically Spaced Vector

14.17.1 Usage

Generates a row vector with the specified number number of elements, with entries logarithmically spaced between two specified endpoints. The syntax for its use is either

```
y = logspace(a,b,count)
```

or, for a default count = 50,

```
y = logspace(a,b)
```

A third special use is when

```
y = logspace(a,pi)
```

where it generates points between 10^a and π

Contributed by Paulo Xavier Candeias under GPL.

14.17.2 Example

Here is an example of the use of `logspace`

```
--> logspace(1,2,3)
```

```
ans =  
    10.0000    31.6228   100.0000
```

14.18 MESHGRID Generate Grid Mesh For Plots

14.18.1 Usage

The `meshgrid` function generates arrays that can be used for the generation of surface plots. The syntax is one of

```
[X,Y] = meshgrid(x)  
[X,Y] = meshgrid(x,y)  
[X,Y,Z] = meshgrid(x,y,z)
```

where x, y, z are vectors, and X, Y, Z are matrices. In the first case $[X, Y] = \text{meshgrid}(x)$, the rows of X and the columns of Y contain copies of the vector x . In the second case $[X, Y] = \text{meshgrid}(x, y)$, the rows of X contain copies of x , and the columns of Y contain copies of y . In the third case, each input is copied along the row, column or slice direction of the corresponding output variable.

14.18.2 Example

In the first example:

```
--> [X,Y] = meshgrid(-2:.4:2)
```

```
X =
```

```
Columns 1 to 8
```

```
-2.0000  -1.6000  -1.2000  -0.8000  -0.4000   0.0000   0.4000   0.8000
-2.0000  -1.6000  -1.2000  -0.8000  -0.4000   0.0000   0.4000   0.8000
-2.0000  -1.6000  -1.2000  -0.8000  -0.4000   0.0000   0.4000   0.8000
-2.0000  -1.6000  -1.2000  -0.8000  -0.4000   0.0000   0.4000   0.8000
-2.0000  -1.6000  -1.2000  -0.8000  -0.4000   0.0000   0.4000   0.8000
-2.0000  -1.6000  -1.2000  -0.8000  -0.4000   0.0000   0.4000   0.8000
-2.0000  -1.6000  -1.2000  -0.8000  -0.4000   0.0000   0.4000   0.8000
-2.0000  -1.6000  -1.2000  -0.8000  -0.4000   0.0000   0.4000   0.8000
-2.0000  -1.6000  -1.2000  -0.8000  -0.4000   0.0000   0.4000   0.8000
-2.0000  -1.6000  -1.2000  -0.8000  -0.4000   0.0000   0.4000   0.8000
```

```
Columns 9 to 11
```

```
 1.2000   1.6000   2.0000
 1.2000   1.6000   2.0000
 1.2000   1.6000   2.0000
 1.2000   1.6000   2.0000
 1.2000   1.6000   2.0000
 1.2000   1.6000   2.0000
 1.2000   1.6000   2.0000
 1.2000   1.6000   2.0000
 1.2000   1.6000   2.0000
 1.2000   1.6000   2.0000
 1.2000   1.6000   2.0000
```

```
Y =
```

```
Columns 1 to 8
```

```
-2.0000  -2.0000  -2.0000  -2.0000  -2.0000  -2.0000  -2.0000  -2.0000
-1.6000  -1.6000  -1.6000  -1.6000  -1.6000  -1.6000  -1.6000  -1.6000
```

```

-1.2000 -1.2000 -1.2000 -1.2000 -1.2000 -1.2000 -1.2000 -1.2000
-0.8000 -0.8000 -0.8000 -0.8000 -0.8000 -0.8000 -0.8000 -0.8000
-0.4000 -0.4000 -0.4000 -0.4000 -0.4000 -0.4000 -0.4000 -0.4000
 0.0000  0.0000  0.0000  0.0000  0.0000  0.0000  0.0000  0.0000
 0.4000  0.4000  0.4000  0.4000  0.4000  0.4000  0.4000  0.4000
 0.8000  0.8000  0.8000  0.8000  0.8000  0.8000  0.8000  0.8000
 1.2000  1.2000  1.2000  1.2000  1.2000  1.2000  1.2000  1.2000
 1.6000  1.6000  1.6000  1.6000  1.6000  1.6000  1.6000  1.6000
 2.0000  2.0000  2.0000  2.0000  2.0000  2.0000  2.0000  2.0000

```

Columns 9 to 11

```

-2.0000 -2.0000 -2.0000
-1.6000 -1.6000 -1.6000
-1.2000 -1.2000 -1.2000
-0.8000 -0.8000 -0.8000
-0.4000 -0.4000 -0.4000
 0.0000  0.0000  0.0000
 0.4000  0.4000  0.4000
 0.8000  0.8000  0.8000
 1.2000  1.2000  1.2000
 1.6000  1.6000  1.6000
 2.0000  2.0000  2.0000

```

Next, we use different vectors for X and for Y:

```

--> [X,Y] = meshgrid([1,2,3,4],[6,7,8])
X =
 1 2 3 4
 1 2 3 4
 1 2 3 4

Y =
 6 6 6 6
 7 7 7 7
 8 8 8 8

```

14.19 NaN Not-a-Number Constant

14.19.1 Usage

Returns a value that represents “not-a-number” for both 32 and 64-bit floating point values. This constant is meant to represent the result of arithmetic operations whose output cannot be meaningfully defined (like zero divided by zero). There are several forms for the NaN function. The first form returns a double precision NaN.

```
y = nan
```

The next form takes a class name that can be either `'double'`

```
y = nan('double')
```

or `'single'`:

```
y = nan('single')
```

With a single parameter it generates a square matrix of `nans`.

```
y = nan(n)
```

Alternatively, you can specify the dimensions of the array via

```
y = nan(m,n,p,...)
```

or

```
y = nan([m,n,p,...])
```

Finally, you can add a classname of either `'single'` or `'double'`.

14.19.2 Example

The following examples demonstrate a few calculations with the not-a-number constant.

```
--> nan*0
```

```
ans =  
NaN
```

```
--> nan-nan
```

```
ans =  
NaN
```

Note that NaNs are preserved under type conversion to floating point types (i.e., `float`, `double`, `complex` and `dcomplex` types), but not integer types.

```
--> uint32(nan)
```

```
ans =  
0
```

```
--> complex(nan)
```

```
ans =  
NaN
```

14.20 NDGRID Generate N-Dimensional Grid

14.20.1 Usage

Generates N-dimensional grids, each of which is constant in all but one dimension. The syntax for its use is either

```
[y1, y2, ..., ym] = ndgrid(x1, x2, ..., xn)
```

where $m \leq n$ or

```
[y1, y2, ..., ym] = ndgrid(x1)
```

which is equivalent to the first form, with $x1=x2=\dots=xn$. Each output y_i is an n-dimensional array, with values such that

$$y_i(d_1, \dots, d_{i-1}, d_i, d_{i+1}, \dots, d_m) = x_i(d_i)$$

`ndgrid` is useful for evaluating multivariate functionals over a range of arguments. It is a generalization of `meshgrid`, except that `meshgrid` transposes the dimensions corresponding to the first two arguments to better fit graphical applications.

14.20.2 Example

Here is a simple `ndgrid` example

```
--> [a,b] = ndgrid(1:2,3:5)
```

```
a =
  1 1 1
  2 2 2
```

```
b =
  3 4 5
  3 4 5
```

```
--> [a,b,c] = ndgrid(1:2,3:5,0:1)
```

```
a =
(:, :, 1) =
  1 1 1
  2 2 2
```

```
(:, :, 2) =
  1 1 1
  2 2 2
```

```
b =
```

```
(:, :, 1) =
  3 4 5
```

```

3 4 5

(:, :, 2) =
3 4 5
3 4 5

```

```
c =
```

```

(:, :, 1) =
0 0 0
0 0 0

```

```

(:, :, 2) =
1 1 1
1 1 1

```

Here we use the second form

```
--> [a,b,c] = ndgrid(1:3)
```

Warning: Newly defined variable nargin shadows a function of the same name. Use clear nargin to recover.

```
a =
```

```

(:, :, 1) =
1 1 1
2 2 2
3 3 3

```

```

(:, :, 2) =
1 1 1
2 2 2
3 3 3

```

```

(:, :, 3) =
1 1 1
2 2 2
3 3 3

```

```
b =
```

```

(:, :, 1) =
1 2 3
1 2 3
1 2 3

```

```

(:, :, 2) =
1 2 3
1 2 3

```

```

1 2 3

(:, :, 3) =
1 2 3
1 2 3
1 2 3

c =

(:, :, 1) =
1 1 1
1 1 1
1 1 1

(:, :, 2) =
2 2 2
2 2 2
2 2 2

(:, :, 3) =
3 3 3
3 3 3
3 3 3

```

14.21 NONZEROS Retrieve Nonzero Matrix Entries

14.21.1 USAGE

Returns a dense column vector containing the nonzero elements of the argument matrix. The syntax for its use is

```
y = nonzeros(x)
```

where x is the argument array. The argument matrix may be sparse as well as dense.

14.21.2 Example

Here is an example of using `nonzeros` on a sparse matrix.

```
--> a = rand(8); a(a>0.2) = 0;
--> A = sparse(a)
```

```
A =
4 1 0.150467
6 1 0.119605
7 1 0.0450479
3 3 0.0592569
```

```
4 3 0.185718
3 4 0.064163
7 4 0.196438
8 4 0.103669
4 5 0.160639
2 6 0.0545093
5 6 0.0226881
8 6 0.0202693
1 7 0.192362
6 8 0.101635
7 8 0.147931
--> nonzeros(A)
```

```
ans =
    0.1505
    0.1196
    0.0450
    0.0593
    0.1857
    0.0642
    0.1964
    0.1037
    0.1606
    0.0545
    0.0227
    0.0203
    0.1924
    0.1016
    0.1479
```

14.22 NORM Norm Calculation

14.22.1 Usage

Calculates the norm of a matrix. There are two ways to use the `norm` function. The general syntax is

```
y = norm(A,p)
```

where `A` is the matrix to analyze, and `p` is the type norm to compute. The following choices of `p` are supported

- `p = 1` returns the 1-norm, or the max column sum of `A`
- `p = 2` returns the 2-norm (largest singular value of `A`)
- `p = inf` returns the infinity norm, or the max row sum of `A`

- `p = 'fro'` returns the Frobenius-norm (vector Euclidean norm, or RMS value)

For a vector, the regular norm calculations are performed:

- `1 <= p < inf` returns `sum(abs(A).^p)^(1/p)`
- `p` unspecified returns `norm(A,2)`
- `p = inf` returns `max(abs(A))`
- `p = -inf` returns `min(abs(A))`

14.22.2 Examples

Here are the various norms calculated for a sample matrix

```
--> A = float(rand(3,4))

A =
    0.0063    0.2224    0.7574    0.9848
    0.7319    0.1965    0.7191    0.7010
    0.8319    0.6392    0.8905    0.9280

--> norm(A,1)

ans =
    2.6138

--> norm(A,2)

ans =
    2.3403

--> norm(A,inf)

ans =
    3.2896

--> norm(A,'fro')

ans =
    2.4353
```

Next, we calculate some vector norms.

```
--> A = float(rand(4,1))

A =
    0.5011
```

```
0.3269
0.8192
0.7321

--> norm(A,1)

ans =
    2.3792

--> norm(A,2)

ans =
    1.2510

--> norm(A,7)

ans =
    0.8671

--> norm(A,inf)

ans =
    0.8192

--> norm(A,-inf)

ans =
    0.3269
```

14.23 NUM2STR Convert Numbers To Strings

14.23.1 Usage

Converts an array into its string representation. The general syntax for this function is

```
s = num2str(X)
```

where **s** is a string (or string matrix) and **X** is an array. By default, the `num2str` function uses 4 digits of precision and an exponent if required. If you want more digits of precision, you can specify the precision via the form

```
s = num2str(X, precision)
```

where **precision** is the number of digits to include in the string representation. For more control over the format of the output, you can also specify a format specifier (see `printf` for more details).

```
s = num2str(X, format)
```

where **format** is the specifier string.

14.24 ONES Array of Ones

14.24.1 Usage

Creates an array of ones of the specified size. Two separate syntaxes are possible. The first syntax specifies the array dimensions as a sequence of scalar dimensions:

```
y = ones(d1,d2,...,dn).
```

The resulting array has the given dimensions, and is filled with all ones. The type of `y` is `float`, a 32-bit floating point array. To get arrays of other types, use the typecast functions (e.g., `uint8`, `int8`, etc.).

The second syntax specifies the array dimensions as a vector, where each element in the vector specifies a dimension length:

```
y = ones([d1,d2,...,dn]).
```

This syntax is more convenient for calling `ones` using a variable for the argument. In both cases, specifying only one dimension results in a square matrix output.

14.24.2 Example

The following examples demonstrate generation of some arrays of ones using the first form.

```
--> ones(2,3,2)
```

```
ans =
```

```
(:,:,1) =
 1 1 1
 1 1 1
```

```
(:,:,2) =
 1 1 1
 1 1 1
```

```
--> ones(1,3)
```

```
ans =
```

```
1 1 1
```

The same expressions, using the second form.

```
--> ones([2,6])
```

```
ans =
```

```
1 1 1 1 1 1
 1 1 1 1 1 1
```

```
--> ones([1,3])
```

```
ans =  
 1 1 1
```

Finally, an example of using the type casting function `uint16` to generate an array of 16-bit unsigned integers with a value of 1.

```
--> uint16(ones(3))
```

```
ans =  
 1 1 1  
 1 1 1  
 1 1 1
```

14.25 PERMUTE Array Permutation Function

14.25.1 Usage

The `permute` function rearranges the contents of an array according to the specified permutation vector. The syntax for its use is

```
y = permute(x,p)
```

where `p` is a permutation vector - i.e., a vector containing the integers `1...ndims(x)` each occurring exactly once. The resulting array `y` contains the same data as the array `x`, but ordered according to the permutation. This function is a generalization of the matrix transpose operation.

14.25.2 Example

Here we use `permute` to transpose a simple matrix (note that `permute` also works for sparse matrices):

```
--> A = [1,2;4,5]
```

```
A =  
 1 2  
 4 5
```

```
--> permute(A,[2,1])
```

```
ans =  
 1 4  
 2 5
```

```
--> A'
```

```
ans =
```

```
1 4
2 5
```

Now we permute a larger n-dimensional array:

```
--> A = randn(13,5,7,2);
--> size(A)

ans =
    13    5    7    2

--> B = permute(A,[3,4,2,1]);
--> size(B)

ans =
     7     2     5    13
```

14.26 PINV Moore-Penrose Pseudoinverse

14.26.1 Usage

Calculates the Moore-Penrose pseudoinverse of a matrix. The general syntax for its use is

```
y = pinv(A,tol)
```

or for a default specification of the tolerance `tol`,

```
y = pinv(A)
```

For any $m \times n$ matrix A , the Moore-Penrose pseudoinverse is the unique $n \times m$ matrix B that satisfies the following four conditions

- $A B A = A$
- $B A B = B$
- $(A B)' = A B$
- $(B A)' = B A$

Also, it is true that $B y$ is the minimum norm, least squares solution to $A x = y$. The Moore-Penrose pseudoinverse is computed from the singular value decomposition of A , with singular values smaller than `tol` being treated as zeros. If `tol` is not specified then it is chosen as

```
tol = max(size(A)) * norm(A) * eps(A).
```

14.26.2 Function Internals

The calculation of the MP pseudo-inverse is almost trivial once the svd of the matrix is available. First, for a real, diagonal matrix with positive entries, the pseudo-inverse is simply

$$(\Sigma^+)_{ii} = \begin{cases} 1/\sigma_{ii} & \sigma_{ii} > 0 \\ 0 & \text{else} \end{cases}$$

One can quickly verify that this choice of matrix satisfies the four properties of the pseudoinverse. Then, the pseudoinverse of a general matrix $A = U S V'$ is defined as

$$A^+ = V S^+ U'$$

and again, using the facts that $U' U = I$ and $V V' = I$, one can quickly verify that this choice of pseudoinverse satisfies the four defining properties of the MP pseudoinverse. Note that in practice, the diagonal pseudoinverse $S^{\{+\}}$ is computed with a threshold (the `tol` argument to `pinv`) so that singular values smaller than `tol` are treated like zeros.

14.26.3 Examples

Consider a simple 1 x 2 matrix example, and note the various Moore-Penrose conditions:

```
--> A = float(rand(1,2))
```

```
A =
    0.9526    0.4847
```

```
--> B = pinv(A)
```

```
B =
    0.8338
    0.4243
```

```
--> A*B*A
```

```
ans =
    0.9526    0.4847
```

```
--> B*A*B
```

```
ans =
    0.8338
    0.4243
```

```
--> A*B
```

```
ans =
```

```
1.0000
```

```
--> B*A
```

```
ans =
    0.7943    0.4042
    0.4042    0.2057
```

To demonstrate that `pinv` returns the least squares solution, consider the following very simple case

```
--> A = float([1;1;1;1])
```

```
A =
    1
    1
    1
    1
```

The least squares solution to $A x = b$ is just $x = \text{mean}(b)$, and computing the `pinv` of A demonstrates this

```
--> pinv(A)
```

```
ans =
    0.2500    0.2500    0.2500    0.2500
```

Similarly, we can demonstrate the minimum norm solution with the following simple case

```
--> A = float([1,1])
```

```
A =
    1 1
```

The solutions of $A x = 5$ are those x_1 and x_2 such that $x_1 + x_2 = 5$. The norm of x is $x_1^2 + x_2^2$, which is $x_1^2 + (5-x_1)^2$, which is minimized for $x_1 = x_2 = 2.5$:

```
--> pinv(A) * 5.0
```

```
ans =
    2.5000
    2.5000
```

14.27 RANK Calculate the Rank of a Matrix

14.27.1 Usage

Returns the rank of a matrix. There are two ways to use the `rank` function is

```
y = rank(A,tol)
```

where `tol` is the tolerance to use when computing the rank. The second form is

```
y = rank(A)
```

in which case the tolerance `tol` is chosen as

```
tol = max(size(A))*max(s)*eps,
```

where `s` is the vector of singular values of `A`. The rank is computed using the singular value decomposition `svd`.

14.27.2 Examples

Some examples of matrix rank calculations

```
--> rank([1,3,2;4,5,6])
```

```
ans =  
2
```

```
--> rank([1,2,3;2,4,6])
```

```
ans =  
1
```

Here we construct an ill-conditioned matrix, and show the use of the `tol` argument.

```
--> A = [1,0;0,eps/2]
```

```
A =  
1.0000      0  
0      0.0000
```

```
--> rank(A)
```

```
ans =  
1
```

```
--> rank(A,eps/8)
```

```
ans =  
2
```

14.28 RCOND Reciprocal Condition Number Estimate

14.28.1 Usage

The `rcond` function is a FreeMat wrapper around LAPACK's function `XGECON`, which estimates the 1-norm condition number (reciprocal). For the details of the algorithm see the LAPACK documentation. The syntax for its use is

```
x = rcond(A)
```

where **A** is a matrix.

14.28.2 Example

Here is the reciprocal condition number for a random square matrix

```
--> A = rand(30);
--> rcond(A)
```

```
ans =
 3.5040e-04
```

And here we calculate the same value using the definition of (reciprocal) condition number

```
--> 1/(norm(A,1)*norm(inv(A),1))
```

```
ans =
 3.5040e-04
```

Note that the values are very similar. LAPACK's `rcond` function is far more efficient than the explicit calculation (which is also used by the `cond` function).

14.29 REPMAT Array Replication Function

14.29.1 Usage

The `repmat` function replicates an array the specified number of times. The source and destination arrays may be multidimensional. There are three distinct syntaxes for the `repmat` function. The first form:

```
y = repmat(x,n)
```

replicates the array **x** on an **n-times-n** tiling, to create a matrix **y** that has **n** times as many rows and columns as **x**. The output **y** will match **x** in all remaining dimensions. The second form is

```
y = repmat(x,m,n)
```

And creates a tiling of **x** with **m** copies of **x** in the row direction, and **n** copies of **x** in the column direction. The final form is the most general

```
y = repmat(x,[m n p...])
```

where the supplied vector indicates the replication factor in each dimension.

14.29.2 Example

Here is an example of using the `repmat` function to replicate a row 5 times. Note that the same effect can be accomplished (although somewhat less efficiently) by a multiplication.

```
--> x = [1 2 3 4]

x =
    1 2 3 4

--> y = repmat(x,[5,1])

y =
    1 2 3 4
    1 2 3 4
    1 2 3 4
    1 2 3 4
    1 2 3 4
```

The `repmat` function can also be used to create a matrix of scalars or to provide replication in arbitrary dimensions. Here we use it to replicate a 2D matrix into a 3D volume.

```
--> x = [1 2;3 4]

x =
    1 2
    3 4

--> y = repmat(x,[1,1,3])

y =

(:, :, 1) =
    1 2
    3 4

(:, :, 2) =
    1 2
    3 4

(:, :, 3) =
    1 2
    3 4
```

14.30 RESHAPE Reshape An Array

14.30.1 Usage

Reshapes an array from one size to another. Two separate syntaxes are possible. The first syntax specifies the array dimensions as a sequence of scalar dimensions:

```
y = reshape(x,d1,d2,...,dn).
```

The resulting array has the given dimensions, and is filled with the contents of **x**. The type of **y** is the same as **x**. The second syntax specifies the array dimensions as a vector, where each element in the vector specifies a dimension length:

```
y = reshape(x,[d1,d2,...,dn]).
```

This syntax is more convenient for calling **reshape** using a variable for the argument. The **reshape** function requires that the length of **x** equal the product of the **di** values. Note that arrays are stored in column format, which means that elements in **x** are transferred to the new array **y** starting with the first column first element, then proceeding to the last element of the first column, then the first element of the second column, etc.

14.30.2 Example

Here are several examples of the use of **reshape** applied to various arrays. The first example reshapes a row vector into a matrix.

```
--> a = uint8(1:6)
```

```
a =
  1 2 3 4 5 6
```

```
--> reshape(a,2,3)
```

```
ans =
  1 3 5
  2 4 6
```

The second example reshapes a longer row vector into a volume with two planes.

```
--> a = uint8(1:12)
```

```
a =
  1  2  3  4  5  6  7  8  9 10 11 12
```

```
--> reshape(a,[2,3,2])
```

```
ans =
```

```
(:,: ,1) =
```

```

1 3 5
2 4 6

(:, :, 2) =
7 9 11
8 10 12

```

The third example reshapes a matrix into another matrix.

```
--> a = [1,6,7;3,4,2]
```

```
a =
1 6 7
3 4 2
```

```
--> reshape(a,3,2)
```

```
ans =
1 4
3 7
6 2
```

14.31 RREF Reduced Row Echelon Form of a Matrix

14.31.1 Usage

Calculates the reduced row echelon form of a matrix using Gauss Jordan elimination with partial pivoting. The generic syntax for `rref` is

```
R = rref(A)
```

A default tolerance of `max(size(A))*eps*norm(A,inf)` is used to detect negligible column elements. The second form of `rref` returns a vector `k` as well as `R`

```
[R,k] = rref(A)
```

where `k` is a vector that corresponds to the columns of `A` used as pivot columns. If you want to control the tolerance used to identify negligible elements, you can use the form

```
[R,k] = rref(A, tolerance)
```

This implementation of `rref` is based on the one from the `matcompat` lib for octave. It is copyright Paul Kienzle, and distributed under the GNU GPL.

14.32 SHIFTDIM Shift Array Dimensions Function

14.32.1 Usage

The `shiftdim` function is used to shift the dimensions of an array. The general syntax for the `shiftdim` function is

```
y = shiftdim(x,n)
```

where x is a multidimensional array, and n is an integer. If n is a positive integer, then `shiftdim` circularly shifts the dimensions of x to the left, wrapping the dimensions around as necessary. If n is a negative integer, then `shiftdim` shifts the dimensions of x to the right, introducing singleton dimensions as necessary. In its second form:

```
[y,n] = shiftdim(x)
```

the `shiftdim` function will shift away (to the left) the leading singleton dimensions of x until the leading dimension is not a singleton dimension (recall that a singleton dimension p is one for which `size(x,p) == 1`).

14.32.2 Example

Here are some simple examples of using `shiftdim` to remove the singleton dimensions of an array, and then restore them:

```
--> x = uint8(10*randn(1,1,1,3,2));
--> [y,n] = shiftdim(x);
--> n
```

```
ans =
     3
```

```
--> size(y)
```

```
ans =
     3 2
```

```
--> c = shiftdim(y,-n);
--> size(c)
```

```
ans =
     1 1 1 3 2
```

```
--> any(c~=x)
```

```
ans =
```

```
(:,:,1,1,1) =
     0
```

```
(:,:,1,1,2) =
     0
```

Note that these operations (where shifting involves only singleton dimensions) do not actually cause data to be resorted, only the size of the arrays change. This is not true for the following example, which triggers a call to `permute`:

```
--> z = shiftdim(x,4);
```

Note that `z` is now the transpose of `x`

```
--> squeeze(x)
```

```
ans =
```

```
0 0
```

```
0 3
```

```
0 0
```

```
--> squeeze(z)
```

```
ans =
```

```
0 0 0
```

```
0 3 0
```

14.33 SORT Sort

14.33.1 Usage

Sorts an `n`-dimensional array along the specified dimension. The first form sorts the array along the first non-singular dimension.

```
B = sort(A)
```

Alternately, the dimension along which to sort can be explicitly specified

```
B = sort(A,dim)
```

FreeMat does not support vector arguments for `dim` - if you need `A` to be sorted along multiple dimensions (i.e., row first, then columns), make multiple calls to `sort`. Also, the direction of the sort can be specified using the `mode` argument

```
B = sort(A,dim,mode)
```

where `mode = 'ascend'` means to sort the data in ascending order (the default), and `mode = 'descend'` means to sort the data into descending order.

When two outputs are requested from `sort`, the indexes are also returned. Thus, for

```
[B,IX] = sort(A)
```

```
[B,IX] = sort(A,dim)
```

```
[B,IX] = sort(A,dim,mode)
```

an array `IX` of the same size as `A`, where `IX` records the indices of `A` (along the sorting dimension) corresponding to the output array `B`.

Two additional issues worth noting. First, a cell array can be sorted if each cell contains a `string`, in which case the strings are sorted by lexical order. The second issue is that FreeMat uses the same method as MATLAB to sort complex numbers. In particular, a complex number `a` is less than another complex number `b` if `abs(a) < abs(b)`. If the magnitudes are the same then we test the angle of `a`, i.e. `angle(a) < angle(b)`, where `angle(a)` is the phase of `a` between `-pi,pi`.

14.33.2 Example

Here are some examples of sorting on numerical arrays.

```
--> A = int32(10*rand(4,3))
```

```
A =  
 8  9  3  
 7  2  1  
 3  0  3  
 3 10  5
```

```
--> [B,IX] = sort(A)
```

```
B =  
 3  0  1  
 3  2  3  
 7  9  3  
 8 10  5
```

```
IX =  
 3 3 2  
 4 2 1  
 2 1 3  
 1 4 4
```

```
--> [B,IX] = sort(A,2)
```

```
B =  
 3  8  9  
 1  2  7  
 0  3  3  
 3  5 10
```

```
IX =  
 3 1 2  
 3 2 1  
 2 1 3  
 1 3 2
```

```
--> [B,IX] = sort(A,1,'descend')
```

```
B =  
 8 10  5  
 7  9  3  
 3  2  3  
 3  0  1
```

```
IX =  
 1 4 4
```

```
2 1 1
3 2 3
4 3 2
```

Here we sort a cell array of strings.

```
--> a = {'hello', 'abba', 'goodbye', 'jockey', 'cake'}
```

```
a =
 [hello] [abba] [goodbye] [jockey] [cake]
```

```
--> b = sort(a)
```

```
b =
 [abba] [cake] [goodbye] [hello] [jockey]
```

14.34 SQUEEZE Remove Singleton Dimensions of an Array

14.34.1 Usage

This function removes the singleton dimensions of an array. The syntax for its use is

```
y = squeeze(x)
```

where x is a multidimensional array. Generally speaking, if x is of size $d_1 \times 1 \times d_2 \times \dots$, then $\text{squeeze}(x)$ is of size $d_1 \times d_2 \times \dots$, i.e., each dimension of x that was singular (size 1) is squeezed out.

14.34.2 Example

Here is a many dimensioned, ungainly array, both before and after squeezing;

```
--> x = zeros(1,4,3,1,1,2);
--> size(x)
```

```
ans =
 1 4 3 1 1 2
```

```
--> y = squeeze(x);
--> size(y)
```

```
ans =
 4 3 2
```

14.35 SUBSREF Array Dereferencing

14.35.1 Usage

This function can be used to index into basic array types (or structures). It provides a functional interface to execute complex indexing expressions such as `a.b(3){5}` at run time (i.e. while executing a script or a function) without resorting to using `eval`. Note that this function should be overloaded for use with user defined classes, and that it cannot be overloaded for base types. The basic syntax of the function is:

```
b = subsref(a,s)
```

where `s` is a structure array with two fields. The first field is

- `type` is a string containing either `'()` or `'{'` or `','` depending on the form of the call.
- `subs` is a cell array or string containing the subscript information.

When multiple indexing expressions are combined together such as `b = a(5).foo{:}`, the `s` array should contain the following entries

```
s(1).type = '()'   s(1).subs = {5}
s(2).type = ','   s(2).subs = 'foo'
s(3).type = '{}'  s(3).subs = ':'
```

14.36 TRANSPOSE Matrix Transpose

14.36.1 Usage

Performs a (nonconjugate) transpose of a matrix. The syntax for its use is

```
y = transpose(x)
```

and is a synonym for `y = x.'`

14.36.2 Example

Here is an example of the transpose of a complex matrix. Note that the entries are not conjugated.

```
--> A = [1+i,2+i;3-2*i,4+2*i]
```

```
A =
 1.0000 + 1.0000i   2.0000 + 1.0000i
 3.0000 - 2.0000i   4.0000 + 2.0000i
```

```
--> transpose(A)
```

```
ans =
 1.0000 + 1.0000i   3.0000 - 2.0000i
 2.0000 + 1.0000i   4.0000 + 2.0000i
```

14.37 UNIQUE Unique

14.37.1 Usage

Returns a vector containing the unique elements of an array. The first form is simply

```
y = unique(x)
```

where x is either a numerical array or a cell-array of strings. The result is sorted in increasing order. You can also retrieve two sets of index vectors

```
[y, m, n] = unique(x)
```

such that $y = x(m)$ and $x = y(n)$. If the argument x is a matrix, you can also indicate that FreeMat should look for unique rows in the matrix via

```
y = unique(x,'rows')
```

and

```
[y, m, n] = unique(x,'rows')
```

14.37.2 Example

Here is an example in row mode

```
--> A = randi(1,3*ones(15,3))
```

```
A =
```

```
1 1 1
3 2 2
1 3 3
2 2 1
1 3 1
2 3 1
2 2 3
1 2 2
3 3 3
2 1 1
2 1 1
2 3 2
1 1 2
2 3 1
3 3 2
```

```
--> unique(A,'rows')
```

```
ans =
```

```
1 1 1
```

```
1 1 2
1 2 2
1 3 1
1 3 3
2 1 1
2 2 1
2 2 3
2 3 1
2 3 2
3 2 2
3 3 2
3 3 3
```

```
--> [b,m,n] = unique(A,'rows');
--> b
```

```
ans =
1 1 1
1 1 2
1 2 2
1 3 1
1 3 3
2 1 1
2 2 1
2 2 3
2 3 1
2 3 2
3 2 2
3 3 2
3 3 3
```

```
--> A(m,:)
```

```
ans =
1 1 1
1 1 2
1 2 2
1 3 1
1 3 3
2 1 1
2 2 1
2 2 3
2 3 1
2 3 2
3 2 2
3 3 2
```

```
3 3 3
--> b(n,:)
```

```
ans =
1 1 1
3 2 2
1 3 3
2 2 1
1 3 1
2 3 1
2 2 3
1 2 2
3 3 3
2 1 1
2 1 1
2 3 2
1 1 2
2 3 1
3 3 2
```

Here is an example in vector mode

```
--> A = randi(1,5*ones(10,1))
```

```
A =
3
5
5
1
3
4
1
5
4
1
```

```
--> unique(A)
```

```
ans =
1
3
4
5
```

```
--> [b,m,n] = unique(A,'rows');
--> b
```

```
ans =  
 1  
 3  
 4  
 5
```

```
--> A(m)
```

```
ans =  
 1  
 3  
 4  
 5
```

```
--> b(n)
```

```
ans =  
 3  
 5  
 5  
 1  
 3  
 4  
 1  
 5  
 4  
 1
```

For cell arrays of strings.

```
--> A = {'hi', 'bye', 'good', 'tell', 'hi', 'bye'}
```

```
A =  
 [hi] [bye] [good] [tell] [hi] [bye]
```

```
--> unique(A)
```

```
ans =  
 [bye] [good] [hi] [tell]
```

14.38 XNRM2 BLAS Norm Calculation

14.38.1 Usage

Calculates the 2-norm of a vector. The syntax for its use is

```
y = xnorm2(A)
```

where **A** is the n-dimensional array to analyze. This form uses the underlying BLAS implementation to compute the 2-norm.

14.39 ZEROS Array of Zeros

14.39.1 Usage

Creates an array of zeros of the specified size. Two separate syntaxes are possible. The first syntax specifies the array dimensions as a sequence of scalar dimensions:

```
y = zeros(d1,d2,...,dn).
```

The resulting array has the given dimensions, and is filled with all zeros. The type of **y** is **double**, a 64-bit floating point array. To get arrays of other types, use the typecast functions (e.g., **uint8**, **int8**, etc.). An alternative syntax is to use the following notation:

```
y = zeros(d1,d2,...,dn,classname)
```

where **classname** is one of 'double', 'single', 'int8', 'uint8', 'int16', 'uint16', 'int32', 'uint32', 'int64', 'uint64', 'float', 'logical'.

The second syntax specifies the array dimensions as a vector, where each element in the vector specifies a dimension length:

```
y = zeros([d1,d2,...,dn]),
```

or

```
y = zeros([d1,d2,...,dn],classname).
```

This syntax is more convenient for calling **zeros** using a variable for the argument. In both cases, specifying only one dimension results in a square matrix output.

14.39.2 Example

The following examples demonstrate generation of some zero arrays using the first form.

```
--> zeros(2,3,2)
```

```
ans =
```

```
(:,:,1) =
 0 0 0
 0 0 0
```

```
(:,:,2) =
 0 0 0
 0 0 0
```

```
--> zeros(1,3)
```

```
ans =  
0 0 0
```

The same expressions, using the second form.

```
--> zeros([2,6])
```

```
ans =  
0 0 0 0 0 0  
0 0 0 0 0 0
```

```
--> zeros([1,3])
```

```
ans =  
0 0 0
```

Finally, an example of using the type casting function `uint16` to generate an array of 16-bit unsigned integers with zero values.

```
--> uint16(zeros(3))
```

```
ans =  
0 0 0  
0 0 0  
0 0 0
```

Here we use the second syntax where the class of the output is specified explicitly

```
--> zeros(3,'int16')
```

```
ans =  
0 0 0  
0 0 0  
0 0 0
```

Chapter 15

Random Number Generation

15.1 RAND Uniform Random Number Generator

15.1.1 Usage

Creates an array of pseudo-random numbers of the specified size. The numbers are uniformly distributed on $[0,1)$. Two separate syntaxes are possible. The first syntax specifies the array dimensions as a sequence of scalar dimensions:

```
y = rand(d1,d2,...,dn).
```

The resulting array has the given dimensions, and is filled with random numbers. The type of `y` is `double`, a 64-bit floating point array. To get arrays of other types, use the typecast functions.

The second syntax specifies the array dimensions as a vector, where each element in the vector specifies a dimension length:

```
y = rand([d1,d2,...,dn]).
```

This syntax is more convenient for calling `rand` using a variable for the argument.

Finally, `rand` supports two additional forms that allow you to manipulate the state of the random number generator. The first retrieves the state

```
y = rand('state')
```

which is a 625 length integer vector. The second form sets the state

```
rand('state',y)
```

or alternately, you can reset the random number generator with

```
rand('state',0)
```

15.1.2 Example

The following example demonstrates an example of using the first form of the `rand` function.

```
--> rand(2,2,2)

ans =

(:, :, 1) =
    0.3478    0.5313
    0.0276    0.9958

(:, :, 2) =
    0.2079    0.7597
    0.4921    0.3365
```

The second example demonstrates the second form of the `rand` function.

```
--> rand([2,2,2])

ans =

(:, :, 1) =
    0.8670    0.2174
    0.2714    0.6897

(:, :, 2) =
    0.2305    0.3898
    0.1721    0.9545
```

The third example computes the mean and variance of a large number of uniform random numbers. Recall that the mean should be $1/2$, and the variance should be $1/12 \sim 0.083$.

```
--> x = rand(1,10000);
--> mean(x)

ans =
    0.5023

--> var(x)

ans =
    0.0840
```

Now, we use the state manipulation functions of `rand` to exactly reproduce a random sequence. Note that unlike using `seed`, we can exactly control where the random number generator starts by saving the state.

```
--> rand('state',0)    % restores us to startup conditions
--> a = rand(1,3)      % random sequence 1
```

```
a =
    0.3759    0.0183    0.9134
```

```
--> b = rand('state'); % capture the state vector
--> c = rand(1,3)      % random sequence 2
```

```
c =
    0.3580    0.7604    0.8077
```

```
--> rand('state',b);  % restart the random generator so...
--> c = rand(1,3)      % we get random sequence 2 again
```

```
c =
    0.3580    0.7604    0.8077
```

15.2 RANDBETA Beta Deviate Random Number Generator

15.2.1 Usage

Creates an array of beta random deviates based on the supplied two parameters. The general syntax for `randbeta` is

```
y = randbeta(alpha, beta)
```

where `alpha` and `beta` are the two parameters of the random deviate. There are three forms for calling `randbeta`. The first uses two vectors `alpha` and `beta` of the same size, in which case the output `y` is the same size as both inputs, and each deviate uses the corresponding values of `alpha` and `beta` from the arguments. In the other forms, either `alpha` or `beta` are scalars.

15.2.2 Function Internals

The probability density function (PDF) of a beta random variable is

$$f(x) = x^{(a-1)} * (1-x)^{(b-1)} / B(a,b)$$

for `x` between 0 and 1. The function `B(a,b)` is defined so that the integral of `f(x)` is 1.


15.2.3 Example

Here is a plot of the PDF of a beta random variable with `a=3`, `b=7`.

```
--> a = 3; b = 7;
--> x = (0:100)/100; t = x.^(a-1).*(1-x).^(b-1);
```

```
--> t = t/(sum(t)*.01);
--> plot(x,t);
```

which is plotted as



If we generate a few random deviates with these values, we see they are distributed around the peak of roughly 0.25.

```
--> randbeta(3*ones(1,5),7*ones(1,5))

ans =
    0.2777    0.0642    0.3305    0.5259    0.4003
```

15.3 RANDBIN Generate Binomial Random Variables

15.3.1 Usage

Generates random variables with a binomial distribution. The general syntax for its use is

```
y = randbin(N,p)
```

where N is a vector representing the number of Bernoulli trials, and p is the success probability associated with each trial.

15.3.2 Function Internals

A Binomial random variable describes the number of successful outcomes from N Bernoulli trials, with the probability of success in each trial being p . The probability distribution is

$$P(n) = \frac{N!}{n!(N-n)!} p^n (1-p)^{N-n}$$

15.3.3 Example

Here we generate 10 binomial random variables, corresponding to $N=100$ trials, each with probability $p=0.1$, using both `randbin` and then again using `rand` (to simulate the trials):

```
--> randbin(100,.1*ones(1,10))
```

```
ans =
    13     6     8     9    11     9     6     9     7    10
```

```
--> sum(rand(100,10)<0.1)
```

```
ans =
     8    12    10     7    12     4    11     8     9     6
```

15.4 RANDCHI Generate Chi-Square Random Variable

15.4.1 Usage

Generates a vector of chi-square random variables with the given number of degrees of freedom. The general syntax for its use is

```
y = randchi(n)
```

where **n** is an array containing the degrees of freedom for each generated random variable.

15.4.2 Function Internals

A chi-square random variable is essentially distributed as the squared Euclidean norm of a vector of standard Gaussian random variables. The number of degrees of freedom is generally the number of elements in the vector. In general, the PDF of a chi-square random variable is


$$f(x) = \frac{x^{r/2-1}e^{-x/2}}{\Gamma(r/2)2^{r/2}}$$

15.4.3 Example

First, a plot of the PDF for a family of chi-square random variables

```
--> f = zeros(7,100);
--> x = (1:100)/10;
--> for n=1:7;t=x.^(n/2-1).*exp(-x/2);f(n,:)=10*t/sum(t);end
--> plot(x,f')
```

The PDF is below:



Here is an example of using `randchi` and `randn` to compute some chi-square random variables with four degrees of freedom.

```
--> randchi(4*ones(1,6))
ans =
    8.9675    4.0015    3.2578    5.5461    2.5090    5.7587

--> sum(randn(4,6).^2)
ans =
    1.1941   10.6441    3.6228    8.4425    2.5031    1.9058
```

15.5 RANDEXP Generate Exponential Random Variable

15.5.1 Usage

Generates a vector of exponential random variables with the specified parameter. The general syntax for its use is

```
y = randexp(lambda)
```

where `lambda` is a vector containing the parameters for the generated random variables.

15.5.2 Function Internals

The exponential random variable is usually associated with the waiting time between events in a Poisson random process. The PDF of an exponential random variable is:

$$f(x) = \lambda e^{-\lambda x}$$

15.5.3 Example

Here is an example of using the `randexp` function to generate some exponentially distributed random variables

```
--> randexp(ones(1,6))
```

```
ans =
    0.0608    0.0019    1.1266    0.2012    0.5079    3.4205
```

15.6 RANDF Generate F-Distributed Random Variable

15.6.1 Usage

Generates random variables with an F-distribution. The general syntax for its use is

```
y = randf(n,m)
```

where **n** and **m** are vectors of the number of degrees of freedom in the numerator and denominator of the chi-square random variables whose ratio defines the statistic.

15.6.2 Function Internals

The statistic $F_{\{n,m\}}$ is defined as the ratio of two chi-square random variables:

$$F_{n,m} = \frac{\chi_n^2/n}{\chi_m^2/m}$$

The PDF is given by

$$f_{n,m} = \frac{m^{m/2} n^{n/2} x^{n/2-1}}{(m+nx)^{(n+m)/2} B(n/2, m/2)},$$

where $B(a,b)$ is the beta function.

15.6.3 Example

Here we use `randf` to generate some F-distributed random variables, and then again using the `randchi` function:

```
--> randf(5*ones(1,9),7)
```

```
ans =
```

```
Columns 1 to 8
```

```
    1.1944    0.9069    0.7558    1.5029    0.0621    1.3860    1.8161    0.3755
```

```
Columns 9 to 9
```

```
    3.5794
```

```
--> randchi(5*ones(1,9))./randchi(7*ones(1,9))
```

```
ans =
```

```
Columns 1 to 8
```

```
1.3085    1.2693    1.0684    0.4377    1.1158    0.7171    0.4151    1.8022
```

```
Columns 9 to 9
```

```
1.4606
```

15.7 RANDGAMMA Generate Gamma-Distributed Random Variable

15.7.1 Usage

Generates random variables with a gamma distribution. The general syntax for its use is

```
y = randgamma(a,r),
```

where **a** and **r** are vectors describing the parameters of the gamma distribution. Roughly speaking, if **a** is the mean time between changes of a Poisson random process, and we wait for the **r** change, the resulting wait time is Gamma distributed with parameters **a** and **r**.

15.7.2 Function Internals

The Gamma distribution arises in Poisson random processes. It represents the waiting time to the occurrence of the **r**-th event in a process with mean time **a** between events. The probability distribution of a Gamma random variable is

$$P(x) = \frac{a^r x^{r-1} e^{-ax}}{\Gamma(r)}.$$

Note also that for integer values of **r** that a Gamma random variable is effectively the sum of **r** exponential random variables with parameter **a**.

15.7.3 Example

Here we use the `randgamma` function to generate Gamma-distributed random variables, and then generate them again using the `randexp` function.

```
--> randgamma(1,15*ones(1,9))
```

```
ans =
```

```
Columns 1 to 8
```

```
22.7804    11.5514    16.8537    12.7457    16.2303    10.7442    19.3942    16.3611
```

```

Columns 9 to 9

    17.4772

--> sum(randexp(ones(15,9)))

ans =

Columns 1 to 8

    14.6404    15.1860    13.3147    11.4380     7.2307    10.8225    14.5271    12.4631

Columns 9 to 9

    11.8753

```

15.8 RANDI Uniformly Distributed Integer

15.8.1 Usage

Generates an array of uniformly distributed integers between the two supplied limits. The general syntax for `randi` is

```
y = randi(low,high)
```

where `low` and `high` are arrays of integers. Scalars can be used for one of the arguments. The output `y` is a uniformly distributed pseudo-random number between `low` and `high` (inclusive).

15.8.2 Example

Here is an example of a set of random integers between zero and 5:

```

--> randi(zeros(1,6),5*ones(1,6))

ans =
    1 0 4 1 5 0

```

15.9 RANDMULTI Generate Multinomial-distributed Random Variables

15.9.1 Usage

This function generates samples from a multinomial distribution given the probability of each outcome. The general syntax for its use is

```
y = randmulti(N,pvec)
```

where N is the number of experiments to perform, and $pvec$ is the vector of probabilities describing the distribution of outcomes.

15.9.2 Function Internals

A multinomial distribution describes the number of times each of m possible outcomes occurs out of N trials, where each outcome has a probability p_i . More generally, suppose that the probability of a Bernoulli random variable X_i is p_i , and that

$$\sum_{i=1}^m p_i = 1.$$

Then the probability that X_i occurs x_i times is

$$P_N(x_1, x_2, \dots, x_n) = \frac{N!}{x_1! \dots x_n!} p_1^{x_1} \dots p_n^{x_n}.$$

15.9.3 Example

Suppose an experiment has three possible outcomes, say heads, tails and edge, with probabilities 0.4999, 0.4999 and 0.0002, respectively. Then if we perform ten thousand coin flips we get

```
--> randmulti(10000, [0.4999, 0.4999, 0.0002])
```

```
ans =
 5026 4973    1
```

15.10 RANDN Gaussian (Normal) Random Number Generator

15.10.1 Usage

Creates an array of pseudo-random numbers of the specified size. The numbers are normally distributed with zero mean and a unit standard deviation (i.e., $\mu = 0$, $\sigma = 1$). Two separate syntaxes are possible. The first syntax specifies the array dimensions as a sequence of scalar dimensions:

```
y = randn(d1, d2, ..., dn).
```

The resulting array has the given dimensions, and is filled with random numbers. The type of y is `double`, a 64-bit floating point array. To get arrays of other types, use the `typecast` functions.

The second syntax specifies the array dimensions as a vector, where each element in the vector specifies a dimension length:

```
y = randn([d1, d2, ..., dn]).
```

This syntax is more convenient for calling `randn` using a variable for the argument.

Finally, `randn` supports two additional forms that allow you to manipulate the state of the random number generator. The first retrieves the state

```
y = randn('state')
```

which is a 625 length integer vector. The second form sets the state

```
randn('state',y)
```

or alternately, you can reset the random number generator with

```
randn('state',0)
```

15.10.2 Function Internals

Recall that the probability density function (PDF) of a normal random variable is

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}.$$

The Gaussian random numbers are generated from pairs of uniform random numbers using a transformation technique.

15.10.3 Example

The following example demonstrates an example of using the first form of the `randn` function.

```
--> randn(2,2,2)

ans =

(:, :, 1) =
   -1.7375   -0.5664
   -0.2634   -1.0112

(:, :, 2) =
   -0.4020    0.0557
   -1.8966    0.2098
```

The second example demonstrates the second form of the `randn` function.

```
--> randn([2,2,2])

ans =

(:, :, 1) =
   -0.7183    1.9415
    0.1010   -1.1747
```

```
(:,:,2) =
    0.3048    3.1685
   -1.4185   -0.6130
```

In the next example, we create a large array of 10000 normally distributed pseudo-random numbers. We then shift the mean to 10, and the variance to 5. We then numerically calculate the mean and variance using `mean` and `var`, respectively.

```
--> x = 10+sqrt(5)*randn(1,10000);
--> mean(x)
```

```
ans =
    10.0135
```

```
--> var(x)
```

```
ans =
    4.9458
```

Now, we use the state manipulation functions of `randn` to exactly reproduce a random sequence. Note that unlike using `seed`, we can exactly control where the random number generator starts by saving the state.

```
--> randn('state',0)    % restores us to startup conditions
--> a = randn(1,3)      % random sequence 1
```

```
a =
   -0.0362   -0.1404    0.6934
```

```
--> b = randn('state'); % capture the state vector
--> c = randn(1,3)      % random sequence 2
```

```
c =
    0.5998    0.7086   -0.9394
```

```
--> randn('state',b); % restart the random generator so...
--> c = randn(1,3)     % we get random sequence 2 again
```

```
c =
    0.5998    0.7086   -0.9394
```

15.11 RANDNBIN Generate Negative Binomial Random Variables

15.11.1 Usage

Generates random variables with a negative binomial distribution. The general syntax for its use is

```
y = randnbin(r,p)
```

where **r** is a vector of integers representing the number of successes, and **p** is the probability of success.

15.11.2 Function Internals

A negative binomial random variable describes the number of failures **x** that occur in **x+r** bernoulli trials, with a success on the **x+r** trial. The pdf is given by

$$P_{r,p}(x) = \binom{x+r-1}{r-1} p^r (1-p)^x.$$

15.11.3 Example

Here we generate some negative binomial random variables:

```
--> randnbin(3*ones(1,4),.01)
```

```
ans =
```

```
150 274 304 159
```

```
--> randnbin(6*ones(1,4),.01)
```

```
ans =
```

```
657 626 357 663
```

15.12 RANDNCHI Generate Noncentral Chi-Square Random Variable

15.12.1 Usage

Generates a vector of non-central chi-square random variables with the given number of degrees of freedom and the given non-centrality parameters. The general syntax for its use is

```
y = randnchi(n,mu)
```

where **n** is an array containing the degrees of freedom for each generated random variable (with each element of **n** $i=1$), and **mu** is the non-centrality shift (must be positive).

15.12.2 Function Internals

A non-central chi-square random variable is the sum of a chisquare deviate with **n-1** degrees of freedom plus the square of a normal deviate with mean **mu** and standard deviation 1.

15.12.3 Examples

Here is an example of a non-central chi-square random variable:

```
--> randnchi(5*ones(1,9),0.3)
```

```
ans =
```

```
Columns 1 to 8
```

```
1.7097    6.6003   14.1463    2.0817    4.4984    5.3132    3.1775    3.5291
```

```
Columns 9 to 9
```

```
5.5185
```

15.13 RANDNF Generate Noncentral F-Distribution Random Variable

15.13.1 Usage

Generates a vector of non-central F-distributed random variables with the specified parameters. The general syntax for its use is

```
y = randnf(n,m,c)
```

where **n** is the number of degrees of freedom in the numerator, and **m** is the number of degrees of freedom in the denominator. The vector **c** determines the non-centrality shift of the numerator.

15.13.2 Function Internals

A non-central F-distributed random variable is the ratio of a non-central chi-square random variable and a central chi-square random variable, i.e.,

$$F_{n,m,c} = \frac{\chi_{n,c}^2/n}{\chi_m^2/m}.$$

15.13.3 Example

Here we use the `randnf` to generate some non-central F-distributed random variables:

```
--> randnf(5*ones(1,9),7,1.34)
```

```
ans =
```

```
Columns 1 to 8
```

```
1.9422    0.5987    0.1890    0.7468    2.3759    8.2553    1.8047    0.2222
```

Columns 9 to 9

```
2.2680
```

15.14 RANDP Generate Poisson Random Variable

15.14.1 Usage

Generates a vector Poisson random variables with the given parameters. The general syntax for its use is

```
y = randp(nu),
```

where `nu` is an array containing the rate parameters for the generated random variables.

15.14.2 Function Internals

A Poisson random variable is generally defined by taking the limit of a binomial distribution as the sample size becomes large, with the expected number of successes being fixed (so that the probability of success decreases as $1/N$). The Poisson distribution is given by

$$P_{\nu}(n) = \frac{\nu^n e^{-\nu}}{n!}.$$

15.14.3 Example

Here is an example of using `randp` to generate some Poisson random variables, and also using `randbin` to do the same using $N=1000$ trials to approximate the Poisson result.

```
--> randp(33*ones(1,10))
```

```
ans =
```

```
31 33 34 44 32 29 34 30 32 32
```

```
--> randbin(1000*ones(1,10),33/1000*ones(1,10))
```

```
ans =
```

```
32 36 36 39 33 34 41 33 42 32
```

15.15 RANDPERM Random permutation

15.15.1 USAGE

```
y = randperm(n)
```

`y` is a random permutation of integers from 1 to `n`. `randperm` calls `rand` and changes its state.

15.15.2 Example

```
--> y = randperm(10)
```

```
y =  
 4 9 6 3 10 7 2 8 1 5
```

15.16 SEED Seed the Random Number Generator

15.16.1 Usage

Seeds the random number generator using the given integer seeds. Changing the seed allows you to choose which pseudo-random sequence is generated. The seed takes two `uint32` values:

```
seed(s,t)
```

where `s` and `t` are the seed values. Note that due to limitations in `ranlib`, the values of `s,t` must be between $0 \leq s,t \leq 2^{30}$.

15.16.2 Example

Here's an example of how the seed value can be used to reproduce a specific random number sequence.

```
--> seed(32,41);  
--> rand(1,5)
```

```
ans =  
 0.8589    0.3727    0.5551    0.9557    0.7367
```

```
--> seed(32,41);  
--> rand(1,5)
```

```
ans =  
 0.8589    0.3727    0.5551    0.9557    0.7367
```

Chapter 16

Input/Ouput Functions

16.1 CSVREAD Read Comma Separated Value (CSV) File

16.1.1 Usage

The `csvread` function reads a text file containing comma separated values (CSV), and returns the resulting numeric matrix (2D). The function supports multiple syntaxes. The first syntax for `csvread` is

```
x = csvread('filename')
```

which attempts to read the entire CSV file into array `x`. The file can contain only numeric values. Each entry in the file should be separated from other entries by a comma. However, FreeMat will attempt to make sense of the entries if the comma is missing (e.g., a space separated file will also parse correctly). For complex values, you must be careful with the spaces). The second form of `csvread` allows you to specify the first row and column (zero-based index)

```
x = csvread('filename',firstrow,firstcol)
```

The last form allows you to specify the range to read also. This form is

```
x = csvread('filename',firstrow,firstcol,readrange)
```

where `readrange` is either a 4-vector of the form `[R1,C1,R2,C2]`, where `R1,C1` is the first row and column to use, and `R2,C2` is the last row and column to use. You can also specify the `readrange` as a spreadsheet range `B12..C34`, in which case the index for the range is 1-based (as in a typical spreadsheet), so that `A1` is the first cell in the upper left corner. Note also that `csvread` is somewhat limited.

16.1.2 Example

Here is an example of a CSV file that we wish to read in

```
sample_data.csv
10, 12, 13, 00, 45, 16
```

```
09, 11, 52, 93, 05, 06
01, 03, 04, 04, 90, -3
14, 17, 13, 67, 30, 43
21, 33, 14, 44, 01, 00
```

We start by reading the entire file

```
--> csvread('sample_data.csv')
```

```
ans =
    10 12 13  0 45 16
     9 11 52 93  5  6
     1  3  4  4 90 -3
    14 17 13 67 30 43
    21 33 14 44  1  0
```

Next, we read everything starting with the second row, and third column

```
--> csvread('sample_data.csv',1,2)
```

```
ans =
    52 93  5  6
     4  4 90 -3
    13 67 30 43
    14 44  1  0
```

Finally, we specify that we only want the 3 x 3 submatrix starting with the second row, and third column

```
--> csvread('sample_data.csv',1,2,[1,2,3,4])
```

```
ans =
    52 93  5
     4  4 90
    13 67 30

    sample_data.csv
    10, 12, 13, 00, 45, 16
    09, 11, 52, 93, 05, 06
    01, 03, 04, 04, 90, -3
    14, 17, 13, 67, 30, 43
    21, 33, 14, 44, 01, 00
```

16.2 CSVWRITE Write Comma Separated Value (CSV) File

16.2.1 Usage

The `csvwrite` function writes a given matrix to a text file using comma separated value (CSV) notation. Note that you can create CSV files with arbitrary sized matrices, but that `csvread` has

limits on line length. If you need to reliably read and write large matrices, use `rawwrite` and `rawread` respectively. The syntax for `csvwrite` is

```
csvwrite('filename',x)
```

where `x` is a numeric array. The contents of `x` are written to `filename` as comma-separated values. You can also specify a row and column offset to `csvwrite` to force `csvwrite` to write the matrix `x` starting at the specified location in the file. This syntax of the function is

```
csvwrite('filename',x,startrow,startcol)
```

where `startrow` and `startcol` are the offsets in zero-based indexing.

16.2.2 Example

Here we create a simple matrix, and write it to a CSV file

```
--> x = [1,2,3;5,6,7]
```

```
x =
  1 2 3
  5 6 7
```

```
--> csvwrite('csvwrite.csv',x)
--> csvread('csvwrite.csv')
```

```
ans =
  1 2 3
  5 6 7
```

Next, we do the same with an offset.

```
--> csvwrite('csvwrite.csv',x,1,2)
--> csvread('csvwrite.csv')
```

```
ans =
  0 0 0 0
  0 1 2 3
  0 5 6 7
```

Note the extra zeros.

16.3 DISP Display a Variable or Expression

16.3.1 Usage

Displays the result of a set of expressions. The `disp` function takes a variable number of arguments, each of which is an expression to output:

```
disp(expr1,expr2,...,exprn)
```

This is functionally equivalent to evaluating each of the expressions without a semicolon after each.

16.3.2 Example

Here are some simple examples of using `disp`.

```
--> a = 32;
--> b = 1:4;
--> disp(a,b,pi)
32

1 2 3 4

3.1416
```

16.4 DLMREAD Read ASCII-delimited File

16.4.1 Usage

Loads a matrix from an ASCII-formatted text file with a delimiter between the entries. This function is similar to the `load -ascii` command, except that it can handle complex data, and it allows you to specify the delimiter. Also, you can read only a subset of the data from the file. The general syntax for the `dmlread` function is

```
y = dlmread(filename)
```

where `filename` is a string containing the name of the file to read. In this form, FreeMat will guess at the type of the delimiter in the file. The guess is made by examining the input for common delimiter characters, which are `,;` or a whitespace (e.g., tab). The text in the file is preprocessed to replace these characters with whitespace and the file is then read in using a whitespace for the delimiter.

If you know the delimiter in the file, you can specify it using this form of the function:

```
y = dlmread(filename, delimiter)
```

where `delimiter` is a string containing the delimiter. If `delimiter` is the empty string, then the delimiter is guessed from the file.

You can also read only a portion of the file by specifying a start row and start column:

```
y = dlmread(filename, delimiter, startrow, startcol)
```

where `startrow` and `startcol` are zero-based. You can also specify the data to read using a range argument:

```
y = dlmread(filename, delimiter, range)
```

where `range` is either a vector `[startrow,startcol,stoprow,stopcol]` or is specified in spreadsheet notation as `B4..ZA5`.

Note that complex numbers can be present in the file if they are encoded without whitespaces inside the number, and use either `i` or `j` as the indicator. Note also that when the delimiter is given, each incidence of the delimiter counts as a separator. Multiple separators generate zeros in the matrix.

16.5 FCLOSE File Close Function

16.5.1 Usage

Closes a file handle, or all open file handles. The general syntax for its use is either

```
fclose(handle)
```

or

```
fclose('all')
```

In the first case a specific file is closed, In the second, all open files are closed. Note that until a file is closed the file buffers are not flushed. Returns a '0' if the close was successful and a '-1' if the close failed for some reason.

16.5.2 Example

A simple example of a file being opened with `fopen` and then closed with `fclose`.

```
--> fp = fopen('test.dat','wb','ieee-le')
```

```
fp =  
8
```

```
--> fclose(fp)
```

16.6 FEOF End Of File Function

16.6.1 Usage

Check to see if we are at the end of the file. The usage is

```
b = feof(handle)
```

The `handle` argument must be a valid and active file handle. The return is true (logical 1) if the current position is at the end of the file, and false (logical 0) otherwise. Note that simply reading to the end of a file will not cause `feof` to return `true`. You must read past the end of the file (which will cause an error anyway). See the example for more details.

16.6.2 Example

Here, we read to the end of the file to demonstrate how `feof` works. At first pass, we force a read of the contents of the file by specifying `inf` for the dimension of the array to read. We then test the end of file, and somewhat counter-intuitively, the answer is `false`. We then attempt to read past the end of the file, which causes an error. An `feof` test now returns the expected value of `true`.

```
--> fp = fopen('test.dat','rb');
--> x = fread(fp,[512,inf],'float');
--> feof(fp)
```

```
ans =
     1
```

```
--> x = fread(fp,[1,1],'float');
--> feof(fp)
```

```
ans =
     1
```

16.7 FFLUSH Force File Flush

16.7.1 Usage

Flushes any pending output to a given file. The general use of this function is

```
fflush(handle)
```

where `handle` is an active file handle (as returned by `fopen`).

16.8 FGETLINE Read a String from a File

16.8.1 Usage

Reads a string from a file. The general syntax for its use is

```
s = fgetline(handle)
```

This function reads characters from the file `handle` into a `string` array `s` until it encounters the end of the file or a newline. The newline, if any, is retained in the output string. If the file is at its end, (i.e., that `feof` would return true on this handle), `fgetline` returns an empty string.

16.8.2 Example

First we write a couple of strings to a test file.

```
--> fp = fopen('testtext','w');
--> fprintf(fp,'String 1\n');
--> fprintf(fp,'String 2\n');
--> fclose(fp);
```

Next, we read then back.

```
--> fp = fopen('testtext','r')

fp =
    10

--> fgetline(fp)

ans =
String 1

--> fgetline(fp)

ans =
String 2

--> fclose(fp);
```

16.9 FOPEN File Open Function

16.9.1 Usage

Opens a file and returns a handle which can be used for subsequent file manipulations. The general syntax for its use is

```
fp = fopen(fname,mode,byteorder)
```

Here **fname** is a string containing the name of the file to be opened. **mode** is the mode string for the file open command. The first character of the mode string is one of the following:

- **'r'** Open file for reading. The file pointer is placed at the beginning of the file. The file can be read from, but not written to.
- **'r+'** Open for reading and writing. The file pointer is placed at the beginning of the file. The file can be read from and written to, but must exist at the outset.
- **'w'** Open file for writing. If the file already exists, it is truncated to zero length. Otherwise, a new file is created. The file pointer is placed at the beginning of the file.
- **'w+'** Open for reading and writing. The file is created if it does not exist, otherwise it is truncated to zero length. The file pointer placed at the beginning of the file.
- **'a'** Open for appending (writing at end of file). The file is created if it does not exist. The file pointer is placed at the end of the file.
- **'a+'** Open for reading and appending (writing at end of file). The file is created if it does not exist. The file pointer is placed at the end of the file.

Starting with FreeMat 4, all files are treated as binary files by default. To invoke the operating systems 'CR/LF' to 'CR' translation (on Win32) add a 't' to the mode string, as in 'rt+'.

Also, you can supply a second argument to **fopen** to retrieve error messages if the **fopen** fails.

```
[fp,messages] = fopen(fname,mode,byteorder)
```

Finally, FreeMat has the ability to read and write files of any byte-sex (endian). The third (optional) input indicates the byte-endianness of the file. If it is omitted, the native endian-ness of the machine running FreeMat is used. Otherwise, the third argument should be one of the following strings:

- 'le', 'ieee-le', 'little-endian', 'littleEndian', 'little', 'l', 'ieee-le.l64', 's'
- 'be', 'ieee-be', 'big-endian', 'bigEndian', 'big', 'b', 'ieee-be.l64', 'a'

If the file cannot be opened, or the file mode is illegal, then an error occurs. Otherwise, a file handle is returned (which is an integer). This file handle can then be used with `fread`, `fwrite`, or `fclose` for file access.

Note that three handles are assigned at initialization time:

- Handle 0 - is assigned to standard input
- Handle 1 - is assigned to standard output
- Handle 2 - is assigned to standard error

These handles cannot be closed, so that user created file handles start at 3.

16.9.2 Examples

Here are some examples of how to use `fopen`. First, we create a new file, which we want to be little-endian, regardless of the type of the machine. We also use the `fwrite` function to write some floating point data to the file.

```
--> fp = fopen('test.dat','w','ieee-le')
```

```
fp =  
8
```

```
--> fwrite(fp,float([1.2,4.3,2.1]))
```

```
ans =  
12
```

```
--> fclose(fp)
```

Next, we open the file and read the data back

```
--> fp = fopen('test.dat','r','ieee-le')
```

```
fp =  
8
```

```
--> fread(fp,[1,3],'float')
```

```
ans =  
    1.2000    4.3000    2.1000
```

```
--> fclose(fp)
```

Now, we re-open the file in append mode and add two additional floats to the file.

```
--> fp = fopen('test.dat','a+','le')
```

```
fp =  
    8
```

```
--> fwrite(fp,float([pi,e]))
```

```
ans =  
    8
```

```
--> fclose(fp)
```

Finally, we read all 5 float values from the file

```
--> fp = fopen('test.dat','r','ieee-le')
```

```
fp =  
    8
```

```
--> fread(fp,[1,5],'float')
```

```
ans =  
    1.2000    4.3000    2.1000    3.1416    2.7183
```

```
--> fclose(fp)
```

16.10 FORMAT Control the Format of Matrix Display

16.10.1 Usage

FreeMat supports several modes for displaying matrices (either through the `disp` function or simply by entering expressions on the command line. There are several options for the `format` command. The default mode is equivalent to

```
format short
```

which generally displays matrices with 4 decimals, and scales matrices if the entries have magnitudes larger than roughly $1e2$ or smaller than $1e-2$. For more information you can use

```
format long
```

which displays roughly 7 decimals for `float` and `complex` arrays, and 14 decimals for `double` and `dcomplex`. You can also use

```
format short e
```

to get exponential format with 4 decimals. Matrices are not scaled for exponential formats. Similarly, you can use

```
format long e
```

which displays the same decimals as `format long`, but in exponential format. You can also use the `format` command to retrieve the current format:

```
s = format
```

where `s` is a string describing the current format.

16.10.2 Example

We start with the short format, and two matrices, one of double precision, and the other of single precision.

```
--> format short
--> a = randn(4)

a =
    1.9819   -2.4403    0.3082   -0.8708
    0.3885    0.6777   -2.1203   -0.6899
   -0.8923   -1.0126   -1.1448   -0.3336
   -0.5528   -0.2117   -0.6066    0.1530

--> b = float(randn(4))

b =
    1.6282   -0.9987   -0.0002    0.7346
   -0.1461   -0.4450    0.3260    0.0591
   -0.1927   -0.2583   -0.3209   -1.7827
   -0.4694   -0.2961   -0.3487   -0.1476
```

Note that in the short format, these two matrices are displayed with the same format. In `long` format, however, they display differently

```
--> format long
--> a

ans =
    1.98194242245660  -2.44033999910309    0.30822105452542  -0.87083520854217
    0.38848412098646    0.67772654122050  -2.12029702950896  -0.68985792035578
   -0.89231827506021  -1.01256221976480  -1.14477420632547  -0.33359041318909
```

```
-0.55283586680694 -0.21170713821002 -0.60660544623052 0.15300924745427
```

```
--> b
```

```
ans =
```

```
1.6282195 -0.9986902 -0.0002282 0.7346091
-0.1460750 -0.4449911 0.3259999 0.0591399
-0.1926918 -0.2583237 -0.3208777 -1.7827009
-0.4693597 -0.2961315 -0.3487136 -0.1476461
```

Note also that we we scale the contents of the matrices, FreeMat rescales the entries with a scale premultiplier.

```
--> format short
```

```
--> a*1e4
```

```
ans =
```

```
1.0e+04 *
1.9819 -2.4403 0.3082 -0.8708
0.3885 0.6777 -2.1203 -0.6899
-0.8923 -1.0126 -1.1448 -0.3336
-0.5528 -0.2117 -0.6066 0.1530
```

```
--> a*1e-4
```

```
ans =
```

```
1.0e-04 *
1.9819 -2.4403 0.3082 -0.8708
0.3885 0.6777 -2.1203 -0.6899
-0.8923 -1.0126 -1.1448 -0.3336
-0.5528 -0.2117 -0.6066 0.1530
```

```
--> b*1e4
```

```
ans =
```

```
1.0e+04 *
1.6282 -0.9987 -0.0002 0.7346
-0.1461 -0.4450 0.3260 0.0591
-0.1927 -0.2583 -0.3209 -1.7827
-0.4694 -0.2961 -0.3487 -0.1476
```

```
--> b*1e-4
```

```
ans =
```

```

1.0e-04 *
 1.6282 -0.9987 -0.0002  0.7346
-0.1461 -0.4450  0.3260  0.0591
-0.1927 -0.2583 -0.3209 -1.7827
-0.4694 -0.2961 -0.3487 -0.1476

```

Next, we use the exponential formats:

```

--> format short e
--> a*1e4

ans =
 1.9819e+04 -2.4403e+04  3.0822e+03 -8.7084e+03
 3.8848e+03  6.7773e+03 -2.1203e+04 -6.8986e+03
-8.9232e+03 -1.0126e+04 -1.1448e+04 -3.3359e+03
-5.5284e+03 -2.1171e+03 -6.0661e+03  1.5301e+03

--> a*1e-4

ans =
 1.9819e-04 -2.4403e-04  3.0822e-05 -8.7084e-05
 3.8848e-05  6.7773e-05 -2.1203e-04 -6.8986e-05
-8.9232e-05 -1.0126e-04 -1.1448e-04 -3.3359e-05
-5.5284e-05 -2.1171e-05 -6.0661e-05  1.5301e-05

--> b*1e4

ans =
 1.6282e+04 -9.9869e+03 -2.2825e+00  7.3461e+03
-1.4608e+03 -4.4499e+03  3.2600e+03  5.9140e+02
-1.9269e+03 -2.5832e+03 -3.2088e+03 -1.7827e+04
-4.6936e+03 -2.9613e+03 -3.4871e+03 -1.4765e+03

--> b*1e-4

ans =
 1.6282e-04 -9.9869e-05 -2.2825e-08  7.3461e-05
-1.4608e-05 -4.4499e-05  3.2600e-05  5.9140e-06
-1.9269e-05 -2.5832e-05 -3.2088e-05 -1.7827e-04
-4.6936e-05 -2.9613e-05 -3.4871e-05 -1.4765e-05

```

Finally, if we assign the `format` function to a variable, we can retrieve the current format:

```

--> format short
--> t = format

```

```
t =  
short
```

16.11 FPRINTF Formated File Output Function (C-Style)

16.11.1 Usage

Prints values to a file. The general syntax for its use is

```
fprintf(fp,format,a1,a2,...).
```

or,

```
fprintf(format,a1,a2,...).
```

Here `format` is the format string, which is a string that controls the format of the output. The values of the variables `ai` are substituted into the output as required. It is an error if there are not enough variables to satisfy the format string. Note that this `fprintf` command is not vectorized! Each variable must be a scalar. The value `fp` is the file handle. If `fp` is omitted, file handle 1 is assumed, and the behavior of `fprintf` is effectively equivalent to `printf`. For more details on the format string, see `printf`.

16.11.2 Examples

A number of examples are present in the Examples section of the `printf` command.

16.12 FREAD File Read Function

16.12.1 Usage

Reads a block of binary data from the given file handle into a variable of a given shape and precision. The general use of the function is

```
A = fread(handle,size,precision)
```

The `handle` argument must be a valid value returned by the `fopen` function, and accessible for reading. The `size` argument determines the number of values read from the file. The `size` argument is simply a vector indicating the size of the array `A`. The `size` argument can also contain a single `inf` dimension, indicating that FreeMat should calculate the size of the array along that dimension so as to read as much data as possible from the file (see the examples listed below for more details). The data is stored as columns in the file, not rows.

Alternately, you can specify two return values to the `fread` function, in which case the second value contains the number of elements read

```
[A,count] = fread(...)
```

where `count` is the number of elements in `A`.

The third argument determines the type of the data. Legal values for this argument are listed below:

- 'uint8','uchar','unsigned char' for an unsigned, 8-bit integer.
- 'int8','char','integer*1' for a signed, 8-bit integer.
- 'uint16','unsigned short' for an unsigned, 16-bit integer.
- 'int16','short','integer*2' for a signed, 16-bit integer.
- 'uint32','unsigned int' for an unsigned, 32-bit integer.
- 'int32','int','integer*4' for a signed, 32-bit integer.
- 'single','float32','float','real*4' for a 32-bit floating point.
- 'double','float64','real*8' for a 64-bit floating point.

Starting with FreeMat 4, the format for the third argument has changed. If you specify only a type, such as 'float', the data is read in as single precision, but the output type is always 'double'. This behavior is consistent with Matlab. If you want the output type to match the input type (as was previous behavior in FreeMat), you must preface the precision string with a '*'. Thus, the precision string '*float' implies that data is read in as single precision, and the output is also single precision.

The third option is to specify the input and output types explicitly. You can do this by specifying a precision string of the form 'type1 => type2', where 'type1' is the input type and 'type2' is the output type. For example, if the input type is 'double' and the output type is to be a 'float', then a type spec of 'double => float' should be used.

16.12.2 Example

First, we create an array of 512 x 512 Gaussian-distributed float random variables, and then writing them to a file called test.dat.

```
--> A = float(randn(512));
--> fp = fopen('test.dat','w');
--> fwrite(fp,A);
--> fclose(fp);
```

Read as many floats as possible into a row vector

```
--> fp = fopen('test.dat','r');
--> x = fread(fp,[1,inf],'float');
--> fclose(fp);
--> who x
```

Variable Name	Type	Flags	Size
x	double		[1 262144]

Note that x is a double array. This behavior is new to FreeMat 4. Read the same floats into a 2-D float array.

```
--> fp = fopen('test.dat','r');
--> x = fread(fp,[512,inf],'float');
--> fclose(fp);
--> who x
Variable Name      Type   Flags      Size
                x    double    [512 512]
```

To read them as a single precision float array, we can use the following form:

```
--> fp = fopen('test.dat','r');
--> x = fread(fp,[512,inf],'*float');
--> fclose(fp);
--> who x
Variable Name      Type   Flags      Size
                x    single    [512 512]
```

16.13 FSCANF Formatted File Input Function (C-Style)

16.13.1 Usage

Reads values from a file. The general syntax for its use is

```
[a,count] = fscanf(handle,format,[size])
```

Here `format` is the format string, which is a string that controls the format of the input, `size` specifies the amount of data to be read. Values that are parsed from the `text` are stored in `a`. Note that `fscanf` is vectorized - the format string is reused as long as there are entries in the `text` string. See `printf` for a description of the format. Note that if the file is at the end-of-file, the `fscanf` will return

16.14 FSEEK Seek File To A Given Position

16.14.1 Usage

Moves the file pointer associated with the given file handle to the specified offset (in bytes). The usage is

```
fseek(handle,offset,style)
```

The `handle` argument must be a value and active file handle. The `offset` parameter indicates the desired seek offset (how much the file pointer is moved in bytes). The `style` parameter determines how the offset is treated. Three values for the `style` parameter are understood:

- string `'bof'` or the value `-1`, which indicate the seek is relative to the beginning of the file. This is equivalent to `SEEK_SET` in ANSI C.
- string `'cof'` or the value `0`, which indicates the seek is relative to the current position of the file. This is equivalent to `SEEK_CUR` in ANSI C.

- string 'eof' or the value 1, which indicates the seek is relative to the end of the file. This is equivalent to SEEK_END in ANSI C.

The offset can be positive or negative.

16.14.2 Example

The first example reads a file and then “rewinds” the file pointer by seeking to the beginning. The next example seeks forward by 2048 bytes from the files current position, and then reads a line of 512 floats.

```
--> % First we create the file
--> fp = fopen('test.dat','wb');
--> fwrite(fp,float(rand(4096,1)));
--> fclose(fp);
--> % Now we open it
--> fp = fopen('test.dat','rb');
--> % Read the whole thing
--> x = fread(fp,[1,inf],'float');
--> % Rewind to the beginning
--> fseek(fp,0,'bof');
--> % Read part of the file
--> y = fread(fp,[1,1024],'float');
--> who x y
Variable Name      Type  Flags      Size
                x   double    [1 4096]
                y   double    [1 1024]
--> % Seek 2048 bytes into the file
--> fseek(fp,2048,'cof');
--> % Read 512 floats from the file
--> x = fread(fp,[512,1],'float');
--> % Close the file
--> fclose(fp);
```

16.15 FTELL File Position Function

16.15.1 Usage

Returns the current file position for a valid file handle. The general use of this function is

```
n = ftell(handle)
```

The `handle` argument must be a valid and active file handle. The return is the offset into the file relative to the start of the file (in bytes).

16.15.2 Example

Here is an example of using `ftell` to determine the current file position. We read 512 4-byte floats, which results in the file pointer being at position $512*4 = 2048$.

```
--> fp = fopen('test.dat','wb');
--> fwrite(fp,randn(512,1));
--> fclose(fp);
--> fp = fopen('test.dat','rb');
--> x = fread(fp,[512,1],'float');
--> ftell(fp)
```

```
ans =
    2048
```

16.16 FWRITE File Write Function

16.16.1 Usage

Writes an array to a given file handle as a block of binary (raw) data. The general use of the function is

```
n = fwrite(handle,A)
```

The `handle` argument must be a valid value returned by the `fopen` function, and accessible for writing. The array `A` is written to the file a column at a time. The form of the output data depends on (and is inferred from) the precision of the array `A`. If the write fails (because we ran out of disk space, etc.) then an error is returned. The output `n` indicates the number of elements successfully written.

Note that unlike MATLAB, FreeMat 4 does not default to `uint8` for writing arrays to files. Alternately, the type of the data to be written to the file can be specified with the syntax

```
n = fwrite(handle,A,type)
```

where `type` is one of the following legal values:

- `'uint8','uchar','unsigned char'` for an unsigned, 8-bit integer.
- `'int8','char','integer*1'` for a signed, 8-bit integer.
- `'uint16','unsigned short'` for an unsigned, 16-bit integer.
- `'int16','short','integer*2'` for a signed, 16-bit integer.
- `'uint32','unsigned int'` for an unsigned, 32-bit integer.
- `'int32','int','integer*4'` for a signed, 32-bit integer.
- `'single','float32','float','real*4'` for a 32-bit floating point.
- `'double','float64','real*8'` for a 64-bit floating point.

16.16.2 Example

Heres an example of writing an array of 512 x 512 Gaussian-distributed `float` random variables, and then writing them to a file called `test.dat`.

```
--> A = float(randn(512));
--> fp = fopen('test.dat','w');
--> fwrite(fp,A,'single');
--> fclose(fp);
```

16.17 GETLINE Get a Line of Input from User

16.17.1 Usage

Reads a line (as a string) from the user. This function has two syntaxes. The first is

```
a = getline(prompt)
```

where `prompt` is a prompt supplied to the user for the query. The second syntax omits the `prompt` argument:

```
a = getline
```

Note that this function requires command line input, i.e., it will only operate correctly for programs or scripts written to run inside the FreeMat GUI environment or from the X11 terminal. If you build a stand-alone application and expect it to operate cross-platform, do not use this function (unless you include the FreeMat console in the final application).

16.18 GETPRINTLIMIT Get Limit For Printing Of Arrays

16.18.1 Usage

Returns the limit on how many elements of an array are printed using either the `disp` function or using expressions on the command line without a semi-colon. The default is set to one thousand elements. You can increase or decrease this limit by calling `setprintlimit`. This function is provided primarily so that you can temporarily change the output truncation and then restore it to the previous value (see the examples).

```
n=getprintlimit
```

where `n` is the current limit in use.

16.18.2 Example

Here is an example of using `getprintlimit` along with `setprintlimit` to temporarily change the output behavior of FreeMat.

```
--> A = randn(100,1);  
--> n = getprintlimit
```

```
n =  
    1000
```

```
--> setprintlimit(5);  
--> A
```

```
ans =  
    0.6674  
   -0.1948  
    0.5671  
   -0.6562  
   -0.4182
```

Print limit has been reached. Use `setprintlimit` function to enable longer printouts
--> `setprintlimit(n)`

16.19 HTMLREAD Read an HTML Document into FreeMat

16.19.1 Usage

Given a filename, reads an HTML document, (attempts to) parse it, and returns the result as a FreeMat data structure. The syntax for its use is:

```
p = htmlread(filename)
```

where `filename` is a string. The resulting object `p` is a data structure containing the information in the document. Note that this function works by internally converting the HTML document into something closer to XHTML, and then using the XML parser to parse it. In some cases, the converted HTML cannot be properly parsed. In such cases, a third party tool such as "tidy" will probably do a better job.

16.20 IMREAD Read Image File To Matrix

16.20.1 Usage

Reads the image data from the given file into a matrix. Note that FreeMat's support for `imread` is not complete. Only some of the formats specified in the MATLAB API are implemented. The syntax for its use is

```
[A,map,alpha] = imread(filename)
```

where `filename` is the name of the file to read from. The returned arrays `A` contain the image data, `map` contains the colormap information (for indexed images), and `alpha` contains the alphas (transparency). The returned values will depend on the type of the original image. Generally you can read images in the `jpg`, `png`, `xpm`, `ppm` and some other formats.

16.21 IMWRITE Write Matrix to Image File

16.21.1 Usage

Write the image data from the matrix into a given file. Note that FreeMat's support for `imwrite` is not complete. You can write images in the `jpg`, `png`, `xpm`, `ppm` and some other formats. The syntax for its use is

```
imwrite(A, filename)
imwrite(A, map, filename)
imwrite(A, map, filename, 'Alpha', alpha)
```

or Octave-style syntax:

```
imwrite(filename, A)
imwrite(filename, A, map)
imwrite(filename, A, map, alpha)
```

where `filename` is the name of the file to write to. The input array `A` contains the image data (2D for gray or indexed, and 3D for color). If `A` is an integer array (`int8`, `uint8`, `int16`, `uint16`, `int32`, `uint32`), the values of its elements should be within 0-255. If `A` is a floating-point array (`float` or `double`), the value of its elements should be in the range `[0,1]`. `map` contains the colormap information (for indexed images), and `alpha` the alphamap (transparency).

16.21.2 Example

Here is a simple example of `imread/imwrite`. First, we generate a grayscale image and save it to an image file.

```
--> a = uint8(255*rand(64));
--> figure(1), image(a), colormap(gray)
--> title('image to save')
```

Warning: Newly defined variable `nargin` shadows a function of the same name. Use `clear nargin`

```
--> imwrite(a, 'test.bmp')
```

Then, we read image file and show it:

```
--> b = imread('test.bmp');
--> figure(2), image(b), colormap(gray)
--> title('loaded image')
```

Warning: Newly defined variable `nargin` shadows a function of the same name. Use `clear nargin`

16.22 INPUT Get Input From User

16.22.1 Usage

The `input` function is used to obtain input from the user. There are two syntaxes for its use. The first is

```
r = input('prompt')
```

in which case, the prompt is presented, and the user is allowed to enter an expression. The expression is evaluated in the current workspace or context (so it can use any defined variables or functions), and returned for assignment to the variable (`r` in this case). In the second form of the `input` function, the syntax is

```
r = input('prompt','s')
```

in which case the text entered by the user is copied verbatim to the output.

16.23 LOAD Load Variables From A File

16.23.1 Usage

Loads a set of variables from a file in a machine independent format. The `load` function takes one argument:

```
load filename,
```

or alternately,

```
load('filename')
```

This command is the companion to `save`. It loads the contents of the file generated by `save` back into the current context. Global and persistent variables are also loaded and flagged appropriately. By default, FreeMat assumes that files that end in a `.mat` or `.MAT` extension are MATLAB-formatted files. Also, FreeMat assumes that files that end in `.txt` or `.TXT` are ASCII files. For other filenames, FreeMat first tries to open the file as a FreeMat binary format file (as created by the `save` function). If the file fails to open as a FreeMat binary file, then FreeMat attempts to read it as an ASCII file.

You can force FreeMat to assume a particular format for the file by using alternate forms of the `load` command. In particular,

```
load -ascii filename
```

will load the data in file `filename` as an ASCII file (space delimited numeric text) loaded into a single variable in the current workspace with the name `filename` (without the extension).

For MATLAB-formatted data files, you can use

```
load -mat filename
```

which forces FreeMat to assume that `filename` is a MAT-file, regardless of the extension on the filename.

You can also specify which variables to load from a file (not from an ASCII file - only single 2-D variables can be successfully saved and retrieved from ASCII files) using the additional syntaxes of the `load` command. In particular, you can specify a set of variables to load by name

```
load filename Var_1 Var_2 Var_3 ...
```

where `Var_n` is the name of a variable to load from the file. Alternately, you can use the regular expression syntax

```
load filename -regex expr_1 expr_2 expr_3 ...
```

where `expr_n` is a regular expression (roughly as expected by `regex`). Note that a simpler regular expression mechanism is used for this syntax than the full mechanism used by the `regex` command.

Finally, you can use `load` to create a variable containing the contents of the file, instead of automatically inserting the variables into the current workspace. For this form of `load` you must use the function syntax, and capture the output:

```
V = load('arg1','arg2',...)
```

which returns a structure `V` with one field for each variable retrieved from the file. For ASCII files, `V` is a double precision matrix.

16.23.2 Example

Here is a simple example of `save/load`. First, we save some variables to a file.

```
--> D = {1,5,'hello'};
--> s = 'test string';
--> x = randn(512,1);
--> z = zeros(512);
--> who
Variable Name      Type   Flags      Size
                D      cell      [1 3]
                s      char      [1 11]
                x      double    [512 1]
                z      double    [512 512]
--> save loadsave.dat
```

Next, we clear the variables, and then load them back from the file.

```
--> clear D s x z
--> who
Variable Name      Type   Flags      Size
                ans     double    [0 0]
--> load loadsave.dat
--> who
Variable Name      Type   Flags      Size
                D      cell      [1 3]
                ans     double    [0 0]
                s      char      [1 11]
                x      double    [512 1]
                z      double    [512 512]
```

16.24 PAUSE Pause Script Execution

16.24.1 Usage

The `pause` function can be used to pause execution of FreeMat scripts. There are several syntaxes for its use. The first form is

```
pause
```

This form of the `pause` function pauses FreeMat until you press any key. The second form of the `pause` function takes an argument

```
pause(p)
```

where `p` is the number of seconds to pause FreeMat for. The pause argument should be accurate to a millisecond on all supported platforms. Alternately, you can control all `pause` statements using:

```
pause on
```

which enables pauses and

```
pause off
```

which disables all `pause` statements, both with and without arguments.

16.25 PRINTF Formated Output Function (C-Style)

16.25.1 Usage

Prints values to the output. The general syntax for its use is

```
printf(format,a1,a2,...)
```

Here `format` is the format string, which is a string that controls the format of the output. The values of the variables `a_i` are substituted into the output as required. It is an error if there are not enough variables to satisfy the format string. Note that this `printf` command is not vectorized! Each variable must be a scalar.

It is important to point out that the `printf` function does not add a newline (or carriage return) to the output by default. That can lead to some confusing behavior if you do not know what to expect. For example, the command `printf('Hello')` does not appear to produce any output. In fact, it does produce the text, but it then gets overwritten by the prompt. To see the text, you need `printf('Hello\n')`. This seems odd, but allows you to assemble a line using multiple `printf` commands, including the `'\n'` when you are done with the line. You can also use the `'\r'` character as an explicit carriage return (with no line feed). This allows you to write to the same line many times (to show a progress string, for example).

16.25.2 Format of the format string

The format string is a character string, beginning and ending in its initial shift state, if any. The format string is composed of zero or more directives: ordinary characters (not unchanged to the output stream; and conversion specifications, each of which results in fetching zero or more subsequent arguments. Each conversion specification is introduced by the character conversion specifier. In between there may be (in this order) zero or more flags, an optional minimum field width, and an optional precision.

The arguments must correspond properly (after type promotion) with the conversion specifier, and are used in the order given.

16.25.3 The flag characters

The character % is followed by zero or more of the following flags:

- `\#` The value should be converted to an “alternate form”. For `o` conversions, the first character of the output string is made zero (by prefixing a `0` if it was not zero already). For `x` and `X` conversions, a nonzero result has the string `'0x'` (or `'0X'` for `X` conversions) prepended to it. For `a`, `A`, `e`, `E`, `f`, `F`, `g`, and `G` conversions, the result will always contain a decimal point, even if no digits follow it (normally, a decimal point appears in the results of those conversions only if a digit follows). For `g` and `G` conversions, trailing zeros are not removed from the result as they would otherwise be. For other conversions, the result is undefined.
- `0` The value should be zero padded. For `d`, `i`, `o`, `u`, `x`, `X`, `a`, `A`, `e`, `E`, `f`, `F`, `g`, and `G` conversions, the converted value is padded on the left with zeros rather than blanks. If the `0` and `-` flags both appear, the `0` flag is ignored. If a precision is given with a numeric conversion (`d`, `i`, `o`, `u`, `x`, and `X`), the `0` flag is ignored. For other conversions, the behavior is undefined.
- `-` The converted value is to be left adjusted on the field boundary. (The default is right justification.) Except for `n` conversions, the converted value is padded on the right with blanks, rather than on the left with blanks or zeros. A `-` overrides a `0` if both are given.
- `' '` (a space) A blank should be left before a positive number (or empty string) produced by a signed conversion.
- `+` A sign (`+` or `-`) always be placed before a number produced by a signed conversion. By default a sign is used only for negative numbers. A `+` overrides a space if both are used.

16.25.4 The field width

An optional decimal digit string (with nonzero first digit) specifying a minimum field width. If the converted value has fewer characters than the field width, it will be padded with spaces on the left (or right, if the left-adjustment flag has been given). A negative field width is taken as a `'-'` flag followed by a positive field width. In no case does a non-existent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is expanded to contain the conversion result.

16.25.5 The precision

An optional precision, in the form of a period (`'.'`) followed by an optional decimal digit string. If the precision is given as just `'.'`, or the precision is negative, the precision is taken to be zero. This gives the minimum number of digits to appear for `d`, `i`, `o`, `u`, `x`, and `X` conversions, the number of digits to appear after the radix character for `a`, `A`, `e`, `E`, `f`, and `F` conversions, the maximum number of significant digits for `g` and `G` conversions, or the maximum number of characters to be printed from a string for `s` conversions.

16.25.6 The conversion specifier

A character that specifies the type of conversion to be applied. The conversion specifiers and their meanings are:

- `d,i` The `int` argument is converted to signed decimal notation. The precision, if any, gives the minimum number of digits that must appear; if the converted value requires fewer digits, it is padded on the left with zeros. The default precision is 1. When 0 is printed with an explicit precision 0, the output is empty.
- `o,u,x,X` The unsigned `int` argument is converted to unsigned octal (`o`), unsigned decimal (`u`), or unsigned hexadecimal (`x` and `X`) notation. The letters `abcdef` are used for `x` conversions; the letters `ABCDEF` are used for `X` conversions. The precision, if any, gives the minimum number of digits that must appear; if the converted value requires fewer digits, it is padded on the left with zeros. The default precision is 1. When 0 is printed with an explicit precision 0, the output is empty.
- `e,E` The double argument is rounded and converted in the style `[-]d.ddde dd` where there is one digit before the decimal-point character and the number of digits after it is equal to the precision; if the precision is missing, it is taken as 6; if the precision is zero, no decimal-point character appears. An `E` conversion uses the letter `E` (rather than `e`) to introduce the exponent. The exponent always contains at least two digits; if the value is zero, the exponent is 00.
- `f,F` The double argument is rounded and converted to decimal notation in the style `[-]ddd.ddd`, where the number of digits after the decimal-point character is equal to the precision specification. If the precision is missing, it is taken as 6; if the precision is explicitly zero, no decimal-point character appears. If a decimal point appears, at least one digit appears before it.
- `g,G` The double argument is converted in style `f` or `e` (or `F` or `E` for `G` conversions). The precision specifies the number of significant digits. If the precision is missing, 6 digits are given; if the precision is zero, it is treated as 1. Style `e` is used if the exponent from its conversion is less than -4 or greater than or equal to the precision. Trailing zeros are removed from the fractional part of the result; a decimal point appears only if it is followed by at least one digit.
- `c` The `int` argument is converted to an unsigned char, and the resulting character is written.
- `s` The string argument is printed.
- `%A '%'` is written. No argument is converted. The complete conversion specification is `'%%'`.

16.25.7 Example

Here are some examples of the use of `printf` with various arguments. First we print out an integer and double value.

```
--> printf('intvalue is %d, floatvalue is %f\n',3,1.53);
intvalue is 3, floatvalue is 1.530000
```

Next, we print out a string value.

```
--> printf('string value is %s\n','hello');
string value is hello
```

Now, we print out an integer using 12 digits, zeros up front.

```
--> printf('integer padded is %012d\n',32);
integer padded is 000000000032
```

Print out a double precision value with a sign, a total of 18 characters (zero prepended if necessary), a decimal point, and 12 digit precision.

```
--> printf('float value is %+018.12f\n',pi);
float value is +0003.141592653590
```

16.26 RAWREAD Read N-dimensional Array From File

16.26.1 Usage

The syntax for `rawread` is

```
function x = rawread(fname,size,precision,byteorder)
```

where `fname` is the name of the file to read from, and `size` is an n-dimensional vector that stores the size of the array in each dimension. The argument `precision` is the type of the data to read in:

- 'uint8','uchar','unsigned char' for unsigned, 8-bit integers
- 'int8','char','integer*1' for signed, 8-bit integers
- 'uint16','unsigned short' for unsigned, 16-bit integers
- 'int16','short','integer*2' for signed, 16-bit integers
- 'uint32','unsigned int' for unsigned, 32-bit integers
- 'int32','int','integer*4' for signed, 32-bit integers
- 'uint64','unsigned int' for unsigned, 64-bit integers
- 'int64','int','integer*8' for signed, 64-bit integers
- 'single','float32','float','real*4' for 32-bit floating point

- 'double', 'float64', 'real*8' for 64-bit floating point
- 'complex', 'complex*8' for 64-bit complex floating point (32 bits for the real and imaginary part).
- 'dcomplex', 'complex*16' for 128-bit complex floating point (64 bits for the real and imaginary part).

As a special feature, one of the size elements can be 'inf', in which case, the largest possible array is read in. If `byteorder` is left unspecified, the file is assumed to be of the same byte-order as the machine `FreeMat` is running on. If you wish to force a particular byte order, specify the `byteorder` argument as

- 'le', 'ieee-le', 'little-endian', 'littleEndian', 'little'
- 'be', 'ieee-be', 'big-endian', 'bigEndian', 'big'

16.27 RAWWRITE Write N-dimensional Array From File

16.27.1 Usage

The syntax for `rawwrite` is

```
function rawwrite(fname,x,byteorder)
```

where `fname` is the name of the file to write to, and the (numeric) array `x` is written to the file in its native type (e.g. if `x` is of type `int16`, then it will be written to the file as 16-bit signed integers. If `byteorder` is left unspecified, the file is assumed to be of the same byte-order as the machine `FreeMat` is running on. If you wish to force a particular byte order, specify the `byteorder` argument as

- 'le', 'ieee-le', 'little-endian', 'littleEndian', 'little'
- 'be', 'ieee-be', 'big-endian', 'bigEndian', 'big'

16.28 SAVE Save Variables To A File

16.28.1 Usage

Saves a set of variables to a file in a machine independent format. There are two formats for the function call. The first is the explicit form, in which a list of variables are provided to write to the file:

```
save filename a1 a2 ...
```

In the second form,

```
save filename
```

all variables in the current context are written to the file. The format of the file is a simple binary encoding (raw) of the data with enough information to restore the variables with the `load` command. The endianness of the machine is encoded in the file, and the resulting file should be portable between machines of similar types (in particular, machines that support IEEE floating point representation).

You can also specify both the filename as a string, in which case you also have to specify the names of the variables to save. In particular

```
save('filename','a1','a2')
```

will save variables `a1` and `a2` to the file.

Starting with version 2.0, FreeMat can also read and write MAT files (the file format used by MATLAB) thanks to substantial work by Thomas Beutlich. Support for MAT files in version 2.1 has improved over previous versions. In particular, classes should be saved properly, as well as a broader range of sparse matrices. Compression is supported for both reading and writing to MAT files. MAT file support is still in the alpha stages, so please be cautious with using it to store critical data. The file format is triggered by the extension. To save files with a MAT format, simply use a filename with a ".mat" ending.

The `save` function also supports ASCII output. This is a very limited form of the save command - it can only save numeric arrays that are 2-dimensional. This form of the `save` command is triggered using

```
save -ascii filename var1 var 2
```

although where `-ascii` appears on the command line is arbitrary (provided it comes after the `save` command, of course). By default, the `save` command uses an 8-digit exponential format notation to save the values to the file. You can specify that you want 16-digits using the

```
save -ascii -double filename var1 var2
```

form of the command. Also, by default, `save` uses spaces as the delimiters between the entries in the matrix. If you want tabs instead, you can use

```
save -ascii -tabs filename var1 var2
```

(you can also use both the `-tabs` and `-double` flags simultaneously).

Finally, you can specify that `save` should only save variables that match a particular regular expression. Any of the above forms can be combined with the `-regexp` flag:

```
save filename -regexp pattern1 pattern2
```

in which case variables that match any of the patterns will be saved.

16.28.2 Example

Here is a simple example of `save/load`. First, we save some variables to a file.

```
--> D = {1,5,'hello'};
--> s = 'test string';
--> x = randn(512,1);
--> z = zeros(512);
```

```
--> who
Variable Name      Type  Flags      Size
      D      cell
      s      char
      x     double
      z     double
[512 512]
```

```
--> save loadsave.dat
```

Next, we clear the variables, and then load them back from the file.

```
--> clear D s x z
--> who
Variable Name      Type  Flags      Size
      ans    double
[0 0]
--> load loadsave.dat
--> who
Variable Name      Type  Flags      Size
      D      cell
      ans    double
      s      char
      x     double
      z     double
[512 512]
```

16.29 SETPRINTLIMIT Set Limit For Printing Of Arrays

16.29.1 Usage

Changes the limit on how many elements of an array are printed using either the `disp` function or using expressions on the command line without a semi-colon. The default is set to one thousand elements. You can increase or decrease this limit by calling

```
setprintlimit(n)
```

where `n` is the new limit to use.

16.29.2 Example

Setting a smaller print limit avoids pages of output when you forget the semicolon on an expression.

```
--> A = randn(512);
--> setprintlimit(10)
--> A
```

```
ans =
```

```
Columns 1 to 8
```

```
-0.2911  -2.2394   0.7388   0.4197   0.4382  -0.1399   0.9364  -1.8721
 0.6263   0.5764
```

Print limit has been reached. Use `setprintlimit` function to enable longer printouts
 --> `setprintlimit(1000)`

16.30 SPRINTF Formated String Output Function (C-Style)

16.30.1 Usage

Prints values to a string. The general syntax for its use is

```
y = sprintf(format,a1,a2,...).
```

Here `format` is the format string, which is a string that controls the format of the output. The values of the variables `a_i` are substituted into the output as required. It is an error if there are not enough variables to satisfy the format string. Note that this `sprintf` command is not vectorized! Each variable must be a scalar. The returned value `y` contains the string that would normally have been printed. For more details on the format string, see `printf`.

16.30.2 Examples

Here is an example of a loop that generates a sequence of files based on a template name, and stores them in a cell array.

```
--> l = {}; for i = 1:5; s = sprintf('file_%d.dat',i); l(i) = {s}; end;
--> l
```

```
ans =
 [file_1.dat] [file_2.dat] [file_3.dat] [file_4.dat] [file_5.dat]
```

16.31 SSCANF Formated String Input Function (C-Style)

16.31.1 Usage

Reads values from a string. The general syntax for its use is

```
[a, count, errmsg, nextind] = sscanf(text,format,[size])
```

Here `format` is the format string, which is a string that controls the format of the input, `size` specifies the amount of data to be read. Values that are parsed from the `text` are stored in `a`. Note that `sscanf` is vectorized - the format string is reused as long as there are entries in the `text` string. See `printf` for a description of the format.

16.32 STR2NUM Convert a String to a Number

16.32.1 Usage

Converts a string to a number. The general syntax for its use is

```
x = str2num(string)
```

Here `string` is the data string, which contains the data to be converted into a number. The output is in double precision, and must be typecasted to the appropriate type based on what you need. Note that by definition, `str2num` is entirely equivalent to `eval(['[' string '']', [])` with all of the associated problems where `string` contains text that causes side effects.

16.33 URLWRITE Retrieve a URL into a File

16.33.1 Usage

Given a URL and a timeout, attempts to retrieve the URL and write the contents to a file. The syntax is

```
f = urlwrite(url,filename,timeout)
```

The `timeout` is in milliseconds. Note that the URL must be a complete spec (i.e., including the name of the resource you wish to retrieve). So for example, you cannot use `http://www.google.com` as a URL, but must instead use `http://www.google.com/index.html`.

16.34 WAVPLAY

16.34.1 Usage

Plays a linear PCM set of samples through the audio system. This function is only available if the `portaudio` library was available when FreeMat was built. The syntax for the command is one of:

```
wavplay(y)
wavplay(y,sampling_rate)
wavplay(...,mode)
```

where `y` is a matrix of audio samples. If `y` has two columns, then the audio playback is in stereo. The `y` input can be of types `float`, `double`, `int32`, `int16`, `int8`, `uint8`. For `float` and `double` types, the sample values in `y` must be between `-1` and `1`. The `sampling_rate` specifies the rate at which the data is recorded. If not specified, the `sampling_rate` defaults to `11025Hz`. Finally, you can specify a playback mode of `'sync'` which is synchronous playback or a playback mode of `'async'` which is asynchronous playback. For `'sync'` playback, the `wavplay` function returns when the playback is complete. For `'async'` playback, the function returns immediately (unless a former playback is still issuing).

16.35 WAVREAD Read a WAV Audio File

16.35.1 Usage

The `wavread` function (attempts) to read the contents of a linear PCM audio WAV file. This function could definitely use improvements - it is based on a very simplistic notion of a WAV file. The simplest form for its use is

```
y = wavread(filename)
```

where `filename` is the name of the WAV file to read. If no extension is provided, FreeMat will add a `.wav` extension. This loads the data from the WAV file into `y`, and returns it in `double` precision, normalized format. If you want additional information on, for example, the WAV sampling rate or bit depth, you can request it via

```
[y, SamplingRate, BitDepth] = wavread(filename)
```

where `SamplingRate` and `BitDepth` are the sampling rate (in Hz) and the bit depth of the original data in the WAV file. If you only want to load part of the WAV file, you can use

```
[...] = wavread(filename, N)
```

where `N` indicates the number of samples to read from the file. Alternately, you can indicate a range of samples to load via

```
[...] = wavread(filename, [N1 N2])
```

which returns only the indicated samples from each channel in the file. By default, the output format is `double` precision. You can control the format of the output by indicating

```
[...] = wavread(filename, format)
```

where `format` is either `'double'` for double precision output, or `'native'` for native precision output (meaning whatever bitdepth that was present in the original file). Finally, you can use the `'size'` flag

```
y_siz = wavread(filename, 'size')
```

which returns a vector `[samples channels]` indicating the size of the data present in the WAV file.

16.36 WAVRECORD

16.36.1 Usage

Records linear PCM sound from the audio system. This function is only available if the `portaudio` library was available when FreeMat was built. The syntax for this command is one of:

```
y = wavrecord(samples,rate)
y = wavrecord(...,channels)
y = wavrecord(...,'datatype')
```

where `samples` is the number of samples to record, and `rate` is the sampling rate. If not specified, the `rate` defaults to 11025Hz. If you want to record in stereo, specify `channels = 2`. Finally, you can specify the type of the recorded data (defaults to `FM_DOUBLE`). Valid choices are `float`, `double`, `int32`, `int16`, `int8`, `uint8`.

16.37 WAVWRITE Write a WAV Audio File

16.37.1 Usage

The `wavwrite` function writes an audio signal to a linear PCM WAV file. The simplest form for its use is

```
wavwrite(y,filename)
```

which writes the data stored in `y` to a WAV file with the name `filename`. By default, the output data is assumed to be sampled at a rate of 8 KHz, and is output using 16 bit integer format. Each column of `y` is written as a separate channel. The data are clipped to the range `[-1,1]` prior to writing them out. If you want the data to be written with a different sampling frequency, you can use the following form of the `wavwrite` command:

```
wavwrite(y,SampleRate,filename)
```

where `SampleRate` is in Hz. Finally, you can specify the number of bits to use in the output form of the file using the form

```
wavwrite(y,SampleRate,NBits,filename)
```

where `NBits` is the number of bits to use. Legal values include `8,16,24,32`. For less than 32 bit output format, the data is truncated to the range `[-1,1]`, and an integer output format is used (type 1 PCM in WAV-speak). For 32 bit output format, the data is written in type 3 PCM as floating point data.

16.38 XMLREAD Read an XML Document into FreeMat

16.38.1 Usage

Given a `filename`, reads an XML document, parses it, and returns the result as a FreeMat data structure. The syntax for its use is:

```
p = xmlread(filename)
```

where `filename` is a `string`. The resulting object `p` is a data structure containing the information in the document. Note that the returned object `p` is not the same object as the one returned by MATLAB's `xmlread`, although the information content is the same. The output is largely compatible with the output of the `parseXML` example in the `xmlread` documentation of the MATLAB API.

Chapter 17

String Functions

17.1 BLANKS Create a blank string

17.1.1 Usage

```
str = blanks(n)
```

Create a string `str` containing `n` blank characters.

17.1.2 Example

A simple example:

```
--> sprintf(['x0123456789y\n', 'x', blanks(10), 'y\n'])
```

```
ans =  
x0123456789y  
x          y
```

17.2 CELLSTR Convert character array to cell array of strings

17.2.1 Usage

The `cellstr` converts a character array matrix into a cell array of individual strings. Each string in the matrix is placed in a different cell, and extra spaces are removed. The syntax for the command is

```
y = cellstr(x)
```

where `x` is an `N x M` array of characters as a string.

17.2.2 Example

Here is an example of how to use `cellstr`

```
--> a = ['quick';'brown';'fox  '; 'is  ']
```

```
a =  
quick  
brown  
fox  
is  
--> cellstr(a)
```

```
ans =  
 [quick]  
 [brown]  
 [fox]  
 [is]
```

17.3 DEBLANK Remove trailing blanks from a string

17.3.1 Usage

The `deblank` function removes spaces at the end of a string when used with the syntax

```
y = deblank(x)
```

where `x` is a string, in which case, all of the extra spaces in `x` are stripped from the end of the string. Alternately, you can call `deblank` with a cell array of strings

```
y = deblank(c)
```

in which case each string in the cell array is deblanked.

17.3.2 Example

A simple example

```
--> deblank('hello  ')
```

```
ans =  
hello
```

and a more complex example with a cell array of strings

```
--> deblank({'hello  ','there ',' is ',' sign  '})
```

```
ans =  
 [hello] [there] [ is] [ sign]
```

17.4 ISALPHA Test for Alpha Characters in a String

17.4.1 Usage

The `isalpha` functions returns a logical array that is 1 for characters in the argument string that are letters, and is a logical 0 for characters in the argument that are not letters. The syntax for its use is

```
x = isalpha(s)
```

where `s` is a `string`. Note that this function is not locale sensitive, and returns a logical 1 for letters in the classic ASCII sense (a through z, and A through Z).

17.4.2 Example

A simple example of `isalpha`:

```
--> isalpha('numb3r5')
```

```
ans =  
1 1 1 1 0 1 0
```

17.5 ISDIGIT Test for Digit Characters in a String

17.5.1 Usage

The `isdigit` functions returns a logical array that is 1 for characters in the argument string that are digits, and is a logical 0 for characters in the argument that are not digits. The syntax for its use is

```
x = isdigit(s)
```

where `s` is a `string`.

17.5.2 Example

A simple example of `isdigit`:

```
--> isdigit('numb3r5')
```

```
ans =  
0 0 0 0 1 0 1
```

17.6 ISSPACE Test for Space Characters in a String

17.6.1 Usage

The `isspace` functions returns a logical array that is 1 for characters in the argument string that are spaces, and is a logical 0 for characters in the argument that are not spaces. The syntax for its use is

```
x = isspace(s)
```

where `s` is a string. A blank character is considered a space, newline, tab, carriage return, formfeed, and vertical tab.

17.6.2 Example

A simple example of `isspace`:

```
--> isspace(' hello there world ')
ans =
 1 1 0 0 0 0 0 1 0 0 0 0 0 1 0 0 0 0 0 1
```

17.7 LOWER Convert strings to lower case

17.7.1 Usage

The `lower` function converts a string to lower case with the syntax

```
y = lower(x)
```

where `x` is a string, in which case all of the upper case characters in `x` (defined as the range 'A'-'Z') are converted to lower case. Alternately, you can call `lower` with a cell array of strings

```
y = lower(c)
```

in which case each string in the cell array is converted to lower case.

17.7.2 Example

A simple example:

```
--> lower('this Is Strange CAPitalizaTion')
ans =
this is strange capitalization
```

and a more complex example with a cell array of strings

```
--> lower({'This', 'Is', 'Strange', 'CAPitalizaTion'})
ans =
 [this] [is] [strange] [capitalization]
```

17.8 REGEXP Regular Expression Matching Function

17.8.1 Usage

Matches regular expressions in the provided string. This function is complicated, and compatibility with MATLABs syntax is not perfect. The syntax for its use is

```
regexp('str','expr')
```

which returns a row vector containing the starting index of each substring of `str` that matches the regular expression described by `expr`. The second form of `regexp` returns six outputs in the following order:

```
[start stop tokenExtents match tokens names] = regexp('str','expr')
```

where the meaning of each of the outputs is defined below.

- `start` is a row vector containing the starting index of each substring that matches the regular expression.
- `stop` is a row vector containing the ending index of each substring that matches the regular expression.
- `tokenExtents` is a cell array containing the starting and ending indices of each substring that matches the `tokens` in the regular expression. A token is a captured part of the regular expression. If the `'once'` mode is used, then this output is a double array.
- `match` is a cell array containing the text for each substring that matches the regular expression. In `'once'` mode, this is a string.
- `tokens` is a cell array of cell arrays of strings that correspond to the tokens in the regular expression. In `'once'` mode, this is a cell array of strings.
- `named` is a structure array containing the named tokens captured in a regular expression. Each named token is assigned a field in the resulting structure array, and each element of the array corresponds to a different match.

If you want only some of the the outputs, you can use the following variant of `regexp`:

```
[o1 o2 ...] = regexp('str','expr', 'p1', 'p2', ...)
```

where `p1` etc. are the names of the outputs (and the order we want the outputs in). As a final variant, you can supply some mode flags to `regexp`

```
[o1 o2 ...] = regexp('str','expr', p1, p2, ..., 'mode1', 'mode2')
```

where acceptable mode flags are:

- `'once'` - only the first match is returned.
- `'matchcase'` - letter case must match (selected by default for `regexp`)
- `'ignorecase'` - letter case is ignored (selected by default for `regexp`)

- 'dotall' - the '.' operator matches any character (default)
- 'dotexceptnewline' - the '.' operator does not match the newline character
- 'stringanchors' - the ^ and \$ operators match at the beginning and end (respectively) of a string.
- 'lineanchors' - the ^ and \$ operators match at the beginning and end (respectively) of a line.
- 'literalspacing' - the space characters and comment characters # are matched as literals, just like any other ordinary character (default).
- 'freespacing' - all spaces and comments are ignored in the regular expression. You must use ' ' and '#' to match spaces and comment characters, respectively.

Note the following behavior differences between MATLABs regexp and FreeMats:

- If you have an old version of pcre installed, then named tokens must use the older <?P<name> syntax, instead of the new <?<name> syntax.
- The pcre library is pickier about named tokens and their appearance in expressions. So, for example, the regexp from the MATLAB manual '(?<first>\\w+)\\s+(?<last>\\w+)(?|last;w+),s+(?|first;w+)'— does not work correctly (as of this writing) because the same named tokens appear multiple times. The workaround is to assign different names to each token, and then collapse the results later.

17.8.2 Example

Some examples of using the `regexp` function

```
--> [start,stop,tokenExtents,match,tokens,named] = regexp('quick down town zoo','(.)own')
start =
    7 12

stop =
   10 15

tokenExtents =
   [1 2 double array] [1 2 double array]

match =
   [down] [town]

tokens =
   [1 1 cell array] [1 1 cell array]

named =
   []
```

17.9 REGEXPREP Regular Expression Replacement Function

17.9.1 Usage

Replaces regular expressions in the provided string. The syntax for its use is

```
outstring = regexprep(instring,pattern,replacement,modes)
```

Here `instring` is the string to be operated on. And `pattern` is a regular expression of the type accepted by `regexp`. For each match, the contents of the matched string are replaced with the replacement text. Tokens in the regular expression can be used in the replacement text using `$N` where `N` is the number of the token to use. You can also specify the same `mode` flags that are used by `regexp`.

17.10 STRCMP String Compare Function

17.10.1 USAGE

Compares two strings for equality. The general syntax for its use is

```
p = strcmp(x,y)
```

where `x` and `y` are two strings. Returns `true` if `x` and `y` are the same size, and are equal (as strings). Otherwise, it returns `false`. In the second form, `strcmp` can be applied to a cell array of strings. The syntax for this form is

```
p = strcmp(cellstra,cellstrb)
```

where `cellstra` and `cellstrb` are cell arrays of a strings to compare. Also, you can also supply a character matrix as an argument to `strcmp`, in which case it will be converted via `cellstr` (so that trailing spaces are removed), before being compared.

17.10.2 Example

The following piece of code compares two strings:

```
--> x1 = 'astring';  
--> x2 = 'bstring';  
--> x3 = 'astring';  
--> strcmp(x1,x2)
```

```
ans =  
0
```

```
--> strcmp(x1,x3)
```

```
ans =  
1
```

Here we use a cell array strings

```
--> x = {'astring', 'bstring', 43, 'astring'}
```

```
x =
  [astring] [bstring] [43] [astring]
```

```
--> p = strcmp(x, 'astring')
```

```
p =
  1 0 0 1
```

Here we compare two cell arrays of strings

```
--> strcmp({'this', 'is', 'a', 'pickle'}, {'what', 'is', 'to', 'pickle'})
```

```
ans =
  0 1 0 1
```

Finally, the case where one of the arguments is a matrix string

```
--> strcmp({'this', 'is', 'a', 'pickle'}, ['peter '; 'piper '; 'hated '; 'pickle'])
```

```
ans =
  0 0 0 0
```

17.11 STRCMPI String Compare Case Insensitive Function

17.11.1 Usage

Compares two strings for equality ignoring case. The general syntax for its use is

```
p = strcmpi(x,y)
```

where *x* and *y* are two strings, or cell arrays of strings. See `strcmp` for more help.

17.12 STRFIND Find Substring in a String

17.12.1 Usage

Searches through a string for a pattern, and returns the starting positions of the pattern in an array. There are two forms for the `strfind` function. The first is for single strings

```
ndx = strfind(string, pattern)
```

the resulting array `ndx` contains the starting indices in `string` for the pattern `pattern`. The second form takes a cell array of strings

```
ndx = strfind(cells, pattern)
```

and applies the search operation to each string in the cell array.

17.12.2 Example

Here we apply `strfind` to a simple string

```
--> a = 'how now brown cow?'
```

```
a =
how now brown cow?
--> b = strfind(a,'ow')
```

```
b =
     2     6    11    16
```

Here we search over multiple strings contained in a cell array.

```
--> a = {'how now brown cow','quick brown fox','coffee anyone?'}
```

```
a =
[how now brown cow] [quick brown fox] [coffee anyone?]
```

```
--> b = strfind(a,'ow')
```

```
b =
[1 4 double array] [9] []
```

17.13 STRNCMP String Compare Function To Length N

17.13.1 USAGE

Compares two strings for equality, but only looks at the first `N` characters from each string. The general syntax for its use is

```
p = strncmp(x,y,n)
```

where `x` and `y` are two strings. Returns `true` if `x` and `y` are each at least `n` characters long, and if the first `n` characters from each string are the same. Otherwise, it returns `false`. In the second form, `strncmp` can be applied to a cell array of strings. The syntax for this form is

```
p = strncmp(cellstra,cellstrb,n)
```

where `cellstra` and `cellstrb` are cell arrays of a strings to compare. Also, you can also supply a character matrix as an argument to `strncmp`, in which case it will be converted via `cellstr` (so that trailing spaces are removed), before being compared.

17.13.2 Example

The following piece of code compares two strings:

```
--> x1 = 'astring';
--> x2 = 'bstring';
--> x3 = 'astring';
--> strncmp(x1,x2,4)
```

```
ans =
    0
```

```
--> strncmp(x1,x3,4)
```

```
ans =
    1
```

Here we use a cell array strings

```
--> x = {'ast','bst',43,'astr'}
```

```
x =
    [ast] [bst] [43] [astr]
```

```
--> p = strncmp(x,'ast',3)
```

```
p =
    1 0 0 1
```

Here we compare two cell arrays of strings

```
--> strncmp({'this','is','a','pickle'},{'think','is','to','pickle'},3)
```

```
ans =
    1 0 0 1
```

Finally, the case where one of the arguments is a matrix string

```
--> strncmp({'this','is','a','pickle'},['peter ','piper ','hated ','pickle'],4);
```

17.14 STRREP String Replace Function

17.14.1 Usage

Replace every occurrence of one string with another. The general syntax for its use is

```
p = strrep(source,find,replace)
```

Every instance of the string `find` in the string `source` is replaced with the string `replace`. Any of `source`, `find` and `replace` can be a cell array of strings, in which case each entry has the replace operation applied.

17.14.2 Example

Here are some examples of the use of `strrep`. First the case where the arguments are simple strings

```
--> strrep('Matlab is great','Matlab','FreeMat')
```

```
ans =  
FreeMat is great
```

And here we have the replace operation for a number of strings:

```
--> strrep({'time is money';'A stitch in time';'No time for games'},'time','money')
```

```
ans =  
[money is money]  
[A stitch in money]  
[No money for games]
```

17.15 STRSTR String Search Function

17.15.1 Usage

Searches for the first occurrence of one string inside another. The general syntax for its use is

```
p = strstr(x,y)
```

where `x` and `y` are two strings. The returned integer `p` indicates the index into the string `x` where the substring `y` occurs. If no instance of `y` is found, then `p` is set to zero.

17.15.2 Example

Some examples of `strstr` in action

```
--> strstr('hello','lo')
```

```
ans =  
4
```

```
--> strstr('quick brown fox','own')
```

```
ans =  
9
```

```
--> strstr('free stuff','lunch')
```

```
ans =  
0
```

17.16 STRTRIM Trim Spaces from a String

17.16.1 Usage

Removes the white-spaces at the beginning and end of a string (or a cell array of strings). See `isspace` for a definition of a white-space. There are two forms for the `strtrim` function. The first is for single strings

```
y = strtrim(strng)
```

where `strng` is a string. The second form operates on a cell array of strings

```
y = strtrim(cellstr)
```

and trims each string in the cell array.

17.16.2 Example

Here we apply `strtrim` to a simple string

```
--> strtrim(' lot of blank spaces ');
```

and here we apply it to a cell array

```
--> strtrim({' space', 'enough ', ' for ', ''})
```

```
ans =
 [space] [enough] [for] []
```

17.17 UPPER Convert strings to upper case

17.17.1 Usage

The `upper` function converts a string to upper case with the syntax

```
y = upper(x)
```

where `x` is a string, in which case all of the lower case characters in `x` (defined as the range `'a'-'z'`) are converted to upper case. Alternately, you can call `upper` with a cell array of strings

```
y = upper(c)
```

in which case each string in the cell array is converted to upper case.

17.17.2 Example

A simple example:

```
--> upper('this Is Strange CAPitalizaTion')
```

```
ans =
 THIS IS STRANGE CAPITALIZATION
```

and a more complex example with a cell array of strings

```
--> upper({'This', 'Is', 'Strange', 'CAPitalizaTion'})
```

```
ans =
```

```
[THIS] [IS] [STRANGE] [CAPITALIZATION]
```


Chapter 18

Transforms/Decompositions

18.1 EIG Eigendecomposition of a Matrix

18.1.1 Usage

Computes the eigendecomposition of a square matrix. The `eig` function has several forms. The first returns only the eigenvalues of the matrix:

```
s = eig(A)
```

The second form returns both the eigenvectors and eigenvalues as two matrices (the eigenvalues are stored in a diagonal matrix):

```
[V,D] = eig(A)
```

where `D` is the diagonal matrix of eigenvalues, and `V` is the matrix of eigenvectors.

Eigenvalues and eigenvectors for asymmetric matrices `A` normally are computed with balancing applied. Balancing is a scaling step that normally improves the quality of the eigenvalues and eigenvectors. In some instances (see the Function Internals section for more details) it is necessary to disable balancing. For these cases, two additional forms of `eig` are available:

```
s = eig(A,'nobalance'),
```

which computes the eigenvalues of `A` only, and does not balance the matrix prior to computation. Similarly,

```
[V,D] = eig(A,'nobalance')
```

recovers both the eigenvectors and eigenvalues of `A` without balancing. Note that the `'nobalance'` option has no effect on symmetric matrices.

FreeMat also provides the ability to calculate generalized eigenvalues and eigenvectors. Similarly to the regular case, there are two forms for `eig` when computing generalized eigenvector (see the Function Internals section for a description of what a generalized eigenvector is). The first returns only the generalized eigenvalues of the matrix pair `A,B`

```
s = eig(A,B)
```

The second form also computes the generalized eigenvectors, and is accessible via

$$[V,D] = \text{eig}(A,B)$$

18.1.2 Function Internals

Recall that v is an eigenvector of A with associated eigenvalue d if

$$Av = dv.$$

This decomposition can be written in matrix form as

$$AV = VD$$

where

$$V = [v_1, v_2, \dots, v_n], D = \text{diag}(d_1, d_2, \dots, d_n).$$

The `eig` function uses the LAPACK class of functions `GEEVX` to compute the eigenvalue decomposition for non-symmetric (or non-Hermitian) matrices A . For symmetric matrices, `SSYEV` and `DSYEV` are used for `float` and `double` matrices (respectively). For Hermitian matrices, `CHEEV` and `ZHEEV` are used for `complex` and `dcomplex` matrices.

For some matrices, the process of balancing (in which the rows and columns of the matrix are pre-scaled to facilitate the search for eigenvalues) is detrimental to the quality of the final solution. This is particularly true if the matrix contains some elements on the order of round off error. See the Example section for an example.

A generalized eigenvector of the matrix pair A, B is simply a vector v with associated eigenvalue d such that

$$Av = dBv,$$

where B is a square matrix of the same size as A . This decomposition can be written in matrix form as

$$AV = BVD$$

where

$$V = [v_1, v_2, \dots, v_n], D = \text{diag}(d_1, d_2, \dots, d_n).$$

For general matrices A and B , the `GGEV` class of routines are used to compute the generalized eigen-decomposition. If however, A and B are both symmetric (or Hermitian, as appropriate), Then FreeMat first attempts to use `SSYGV` and `DSYGV` for `float` and `double` arguments and `CHEGV` and `ZHEGV` for `complex` and `dcomplex` arguments (respectively). These routines requires that B also be positive definite, and if it fails to be, FreeMat will revert to the routines used for general arguments.

18.1.3 Example

Some examples of eigenvalue decompositions. First, for a diagonal matrix, the eigenvalues are the diagonal elements of the matrix.

```
--> A = diag([1.02f,3.04f,1.53f])
```

```
A =
  1.0200      0      0
      0  3.0400      0
      0      0  1.5300
```

```
--> eig(A)
```

```
ans =
  1.0200
  1.5300
  3.0400
```

Next, we compute the eigenvalues of an upper triangular matrix, where the eigenvalues are again the diagonal elements.

```
--> A = [1.0f,3.0f,4.0f;0,2.0f,6.7f;0.0f,0.0f,1.0f]
```

```
A =
  1.0000    3.0000    4.0000
      0    2.0000    6.7000
      0      0    1.0000
```

```
--> eig(A)
```

```
ans =
  1
  2
  1
```

Next, we compute the complete eigenvalue decomposition of a random matrix, and then demonstrate the accuracy of the solution

```
--> A = float(randn(2))
```

```
A =
  0.0501    0.1608
  2.1808   -2.4972
```

```
--> [V,D] = eig(A)
```

```
V =
  0.7754   -0.0599
  0.6314    0.9982
```

```
D =
  0.1811      0
```

```

0 -2.6282

--> A*V - V*D

ans =

1.0e-07 *
1.0431 -0.2980
-1.5646 0

Now, we consider a matrix that requires the nobalance option to compute the eigenvalues and
eigenvectors properly. Here is an example from MATLAB's manual.

--> B = [3,-2,-.9,2*eps;-2,4,1,-eps;-eps/4,eps/2,-1,0;-.5,-.5,.1,1]

B =
3.0000 -2.0000 -0.9000 0.0000
-2.0000 4.0000 1.0000 -0.0000
-0.0000 0.0000 -1.0000 0
-0.5000 -0.5000 0.1000 1.0000

--> [VB,DB] = eig(B)
VB =
0.6153 -0.4176 -0.0000 -0.1530
-0.7881 -0.3261 -0.0000 0.1346
-0.0000 -0.0000 0.0000 -0.9790
0.0189 0.8481 1.0000 -0.0097

DB =
5.5616 0 0 0
0 1.4384 0 0
0 0 1.0000 0
0 0 0 -1.0000

--> B*VB - VB*DB

ans =
0.0000 -0.0000 -0.0000 0.0000
-0.0000 -0.0000 0.0000 0.0000
-0.0000 -0.0000 -0.0000 0
-0.0000 0.0000 0 -0.1081

--> [VN,DN] = eig(B,'nobalance')
VN =
0.6153 -0.4176 0.0000 -0.1528
-0.7881 -0.3261 0.0000 0.1345
-0.0000 -0.0000 -0.0000 -0.9781

```

```

0.0189    0.8481   -1.0000    0.0443

DN =
 5.5616         0         0         0
  0    1.4384         0         0
  0         0    1.0000         0
  0         0         0   -1.0000

--> B*VN - VN*DN

ans =

1.0e-16 *
 8.8818   -1.1102   -1.6722   -1.1102
-8.8818    2.7756    0.1811    0.8327
 0.1718    0.0154    0.0663         0
-0.6939         0         0    0.8327

```

18.2 FFT (Inverse) Fast Fourier Transform Function

18.2.1 Usage

Computes the Discrete Fourier Transform (DFT) of a vector using the Fast Fourier Transform technique. The general syntax for its use is

```
y = fft(x,n,d)
```

where \mathbf{x} is an n -dimensional array of numerical type. Integer types are promoted to the `double` type prior to calculation of the DFT. The argument \mathbf{n} is the length of the FFT, and \mathbf{d} is the dimension along which to take the DFT. If $-\mathbf{n}-$ is larger than the length of \mathbf{x} along dimension \mathbf{d} , then \mathbf{x} is zero-padded (by appending zeros) prior to calculation of the DFT. If \mathbf{n} is smaller than the length of \mathbf{x} along the given dimension, then \mathbf{x} is truncated (by removing elements at the end) to length \mathbf{n} .

If \mathbf{d} is omitted, then the DFT is taken along the first non-singleton dimension of \mathbf{x} . If \mathbf{n} is omitted, then the DFT length is chosen to match of the length of \mathbf{x} along dimension \mathbf{d} .

Note that FFT support on Linux builds requires availability of the FFTW libraries at compile time. On Windows and Mac OS X, single and double precision FFTs are available by default.

18.2.2 Function Internals

The output is computed via

$$y(m_1, \dots, m_{d-1}, l, m_{d+1}, \dots, m_p) = \sum_{k=1}^n x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p) e^{-\frac{2\pi(k-1)l}{n}}.$$

For the inverse DFT, the calculation is similar, and the arguments have the same meanings as

the DFT:

$$y(m_1, \dots, m_{d-1}, l, m_{d+1}, \dots, m_p) = \frac{1}{n} \sum_{k=1}^n x(m_1, \dots, m_{d-1}, k, m_{d+1}, \dots, m_p) e^{\frac{2\pi(k-1)l}{n}}.$$

The FFT is computed using the FFTPack library, available from netlib at <http://www.netlib.org>. Generally speaking, the computational cost for a FFT is (in worst case) $O(n^2)$. However, if n is composite, and can be factored as

$$n = \prod_{k=1}^p m_k,$$

then the DFT can be computed in

$$O\left(n \sum_{k=1}^p m_k\right)$$


operations. If n is a power of 2, then the FFT can be calculated in $O(n \log_2 n)$. The calculations for the inverse FFT are identical.

18.2.3 Example

The following piece of code plots the FFT for a sinusoidal signal:

```
--> t = linspace(0,2*pi,128);
--> x = cos(15*t);
--> y = fft(x);
--> plot(t,abs(y));
```

The resulting plot is:



The FFT can also be taken along different dimensions, and with padding and/or truncation. The following example demonstrates the Fourier Transform being computed along each column, and then along each row.

```
--> A = [2,5;3,6]
```

```
A =
 2 5
```

```
3 6
```

```
--> real(fft(A,[],1))
```

```
ans =
    5 11
   -1 -1
```


```
--> real(fft(A,[],2))
```

```
ans =
    7 -3
    9 -3
```

Fourier transforms can also be padded using the `n` argument. This pads the signal with zeros prior to taking the Fourier transform. Zero padding in the time domain results in frequency interpolation. The following example demonstrates the FFT of a pulse (consisting of 10 ones) with (red line) and without (green circles) padding.

```
--> delta(1:10) = 1;
--> plot((0:255)/256*pi*2,real(fft(delta,256)),'r-');
--> hold on
--> plot((0:9)/10*pi*2,real(fft(delta)),'go');
```

The resulting plot is:



18.3 FFTN N-Dimensional Forward FFT

18.3.1 Usage

Computes the DFT of an N-dimensional numerical array along all dimensions. The general syntax for its use is

```
y = fftn(x)
```

which computes the same-size FFTs for each dimension of `x`. Alternately, you can specify the size vector

```
y = fftn(x,dims)
```

where `dims` is a vector of sizes. The array `x` is zero padded or truncated as necessary in each dimension so that the output is of size `dims`. The `fftn` function is implemented by a sequence of calls to `fft`.

18.4 FFTSHIFT Shift FFT Output

18.4.1 Usage

The `fftshift` function shifts the DC component (zero-frequency) of the output from an FFT to the center of the array. For vectors this means swapping the two halves of the vector. For matrices, the first and third quadrants are swapped. So on for N-dimensional arrays. The syntax for its use is

```
y = fftshift(x).
```

Alternately, you can specify that only one dimension be shifted

```
y = fftshift(x,dim).
```

18.5 HILBERT Hilbert Transform

18.5.1 Usage

The `hilbert` function computes the hilbert transform of the argument vector or matrix. The FreeMat `hilbert` function is compatible with the one from the MATLAB API. This means that the output of the `hilbert` function is the sum of the original function and an imaginary signal containing the Hilbert transform of it. There are two syntaxes for the `hilbert` function. The first is

```
y = hilbert(x)
```

where `x` is real vector or matrix. If `x` is a matrix, then the Hilbert transform is computed along the columns of `x`. The second syntax provides a dimension along which to take the transform.

```
y = hilbert(x,n)
```

where `n` is the dimension along which to apply the transformation.

18.6 IFFTN N-Dimensional Inverse FFT

18.6.1 Usage

Computes the inverse DFT of an N-dimensional numerical array along all dimensions. The general syntax for its use is

```
y = ifftn(x)
```

which computes the same-size inverse FFTs for each dimension of `x`. Alternately, you can specify the size vector

```
y = ifftn(x,dims)
```

where `dims` is a vector of sizes. The array `x` is zero padded or truncated as necessary in each dimension so that the output is of size `dims`. The `ifftn` function is implemented by a sequence of calls to `ifft`.

18.7 IFFTSHIFT Inverse Shift FFT Output

18.7.1 Usage

The `ifftshift` function shifts the DC component (zero-frequency) of the output from the center of the array back to the first position and iseffectively the inverse of `fftshift`. For vectors this means swapping the two halves of the vector. For matrices, the first and third quadrants are swapped. So on for N-dimensional arrays. The syntax for its use is

```
y = ifftshift(x).
```

Alternately, you can specify that only one dimension be shifted

```
y = ifftshift(x,dim).
```

18.8 INV Invert Matrix

18.8.1 Usage

Inverts the argument matrix, provided it is square and invertible. The syntax for its use is

```
y = inv(x)
```

Internally, the `inv` function uses the matrix divide operators. For sparse matrices, a sparse matrix solver is used.

18.8.2 Example

Here we invert some simple matrices

```
--> a = randi(zeros(3),5*ones(3))
```

```
a =
```

```
1 4 3
1 1 3
1 4 5
```

```
--> b = inv(a)
```

```
b =
```

```
1.1667    1.3333   -1.5000
0.3333   -0.3333         0
```

```

-0.5000      0      0.5000

--> a*b

ans =
  1.0000      0      0
  0.0000  1.0000      0
      0      0  1.0000

--> b*a

ans =
  1.0000  0.0000      0
      0  1.0000  0.0000
      0      0  1.0000

```

18.9 LU LU Decomposition for Matrices

18.9.1 Usage

Computes the LU decomposition for a matrix. The form of the command depends on the type of the argument. For full (non-sparse) matrices, the primary form for `lu` is

$$[L,U,P] = \text{lu}(A),$$

where L is lower triangular, U is upper triangular, and P is a permutation matrix such that $L*U = P*A$. The second form is

$$[V,U] = \text{lu}(A),$$

where V is $P'*L$ (a row-permuted lower triangular matrix), and U is upper triangular. For sparse, square matrices, the LU decomposition has the following form:

$$[L,U,P,Q,R] = \text{lu}(A),$$

where A is a sparse matrix of either `double` or `dcomplex` type. The matrices are such that $L*U=P*R*A*Q$, where L is a lower triangular matrix, U is upper triangular, P and Q are permutation vectors and R is a diagonal matrix of row scaling factors. The decomposition is computed using UMFPACK for sparse matrices, and LAPACK for dense matrices.

18.9.2 Example

First, we compute the LU decomposition of a dense matrix.

```

--> a = float([1,2,3;4,5,8;10,12,3])

a =
  1  2  3

```

```

 4 5 8
10 12 3

--> [l,u,p] = lu(a)
l =
 1.0000      0      0
 0.1000  1.0000      0
 0.4000  0.2500  1.0000

u =
10.0000 12.0000  3.0000
      0  0.8000  2.7000
      0      0  6.1250

p =
0 0 1
1 0 0
0 1 0

--> l*u

ans =
10 12 3
 1  2 3
 4  5 8

--> p*a

ans =
10 12 3
 1  2 3
 4  5 8

```

Now we repeat the exercise with a sparse matrix, and demonstrate the use of the permutation vectors.

```

--> a = sparse([1,0,0,4;3,2,0,0;0,0,0,1;4,3,2,4])

a =
 1 1 1
 2 1 3
 4 1 4
 2 2 2
 4 2 3
 4 3 2
 1 4 4
 3 4 1

```

```
4 4 4
--> [l,u,p,q,r] = lu(a)
l =
 1 1 1
 2 2 1
 3 3 1
 4 4 1
u =
 1 1 0.153846
 1 2 0.230769
 2 2 0.4
 1 3 0.307692
 2 3 0.6
 3 3 0.2
 1 4 0.307692
 3 4 0.8
 4 4 1
p =
 4
 2
 1
 3

q =
 3
 2
 1
 4

r =
 1 1 0.2
 2 2 0.2
 3 3 1
 4 4 0.0769231
--> full(l*a)

ans =
 1 0 0 4
 3 2 0 0
 0 0 0 1
 4 3 2 4

--> b = r*a

b =
 1 1 0.2
```

```

2 1 0.6
3 1 0
4 1 0.307692
1 2 0
2 2 0.4
3 2 0
4 2 0.230769
1 3 0
2 3 0
3 3 0
4 3 0.153846
1 4 0.8
2 4 0
3 4 1
4 4 0.307692
--> full(b(p,q))

```

```

ans =
    0.1538    0.2308    0.3077    0.3077
         0    0.4000    0.6000         0
         0         0    0.2000    0.8000
         0         0         0    1.0000

```

18.10 QR QR Decomposition of a Matrix

18.10.1 Usage

Computes the QR factorization of a matrix. The `qr` function has multiple forms, with and without pivoting. The non-pivot version has two forms, a compact version and a full-blown decomposition version. The compact version of the decomposition of a matrix of size $M \times N$ is

$$[q,r] = qr(a,0)$$

where q is a matrix of size $M \times L$ and r is a matrix of size $L \times N$ and $L = \min(N,M)$, and $q*r = a$. The QR decomposition is such that the columns of Q are orthonormal, and R is upper triangular. The decomposition is computed using the LAPACK routine `xgeqrf`, where x is the precision of the matrix. FreeMat supports decompositions of `single` and `double` types.

The second form of the non-pivot decomposition omits the second 0 argument:

$$[q,r] = qr(a)$$

This second form differs from the previous form only for matrices with more rows than columns ($M > N$). For these matrices, the full decomposition is of a matrix Q of size $M \times M$ and a matrix R of size $M \times N$. The full decomposition is computed using the same LAPACK routines as the compact decomposition, but on an augmented matrix $[a \ 0]$, where enough columns are added to form a square matrix.

Generally, the QR decomposition will not return a matrix R with diagonal elements in any specific order. The remaining two forms of the `qr` command utilize permutations of the columns of a so that the diagonal elements of r are in decreasing magnitude. To trigger this form of the decomposition, a third argument is required, which records the permutation applied to the argument a . The compact version is

$$[q,r,e] = qr(a,0)$$

where e is an integer vector that describes the permutation of the columns of a necessary to reorder the diagonal elements of r . This result is computed using the LAPACK routines `(s,d)geqp3`. In the non-compact version of the QR decomposition with pivoting,

$$[q,r,e] = qr(a)$$

the returned matrix e is a permutation matrix, such that $q*r*e' = a$.

18.11 SVD Singular Value Decomposition of a Matrix

18.11.1 Usage

Computes the singular value decomposition (SVD) of a matrix. The `svd` function has three forms. The first returns only the singular values of the matrix:

$$s = svd(A)$$

The second form returns both the singular values in a diagonal matrix S , as well as the left and right eigenvectors.

$$[U,S,V] = svd(A)$$

The third form returns a more compact decomposition, with the left and right singular vectors corresponding to zero singular values being eliminated. The syntax is

$$[U,S,V] = svd(A,0)$$

18.11.2 Function Internals

Recall that σ_i is a singular value of an $M \times N$ matrix A if there exists two vectors u_i , v_i where u_i is of length M , and v_i is of length N and

$$Av_i = \sigma_i u_i$$

and generally

$$A = \sum_{i=1}^K \sigma_i u_i v_i'$$

where K is the rank of A . In matrix form, the left singular vectors u_i are stored in the matrix U as

$$U = [u_1, \dots, u_m], V = [v_1, \dots, v_n]$$

The matrix S is then of size $M \times N$ with the singular values along the diagonal. The SVD is computed using the LAPACK class of functions `GESVD` (Note that this has changed. Previous versions of FreeMat used `GESDD`, which yields a valid, but slightly different choice of the decomposition. Starting in version 4, it was changed to `GESVD` to improve compatibility.

18.11.3 Examples

Here is an example of a partial and complete singular value decomposition.

```
--> A = float(randn(2,3))
```

```
A =  
 -0.9542    1.2478   -0.2295  
  0.3075    1.0686   -0.4849
```

```
--> [U,S,V] = svd(A)
```

```
U =  
 -0.8410   -0.5411  
 -0.5411    0.8410
```

```
S =  
  1.8058     0     0  
     0    0.8549     0
```

```
V =  
  0.3522    0.9064    0.2331  
 -0.9013    0.2614    0.3454  
  0.2521   -0.3317    0.9091
```

```
--> U*S*V'
```

```
ans =  
 -0.9542    1.2478   -0.2295  
  0.3075    1.0686   -0.4849
```

```
--> svd(A)
```

```
ans =  
  1.8058  
  0.8549
```


Chapter 19

Signal Processing Functions

19.1 CONV Convolution Function

19.1.1 Usage

The `conv` function performs a one-dimensional convolution of two vector arguments. The syntax for its use is

$$z = \text{conv}(x,y)$$

where `x` and `y` are vectors. The output is of length `nx + ny - 1`. The `conv` function calls `conv2` to do the calculation. See its help for more details.

19.2 CONV2 Matrix Convolution

19.2.1 Usage

The `conv2` function performs a two-dimensional convolution of matrix arguments. The syntax for its use is

$$Z = \text{conv2}(X,Y)$$

which performs the full 2-D convolution of `X` and `Y`. If the input matrices are of size `[xm,xn]` and `[ym,yn]` respectively, then the output is of size `[xm+ym-1,xn+yn-1]`. Another form is

$$Z = \text{conv2}(\text{hcol},\text{hrow},X)$$

where `hcol` and `hrow` are vectors. In this form, `conv2` first convolves `Y` along the columns with `hcol`, and then convolves `Y` along the rows with `hrow`. This is equivalent to `conv2(hcol(:)*hrow(:)',Y)`.

You can also provide an optional `shape` argument to `conv2` via either

$$\begin{aligned} Z &= \text{conv2}(X,Y,'shape') \\ Z &= \text{conv2}(\text{hcol},\text{hrow},X,'shape') \end{aligned}$$

where `shape` is one of the following strings

- 'full' - compute the full convolution result - this is the default if no `shape` argument is provided.
- 'same' - returns the central part of the result that is the same size as `X`.
- 'valid' - returns the portion of the convolution that is computed without the zero-padded edges. In this situation, `Z` has size `[xm-ym+1,xn-yn+1]` when `xm>=ym` and `xn>=yn`. Otherwise `conv2` returns an empty matrix.

19.2.2 Function Internals

The convolution is computed explicitly using the definition:

$$Z(m, n) = \sum_k \sum_j X(k, j) Y(m - k, n - j)$$

If the full output is requested, then `m` ranges over $0 \leq m < x_m + y_m - 1$ and `n` ranges over $0 \leq n < x_n + y_n - 1$. For the case where `shape` is 'same', the output ranges over $(y_m - 1)/2 \leq m < x_m + (y_m - 1)/2$ and $(y_n - 1)/2 \leq n < x_n + (y_n - 1)/2$.

Chapter 20

Numerical Methods

20.1 ODE45 Numerical Solution of ODEs

20.1.1 Usage

function [t,y] = ode45(f,tspan,y0,options,varargin) function SOL = ode45(f,tspan,y0,options,varargin)
ode45 is a solver for ordinary differential equations and initial value problems. To solve the ODE

$$\begin{aligned}y'(t) &= f(t,y) \\ y(0) &= y_0\end{aligned}$$

over the interval tspan=[t0 t1], you can use ode45. For example, to solve the ode

$$y' = y \quad y(0) = 1$$

whose exact solution is $y(t)=\exp(t)$, over the interval t0=0, t1=3, do

```
--> [t,y]=ode45(@(t,y) y,[0 3],1)
```

Warning: Newly defined variable error shadows a function of the same name. Use clear error to recover a

```
k =  
2
```

```
y =  
1.0000 1.0030
```

```
t =  
1.0e-03 *  
0 3.0000
```

```
k =
```

3

y =
1.0000 1.0030 1.0182

t =
1.0e-02 *
0 0.3000 1.8000

k =
4

y =
1.0000 1.0030 1.0182 1.0975

t =
1.0e-02 *
0 0.3000 1.8000 9.3000

k =
5

y =
1.0000 1.0030 1.0182 1.0975 1.4814

t =
0 0.0030 0.0180 0.0930 0.3930

k =
6

y =
1.0000 1.0030 1.0182 1.0975 1.4814 1.9997

t =
 0 0.0030 0.0180 0.0930 0.3930 0.6930

k =
 7

y =
 1.0000 1.0030 1.0182 1.0975 1.4814 1.9997 2.6993

t =
 0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930

k =
 8

y =
 1.0000 1.0030 1.0182 1.0975 1.4814 1.9997 2.6993 3.6437

t =
 0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930 1.2930

k =
 9

y =

Columns 1 to 8

1.0000 1.0030 1.0182 1.0975 1.4814 1.9997 2.6993 3.6437

Columns 9 to 9

4.9185

t =

Columns 1 to 8

0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930 1.2930

Columns 9 to 9

1.5930

k =

10

y =

Columns 1 to 8

1.0000 1.0030 1.0182 1.0975 1.4814 1.9997 2.6993 3.6437

Columns 9 to 10

4.9185 6.6392

t =

Columns 1 to 8

0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930 1.2930

Columns 9 to 10

1.5930 1.8930

k =

11

y =

Columns 1 to 8

1.0000 1.0030 1.0182 1.0975 1.4814 1.9997 2.6993 3.6437

Columns 9 to 11

4.9185 6.6392 8.9620

t =

Columns 1 to 8

0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930 1.2930

Columns 9 to 11

1.5930 1.8930 2.1930

k =

12

y =

Columns 1 to 8

1.0000 1.0030 1.0182 1.0975 1.4814 1.9997 2.6993 3.6437

Columns 9 to 12

4.9185 6.6392 8.9620 12.0975

t =

Columns 1 to 8

0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930 1.2930

Columns 9 to 12

1.5930 1.8930 2.1930 2.4930

k =

13

y =

Columns 1 to 8

1.0000	1.0030	1.0182	1.0975	1.4814	1.9997	2.6993	3.6437
--------	--------	--------	--------	--------	--------	--------	--------

Columns 9 to 13

4.9185	6.6392	8.9620	12.0975	16.3299
--------	--------	--------	---------	---------

t =

Columns 1 to 8

0	0.0030	0.0180	0.0930	0.3930	0.6930	0.9930	1.2930
---	--------	--------	--------	--------	--------	--------	--------

Columns 9 to 13

1.5930	1.8930	2.1930	2.4930	2.7930
--------	--------	--------	--------	--------

k =

14

y =

Columns 1 to 8

1.0000	1.0030	1.0182	1.0975	1.4814	1.9997	2.6993	3.6437
--------	--------	--------	--------	--------	--------	--------	--------

Columns 9 to 14

4.9185	6.6392	8.9620	12.0975	16.3299	20.0854
--------	--------	--------	---------	---------	---------

t =

Columns 1 to 8

0	0.0030	0.0180	0.0930	0.3930	0.6930	0.9930	1.2930
---	--------	--------	--------	--------	--------	--------	--------

Columns 9 to 14

1.5930	1.8930	2.1930	2.4930	2.7930	3.0000
--------	--------	--------	--------	--------	--------

t =

Columns 1 to 8

0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930 1.2930

Columns 9 to 14

1.5930 1.8930 2.1930 2.4930 2.7930 3.0000

y =

1.0000
1.0030
1.0182
1.0975
1.4814
1.9997
2.6993
3.6437
4.9185
6.6392
8.9620
12.0975
16.3299
20.0854

If you want a dense output (i.e., an output that also contains an interpolating spline), use instead

```
--> SOL=ode45(@(t,y) y,[0 3],1)
```

Warning: Newly defined variable error shadows a function of the same name. Use clear error to recover a

k =

2

y =

1.0000 1.0030

t =

1.0e-03 *
0 3.0000

k =

3

y =
1.0000 1.0030 1.0182

t =
1.0e-02 *
0 0.3000 1.8000

k =
4

y =
1.0000 1.0030 1.0182 1.0975

t =
1.0e-02 *
0 0.3000 1.8000 9.3000

k =
5

y =
1.0000 1.0030 1.0182 1.0975 1.4814

t =
0 0.0030 0.0180 0.0930 0.3930

k =
6

y =
1.0000 1.0030 1.0182 1.0975 1.4814 1.9997

t =
 0 0.0030 0.0180 0.0930 0.3930 0.6930

k =
 7

y =
 1.0000 1.0030 1.0182 1.0975 1.4814 1.9997 2.6993

t =
 0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930

k =
 8

y =
 1.0000 1.0030 1.0182 1.0975 1.4814 1.9997 2.6993 3.6437

t =
 0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930 1.2930

k =
 9

y =

Columns 1 to 8

1.0000 1.0030 1.0182 1.0975 1.4814 1.9997 2.6993 3.6437

Columns 9 to 9

4.9185

t =

Columns 1 to 8

0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930 1.2930

Columns 9 to 9

1.5930

k =

10

y =

Columns 1 to 8

1.0000 1.0030 1.0182 1.0975 1.4814 1.9997 2.6993 3.6437

Columns 9 to 10

4.9185 6.6392

t =

Columns 1 to 8

0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930 1.2930

Columns 9 to 10

1.5930 1.8930

k =

11

y =

Columns 1 to 8

1.0000 1.0030 1.0182 1.0975 1.4814 1.9997 2.6993 3.6437

Columns 9 to 11

4.9185 6.6392 8.9620

t =

Columns 1 to 8

0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930 1.2930

Columns 9 to 11

1.5930 1.8930 2.1930

k =

12

y =

Columns 1 to 8

1.0000 1.0030 1.0182 1.0975 1.4814 1.9997 2.6993 3.6437

Columns 9 to 12

4.9185 6.6392 8.9620 12.0975

t =

Columns 1 to 8

0 0.0030 0.0180 0.0930 0.3930 0.6930 0.9930 1.2930

Columns 9 to 12

1.5930 1.8930 2.1930 2.4930

k =

13

y =

Columns 1 to 8

1.0000	1.0030	1.0182	1.0975	1.4814	1.9997	2.6993	3.6437
--------	--------	--------	--------	--------	--------	--------	--------

Columns 9 to 13

4.9185	6.6392	8.9620	12.0975	16.3299
--------	--------	--------	---------	---------

t =

Columns 1 to 8

0	0.0030	0.0180	0.0930	0.3930	0.6930	0.9930	1.2930
---	--------	--------	--------	--------	--------	--------	--------

Columns 9 to 13

1.5930	1.8930	2.1930	2.4930	2.7930
--------	--------	--------	--------	--------

k =

14

y =

Columns 1 to 8

1.0000	1.0030	1.0182	1.0975	1.4814	1.9997	2.6993	3.6437
--------	--------	--------	--------	--------	--------	--------	--------

Columns 9 to 14

4.9185	6.6392	8.9620	12.0975	16.3299	20.0854
--------	--------	--------	---------	---------	---------

t =

Columns 1 to 8

0	0.0030	0.0180	0.0930	0.3930	0.6930	0.9930	1.2930
---	--------	--------	--------	--------	--------	--------	--------

Columns 9 to 14

1.5930	1.8930	2.1930	2.4930	2.7930	3.0000
--------	--------	--------	--------	--------	--------

```
SOL =
  x: 1 14 double array
  y: 1 14 double array
  xe:
  ye:
  ie:
  solver: generic_ode_solver
  interpolant: 1 1 functionpointer array
  idata: 1 1 struct array
```

You can view the result using

```
plot(0:0.01:3,deval(SOL,0:0.01:3))
```

You will notice that this function is available for "every" value of t, while

```
plot(t,y,'o')
```

is only available at a few points.

The optional argument 'options' is a structure. It may contain any of the following fields:

'AbsTol' - Absolute tolerance, default is 1e-6. 'RelTol' - Relative tolerance, default is 1e-3.
 'MaxStep' - Maximum step size, default is (tspan(2)-tspan(1))/10 'InitialStep' - Initial step size,
 default is maxstep/100 'Stepper' - To override the default Fehlberg integrator 'Events' - To provide
 an event function 'Projection' - To provide a projection function

The varargin is ignored by this function, but is passed to all your callbacks, i.e., f, the event function and the projection function.

==Event Function==

The event function can be used to detect situations where the integrator should stop, possibly because the right-hand-side has changed, because of a collision, etc...

An event function should look like

```
function [val,isterminal,direction]=event(t,y,...)
```

The return values are:

val - the value of the event function. isterminal - whether or not this event should cause termination of the integrator. direction - 1=upcrossings only matter, -1=downcrossings only, 0=both.

== Projection function ==

For geometric integration, you can provide a projection function which will be called after each time step. The projection function has the following signature:

```
function yn=project(t,yn,...);
```

If the output yn is very different from the input yn, the quality of interpolation may decrease.

Chapter 21

Operating System Functions

21.1 CD Change Working Directory Function

21.1.1 Usage

Changes the current working directory to the one specified as the argument. The general syntax for its use is

```
cd('dirname')
```

but this can also be expressed as

```
cd 'dirname'
```

or

```
cd dirname
```

Examples of all three usages are given below. Generally speaking, `dirname` is any string that would be accepted by the underlying OS as a valid directory name. For example, on most systems, `'.'` refers to the current directory, and `'..'` refers to the parent directory. Also, depending on the OS, it may be necessary to “escape” the directory separators. In particular, if directories are separated with the backwards-slash character `'\'`, then the path specification must use double-slashes `'\\'`. Note: to get file-name completion to work at this time, you must use one of the first two forms of the command.

21.1.2 Example

The `pwd` command returns the current directory location. First, we use the simplest form of the `cd` command, in which the directory name argument is given unquoted.

```
--> pwd
```

```
ans =  
/home/basu/dev/branches/FreeMat4/help/tmp
```

```
--> cd ..
--> pwd
```

```
ans =
/home/basu/dev/branches/FreeMat4/help
```

Next, we use the “traditional” form of the function call, using both the parenthesis and a variable to store the quoted string.

```
--> a = pwd;
--> cd(a)
--> pwd
```

```
ans =
/home/basu/dev/branches/FreeMat4/help/tmp
```

21.2 COPYFILE Copy Files

21.2.1 Usage

Copies a file or files from one location to another. There are several syntaxes for this function that are acceptable:

```
copyfile(file_in,file_out)
```

copies the file from `file_in` to `file_out`. Also, the second argument can be a directory name:

```
copyfile(file_in,directory_out)
```

in which case `file_in` is copied into the directory specified by `directory_out`. You can also use `copyfile` to copy entire directories as in

```
copyfile(dir_in,dir_out)
```

in which case the directory contents are copied to the destination directory (which is created if necessary). Finally, the first argument to `copyfile` can contain wildcards

```
copyfile(pattern,directory_out)
```

in which case all files that match the given pattern are copied to the output directory. Note that to remain compatible with the MATLAB API, this function will delete/replace destination files that already exist, unless they are marked as read-only. If you want to force the copy to succeed, you can append a `'f'` argument to the `copyfile` function:

```
copyfile(arg1,arg2,'f')
```

or equivalently

```
copyfile arg1 arg2 f
```

21.3 DELETE Delete a File

21.3.1 Usage

Deletes a file. The general syntax for its use is

```
delete('filename')
```

or alternately

```
delete filename
```

which removes the file described by `filename` which must be relative to the current path.

21.4 DIR List Files Function

21.4.1 Usage

In some versions of FreeMat (pre 3.1), the `dir` function was aliased to the `ls` function. Starting with version 3.1, the `dir` function has been rewritten to provide compatibility with MATLAB. The general syntax for its use is

```
dir
```

in which case, a listing of the files in the current directory are output to the console. Alternately, you can specify a target via

```
dir('name')
```

or using the string syntax

```
dir name
```

If you want to capture the output of the `dir` command, you can assign the output to an array

```
result = dir('name')
```

(or you can omit `'name'` to get a directory listing of the current directory. The resulting array `result` is a structure array containing the fields:

- `name` the filename as a string
- `date` the modification date and time stamp as a string
- `bytes` the size of the file in bytes as a `uint64`
- `isdir` a logical that is 1 if the file corresponds to a directory.

Note that `'name'` can also contain wildcards (e.g., `dir *.m` to get a listing of all FreeMat scripts in the current directory.

21.5 DIRSEP Director Seperator

21.5.1 Usage

Returns the directory separator character for the current platform. The general syntax for its use is

```
y = dirsep
```

This function can be used to build up paths (or see `fullfile` for another way to do this).

21.6 FILEPARTS Extract Filename Parts

21.6.1 Usage

The `fileparts` takes a filename, and returns the path, filename, extension, and (for MATLAB-compatibility) an empty version number of the file. The syntax for its use is

```
[path,name,extension,version] = fileparts(filename)
```

where `filename` is a string containing the description of the file, and `path` is the path to the file,

21.7 FULLFILE Build a Full Filename From Pieces

21.7.1 Usage

The `fullfile` routine constructs a full filename from a set of pieces, namely, directory names and a filename. The syntax is:

```
x = fullfile(dir1,dir2,...,dirn,filename)
```

where each of the arguments are strings. The `fullfile` function is equivalent to `[dir1 dirsep dir2 dirsep`

21.7.2 Example

```
--> fullfile('path','to','my','file.m')
```

```
ans =  
path/to/my/file.m
```

21.8 GETPATH Get Current Search Path

21.8.1 Usage

Returns a `string` containing the current FreeMat search path. The general syntax for its use is

```
y = getpath
```

The delimiter between the paths depends on the system being used. For Win32, the delimiter is a semicolon. For all other systems, the delimiter is a colon.

21.8.2 Example

The `getpath` function is straightforward.

```
--> getpath
```

```
ans =  
    Empty array 1 0
```

21.9 LS List Files Function

21.9.1 Usage

Lists the files in a directory or directories. The general syntax for its use is

```
ls('dirname1','dirname2',...,'dirnameN')
```

but this can also be expressed as

```
ls 'dirname1' 'dirname2' ... 'dirnameN'
```

or

```
ls dirname1 dirname2 ... dirnameN
```

For compatibility with some environments, the function `dir` can also be used instead of `ls`. Generally speaking, `dirname` is any string that would be accepted by the underlying OS as a valid directory name. For example, on most systems, `'.'` refers to the current directory, and `'..'` refers to the parent directory. Also, depending on the OS, it may be necessary to “escape” the directory separators. In particular, if directories are separated with the backwards-slash character `'\'`, then the path specification must use double-slashes `'\\'`. Two points worth mentioning about the `ls` function:

- To get file-name completion to work at this time, you must use one of the first two forms of the command.
- If you want to capture the output of the `ls` command, use the `system` function instead.

21.9.2 Example

First, we use the simplest form of the `ls` command, in which the directory name argument is given unquoted.

```
--> ls m*.m
```

Next, we use the “traditional” form of the function call, using both the parenthesis and the quoted string.

```
--> ls('m*.m')
```

In the third version, we use only the quoted string argument without parenthesis.

```
--> ls 'm*.m'
```

21.10 MKDIR Make Directory

21.10.1 Usage

Creates a directory. The general syntax for its use is

```
mkdir('dirname')
```

which creates the directory `dirname` if it does not exist. The argument `dirname` can be either a relative path or an absolute path. For compatibility with MATLAB, the following syntax is also allowed

```
mkdir('parentdir','dirname')
```

which attempts to create a directory `dirname` in the directory given by `parentdir`. However, this simply calls `mkdir([parentdir dirsep dirname])`, and if this is not the required behavior, please file an enhancement request to have it changed. Note that `mkdir` returns a logical 1 if the call succeeded, and a logical 0 if not.

21.11 PWD Print Working Directory Function

21.11.1 Usage

Returns a `string` describing the current working directory. The general syntax for its use is

```
y = pwd
```

21.11.2 Example

The `pwd` function is fairly straightforward.

```
--> pwd
```

```
ans =
```

```
/home/basu/dev/branches/FreeMat4/help/tmp
```

21.12 RMDIR Remove Directory

21.12.1 Usage

Deletes a directory. The general syntax for its use is

```
rmdir('dirname')
```

which removes the directory `dirname` if it is empty. If you want to delete the directory and all subdirectories and files in it, use the syntax

```
rmdir('dirname','s')
```

21.13 SETPATH Set Current Search Path

21.13.1 Usage

Changes the current FreeMat search path. The general syntax for its use is

```
setpath(y)
```

where `y` is a **string** containing a delimited list of directories to be searched for M files and libraries. The delimiter between the paths depends on the system being used. For Win32, the delimiter is a semicolon. For all other systems, the delimiter is a colon.

@Example The `setpath` function is straightforward.

```
--> getpath
```

```
ans =
```

```
Empty array 1 0
```

```
--> setpath('/usr/local/FreeMat/MFiles:/localhome/basu/MFiles')
```

```
--> getpath
```

```
ans =
```

```
/usr/local/FreeMat/MFiles:/localhome/basu/MFiles
```

21.14 SYSTEM Call an External Program

21.14.1 Usage

The `system` function allows you to call an external program from within FreeMat, and capture the output. The syntax of the `system` function is

```
y = system(cmd)
```

where `cmd` is the command to execute. The return array `y` is of type **cell-array**, where each entry in the array corresponds to a line from the output.

21.14.2 Example

Here is an example of calling the `ls` function (the list files function under Un*x-like operating system).

```
--> y = system('ls')
```

```
y =
```

```
Columns 1 to 3
```

```
[chain1.m] [chain2.m] [chain3.m]
```

Columns 4 to 6

[conv2_eqn1.png] [conv2_eqn.aux] [conv2_eqn.dvi]

Columns 7 to 9

[conv2_eqn.log] [conv2_eqn.tex] [cumprod_eqn1.png]

Columns 10 to 12

[cumprod_eqn.aux] [cumprod_eqn.dvi] [cumprod_eqn.log]

Columns 13 to 15

[cumprod_eqn.tex] [cumsum_eqn1.png] [cumsum_eqn.aux]

Columns 16 to 18

[cumsum_eqn.dvi] [cumsum_eqn.log] [cumsum_eqn.tex]

Columns 19 to 21

[do_jit_test.m] [eig_eqn1.png] [eig_eqn2.png]

Columns 22 to 24

[eig_eqn3.png] [eig_eqn4.png] [eig_eqn5.png]

Columns 25 to 27

[eig_eqn6.png] [eig_eqn.aux] [eig_eqn.dvi]

Columns 28 to 30

[eig_eqn.log] [eig_eqn.tex] [evenoddttest.m]

Columns 31 to 33

[fft1.jpg] [fft1.png] [fft2.jpg]

Columns 34 to 36

[fft2.png] [fft_eqn1.png] [fft_eqn2.png]

Columns 37 to 39

[fft_eqn3.png] [fft_eqn4.png] [fft_eqn.aux]

Columns 40 to 42

[fft_eqn.dvi] [fft_eqn.log] [fft_eqn.tex]

Columns 43 to 45

[find_eqn1.png] [find_eqn.aux] [find_eqn.dvi]

Columns 46 to 48

[find_eqn.log] [find_eqn.tex] [help]

Columns 49 to 51

[jit_test001.m] [jit_test002.m] [jit_test003.m]

Columns 52 to 54

[jit_test004.m] [jit_test005.m] [jit_test006.m]

Columns 55 to 57

[jit_test007.m] [jit_test008.m] [jit_test009.m]

Columns 58 to 60

[jit_test010.m] [jit_test011.m] [jit_test012.m]

Columns 61 to 63

[jit_test013.m] [jit_test014.m] [jit_test015.m]

Columns 64 to 66

[jit_test016.m] [jit_test017.m] [jit_test018.m]

Columns 67 to 69

[jit_test019.m] [jit_test020.m] [jit_test021.m]

Columns 70 to 72

[jit_test022.m] [jit_test023.m] [jit_test024.m]

Columns 73 to 75

[loadsave.dat] [max_eqn1.png] [max_eqn2.png]

Columns 76 to 78

[max_eqn3.png] [max_eqn.aux] [max_eqn.dvi]

Columns 79 to 81

[max_eqn.log] [max_eqn.tex] [mean_eqn1.png]

Columns 82 to 84

[mean_eqn.aux] [mean_eqn.dvi] [mean_eqn.log]

Columns 85 to 87

[mean_eqn.tex] [min_eqn1.png] [min_eqn2.png]

Columns 88 to 90

[min_eqn3.png] [min_eqn.aux] [min_eqn.dvi]

Columns 91 to 93

[min_eqn.log] [min_eqn.tex] [nargintest.m]

Columns 94 to 96

[nargouttest.m] [source_test] [source_test_script.m]

Columns 97 to 99

[test_assignin1.m] [test_assignin2.m] [test_bin2int1.m]

Columns 100 to 102

[test_bitand1.m] [test_bitor1.m] [test_bitxor1.m]

Columns 103 to 105

[test.bmp] [test_builtin1.m] [test_cell1.m]

Columns 106 to 108

[test_clear1.m] [test_clear2.m] [test_conv2_1.m]

Columns 109 to 111

[test_ctype1.m] [test.dat] [test_dcomplex1.m]

Columns 112 to 114

[test_diag1.m] [test_diag2.m] [test_diag3.m]

Columns 115 to 117

[test_diag4.m] [test_diag5.m] [test_dlmread1.m]

Columns 118 to 120

[test_eig1.m] [test_eig2.m] [test_eig3.m]

Columns 121 to 123

[test_eig4.m] [test_eig5.m] [test_eig6.m]

Columns 124 to 126

[test_empty.m] [test_eps1.m] [test_error1.m]

Columns 127 to 129

[test_eval1.m] [test_eval2.m] [test_eval3.m]

Columns 130 to 132

[test_evalin1.m] [test_evalin2.m] [test_exist1.m]

Columns 133 to 135

[test_exist2.m] [test_feval1.m] [test_fieldnames1.m]

Columns 136 to 138

[test_imwrite_imread.m] [test_int2bin1.m] [test_isset1.m]

Columns 139 to 141

[test_jit.m] [test_lasterr1.m] [test_load1.m]

Columns 142 to 144

[test_lu1.m] [test_lu2.m] [test_nargin1.m]

Columns 145 to 147

[test_nargin2.m] [test_permute1.m] [test_permute2.m]

Columns 148 to 150

[test_repmat1.m] [test_repmat2.m] [test_repmat3.m]

Columns 151 to 153

[test_save1.m] [test_source.m] [test_sparse20.m]

Columns 154 to 156

[test_sparse21.m] [test_sparse45.m] [test_sparse58.m]

Columns 157 to 159

[test_sparse68.m] [test_sparse70.m] [test_sparse74.m]

Columns 160 to 162

[test_sparse75.m] [testtext] [test_typeof1.m]

Columns 163 to 165

[test_typeof2.m] [test_typeof3.m] [test_typeof4.m]

Columns 166 to 168

[test_typeof5.m] [test_typeof6.m] [test_uint64_1.m]

Columns 169 to 169

[toolbox]

--> y{1}

ans =
chain1.m

Chapter 22

Optimization and Curve Fitting

22.1 FITFUN Fit a Function

22.1.1 Usage

Fits n (non-linear) functions of m variables using least squares and the Levenberg-Marquardt algorithm. The general syntax for its usage is

```
[xopt,yopt] = fitfun(fcn,xinit,y,weights,tol,params...)
```

Where `fcn` is the name of the function to be fit, `xinit` is the initial guess for the solution (required), `y` is the right hand side, i.e., the vector `y` such that:

$$xopt = \arg \min_x \|\text{diag}(weights) * (f(x) - y)\|_2^2,$$

the output `yopt` is the function `fcn` evaluated at `xopt`. The vector `weights` must be the same size as `y`, and contains the relative weight to assign to an error in each output value. Generally, the i th weight should reflect your confidence in the i th measurement. The parameter `tol` is the tolerance used for convergence. The function `fcn` must return a vector of the same size as `y`, and `params` are passed to `fcn` after the argument `x`, i.e.,

$$y = fcn(x, param1, param2, \dots).$$

Note that both `x` and `y` (and the output of the function) must all be real variables. Complex variables are not handled yet.

```
fitfunc_func1.m
function y = fitfunc_func1(init,junk)
c = 2;
d = 0;
init2 = [c,d];
y2 = [1:100]*3;
weights2 = y2*0+1;
tol2 = 1.e-08;
```

```

junk2 = rand(100);
[xopt2,yopt2] = fitfun('fitfunc_func2',init2,y2,weights2,tol2,junk2);
a = init(1);
b = init(2);
x = [1:100];
y = a*x + b;

```

```

fitfunc_func2.m
function y = fitfunc_func2(init,junk);
a = init(1);
b = init(2);
x = [1:100];
y = a*x+b;

```

22.2 GAUSFIT Gaussian Curve Fit

22.2.1 Usage

The `gausfit` routine has the following syntax

```
[mu,sigma,dc,gain,yhat] = gausfit(t,y,w,mug,sigmag,dcg,gaing).
```

where the required inputs are

- `t` - the values of the independent variable (e.g., time samples)
- `y` - the values of the dependent variable (e.g., $f(t)$)

The following inputs are all optional, and default values are available for each of them.

- `w` - the weights to use in the fitting (set to ones if omitted)
- `mug` - initial estimate of the mean
- `sigmag` - initial estimate of the sigma (standard deviation)
- `dcg` - initial estimate of the DC value
- `gaing` - initial estimate of the gain

The fit is of the form $yhat = gain * \exp((t - mu)^2 / (2 * sigma^2)) + dc$. The outputs are

- `mu` - the mean of the fit
- `sigma` - the sigma of the fit
- `dc` - the dc term of the fit
- `gain` - the gain of the gaussian fit
- `yhat` - the output samples (the Gaussian fits)


Because the fit is nonlinear, a good initial guess is critical to convergence of the solution. Thus, you can supply initial guesses for each of the parameters using the `mu`, `sigmag`, `dcg`, `gaing` arguments. Any arguments not supplied are estimated using a simple algorithm. In particular, the DC value is estimated by taking the minimum value from the vector `y`. The gain is estimated from the range of `y`. The mean and standard deviation are estimated using the first and second order moments of `y`. This function uses `fitfun`.

22.2.2 Example

Suppose we want to fit a cycle of a cosine using a Gaussian shape.

```
--> t = linspace(-pi,pi);
--> y = cos(t);
--> [mu,sigma,dc,gaing,yhat] = gausfit(t,y);
--> plot(t,y,'rx',t,yhat,'g-');
```

Which results in the following plot



22.3 INTERPLIN1 Linear 1-D Interpolation

22.3.1 Usage

Given a set of monotonically increasing `x` coordinates and a corresponding set of `y` values, performs simple linear interpolation to a new set of `x` coordinates. The general syntax for its usage is

```
yi = interplin1(x1,y1,xi)
```

where `x1` and `y1` are vectors of the same length, and the entries in `x1` are monotonically increasing. The output vector `yi` is the same size as the input vector `xi`. For each element of `xi`, the values in `y1` are linearly interpolated. For values in `xi` that are outside the range of `x1` the default value returned is `nan`. To change this behavior, you can specify the extrapolation flag:

```
yi = interplin1(x1,y1,xi,extrapflag)
```

Valid options for `extrapflag` are:

- `'nan'` - extrapolated values are tagged with `nans`

- 'zero' - extrapolated values are set to zero
- 'endpoint' - extrapolated values are set to the endpoint values
- 'extrap' - linear extrapolation is performed


The `x1` and `xi` vectors must be real, although complex types are allowed for `y1`.

22.3.2 Example

Here is an example of simple linear interpolation with the different extrapolation modes. We start with a fairly coarse sampling of a cosine.

```
--> x = linspace(-pi*7/8,pi*7/8,15);
--> y = cos(x);
--> plot(x,y,'ro');
```


which is shown here



Next, we generate a finer sampling over a slightly broader range (in this case $[-\pi, \pi]$). First, we demonstrate the 'nan' extrapolation method

```
--> xi = linspace(-4,4,100);
--> yi_nan = interplin1(x,y,xi,'nan');
--> yi_zero = interplin1(x,y,xi,'zero');
--> yi_endpoint = interplin1(x,y,xi,'endpoint');
--> yi_extrap = interplin1(x,y,xi,'extrap');
--> plot(x,y,'ro',xi,yi_nan,'g-x',xi,yi_zero,'g-x',xi,yi_endpoint,'g-x',xi,yi_extrap,'g-x');
```

which is shown here



22.4 POLY Convert Roots To Polynomial Coefficients

22.4.1 Usage

This function calculates the polynomial coefficients for given roots

```
p = poly(r)
```

when **r** is a vector, is a vector whose elements are the coefficients of the polynomial whose roots are the elements of **r**. Alternately, you can provide a matrix

```
p = poly(A)
```

when **A** is an $N \times N$ square matrix, is a row vector with $N+1$ elements which are the coefficients of the characteristic polynomial, $\det(\lambda \text{eye}(\text{size}(A)) - A)$.

Contributed by Paulo Xavier Candeias under GPL.

22.4.2 Example

Here are some examples of the use of `poly`

```
--> A = [1,2,3;4,5,6;7,8,0]
```

```
A =
 1 2 3
 4 5 6
 7 8 0
```

```
--> p = poly(A)
```

```
p =
 1.0000 -6.0000 -72.0000 -27.0000
```

```
--> r = roots(p)
```

```
r =
 12.1229
 -5.7345
 -0.3884
```

22.5 POLYDER Polynomial Coefficient Differentiation

22.5.1 Usage

The `polyder` function returns the polynomial coefficients resulting from differentiation of polynomial `p`. The syntax for its use is either

```
pder = polyder(p)
```

for the derivitave of polynomial `p`, or

```
convp1p2der = polyder(p1,p2)
```

for the derivitave of polynomial `conv(p1,p2)`, or still

```
[nder,dder] = polyder(n,d)
```

for the derivative of polynomial `n/d` (`nder` is the numerator and `dder` is the denominator). In all cases the polynomial coefficients are assumed to be in decreasing degree. Contributed by Paulo Xavier Candeias under GPL

22.5.2 Example

Here are some examples of the use of `polyder`

```
--> polyder([2,3,4])
```

```
ans =
 4 3
```

```
--> polyder([2,3,4],7)
```

```
ans =
 28 21
```

```
--> [n,d] = polyder([2,3,4],5)
```

```
n =
 -20 -15
```

```
d =
 25
```

22.6 POLYFIT Fit Polynomial To Data

22.6.1 Usage

The `polyfit` routine has the following syntax

```
p = polyfit(x,y,n)
```

where `x` and `y` are vectors of the same size, and `n` is the degree of the approximating polynomial. The resulting vector `p` forms the coefficients of the optimal polynomial (in descending degree) that fit `y` with `x`.

22.6.2 Function Internals

The `polyfit` routine finds the approximating polynomial

$$p(x) = p_1x^n + p_2x^{n-1} + \cdots + p_nx + p_{n+1}$$

such that

$$\sum_i (p(x_i) - y_i)^2$$


is minimized. It does so by forming the Vandermonde matrix and solving the resulting set of equations using the backslash operator. Note that the Vandermonde matrix can become poorly conditioned with large `n` quite rapidly.

22.6.3 Example

A classic example from Edwards and Penny, consider the problem of approximating a sinusoid with a polynomial. We start with a vector of points evenly spaced on the unit interval, along with a vector of the sine of these points.

```
--> x = linspace(0,1,20);
--> y = sin(2*pi*x);
--> plot(x,y,'r-')
```

The resulting plot is shown here




Next, we fit a third degree polynomial to the sine, and use `polyval` to plot it

```
--> p = polyfit(x,y,3)

p =
    21.9170   -32.8756    11.1897   -0.1156

--> f = polyval(p,x);
--> plot(x,y,'r-',x,f,'ko');
```

The resulting plot is shown here



Increasing the order improves the fit, as

```
--> p = polyfit(x,y,11)


p =

Columns 1 to 8
    12.4644   -68.5541   130.0555   -71.0940   -38.2814   -14.1222    85.1018   -0.5642

Columns 9 to 12
   -41.2861   -0.0029    6.2832   -0.0000

--> f = polyval(p,x);
--> plot(x,y,'r-',x,f,'ko');
```

The resulting plot is shown here



22.7 POLYINT Polynomial Coefficient Integration

22.7.1 Usage

The `polyint` function returns the polynomial coefficients resulting from integration of polynomial `p`. The syntax for its use is either

```
pint = polyint(p,k)
```

or, for a default `k = 0`,

```
pint = polyint(p);
```

where `p` is a vector of polynomial coefficients assumed to be in decreasing degree and `k` is the integration constant. Contributed by Paulo Xavier Candeias under GPL

22.7.2 Example

Here is are some examples of the use of `polyint`.

```
--> polyint([2,3,4])
```

```
ans =
    0.6667    1.5000    4.0000         0
```

And

```
--> polyint([2,3,4],5)
```

```
ans =
    0.6667    1.5000    4.0000    5.0000
```

22.8 POLYVAL Evaluate Polynomial Fit at Selected Points

22.8.1 Usage

The `polyval` routine has the following syntax

```
y = polyval(p,x)
```

where **p** is a vector of polynomial coefficients, in decreasing degree (as generated by `polyfit`, for example). If **x** is a matrix, the polynomial is evaluated in the matrix sense (in which case **x** must be square).

22.8.2 Function Internals

The polynomial is evaluated using a recursion method. If the polynomial is

$$p(x) = p_1x^n + p_2x^{n-1} + \cdots + p_nx + p_{n+1}$$

then the calculation is performed as

$$p(x) = ((p_1)x + p_2)x + p_3$$

22.8.3 Example

Here is a plot of x^3 generated using `polyval`

```
--> p = [1 0 0 0]
```


```
p =
 1 0 0 0
```

```
--> x = linspace(-1,1);
```

```
--> y = polyval(p,x);
```

```
--> plot(x,y,'r-')
```

Here is the resulting plot



22.9 ROOTS Find Roots of Polynomial

22.9.1 Usage

The `roots` routine will return a column vector containing the roots of a polynomial. The general syntax is

```
z = roots(p)
```

where `p` is a vector containing the coefficients of the polynomial ordered in descending powers.

22.9.2 Function Internals

Given a vector

$$[p_1, p_2, \dots, p_n]$$

which describes a polynomial

$$p_1x^{n-1} + p_2x^{n-2} + \dots + p_n$$

we construct the companion matrix (which has a characteristic polynomial matching the polynomial described by `p`), and then find the eigenvalues of it (which are the roots of its characteristic polynomial), and which are also the roots of the polynomial of interest. This technique for finding the roots is described in the help page for `roots` on the Mathworks website.

22.9.3 Example

Here is an example of finding the roots to the polynomial

$$x^3 - 6x^2 - 72x - 27$$

```
--> roots([1 -6 -72 -27])
```

```
ans =  
    12.1229  
    -5.7345  
    -0.3884
```


Chapter 23

Handle-Based Graphics

23.1 AXES Create Handle Axes

23.1.1 Usage

This function has three different syntaxes. The first takes no arguments,

```
h = axes
```

and creates a new set of axes that are parented to the current figure (see `gcf`). The newly created axes are made the current axes (see `gca`) and are added to the end of the list of children for the current figure. The second form takes a set of property names and values

```
h = axes(propertyname,value,propertyname,value,...)
```

Creates a new set of axes, and then sets the specified properties to the given value. This is a shortcut for calling `set(h,propertyname,value)` for each pair. The third form takes a handle as an argument

```
axes(handle)
```

and makes `handle` the current axes, placing it at the head of the list of children for the current figure.

23.2 AXIS Setup Axis Behavior

23.2.1 Usage

Control the axis behavior. There are several versions of the axis command based on what you would like the axis command to do. The first versions set scalings for the current plot. The general syntax for its use is

```
axis([xmin xmax ymin ymax zmin zmax cmin cmax])
```

which sets the limits in the X, Y, Z and color axes. You can also set only the X, Y and Z axes:

```
axis([xmin xmax ymin ymax zmin zmax])
```

or only the X and Y axes:

```
axis([xmin xmax ymin ymax])
```

To retrieve the current axis limits, use the syntax

```
x = axis
```

where **x** is a 4-vector for 2D plots, and a 6-vector for 3D plots.

There are a number of axis options supported by FreeMat. The first version sets the axis limits to be automatically selected by FreeMat for each dimension. This state is the default one for new axes created by FreeMat.

```
axis auto
```

The next option sets all of the axis limits to **manual** mode. This state turns off automatic scaling of the axis based on the children of the current axis object.

```
axis manual
```

The next option sets the axis limits to fit tightly around the data.

```
axis tight
```

The next option adjusts the axis limits and plotbox aspect ratio so that the axis fills the position rectangle.

```
axis fill
```

The next option puts the axis in matrix mode. This mode is equivalent to the standard mode, but with the vertical axis reversed. Thus, the origin of the coordinate system is at the top left corner of the plot. This mode makes plots of matrix elements look normal (i.e., an identity matrix goes from upper left to lower right).

```
axis ij
```

The next option puts the axis in normal mode, with the origin at the lower left corner.

```
axis xy
```

The next option sets the axis parameters (specifically the data aspect ratio) so that equal ticks on each axis represent equal length. In this mode, spheres look spherical instead of ellipsoidal.

```
axis equal
```

The next option is the same as **axis equal**, but sets the plot box to fit tightly around the data (so no background shows through). It is the best option to use when displaying images.

```
axis image
```

The next option makes the axis box square.

`axis square`

The next option restores many of the normal characteristics of the axis. In particular, it undoes the effects of `square image` and `equal` modes.

`axis normal`

The next mode freezes axis properties so that 3D objects can be rotated properly.

`axis vis3d`

The next mode turns off all labels, tick marks and background.

`axis on`

The next mode turns on all labels, tick marks and background.

`axis off`

The next mode is similar to `axis off`, but also repacks the figure as tightly as possible. The result is a plot box that takes up the entire `outerposition` vector.

`axis maximal`

The `axis` command can also be applied to a particular axis (as opposed to the current axis as returned by `gca`) handle

`axis(M,...)`

23.3 AXISPROPERTIES Axis Object Properties

23.3.1 Usage

Below is a summary of the properties for the axis.

- `activepositionproperty` - four vector - Not used.
- `alim` - two vector - Controls the mapping of transparency. The vector `[a_1,a_2]`@ defines the scale for transparency. Plots then map `a_1` to a completely opaque value, and `a_2` to a completely transparent value. This mapping is applied to the alpha data of the plot data.
- `alimmode` - {'auto','manual'} - For `auto` mode, we map the alpha ranges of all objects in the plot to a full scale. For `manual` mode, we use the `alim` vector.
- `ambientlightcolor` - colorspec - Not used.
- `box` - On/Off - Not used.
- `cameraposition` - three vector - Set the position for the camera in axis space.
- `camerapositionmode` - {'auto','manual'} - For `manual` mode, the camera position is picked up from the `cameraposition` vector. For `auto` mode, the camera position is set to be centered on the `x` and `y` axis limits, and beyond the `z` maximum limit.

- `cameratarget` - `three vector` - Defines the point in axis space that the camera is targetted at.
- `cameratargetmode` - `{'auto', 'manual'}` - For `manual` mode the camera target is picked up from the `cameratarget` vector. For `auto` mode, the camera target is chosen to be the center of the three axes.
- `cameraupvector` - `three vector` - Defines the upwards vector for the camera (what is ultimately mapped to the vertical axis of the plot or screen). This vector must not be parallel to the vector that is defined by the optical axis (i.e., the one connecting the target to the camera position).
- `cameraupvectormode` - `{'auto', 'manual'}` - For `manual` mode, the camera up vector is picked up from the `cameraupvector`. The `auto` mode chooses the up vector to point along the positive y axis.
- `cameraviewangle` - `scalar` - Not used.
- `cameraviewanglemode` - `{'auto', 'manual'}` - Not used.
- `children` - `vector of handles` - A vector containing handles to children of the current axis. Be careful as to how you manipulate this vector. FreeMat uses a reference counting mechanism for graphics objects, so if you remove a handle from the `children` property of an axis, and you have not added it to the `children` property of another object, it will be deleted.
- `clim` - `two vector` - The color range vector. This vector contains two values that dictate how children of this axis get mapped to the colormap. Values between the two endpoints of this vector are mapped to the extremes of the colormap.
- `climmode` - `{'auto', 'manual'}` - For `auto` mode, the color limits are chosen to span the colordata for all of the children objects. For `manual` mode, the color mapping is based on `clim`.
- `clipping` - `{'on', 'off'}` - Not used.
- `color` - `colormap` - The color used to draw the background box for the axes. Defaults to white.
- `colororder` - `color vector` - A vector of color specs (in RGB) that are cycled between when drawing line plots into this axis. The default is `order blue,green,red,cyan,magenta,yellow,black`.
- `datalimits` - `six vector` - A vector that contains the x, y and z limits of the data for children of the current axis. Changes to this property are ignored - it is calculated by FreeMat based on the datasets.
- `dataaspectratio` - `three vector` - A vector that describes the aspect ratio of the data. You can think of this as the relative scaling of units for each axis. For example, if one unit along the x axis is twice as long as one unit along the y axis, you would specify a data aspect ratio of `[2,1,1]`.

- `dataaspectratiomode` - `{'auto','manual'}` - When the data aspect ratio is set to `manual`, the data is scaled by the data aspect ratio before being plotted. When the data aspect ratio mode is `auto` a complex set of rules are applied to determine how the data should be scaled. If `dataaspectratio` mode is `auto` and `plotboxaspectratio` is `auto`, then the default data aspect ratio is set to `[1,1,1]` and the default plot box aspect ratio is chosen proportional to `[xrange,yrange,zrange]`, where `xrange` is the span of data along the `x` axis, and similarly for `yrange` and `zrange`. If `plotboxaspectratio` is set to `[px,py,pz]`, then the `dataaspectratio` is set to `[xrange/px,yrange/py,zrange/pz]`. If one of the axes has been specified manually, then the data will be scaled to fit the axes as well as possible.
- `fontangle` - `{'normal','italic','oblique'}` - The angle of the fonts used for text labels (e.g., tick labels).
- `fontsize` - `scalar` - The size of fonts used for text labels (tick labels).
- `fontunits` - Not used.
- `fontweight` - `{'normal','bold','light','demi'}` - The weight of the font used for tick labels.
- `gridlinestyle` - `{'-' , '--' , ':' , '-.' , 'none'}` - The line style to use for drawing the grid lines. Defaults to `':'`.
- `handlevisibility` - Not used.
- `hitest` - Not used.
- `interruptible` - Not used.
- `layer` - Not used.
- `linestyleorder` - `linestyle vector` - A vector of linestyles that are cycled through when plotted line series.
- `linewidth` - `scalar` - The width of line used to draw grid lines, axis lines, and other lines.
- `minorgridlinestyle` - `{'-' , '--' , ':' , '-.' , 'none'}` - The line style used for drawing grid lines through minor ticks.
- `nextplot` - `{'add','replace','replacechildren'}` - Controls how the next plot interacts with the axis. If it is set to `'add'` the next plot will be added to the current axis. If it is set to `'replace'` the new plot replaces all of the previous children.
- `outerposition` - `four vector` - Specifies the coordinates of the outermost box that contains the axis relative to the containing figure. This vector is in normalized coordinates and corresponds to the `x`, `y`, `width`, `height` coordinates of the box.
- `parent` - `handle` - The handle for the containing object (a figure).
- `plotboxaspectratio` - `three vector` - Controls the aspect ratio of the plot box. See the entry under `dataaspectratio` for details on how FreeMat uses this vector in combination with the axis limits and the `plotboxaspectratio` to determine how to scale the data.

- `plotboxaspectratiomode` - {'auto', 'manual'} - The plot box aspect ratio mode interacts with the `dataaspectratiomode` and the axis limits.
- `position` - `fourvector` - The normalized coordinates of the plot box space. Should be inside the rectangle defined by `outerposition`.
- `positionmode` - {'auto', 'manual'} - the position mode is normally 'auto' which means that FreeMat computes the position vector to fit the plot inside the `outerposition` vector. If you set the `position` vector manually (using a `set` command), this mode flag becomes 'manual' and remains that way until reset to '@—'auto'.
- `projection` - Not used.
- `selected` - Not used.
- `selectionhighlight` - Not used.
- `tag` - A string that can be set to tag the axes with a name.
- `textheight` - `scalar` - This value is set by FreeMat to the height of the current font in pixels.
- `tickdir` - {'in', 'out'} - The direction of ticks. Defaults to 'in' for 2D plots, and 'out' for 3D plots if `tickdirmode` is `auto`.
- `tickdirmode` - {'auto', 'manual'} - When set to 'auto' the `tickdir` defaults to 'in' for 2D plots, and 'out' for 3D plots.
- `ticklength` - `two vector` - The first element is the length of the tick in 2D plots, and the second is the length of the tick in the 3D plots. The lengths are described as fractions of the longer dimension (width or height).
- `tightinset` - Not used.
- `title` - `handle` - The handle of the label used to represent the title of the plot.
- `type` - `string` - Takes the value of 'axes' for objects of the axes type.
- `units` - Not used.
- `userdata` - `array` - An arbitrary array you can set to anything you want.
- `visible` - {'on', 'off'} - If set to 'on' the axes are drawn as normal. If set to 'off', only the children of the axes are drawn. The plot box, axis lines, and tick labels are not drawn.
- `xaxislocation` - {'top', 'bottom'} - Controls placement of the x axis.
- `yaxislocation` - {'left', 'right'} - Controls placement of the y axis.
- `xcolor` - `colorspec` - The color of the x elements including the the x axis line, ticks, grid lines and tick labels
- `ycolor` - `colorspec` - The color of the y elements including the the y axis line, ticks, grid lines and tick labels.

- **zcolor** - **colspec** - The color of the z elements including the the z axis line, ticks, grid lines and tick labels.
- **xdir** - {'normal', 'reverse'} - For **normal**, axes are drawn as you would expect (e.g, in default 2D mode, the x axis has values increasing from left to right. For **reverse**, the x axis has values increasing from right to left.
- **ydir** - {'normal', 'reverse'} - For **normal**, axes are drawn as you would expect (e.g, in default 2D mode, the y axis has values increasing from bottom to top. For **reverse**, the y axis has values increasing from top to bottom.
- **zdir** - {'normal', 'reverse'} - For **normal**, axes are drawn as you would expect. In default 3D mode, the z axis has values increasing in some direction (usually up). For **reverse** the z axis increases in the opposite direction.
- **xgrid** - {'on', 'off'} - Set to **on** to draw grid lines from ticks on the x axis.
- **ygrid** - {'on', 'off'} - Set to **on** to draw grid lines from ticks on the y axis.
- **zgrid** - {'on', 'off'} - Set to **on** to draw grid lines from ticks on the z axis.
- **xlabel** - **handle** - The handle of the text label attached to the x axis. The position of that label and the rotation angle is computed automatically by FreeMat.
- **ylabel** - **handle** - The handle of the text label attached to the y axis. The position of that label and the rotation angle is computed automatically by FreeMat.
- **zlabel** - **handle** - The handle of the text label attached to the z axis. The position of that label and the rotation angle is computed automatically by FreeMat.
- **xlim** - **two vector** - Contains the limits of the data along the x axis. These are set automatically for **xlimmode**. When manually set it allows you to zoom into the data. The first element of this vector should be the smallest x value you want mapped to the axis, and the second element should be the largest.
- **ylim** - **two vector** - Contains the limits of the data along the y axis. These are set automatically for **ylimmode**. When manually set it allows you to zoom into the data. The first element of this vector should be the smallest y value you want mapped to the axis, and the second element should be the largest.
- **zlim** - **two vector** - Contains the limits of the data along the z axis. These are set automatically for **zlimmode**. When manually set it allows you to zoom into the data. The first element of this vector should be the smallest z value you want mapped to the axis, and the second element should be the largest.
- **xlimmode** - {'auto', 'manual'} - Determines if **xlim** is determined automatically or if it is determined manually. When determined automatically, it is chosen to span the data range (at least).

- `ylimmode` - {'auto', 'manual'} - Determines if `ylim` is determined automatically or if it is determined manually. When determined automatically, it is chosen to span the data range (at least).
- `zlimmode` - {'auto', 'manual'} - Determines if `zlim` is determined automatically or if it is determined manually. When determined automatically, it is chosen to span the data range (at least).
- `xminorgrid` - {'on', 'off'} - Set to `on` to draw grid lines from minor ticks on the x axis.
- `yminorgrid` - {'on', 'off'} - Set to `on` to draw grid lines from minor ticks on the y axis.
- `zminorgrid` - {'on', 'off'} - Set to `on` to draw grid lines from minor ticks on the z axis.
- `xscale` - {'linear', 'log'} - Determines if the data on the x axis is linear or logarithmically scaled.
- `yscale` - {'linear', 'log'} - Determines if the data on the y axis is linear or logarithmically scaled.
- `zscale` - {'linear', 'log'} - Determines if the data on the z axis is linear or logarithmically scaled.
- `xtick` - vector - A vector of x coordinates where ticks are placed on the x axis. Setting this vector allows you complete control over the placement of ticks on the axis.
- `ytick` - vector - A vector of y coordinates where ticks are placed on the y axis. Setting this vector allows you complete control over the placement of ticks on the axis.
- `ztick` - vector - A vector of z coordinates where ticks are placed on the z axis. Setting this vector allows you complete control over the placement of ticks on the axis.
- `xticklabel` - string vector - A string vector, of the form 'stringstring—string'— that contains labels to assign to the labels on the axis. If this vector is shorter than `xtick`, then FreeMat will cycle through the elements of this vector to fill out the labels.
- `yticklabel` - string vector - A string vector, of the form 'stringstring—string'— that contains labels to assign to the labels on the axis. If this vector is shorter than `ytick`, then FreeMat will cycle through the elements of this vector to fill out the labels.
- `zticklabel` - string vector - A string vector, of the form 'stringstring—string'— that contains labels to assign to the labels on the axis. If this vector is shorter than `ztick`, then FreeMat will cycle through the elements of this vector to fill out the labels.
- `xtickmode` - {'auto', 'manual'} - Set to 'auto' if you want FreeMat to calculate the tick locations. Setting 'xtick' will cause this property to switch to 'manual'.
- `ytickmode` - {'auto', 'manual'} - Set to 'auto' if you want FreeMat to calculate the tick locations. Setting 'ytick' will cause this property to switch to 'manual'.
- `ztickmode` - {'auto', 'manual'} - Set to 'auto' if you want FreeMat to calculate the tick locations. Setting 'ztick' will cause this property to switch to 'manual'.

- `xticklabelmode` - {'auto', 'manual'} - Set to 'auto' if you want FreeMat to set the tick labels. This will be based on the vector `xtick`.
- `yticklabelmode` - {'auto', 'manual'} - Set to 'auto' if you want FreeMat to set the tick labels. This will be based on the vector `ytick`.
- `zticklabelmode` - {'auto', 'manual'} - Set to 'auto' if you want FreeMat to set the tick labels. This will be based on the vector `ztick`.

23.4 CLA Clear Current Axis

23.4.1 Usage

Clears the current axes. The syntax for its use is

```
cla
```

23.5 CLABEL Add Labels To Contour Plot

23.5.1 Usage

The `clabel` function adds labels to a contour plot. Generate contour labels for a contour plot. The syntax for its use is either:

```
handles = clabel(contourhandle,property,value,property,value,...)
```

which labels all of the contours in the plot, or


```
handles = clabel(contourhandle,vals,property,value,property,value,...)
```

which only labels those contours indicated by the vector `vals`. The `contourhandle` must be the handle to a contour plot object. The remaining property/value pairs are passed to the `text` function to control the parameters of the generated text labels. See the `text properties` for the details on what can be used in those labels.

23.5.2 Example

```
--> [x,y] = meshgrid(linspace(-1,1,50));  
--> z = x.*exp(-(x.^2+y.^2));  
--> h = contour(z);  
--> clabel(h,'backgroundcolor',[1,1,.6],'edgecolor',[.7,.7,.7]);
```


which results in



Alternately, we can just label a subset of the contours

```
--> h = contour(z);
--> clabel(h, [-.2,0,.3]);
```

which results in



23.6 CLF Clear Figure

23.6.1 Usage

This function clears the contents of the current figure. The syntax for its use is

```
clf
```

23.7 CLIM Adjust Color limits of plot

23.7.1 Usage

There are several ways to use `clim` to adjust the color limits of a plot. The various syntaxes are

```
clim
clim([lo,hi])
clim('auto')
clim('manual')
```

```

clim('mode')
clim(handle,...)

```

The first form (without arguments), returns a 2-vector containing the current limits. The second form sets the limits on the plot to `[lo,hi]`. The third and fourth form set the mode for the limit to `auto` and `manual` respectively. In `auto` mode, FreeMat chooses the range for the axis automatically. The `clim('mode')` form returns the current mode for the axis (either `'auto'` or `'manual'`).

Switching to `manual` mode does not change the limits, it simply allows you to modify them (and disables the automatic adjustment of the limits as more objects are added to the plot). Also, if you specify a set of limits explicitly, the mode is set to `manual`.

Finally, you can specify the handle of an axis to manipulate instead of using the current one.

23.7.2 Example

Here is an example of using `clim` to change the effective window and level onto an image. First, the image with default limits

```

--> x = repmat(linspace(-1,1),[100,1]); y = x';
--> z = exp(-x.^2-y.^2);
--> image(z);
--> min(z(:))

```

```

ans =
    0.1353

```

```

--> max(z(:))


```

```

ans =
    0.9998

```

which results in




Next, we change the colorscale of the image using the `clim` function

```

--> image(z);
--> clim([0,0.2]);

```

which results in



23.8 CLOSE Close Figure Window

23.8.1 Usage

Closes a figure window, either the currently active window, a window with a specific handle, or all figure windows. The general syntax for its use is

```
close(handle)
```

in which case the figure window with the specified `handle` is closed. Alternately, issuing the command with no argument

```
close
```

is equivalent to closing the currently active figure window. Finally the command

```
close('all')
```

closes all figure windows currently open.

23.9 COLORBAR Add Colorbar to Current Plot

23.9.1 Usage

There are a number of syntaxes for the `colorbar` command. The first takes no arguments, and adds a vertical colorbar to the right of the current axes.

```
colorbar
```

You can also pass properties to the newly created axes object using the second syntax for colorbar

```
colorbar(properties...)
```

23.10 COLORMAP Image Colormap Function

23.10.1 Usage

Changes the colormap for the current figure. The generic syntax for its use is

```
colormap(map)
```

where `map` is a an array organized as `3 \times N`, which defines the RGB (Red Green Blue) coordinates for each color in the colormap. You can also use the function with no arguments to recover the current colormap

```
map = colormap
```

23.10.2 Function Internals

Assuming that the contents of the colormap function argument `c` are labeled as:

$$c = \begin{bmatrix} r_1 & g_1 & b_1 \\ r_1 & g_2 & b_2 \\ r_1 & g_3 & b_3 \\ \vdots & \vdots & \vdots \end{bmatrix}$$

then these columns for the RGB coordinates of pixel in the mapped image. Assume that the image occupies the range $[a, b]$. Then the RGB color of each pixel depends on the value x via the following integer

$$k = 1 + \lfloor 256 \frac{x - a}{b - a} \rfloor,$$


so that a pixel corresponding to image value x will receive RGB color $[r_k, g_k, b_k]$. Colormaps are generally used to pseudo color images to enhance visibility of features, etc.

23.10.3 Examples

We start by creating a smoothly varying image of a 2D Gaussian pulse.

```
--> x = linspace(-1,1,512)';
--> y = x';
--> Z = exp(-(x.^2+y.^2)/0.3);
--> image(Z);
```


which we display with the default (grayscale) colormap here.



Next we switch to the `copper` colormap, and redisplay the image.


```
--> colormap(copper);
--> image(Z);
```

which results in the following image.



If we capture the output of the `copper` command and plot it, we obtain the following result:


```
--> a = copper;
--> plot(a);
```



Note that in the output that each of the color components are linear functions of the index, with the ratio between the red, blue and green components remaining constant as a function of index. The result is an intensity map with a copper tint. We can similarly construct a colormap of our own


by defining the three components separately. For example, suppose we take three gaussian curves, one for each color, centered on different parts of the index space:

```
--> t = linspace(0,1,256);
--> A = [exp(-(t-1.0).^2/0.1);exp(-(t-0.5).^2/0.1);exp(-t.^2/0.1)]';
--> plot(A);
```




The resulting image has dark bands in it near the color transitions.

```
--> image(Z);
--> colormap(A);
```




These dark bands are a result of the nonuniform color intensity, which we can correct for by renormalizing each color to have the same norm.

```
--> w = sqrt(sum(A'.^2));
--> sA = diag(1./w)*A;
--> plot(A);
```



The resulting image has no more dark bands.

```
--> image(Z);
--> colormap(A);
```



23.11 COLORSPEC Color Property Description

23.11.1 Usage

There are a number of ways of specifying a color value for a color-based property. Examples include line colors, marker colors, and the like. One option is to specify color as an RGB triplet

```
set(h, 'color', [r,g,b])
```

where r, g, b are between $@[0,1]@$. Alternately, you can use color names to specify a color.

- 'none' - No color.
- 'y', 'yellow' - The color $@[1,1,0]@$ in RGB space.
- 'm', 'magenta' - The color $@[1,0,1]@$ in RGB space.
- 'c', 'cyan' - The color $@[0,1,1]@$ in RGB space.
- 'r', 'red' - The color $@[1,0,0]@$ in RGB space.

- 'g', 'green' - The color @[0,1,0]@ in RGB space.
- 'b', 'blue' - The color @[0,0,1]@ in RGB space.
- 'w', 'white' - The color @[1,1,1]@ in RGB space.
- 'k', 'black' - The color @[0,0,0]@ in RGB space.

23.12 CONTOUR Contour Plot Function

23.12.1 Usage

This command generates contour plots. There are several syntaxes for the command. The simplest is

```
contour(Z)
```

which generates a contour plot of the data in matrix *Z*, and will automatically select the contour levels. The *x,y* coordinates of the contour default to *1:n* and *1:m*, where *n* is the number of columns and *m* is the number of rows in the *Z* matrix. Alternately, you can specify a scalar *n*

```
contour(Z,n)
```

which indicates that you want *n* contour levels. For more control, you can provide a vector *v* containing the levels to contour. If you want to generate a contour for a particular level, you must pass a vector [*t,t*] where *t* is the level you want to contour. If you have data that lies on a particular *X,Y* grid, you can pass either vectors *x,y* or matrices *X,Y* to the contour function via

```
contour(X,Y,Z)
contour(X,Y,Z,n)
contour(X,Y,Z,v)
```

Each form of `contour` can optionally take a line spec to indicate the color and linestyle of the contours to draw:

```
contour(...,linespec)
```

or any of the other forms of `contour`. Furthermore, you can supply an axis to target the contour plot to (so that it does not get added to the current axis, which is the default):

```
contour(axis_handle,...)
```

Finally, the `contour` command returns a handle to the newly returned contour plot.

```
handle = contour(...)
```


To place labels on the contour plot, use the `clabel` function.

23.12.2 Example

Here is a simple example of a contour plot with the default x,y coordinates:

```
--> [x,y] = meshgrid(linspace(-1,1,25),linspace(-2,2,30));
--> z = x.*exp(-x.^2-y.^2);
--> contour(z)
```


which results in the following plot



Here, we specify the x and y coordinates explicitly

```
--> contour(x,y,z)
```


note that the axis limits have changed appropriately



By default, contours are created at values selected by FreeMat. To provide our own set of contour values (asymmetrically about zero in this case), we supply them as

```
--> contour(x,y,z,[-.4,0.,3])
```

which is here




Also by default, `contour` uses the current color map and `clim` limits to determine the coloring of the contours. Here, we override the color spec so that we have all black contour

```
--> contour(x,y,z,'b-')
```

Warning: Newly defined variable `margin` shadows a function of the same name. Use `clear margin` to recover

which is here



23.13 CONTOUR3 3D Contour Plot Function

23.13.1 Usage

This command generates contour plots where the lines are plotted in 3D. The syntax for its use is identical to the `contour` function. Please see its help for details.

23.13.2 Example

Here is a simple example of a 3D contour plot.

```
--> [x,y] = meshgrid([-2:.25:2]);
```

```
--> z=x.*exp(-x.^2-y.^2);
```

```
--> contour3(x,y,z,30);
```


Warning: Newly defined variable `xlim` shadows a function of the same name. Use `clear xlim` to recover ac

Warning: Newly defined variable `ylim` shadows a function of the same name. Use `clear ylim` to recover ac

Warning: Newly defined variable `zlim` shadows a function of the same name. Use `clear zlim` to recover ac

```
--> axis square;  
--> view(-15,25)  
Warning: Newly defined variable xlim shadows a function of the same name. Use clear xlim  
Warning: Newly defined variable ylim shadows a function of the same name. Use clear ylim  
Warning: Newly defined variable zlim shadows a function of the same name. Use clear zlim
```

The resulting plot



23.14 COPPER Copper Colormap

23.14.1 Usage

Returns a copper colormap. The syntax for its use is


```
y = copper
```

23.14.2 Example

Here is an example of an image displayed with the `copper` colormap

```
--> x = linspace(-1,1,512)'*ones(1,512);  
--> y = x';  
--> Z = exp(-(x.^2+y.^2)/0.3);  
--> image(Z);  
--> colormap(copper);
```

which results in the following image



23.15 COPY Copy Figure Window

23.15.1 Usage

Copies the currently active figure window to the clipboard. The syntax for its use is:

```
copy
```

The resulting figure is copied as a bitmap to the clipboard, and can then be pasted into any suitable application.

23.16 COUNTOUR Contour Object Properties

23.16.1 Usage

Below is a summary of the properties for a line series.

- **contourmatrix** - **array** - the matrix containing contour data for the plot. This is a 2 x N matrix containing x and y coordinates for points on the contours. In addition, each contour line starts with a column containing the number of points and the contour value.
- **displayname** - **string** - The name of this line series as it appears in a legend.
- **floating** - {'on', 'off'} - set to on to have floating (3D) contours
- **levellist** - **vector** - a vector of Z-values for the contours
- **levellistmode** - {'auto', 'manual'} - set to auto for automatic selection of Z-values of the contours.
- **linecolor** - color of the contour lines.
- **linestyle** - {'-', '--', ':', '-.', 'none'} - the line style to draw the contour in.
- **linewidth** - **scalar** - the width of the lines
- **parent** - **handle** - The axis that contains this object

- **tag** - **string** - A string that can be used to tag the object.
- **type** - **string** - Returns the string 'contour'.
- **userdata** - **array** - Available to store any variable you want in the handle object.
- **visible** - {'on', 'off'} - Controls visibility of the the line.
- **xdata** - **matrix** - Contains the x coordinates of the surface on which the zdata is defined. This can either be a monotonic vector of the same number of columns as **zdata**, or a 2D matrix that is the same size as **zdata**.
- **xdatamode** - {'auto', 'manual'} - When set to 'auto' FreeMat will autogenerate the x coordinates for the contours. These values will be 1, . . . ,N where N is the number of columns of **zdata**.
- **ydata** - **matrix** - Contains the y coordinates of the surface on which the zdata is defined. This can either be a monotonic vector of the same number of rows as **zdata** or a 2D matrix that is the same size as **zdata**.
- **ydatamode** - {'auto', 'manual'} - When set to 'auto' FreeMat will autogenerate the y coordinates for the contour data.
- **zdata** - **matrix** - The matrix of z values that are to be contoured.

23.17 DRAWNOW Flush the Event Queue

23.17.1 Usage

The `drawnow` function can be used to process the events in the event queue of the FreeMat application. The syntax for its use is

```
drawnow
```

Now that FreeMat is threaded, you do not generally need to call this function, but it is provided for compatibility.

23.18 FIGLOWER Lower a Figure Window

23.18.1 Usage

Lowers a figure window indicated by the figure number. The syntax for its use is

```
figlower(fignum)
```

where `fignum` is the number of the figure to lower. The figure will be lowered to the bottom of the GUI stack (meaning that it will be behind other windows). Note that this function does not cause `fignum` to become the current figure, you must use the `figure` command for that. Similarly, if `fignum` is the current figure, it will remain the current figure (even though the figure is now behind others).

23.19 FIGRAISE Raise a Figure Window

23.19.1 Usage

Raises a figure window indicated by the figure number. The syntax for its use is

```
figraise(fignum)
```

where `fignum` is the number of the figure to raise. The figure will be raised to the top of the GUI stack (meaning that it will be visible). Note that this function does not cause `fignum` to become the current figure, you must use the `figure` command for that.

23.20 FIGURE Figure Window Select and Create Function

23.20.1 Usage

Changes the active figure window to the specified figure number. The general syntax for its use is

```
figure(number)
```

where `number` is the figure number to use. If the figure window corresponding to `number` does not already exist, a new window with this number is created. If it does exist then it is brought to the forefront and made active. You can use `gcf` to obtain the number of the current figure.

Note that the figure number is also the handle for the figure. While for most graphical objects (e.g., axes, lines, images), the handles are large integers, for figures, the handle is the same as the figure number. This means that the figure number can be passed to `set` and `get` to modify the properties of the current figure, (e.g., the colormap). So, for figure 3, for example, you can use `get(3, 'colormap')` to retrieve the colormap for the current figure.

23.21 FIGUREPROPERTIES Figure Object Properties

23.21.1 Usage

Below is a summary of the properties for the axis.

- **alphamap** - **vector** - Contains the alpha (transparency) map for the figure. If this is set to a scalar, then all values are mapped to the same transparency. It defaults to 1, which is all values being fully opaque. If you set this to a vector, the values of graphics objects will be mapped to different transparency values, based on the setting of their `alphadatamapping` property.
- **color** - **colormap** - The background color of the figure (defaults to a gray [0.6,0.6,0.6]). During printing, this color is set to white, and then is restored.
- **colormap** - **color vector** - an $N \times 3$ matrix of RGB values that specifies the colormap for the figure. Defaults to an HSV map.
- **children** - **handle vector** - the handles for objects that are children of this figure. These should be axis objects.

- `currentaxes` - `handle` - the handle for the current axes. Also returned by `gca`.
- `doublebuffer` - Not used.
- `parent` - Not used.
- `position` - Not used.
- `type` - `string` - returns the string `'figure'`.
- `userdata` - `array` - arbitrary array you can use to store data associated with the figure.
- `nextplot` - `{'add', 'replace', 'replacechildren'}` - If set to `'add'` then additional axes are added to the list of children for the current figure. If set to `'replace'`, then a new axis replaces all of the existing children.
- `figsize` - `two vector` - the size of the figure window in pixels (width x height).
- `renderer` - `{'painters', 'opengl'}` - When set to `'painters'` drawing is based on the Qt drawing methods (which can handle flat shading of surfaces with transparency). If you set the renderer to `'opengl'` then OpenGL is used for rendering. Support for OpenGL is currently in the alpha stage, and FreeMat does not enable it automatically. You can set the renderer mode to `'opengl'` manually to experiment. Also, OpenGL figures cannot be printed yet.

23.22 GCA Get Current Axis

23.22.1 Usage

Returns the handle for the current axis. The syntax for its use is

```
handle = gca
```

where `handle` is the handle of the active axis. All object creation functions will be children of this axis.

23.23 GCF Get Current Figure

23.23.1 Usage

Returns the figure number for the current figure (which is also its handle, and can be used to set properties of the current figure using `set`). The syntax for its use is

```
figure_number = gcf
```

where `figure_number` is the number of the active figure (also the handle of the figure).

Note that figures have handles, just like axes, images, plots, etc. However the handles for figures match the figure number (while handles for other graphics objects tend to be large, somewhat arbitrary integers). So, to retrieve the colormap of the current figure, you could use `get(gcf, 'colormap')`, or to obtain the colormap for figure 3, use `get(3, 'colormap')`.

23.24 GET Get Object Property

23.24.1 Usage

This function allows you to retrieve the value associated with a property. The syntax for its use is

```
value = get(handle,property)
```

where `property` is a string containing the name of the property, and `value` is the value for that property. The type of the variable `value` depends on the property being set. See the help for the properties to see what values you can set.

23.25 GLSHOW Show a GL Assembly in a GL Window

23.25.1 Usage

Shows a GL Assembly in a GL Window. The syntax for its use is

```
glshow(name,scale)
```

which shows the `glassembly` named `name` in a new GL window, with the scale set to `scale`. Roughly speaking `scale` should represent the radial size of the object that you want to see in the window.

23.26 GRAY Gray Colormap

23.26.1 Usage

Returns a gray colormap. The syntax for its use is


```
y = gray
```

23.26.2 Example

Here is an example of an image displayed with the `gray` colormap

```
--> x = linspace(-1,1,512)*ones(1,512);  
--> y = x';  
--> Z = exp(-(x.^2+y.^2)/0.3);  
--> image(Z);  
--> colormap(gray);
```

which results in the following image



23.27 GRID Plot Grid Toggle Function

23.27.1 Usage

Toggles the drawing of grid lines on the currently active plot. The general syntax for its use is

```
grid(state)
```

where `state` is either

```
grid('on')
```

to activate the grid lines, or

```
grid('off')
```

to deactivate the grid lines. If you specify no argument then `grid` toggles the state of the grid:

```
grid
```

You can also specify a particular axis to the grid command


```
grid(handle,...)
```

where `handle` is the handle for a particular axis.

23.27.2 Example


Here is a simple plot without grid lines.

```
--> x = linspace(-1,1);  
--> y = cos(3*pi*x);  
--> plot(x,y,'r-');
```



Next, we activate the grid lines.

```
--> plot(x,y,'r-');
--> grid on
```



23.28 HCONTOUR Create a contour object

23.28.1 Usage

Creates a contour object and parents it to the current axis. The syntax for its use is

```
handle = hcontour(property,value,property,value,...)
```

where `property` and `value` are set. The handle ID for the resulting object is returned. It is automatically added to the children of the current axis.

23.29 HIMAGE Create a image object

23.29.1 Usage

Creates a image object and parents it to the current axis. The syntax for its use is

```
handle = himage(property,value,property,value,...)
```

where `property` and `value` are set. The handle ID for the resulting object is returned. It is automatically added to the children of the current axis.

23.30 HIST Histogram Function

23.30.1 Usage

```
n=hist (y)
n=hist (y,x)
n=hist (y,x,norm)
```

Produce histogram counts or plots.

With one vector input argument, plot a histogram of the values with 10 bins. The range of the histogram bins is determined by the range of the data.

Given a second scalar argument, use that as the number of bins.

Given a second vector argument, use that as the centers of the bins, with the width of the bins determined from the adjacent values in the vector.

If third argument is provided, the histogram is normalised such that the sum of the bars is equal to `norm`.

Extreme values are lumped in the first and last bins.

23.31 HLINE Create a line object

23.31.1 Usage

Creates a line object and parents it to the current axis. The syntax for its use is

```
handle = hline(property,value,property,value,...)
```

where `property` and `value` are set. The handle ID for the resulting object is returned. It is automatically added to the children of the current axis.

23.32 HOLD Plot Hold Toggle Function

23.32.1 Usage

Toggles the hold state on the currently active plot. The general syntax for its use is

```
hold(state)
```

where `state` is either

```
hold('on')
```

to turn hold on, or

```
hold('off')
```

to turn hold off. If you specify no argument then `hold` toggles the state of the hold:

```
hold
```

You can also specify a particular axis to the hold command

```
hold(handle,...)
```

where `handle` is the handle for a particular axis.


23.32.2 Function Internals

The `hold` function allows one to construct a plot sequence incrementally, instead of issuing all of the plots simultaneously using the `plot` command.

23.32.3 Example


Here is an example of using both the `hold` command and the multiple-argument `plot` command to construct a plot composed of three sets of data. The first is a plot of a modulated Gaussian.

```
--> x = linspace(-5,5,500);
--> t = exp(-x.^2);
--> y = t.*cos(2*pi*x*3);
--> plot(x,y);
```



We now turn the hold state to 'on', and add another plot sequence, this time composed of the top and bottom envelopes of the modulated Gaussian. We add the two envelopes simultaneously using a single `plot` command. The fact that `hold` is 'on' means that these two envelopes are added to (instead of replace) the current contents of the plot.

```
--> plot(x,y);
--> hold on
--> plot(x,t,'g-',x,-t,'b-')
```



23.33 HPATCH Create a patch object

23.33.1 Usage

Creates a patch object and parents it to the current axis. The syntax for its use is

```
handle = hpatch(property,value,property,value,...)
```

where `property` and `value` are set. The handle ID for the resulting object is returned. It is automatically added to the children of the current axis.

23.34 HPOINT Get Point From Window

23.34.1 Usage

This function waits for the user to click on the current figure window, and then returns the coordinates of that click. The generic syntax for its use is

```
[x,y] = hpoint
```

23.35 HRAWPLOT Generate a Raw Plot File

23.35.1 Usage

This function takes a sequence of commands, and generates a raw plot (to a file) that renders the commands. It is a useful tool for creating high quality fully customized PDF plots from within FreeMat scripts that are portable. The syntax for its use

```
hrawplot(filename,commands)
```

where `filename` is the name of the file to plot to, and `commands` is a cell array of strings. Each entry in the cell array contains a string with a command text. The commands describe a simple mini-language for describing plots. The complete dictionary of commands is given

- `LINE x1 y1 x2 y2` – draw a line
- `FONT name size` – select a font of the given name and size
- `TEXT x1 y1 string` – draw the given text string at the given location
- `STYLE style` – select line style ('solid' or 'dotted')
- `PAGE` – force a new page
- `SIZE x1 y1` – Set the page mapping

23.36 **HSURFACE** Create a surface object

23.36.1 Usage

Creates a surface object and parents it to the current axis. The syntax for its use is

```
handle = hsurface(property,value,property,value,...)
```

where `property` and `value` are set. The handle ID for the resulting object is returned. It is automatically added to the children of the current axis.

23.37 **HTEXT** Create a text object

23.37.1 Usage

Creates a text object and parents it to the current axis. The syntax for its use is

```
handle = htext(property,value,property,value,...)
```

where `property` and `value` are set. The handle ID for the resulting object is returned. It is automatically added to the children of the current axis.

23.38 **HTEXTBITMAP** Get Text Rendered as a Bitmap

23.38.1 Usage

This function takes a fontname, a size, and a text string and returns a bitmap containing the text. The generic syntax for its use is

```
bitmap = htextbitmap(fontname,size,text)
```

where the output bitmap contains the text rendered into a matrix.

23.39 **IMAGE** Image Display Function

23.39.1 Usage

The `image` command has the following general syntax

```
handle = image(x,y,C,properties...)
```

where `x` is a two vector containing the `x` coordinates of the first and last pixels along a column, and `y` is a two vector containing the `y` coordinates of the first and last pixels along a row. The matrix `C` constitutes the image data. It must either be a scalar matrix, in which case the image is colormapped using the `colormap` for the current figure. If the matrix is $M \times N \times 3$, then `C` is interpreted as RGB data, and the image is not colormapped. The `properties` argument is a set of `property/value` pairs that affect the final image. You can also omit the `x` and `y`,

```
handle = image(C, properties...)
```

in which case they default to $x = [1, \text{size}(C, 2)]$ and $y = [1, \text{size}(C, 1)]$. Finally, you can use the `image` function with only formal arguments

```
handle = image(properties...)
```

To support legacy FreeMat code, you can also use the following form of `image`

```
image(C, zoomfactor)
```

which is equivalent to `image(C)` with the axes removed so that the image takes up the full figure window, and the size of the figure window adjusted to achieve the desired zoom factor using the `zoom` command.

23.39.2 Example


In this example, we create an image that is 512 x 512 pixels square, and set the background to a noise pattern. We set the central 128 x 256 pixels to be white.

```
--> x = rand(512);
--> x((-64:63)+256, (-128:127)+256) = 1.0;
--> figure
```

```
ans =
     1
```

```
--> image(x)
--> colormap(gray)
```


The resulting image looks like:



Here is an example of an RGB image

```
--> t = linspace(0,1);
--> red = t'*t;
--> green = t'*(t.^2);
--> blue = t'*(0*t+1);
--> A(:,:,1) = red;
--> A(:,:,2) = green;
--> A(:,:,3) = blue;
--> image(A);
```

The resulting image looks like:



23.40 IMAGEPROPERTIES Image Object Properties

23.40.1 Usage

Below is a summary of the properties for the axis.

- **alphadata** - **vector** - This is a vector that should contain as many elements as the image data itself **cdata**, or a single scalar. For a single scalar, all values of the image take on the same transparency. Otherwise, the transparency of each pixel is determined by the corresponding value from the **alphadata** vector.
- **alphadatamapping** - **{'scaled', 'direct', 'none'}** - For **none** mode (the default), no transparency is applied to the data. For **direct** mode, the vector **alphadata** contains values between $@[0,M-1]$ — where **M** is the length of the alpha map stored in the figure. For **scaled** mode, the **alim** vector for the figure is used to linearly rescale the alpha data prior to lookup in the alpha map.
- **cdata** - **array** - This is either a **M x N** array or an **M x N x 3** array. If the data is **M x N** the image is a scalar image (indexed mode), where the color associated with each image pixel is computed using the colormap and the **cdatamapping** mode. If the data is **M x N x 3** the image is assumed to be in RGB mode, and the colorpanes are taken directly from **cdata** (the colormap is ignored). Note that in this case, the data values must be between $@[0,1]$ — for each color channel and each pixel.
- **cdatamapping** - **{'scaled', 'direct'}** - For **scaled** (the default), the pixel values are scaled using the **clim** vector for the figure prior to looking up in the colormap. For **direct** mode, the pixel values must be in the range $[0,N-1]$ where **N** is the number of colors in the colormap if the data is integer type. For floating point data types, values must be in the range $[1,N]$.
- **children** - Not used.
- **parent** - **handle** - The axis containing the image.
- **tag** - **string** - You can set this to any string you want.

- `type` - string - Set to the string `'image'`.
- `xdata` - two vector - contains the x coordinates of the first and last column (respectively). Defaults to `[1,C]` where `C` is the number of columns in the image.
- `ydata` - two vector - contains the y coordinates of the first and last row (respectively). Defaults to `[1,R]` where `R` is the number of rows in the image.
- `userdata` - array - Available to store any variable you want in the handle object.
- `visible` - `{'on','off'}` - Controls whether the image is visible or not.

23.41 IMAGESC Image Display Function

23.41.1 Usage

The `imagesc` command has the following general syntax

```
handle = imagesc(x,y,C,clim)
```

where `x` is a two vector containing the x coordinates of the first and last pixels along a column, and `y` is a two vector containing the y coordinates of the first and last pixels along a row. The matrix `C` constitutes the image data. It must either be a scalar matrix, in which case the image is colormapped using the `colormap` for the current figure. If the matrix is `M x N x 3`, then `C` is interpreted as RGB data, and the image is not colormapped. The `clim` argument is a pairs [low high] that specifies scaling. You can also omit the `x` and `y`,

```
handle = imagesc(C, clim)
```

in which case they default to `x = [1,size(C,2)]` and `y = [1,size(C,1)]`. Finally, you can use the `image` function with only formal arguments

```
handle = imagesc(properties...)
```

23.41.2 Example

In this example, we create an image that is 512 x 512 pixels square, and set the background to a noise pattern. We set the central 128 x 256 pixels to be white.

```
--> x = rand(512);
--> x((-64:63)+256, (-128:127)+256) = 1.0;
--> figure
```

```
ans =
     1
```

```
--> imagesc(x, [0 .5])
```

Warning: Newly defined variable `nargin` shadows a function of the same name. Use `clear nargin`

```
--> colormap(gray)
```

23.42 IS2DVIEW Test Axes For 2D View

23.42.1 Usage

This function returns `true` if the current axes are in a 2-D view, and `false` otherwise. The generic syntax for its use is

```
y = is2dview(x)
```

where `x` is the handle of an axes object.

23.43 ISHOLD Test Hold Status

23.43.1 Usage

Returns the state of the `hold` flag on the currently active plot. The general syntax for its use is

```
ishold
```

and it returns a logical 1 if `hold` is on, and a logical 0 otherwise.

23.44 LEGEND Add Legent to Plot

23.44.1 Usage

This command adds a legend to the current plot. Currently, the following forms of the `legend` command are supported. The first form creates a legend with the given labels for the data series:

```
legend('label1','label2',...)
```

where `'label1'` is the text label associated with data plot 1 and so on. You can also use the `legend` command to control the appearance of the legend in the current plot. To remove the legend from the current plot, use

```
legend('off')
```

To hide the legend for the current plot (but do not remove it)

```
legend('hide')
```

And to show the legend that has been hidden, use

```
legend('show')
```

You can also toggle the display of the box surrounding the legend. Use

```
legend('boxoff')
```

or

```
legend('boxon')
```

to turn the legend box off or on, respectively. To toggle the visible state of the current legend, use

```
legend('toggle')
```

Specifying no arguments at all (apart from an optional location argument as specified below) results in the legend being rebuilt. This form is useful for picking up font changes or relocating the legend.

```
legend
```

By default, the `legend` command places the new legend in the upper right corner of the current plot. To change this behavior, use the `'location'` specifier (must be the last two options to the command)

```
legend(...,'location',option)
```

where `option` takes on the following possible values

- `north,N` - top center of plot
- `south,S` - bottom center of plot
- `east,E` - middle right of plot
- `west,W` - middle left of plot
- `northeast,NE` - top right of plot (default behavior)
- `northwest,NW` - top left of plot
- `southeast,SE` - bottom right of plot
- `southwest,SW` - bottom left of plot

This implementation of `legend` is incomplete relative to the MATLAB API. The functionality will be improved in future versions of FreeMat.

23.45 LINE Line Display Function

23.45.1 Usage

The `line` command has the following general syntax

```
handle = line(x,y,z,properties...)
```

where...

23.46 LINEPROPERTIES Line Series Object Properties

23.46.1 Usage

Below is a summary of the properties for a line series.

- **color** - **colormap** - The color that is used to draw the line.
- **children** - Not used.
- **displayname** - The name of this line series as it appears in a legend.
- **linestyle** - {'-', '--', ':', '-.', 'none'} - The style of the line.
- **linewidth** - **scalar** - The width of the line.
- **marker** - {'+', 'o', '*', '.', 'x', 'square', 's', 'diamond', 'd', '^', 'v', '>', '<'} - The marker for data points on the line. Some of these are redundant, as 'square' 's' are synonyms, and 'diamond' and 'd' are also synonyms.
- **markeredgecolor** - **colormap** - The color used to draw the marker. For some of the markers (circle, square, etc.) there are two colors used to draw the marker. This property controls the edge color (which for unfilled markers) is the primary color of the marker.
- **markerfacecolor** - **colormap** - The color used to fill the marker. For some of the markers (circle, square, etc.) there are two colors used to fill the marker.
- **markersize** - **scalar** - Control the size of the marker. Defaults to 6, which is effectively the radius (in pixels) of the markers.
- **parent** - **handle** - The axis that contains this object.
- **tag** - **string** - A string that can be used to tag the object.
- **type** - **string** - Returns the string 'line'.
- **visible** - {'on', 'off'} - Controls visibility of the the line.
- **xdata** - **vector** - Vector of x coordinates of points on the line. Must be the same size as the **ydata** and **zdata** vectors.
- **ydata** - **vector** - Vector of y coordinates of points on the line. Must be the same size as the **xdata** and **zdata** vectors.
- **zdata** - **vector** - Vector of z coordinates of points on the line. Must be the same size as the **xdata** and **ydata** vectors.
- **xdatamode** - {'auto', 'manual'} - When set to 'auto' FreeMat will autogenerate the x coordinates for the points on the line. These values will be 1, . . . ,N where N is the number of points in the line.
- **userdata** - **array** - Available to store any variable you want in the handle object.

23.47 LOGLOG Log-Log Plot Function

23.47.1 Usage

This command has the exact same syntax as the `plot` command:


```
loglog(<data 1>,{linespec 1},<data 2>,{linespec 2}...,properties...)
```

in fact, it is a simple wrapper around `plot` that sets the x and y axis to have a logarithmic scale.

23.47.2 Example


Here is an example of a doubly exponential signal plotted first on a linear plot:

```
--> x = linspace(1,100);
--> y = x;
--> plot(x,y,'r-');
```



and now on a log-log plot

```
--> loglog(x,y,'r-');
```



23.48 NEWPLOT Get Handle For Next Plot

23.48.1 Usage

Returns the handle for the next plot operation. The general syntax for its use is

```
h = newplot
```

This routine checks the `nextplot` properties of the current figure and axes to see if they are set to `replace` or not. If the figure's `nextplot` property is set to `replace`, the current figure is cleared. If the axes' `nextplot` property is set to `replace` then the axes are cleared for the next operation.

23.49 PATCH Patch Graphics Function

23.49.1 Usage

This routine is used to create a patch object that can be plotting 2D and 3D surfaces. A patch is a polygon defined by the xyz coordinates of its vertices and optionally by the color at the vertices. There are several forms for the `patch` function:

```
h = patch(X,Y,C,properties...)
h = patch(X,Y,Z,C,properties...)
h = patch(properties...)
h = patch(V)
```

Where `X`, `Y` and `Z` are matrices or vectors of `x`, `y` or `z` coordinates and `C` is a matrix or vector of color values (the colormap for the current fig is applied).

23.49.2 Example


Here we generate a surface specifying all four components.

```
--> x = [ 0 1 0 1];
--> y = [ 0 0 1 1];
--> c = [ 1 1 1 ];
--> patch(x,y,c)
--> axis equal
--> view(3)
```

Warning: Newly defined variable `xlim` shadows a function of the same name. Use `clear xlim` to recover ac

Warning: Newly defined variable `ylim` shadows a function of the same name. Use `clear ylim` to recover ac

Warning: Newly defined variable `zlim` shadows a function of the same name. Use `clear zlim` to recover ac



23.50 PCOLOR Pseudocolor Plot

23.50.1 Usage

This routine is used to create a pseudocolor plot of the data. A pseudocolor plot is essentially a surface plot seen from above. There are two forms for the `pcolor` command:

```
pcolor(C)
```

which uses a rectangular grid. Alternately, you can specify `X,Y` matrices or vectors.

```
pcolor(X,Y,C)
```

23.51 PLOT Plot Function

23.51.1 Usage

This is the basic plot command for FreeMat. The general syntax for its use is

```
plot(<data 1>,{linespec 1},<data 2>,{linespec 2}...,properties...)
```

where the `<data>` arguments can have various forms, and the `linespec` arguments are optional. We start with the `<data>` term, which can take on one of multiple forms:

- **Vector Matrix Case** – In this case the argument data is a pair of variables. A set of `x` coordinates in a numeric vector, and a set of `y` coordinates in the columns of the second, numeric matrix. `x` must have as many elements as `y` has columns (unless `y` is a vector, in which case only the number of elements must match). Each column of `y` is plotted sequentially against the common vector `x`.
- **Unpaired Matrix Case** – In this case the argument data is a single numeric matrix `y` that constitutes the `y`-values of the plot. An `x` vector is synthesized as `x = 1:length(y)`, and each column of `y` is plotted sequentially against this common `x` axis.
- **Complex Matrix Case** – Here the argument data is a complex matrix, in which case, the real part of each column is plotted against the imaginary part of each column. All columns receive the same line styles.

Multiple data arguments in a single plot command are treated as a *sequence*, meaning that all of the plots are overlapped on the same set of axes. The `linespec` is a string used to change the characteristics of the line. In general, the `linespec` is composed of three optional parts, the `colourspec`, the `symbolspec` and the `linestylespec` in any order. Each of these specifications is a single character that determines the corresponding characteristic. First, the `colourspec`:

- `'b'` - Color Blue
- `'g'` - Color Green
- `'r'` - Color Red
- `'c'` - Color Cyan

- 'm' - Color Magenta
- 'y' - Color Yellow
- 'k' - Color Black

The `symbolspec` specifies the (optional) symbol to be drawn at each data point:

- '.' - Dot symbol
- 'o' - Circle symbol
- 'x' - Times symbol
- '+' - Plus symbol
- '*' - Asterisk symbol
- 's' - Square symbol
- 'd' - Diamond symbol
- 'v' - Downward-pointing triangle symbol
- '^' - Upward-pointing triangle symbol
- '<' - Left-pointing triangle symbol
- '>' - Right-pointing triangle symbol

The `linespec` specifies the (optional) line style to use for each data series:

- '-' - Solid line style
- ':' - Dotted line style
- '-.' - Dot-Dash-Dot-Dash line style
- '--' - Dashed line style

For sequences of plots, the `linespec` is recycled with color order determined by the properties of the current axes. You can also use the `properties` argument to specify handle properties that will be inherited by all of the plots generated during this event. Finally, you can also specify the handle for the axes that are the target of the `plot` operation.


```
handle = plot(handle,...)
```

23.51.2 Example

The most common use of the `plot` command probably involves the vector-matrix paired case. Here, we generate a simple cosine, and plot it using a red line, with no symbols (i.e., a `linespec` of `'r-'`).

```
--> x = linspace(-pi,pi);
--> y = cos(x);
--> plot(x,y,'r-');
```


which results in the following plot.



Next, we plot multiple sinusoids (at different frequencies). First, we construct a matrix, in which each column corresponds to a different sinusoid, and then plot them all at once.


```
--> x = linspace(-pi,pi);
--> y = [cos(x(:)), cos(3*x(:)), cos(5*x(:))];
--> plot(x,y);
```

In this case, we do not specify a `linespec`, so that we cycle through the colors automatically (in the order listed in the previous section).



This time, we produce the same plot, but as we want to assign individual `linespecs` to each line, we use a sequence of arguments in a single plot command, which has the effect of plotting all of the data sets on a common axis, but which allows us to control the `linespec` of each plot. In the following example, the first line (harmonic) has red, solid lines with times symbols marking the data points, the second line (third harmonic) has blue, solid lines with right-pointing triangle symbols, and the third line (fifth harmonic) has green, dotted lines with asterisk symbols.


```
--> plot(x,y(:,1),'rx-',x,y(:,2),'b>-',x,y(:,3),'g*:')
```



The second most frequently used case is the unpaired matrix case. Here, we need to provide only one data component, which will be automatically plotted against a vector of natural number of the appropriate length. Here, we use a plot sequence to change the style of each line to be dotted, dot-dashed, and dashed.


```
--> plot(y(:,1),'r:',y(:,2),'b;',y(:,3),'g|')
```

Note in the resulting plot that the x-axis no longer runs from $[-\pi, \pi]$, but instead runs from $[1, 100]$.




The final case is for complex matrices. For complex arguments, the real part is plotted against the imaginary part. Hence, we can generate a 2-dimensional plot from a vector as follows.

```
--> y = cos(2*x) + i * cos(3*x);
--> plot(y);
```



Here is an example of using the handle properties to influence the behavior of the generated lines.

```
--> t = linspace(-3,3);
--> plot(cos(5*t).*exp(-t),'r-','linewidth',3);
```



23.52 PLOT3 Plot 3D Function

23.52.1 Usage

This is the 3D plot command. The general syntax for its use is

```
plot3(X,Y,Z,{linespec 1},X,Y,Z,{linespec 2},...,properties...)
```

where X Y and Z are the coordinates of the points on the 3D line. Note that in general, all three should be vectors. If some or all of the quantities are matrices, then FreeMat will attempt to expand the vector arguments to the same size, and then generate multiple plots, one for each column of the matrices. The linespec is optional, see `plot` for details. You can specify `properties` for the generated line plots. You can also specify a handle as an axes to target

```
plot3(handle,...)
```

23.52.2 Example


Here is a simple example of a 3D helix.

```

--> t = linspace(0,5*pi,200);
--> x = cos(t); y = sin(t); z = t;
--> plot3(x,y,z);
Warning: Newly defined variable xlim shadows a function of the same name. Use clear xlim to recover ac
Warning: Newly defined variable ylim shadows a function of the same name. Use clear ylim to recover ac
Warning: Newly defined variable zlim shadows a function of the same name. Use clear zlim to recover ac
--> view(3);
Warning: Newly defined variable xlim shadows a function of the same name. Use clear xlim to recover ac
Warning: Newly defined variable ylim shadows a function of the same name. Use clear ylim to recover ac
Warning: Newly defined variable zlim shadows a function of the same name. Use clear zlim to recover ac

```

Shown here



23.53 POINT Get Axis Position From Mouse Click

23.53.1 Usage

Returns information about the currently displayed image based on a user supplied mouse-click. The general syntax for its use is

```
t = point
```

The returned vector y has two elements:

$$t = [x, y]$$

where x, y are the coordinates in the current axes of the click. This function has changed since FreeMat 1.10. If the click is not inside the active area of any set of axes, a pair of NaNs are returned.

23.54 PRINT Print a Figure To A File

23.54.1 Usage

This function “prints” the currently active fig to a file. The generic syntax for its use is

```
print(filename)
```

or, alternately,

```
print filename
```

where `filename` is the (string) filename of the destined file. The current fig is then saved to the output file using a format that is determined by the extension of the filename. The exact output formats may vary on different platforms, but generally speaking, the following extensions should be supported cross-platform:

- `jpg, jpeg` – JPEG file
- `pdf` – Portable Document Format file
- `png` – Portable Net Graphics file
- `svg` – Scalable Vector Graphics file


Postscript (PS, EPS) is supported on non-Mac-OSX Unix only. Note that only the fig is printed, not the window displaying the fig. If you want something like that (essentially a window-capture) use a separate utility or your operating system's built in screen capture ability.

23.54.2 Example

Here is a simple example of how the figures in this manual are generated.

```
--> x = linspace(-1,1);
--> y = cos(5*pi*x);
--> plot(x,y,'r-');
--> print('printfig1.jpg')
--> print('printfig1.png')
```

which creates two plots `printfig1.png`, which is a Portable Net Graphics file, and `printfig1.jpg` which is a JPEG file.



23.55 PVALID Validate Property Name

23.55.1 Usage

This function checks to see if the given string is a valid property name for an object of the given type. The syntax for its use is

```
b = pvalid(type,propertyname)
```

where `string` is a string that contains the name of a valid graphics object type, and `propertyname` is a string that contains the name of the property to test for.

23.55.2 Example

Here we test for some properties on an `axes` object.

```
--> pvalid('axes','type')

ans =
     1

--> pvalid('axes','children')

ans =
     1

--> pvalid('axes','foobar')

ans =
     0
```

23.56 SEMILOGX Semilog X Axis Plot Function

23.56.1 Usage

This command has the exact same syntax as the `plot` command:


```
semilogx(<data 1>,{linespec 1},<data 2>,{linespec 2}...,properties...)
```

in fact, it is a simple wrapper around `plot` that sets the x axis to have a logarithmic scale.

23.56.2 Example


Here is an example of an exponential signal plotted first on a linear plot:

```
--> y = linspace(0,2);
--> x = (10).^y;
--> plot(x,y,'r-');
```



and now with a logarithmic x axis

```
--> semilogx(x,y,'r-');
```



23.57 SEMILOGY Semilog Y Axis Plot Function

23.57.1 Usage

This command has the exact same syntax as the `plot` command:


```
semilogy(<data 1>,{linespec 1},<data 2>,{linespec 2}...,properties...)
```

in fact, it is a simple wrapper around `plot` that sets the y axis to have a logarithmic scale.

23.57.2 Example


Here is an example of an exponential signal plotted first on a linear plot:

```
--> x = linspace(0,2);
--> y = 10.0.^x;
--> plot(x,y,'r-');
```



and now with a logarithmic y axis

```
--> semilogy(x,y,'r-');
```



23.58 SET Set Object Property

23.58.1 Usage

This function allows you to change the value associated with a property. The syntax for its use is

```
set(handle,property,value,property,value,...)
```

where `property` is a string containing the name of the property, and `value` is the value for that property. The type of the variable `value` depends on the property being set. See the help for the properties to see what values you can set.

23.59 SIZEFIG Set Size of Figure

23.59.1 Usage

The `sizefig` function changes the size of the currently selected fig window. The general syntax for its use is

```
sizefig(width,height)
```

where `width` and `height` are the dimensions of the fig window.

23.60 SUBPLOT Subplot Function

23.60.1 Usage

This function divides the current figure into a 2-dimensional grid, each of which can contain a plot of some kind. The function has a number of syntaxes. The first version

```
subplot(row,col,num)
```

which either activates subplot number `num`, or sets up a subplot grid of size `row` x `col`, and then activates `num`. You can also set up subplots that cover multiple grid elements

```
subplot(row,col,[vec])
```

where `vec` is a set of indexes covered by the new subplot. Finally, as a shortcut, you can specify a string with three components

```
subplot('mnp')
```

or using the alternate notation

```
subplot mnp
```


where `m` is the number of rows, `n` is the number of columns and `p` is the index. You can also specify the location of the subplot explicitly using the syntax

```
subplot('position',[left bottom width height])
```

23.60.2 Example


Here is the use of `subplot` to set up a 2 x 2 grid of plots

```
--> t = linspace(-pi,pi);
--> subplot(2,2,1)
--> plot(t,cos(t).*exp(-2*t));
--> subplot(2,2,2);
--> plot(t,cos(t*2).*exp(-2*t));
--> subplot(2,2,3);
--> plot(t,cos(t*3).*exp(-2*t));
--> subplot(2,2,4);
--> plot(t,cos(t*4).*exp(-2*t));
```



Here we use the second form of `subplot` to generate one subplot that is twice as large.

```
--> t = linspace(-pi,pi);
--> subplot(2,2,[1,2])
--> plot(t,cos(t).*exp(-2*t));
--> subplot(2,2,3);
--> plot(t,cos(t*3).*exp(-2*t));
--> subplot(2,2,4);
--> plot(t,cos(t*4).*exp(-2*t));
```




Note that the subplots can contain any handle graphics objects, not only simple plots.

```
--> t=0:(2*pi/100):(2*pi);
--> x=cos(t*2).*(2+sin(t*3)*.3);
--> y=sin(t*2).*(2+sin(t*3)*.3);
--> z=cos(t*3)*.3;
--> subplot(2,2,1)
--> plot(t,x);
--> subplot(2,2,2);
--> plot(t,y);
--> subplot(2,2,3);
--> plot(t,z);
--> subplot(2,2,4);
--> tubeplot(x,y,z,0.14*sin(t*5)+.29,t,10)
--> axis equal
--> view(3)
```

Warning: Newly defined variable xlim shadows a function of the same name. Use `clear xlim` to recover ac

Warning: Newly defined variable ylim shadows a function of the same name. Use `clear ylim` to recover ac

Warning: Newly defined variable zlim shadows a function of the same name. Use `clear zlim` to recover ac



23.61 SURF Surface Plot Function

23.61.1 Usage

This routine is used to create a surface plot of data. A surface plot is a 3D surface defined by the xyz coordinates of its vertices and optionally by the color at the vertices. The most general syntax for the `surf` function is

```
h = surf(X,Y,Z,C,properties...)
```

Where `X` is a matrix or vector of `x` coordinates, `Y` is a matrix or vector of `y` coordinates, `Z` is a 2D matrix of coordinates, and `C` is a 2D matrix of color values (the colormap for the current fig is applied). In general, `X` and `Y` should be the same size as `Z`, but FreeMat will expand vectors to match the matrix if possible. If you want the color of the surface to be defined by the height of the surface, you can omit `C`

```
h = surf(X,Y,Z,properties...)
```

in which case `C=Z`. You can also eliminate the `X` and `Y` matrices in the specification

```
h = surf(Z,properties)
```

in which case they are set to `1:size(Z,2)` and `1:size(Y,2)` respectively. You can also specify a handle as the target of the `surf` command via

```
h = surf(handle,...)
```

23.61.2 Example

Here we generate a surface specifying all four components.


```
--> x = repmat(linspace(-1,1),[100,1]);
--> y = x';
--> r = x.^2+y.^2;
--> z = exp(-r*3).*cos(5*r);
--> c = r;
```

```
--> surf(x,y,z,c)
--> axis equal
--> view(3)
```

Warning: Newly defined variable xlim shadows a function of the same name. Use clear xlim to recover ac

Warning: Newly defined variable ylim shadows a function of the same name. Use clear ylim to recover ac

Warning: Newly defined variable zlim shadows a function of the same name. Use clear zlim to recover ac




If we allow FreeMat to specify the color component, we see that the colorfield is the same as the height

```
--> surf(x,y,z)
--> axis equal
--> view(3)
```

Warning: Newly defined variable xlim shadows a function of the same name. Use clear xlim to recover ac

Warning: Newly defined variable ylim shadows a function of the same name. Use clear ylim to recover ac

Warning: Newly defined variable zlim shadows a function of the same name. Use clear zlim to recover ac



23.62 SURFACEPROPERTIES Surface Object Properties

23.62.1 Usage

Below is a summary of the properties for the axis.

- **alphadata** - vector - This is a vector that should contain as many elements as the surface data itself **cdata**, or a single scalar. For a single scalar, all values of the surface take on the same transparency. Otherwise, the transparency of each pixel is determined by the corresponding value from the **alphadata** vector.
- **alphadatamapping** - {'scaled', 'direct', 'none'} - For **none** mode (the default), no transparency is applied to the data. For **direct** mode, the vector **alphadata** contains values between $@[0,M-1]$ — where **M** is the length of the alpha map stored in the figure. For **scaled** mode, the **alim** vector for the figure is used to linearly rescale the alpha data prior to lookup in the alpha map.
- **ambientstrength** - Not used.
- **backfacelighting** - Not used.
- **cdata** - array - This is either a **M x N** array or an **M x N x 3** array. If the data is **M x N** the surface is a scalar surface (indexed mode), where the color associated with each surface pixel is computed using the colormap and the **cdatamapping** mode. If the data is **M x N x 3** the surface is assumed to be in RGB mode, and the colorpanes are taken directly from **cdata** (the colormap is ignored). Note that in this case, the data values must be between $@[0,1]$ — for each color channel and each point on the surface.
- **cdatamapping** - {'scaled', 'direct'} - For **scaled** (the default), the pixel values are scaled using the **clim** vector for the figure prior to looking up in the colormap. For **direct** mode, the pixel values must be in the range $[0,N-1]$ where **N** is the number of colors in the colormap.
- **children** - Not used.
- **diffusestrength** - Not used.
- **edgealpha** - {'flat', 'interp', 'scalar'} - Controls how the transparency is mapped for the edges of the surface.
- **edgecolor** - {'flat', 'interp', 'none', 'colormap'} - Specifies how the edges are colored. For **'flat'** the edges are flat colored, meaning that the line segments that make up the edges are not shaded. The color for the line is determined by the first edge point it is connected to.
- **edgelighting** - Not used.
- **facealpha** - {'flat', 'interp', 'texturemap', 'scalar'} - Controls how the transparency of the faces of the surface are controlled. For flat shading, the faces are constant transparency. For **interp** mode, the faces are smoothly transparently mapped. If set to a scalar, all faces have the same transparency.
- **facecolor** - {'none', 'flat', 'interp', 'colormap'} - Controls how the faces are colored. For **'none'** the faces are uncolored, and the surface appears as a mesh without hidden lines removed. For **'flat'** the surface faces have a constant color. For **'interp'** smooth shading is applied to the surface. And if a colormap is provided, then the faces all have the same color.
- **facelighting** - Not used.

- `linestyle` - `{'-' , '--' , ':' , '-.' , 'none'}` - The style of the line used to draw the edges.
- `linewidth` - `scalar` - The width of the line used to draw the edges.
- `marker` - `{'+', 'o', '*', '.', 'x', 'square', 's', 'diamond', 'd', '^', 'v', '>', '<'}` - The marker for data points on the line. Some of these are redundant, as `'square'` `'s'` are synonyms, and `'diamond'` and `'d'` are also synonyms.
- `markeredgecolor` - `colorspec` - The color used to draw the marker. For some of the markers (circle, square, etc.) there are two colors used to draw the marker. This property controls the edge color (which for unfilled markers) is the primary color of the marker.
- `markerfacecolor` - `colorspec` - The color used to fill the marker. For some of the markers (circle, square, etc.) there are two colors used to fill the marker.
- `markersize` - `scalar` - Control the size of the marker. Defaults to 6, which is effectively the radius (in pixels) of the markers.
- `meshstyle` - `{'both', 'rows', 'cols'}` - This property controls how the mesh is drawn for the surface. For `rows` and `cols` modes, only one set of edges is drawn.
- `normalmode` - Not used.
- `parent` - `handle` - The axis containing the surface.
- `specularcolorreflectance` - Not used.
- `specularexponent` - Not used.
- `specularstrength` - Not used.
- `tag` - `string` - You can set this to any string you want.
- `type` - `string` - Set to the string `'surface'`.
- `userdata` - `array` - Available to store any variable you want in the handle object.
- `vertexnormals` - Not used.
- `xdata` - `array` - Must be a numeric array of size `M x N` which contains the x location of each point in the defined surface. Must be the same size as `ydata` and `zdata`. Alternately, you can specify an array of size `1 x N` in which case FreeMat replicates the vector to fill out an `M x N` matrix.
- `xdatamode` - `{'auto', 'manual'}` - When set to `auto` then FreeMat will automatically generate the x coordinates.
- `ydata` - `array` - Must be a numeric array of size `M x N` which contains the y location of each point in the defined surface. Must be the same size as `xdata` and `zdata`. Alternately, you can specify an array of size `M x 1` in which case FreeMat replicates the vector to fill out an `M x N` matrix.

- **ydatamode** - {'auto', 'manual'} - When set to `auto` then FreeMat will automatically generate the y coordinates.
- **zdata** - **array** - Must be a numeric array of size $M \times N$ which contains the y location of each point in the defined surface. Must be the same size as **xdata** and **ydata**.
- **visible** - {'on', 'off'} - Controls whether the surface is visible or not.

23.63 SURFACEPROPERTIES Surface Object Properties

23.63.1 Usage

Below is a summary of the properties for the axis.

- **alphadata** - **vector** - This is a vector that should contain as many elements as the surface data itself **cdata**, or a single scalar. For a single scalar, all values of the surface take on the same transparency. Otherwise, the transparency of each pixel is determined by the corresponding value from the **alphadata** vector.
- **alphadatamapping** - {'scaled', 'direct', 'none'} - For `none` mode (the default), no transparency is applied to the data. For `direct` mode, the vector **alphadata** contains values between $@[0, M-1]$ — where M is the length of the alpha map stored in the figure. For `scaled` mode, the **alim** vector for the figure is used to linearly rescale the alpha data prior to lookup in the alpha map.
- **ambientstrength** - Not used.
- **backfacelighting** - Not used.
- **cdata** - **array** - This is either a $M \times N$ array or an $M \times N \times 3$ array. If the data is $M \times N$ the surface is a scalar surface (indexed mode), where the color associated with each surface pixel is computed using the colormap and the **cdatamapping** mode. If the data is $M \times N \times 3$ the surface is assumed to be in RGB mode, and the colorpanes are taken directly from **cdata** (the colormap is ignored). Note that in this case, the data values must be between $@[0, 1]$ — for each color channel and each point on the surface.
- **cdatamapping** - {'scaled', 'direct'} - For `scaled` (the default), the pixel values are scaled using the **clim** vector for the figure prior to looking up in the colormap. For `direct` mode, the pixel values must be in the range $[0, N-1]$ where N is the number of colors in the colormap.
- **children** - Not used.
- **diffusestrength** - Not used.
- **edgealpha** - {'flat', 'interp', 'scalar'} - Controls how the transparency is mapped for the edges of the surface.
- **edgecolor** - {'flat', 'interp', 'none', 'colorspec'} - Specifies how the edges are colored. For `'flat'` the edges are flat colored, meaning that the line segments that make up the edges are not shaded. The color for the line is determined by the first edge point it is connected to.

- `edgelighting` - Not used.
- `facealpha` - `{'flat','interp','texturemap',scalar}` - Controls how the transparency of the faces of the surface are controlled. For flat shading, the faces are constant transparency. For `interp` mode, the faces are smoothly transparently mapped. If set to a scalar, all faces have the same transparency.
- `facecolor` - `{'none','flat','interp',colorspec}` - Controls how the faces are colored. For `'none'` the faces are uncolored, and the surface appears as a mesh without hidden lines removed. For `'flat'` the surface faces have a constant color. For `'interp'` smooth shading is applied to the surface. And if a `colorspec` is provided, then the faces all have the same color.
- `facelighting` - Not used.
- `linestyle` - `{'-' , '--' , ':' , '-.' , 'none'}` - The style of the line used to draw the edges.
- `linewidth` - `scalar` - The width of the line used to draw the edges.
- `marker` - `{'+', 'o', '*', '.', 'x', 'square', 's', 'diamond', 'd', '^', 'v', '>', '<'}` - The marker for data points on the line. Some of these are redundant, as `'square'` `'s'` are synonyms, and `'diamond'` and `'d'` are also synonyms.
- `markeredgecolor` - `colorspec` - The color used to draw the marker. For some of the markers (circle, square, etc.) there are two colors used to draw the marker. This property controls the edge color (which for unfilled markers) is the primary color of the marker.
- `markerfacecolor` - `colorspec` - The color used to fill the marker. For some of the markers (circle, square, etc.) there are two colors used to fill the marker.
- `markersize` - `scalar` - Control the size of the marker. Defaults to 6, which is effectively the radius (in pixels) of the markers.
- `meshstyle` - `{'both', 'rows', 'cols'}` - This property controls how the mesh is drawn for the surface. For `rows` and `cols` modes, only one set of edges is drawn.
- `normalmode` - Not used.
- `parent` - `handle` - The axis containing the surface.
- `specularcolorreflectance` - Not used.
- `specularexponent` - Not used.
- `specularstrength` - Not used.
- `tag` - `string` - You can set this to any string you want.
- `type` - `string` - Set to the string `'surface'`.
- `userdata` - `array` - Available to store any variable you want in the handle object.
- `vertexnormals` - Not used.

- **xdata** - **array** - Must be a numeric array of size $M \times N$ which contains the x location of each point in the defined surface. Must be the same size as **ydata** and **zdata**. Alternately, you can specify an array of size $1 \times N$ in which case FreeMat replicates the vector to fill out an $M \times N$ matrix.
- **xdatamode** - {'auto', 'manual'} - When set to **auto** then FreeMat will automatically generate the x coordinates.
- **ydata** - **array** - Must be a numeric array of size $M \times N$ which contains the y location of each point in the defined surface. Must be the same size as **xdata** and **zdata**. Alternately, you can specify an array of size $M \times 1$ in which case FreeMat replicates the vector to fill out an $M \times N$ matrix.
- **ydatamode** - {'auto', 'manual'} - When set to **auto** then FreeMat will automatically generate the y coordinates.
- **zdata** - **array** - Must be a numeric array of size $M \times N$ which contains the z location of each point in the defined surface. Must be the same size as **xdata** and **ydata**.
- **visible** - {'on', 'off'} - Controls whether the surface is visible or not.

23.64 TEXT Add Text Label to Plot

23.64.1 Usage

Adds a text label to the currently active plot. The general syntax for its use is either

```
text(x,y,'label')
```

where **x** and **y** are both vectors of the same length, in which case the text **'label'** is added to the current plot at each of the coordinates **x(i),y(i)** (using the current axis to map these to screen coordinates). The second form supplies a cell-array of strings as the second argument, and allows you to place many labels simultaneously

```
text(x,y,{'label1','label2',...})
```


where the number of elements in the cell array must match the size of vectors **x** and **y**. You can also specify properties for the labels via

```
handles = text(x,y,{labels},properties...)
```

23.64.2 Example


Here is an example of a few labels being added to a random plot:

```
--> plot(rand(1,4))
--> text([2,3],[0.5,0.5],{'hello','there'})
```



Here is the same example, but with larger labels:

```
--> plot(rand(1,4))
--> text([2,3],[0.5,0.5],{'hello','there'},'fontsize',20)
```



23.65 TEXTPROPERTIES Text Object Properties

23.65.1 Usage

Below is a summary of the properties for a text object.

- **boundingbox** - **four vector** - The size of the bounding box containing the text (in pixels). May contain negative values if the text is slanted.
- **children** - Not used.
- **string** - **string** - The text contained in the label.
- **extent** - Not used.
- **horizontalalignment** - **{'left','center','right'}** - Controls the alignment of the text relative to the specified position point.
- **position** - **three vector** - The position of the label in axis coordinates.
- **rotation** - **scalar** - The rotation angle (in degrees) of the label.

- `units` - Not used.
- `verticalalignment` - `{'top', 'bottom', 'middle'}` - Controls the alignment fo the text relative to the specified position point in the vertical position.
- `backgroundcolor` - `colormap` - The color used to fill in the background rectangle for the label. Normally this is `none`.
- `edgecolor` - `colormap` - The color used to draw the bounding rectangle for the label. Normally this is `none`.
- `linewidth` - `scalar` - The width of the line used to draw the border.
- `linestyle` - `{'-' , '--' , ':' , '-.' , 'none'}` - The style of the line used to draw the border.
- `margin` - `scalar` - The amount of spacing to place around the text as padding when drawing the rectangle.
- `fontangle` - `{'normal', 'italic', 'oblique'}` - The angle of the fonts used for the labels.
- `fontsize` - `scalar` - The size of fonts used for the text.
- `fontunits` - Not used.
- `fontweight` - `{'normal', 'bold', 'light', 'demi'}` - The weight of the font used for the label
- `visible` - `{'on', 'off'}` - Controls visibility of the the line.
- `color` - `colormap` - The color of the text of the label.
- `children` - Not used.
- `parent` - The handle of the axis that owns this label.
- `tag` - `string` - A string that can be used to tag the object.
- `type` - `string` - Returns the string `'text'`.
- `userdata` - `array` - Available to store any variable you want in the handle object.

23.66 TITLE Plot Title Function

23.66.1 Usage

This command adds a title to the plot. The general syntax for its use is

```
title('label')
```

or in the alternate form

```
title 'label'
```

or simply

```
title label
```

Here `label` is a string variable. You can also specify properties for the label, and a handle to serve as a target for the operation

```
title(handle,'label',properties...)
```


23.66.2 Example

Here is an example of a simple plot with a title.

```
--> x = linspace(-1,1);  
--> y = cos(2*pi*x);  
--> plot(x,y,'r-');  
--> title('cost over time');
```

Warning: Newly defined variable `nargin` shadows a function of the same name. Use `clear nargin` to recover.


which results in the following plot.



We now increase the size of the font using the properties of the `label`

```
--> title('cost over time','fontsize',20);
```

Warning: Newly defined variable `nargin` shadows a function of the same name. Use `clear nargin` to recover.



23.67 TUBE PLOT Creates a Tubeplot

23.67.1 Usage

This `tubeplot` function is from the `tubeplot` package written by Anders Sandberg. The simplest syntax for the `tubeplot` routine is

```
tubeplot(x,y,z)
```

plots the basic tube with radius 1, where `x,y,z` are vectors that describe the tube. If the radius of the tube is to be varied, use the second form

```
tubeplot(x,y,z,r)
```

which plots the basic tube with variable radius `r` (either a vector or a scalar value). The third form allows you to specify the coloring using a vector of values:

```
tubeplot(x,y,z,r,v)
```

where the coloring is now dependent on the values in the vector `v`. If you want to create a tube plot with a greater degree of tangential subdivisions (i.e., the tube is more circular, use the form

```
tubeplot(x,y,z,r,v,s)
```

where `s` is the number of tangential subdivisions (default is 6) You can also use `tubeplot` to calculate matrices to feed to `mesh` and `surf`.

```
[X,Y,Z]=tubeplot(x,y,z)
```

returns `N x 3` matrices suitable for `mesh` or `surf`.

Note that the tube may pinch at points where the normal and binormal misbehaves. It is suitable for general space curves, not ones that contain straight sections. Normally the tube is calculated using the Frenet frame, making the tube minimally twisted except at inflexion points.

To deal with this problem there is an alternative frame:


```
tubeplot(x,y,z,r,v,s,vec)
```

calculates the tube by setting the normal to the cross product of the tangent and the vector `vec`. If it is chosen so that it is always far from the tangent vector the frame will not twist unduly.

23.67.2 Example

Here is an example of a `tubeplot`.

```
--> t=0:(2*pi/100):(2*pi);
--> x=cos(t*2).*(2+sin(t*3)*.3);
--> y=sin(t*2).*(2+sin(t*3)*.3);
--> z=cos(t*3)*.3;
--> tubeplot(x,y,z,0.14*sin(t*5)+.29,t,10);
```



Written by Anders Sandberg, asa@nada.kth.se, 2005 Website says the package is free for anybody to use. www.nada.kth.se/asa/Ray/Tubeplot/tubeplot.html

23.68 UICONTROL Create a UI Control object

23.68.1 Usage

Creates a UI control object and parents it to the current figure. The syntax for its use is

```
handle = uicontrol(property,value,property,value,...)
```

where `property` and `value` are set. The handle ID for the resulting object is returned. It is automatically added to the children of the current figure.

23.69 UICONTROLPROPERTIES UI Control Properties

23.69.1 Usage

Below is a summary of the properties for user interface controls.

- `backgroundcolor` - `colspec` - The background color for the widget.
- `busyaction` - Not used.
- `buttondownfcn` - Not used.
- `callback` - `string` - the callback to execute when the GUI control does its action. Clicking a button or moving a scroller will cause the callback to be executed. Also, pressing enter in a text box causes the callback to be executed.
- `cdata` - an $M \times N \times 3$ array that represents an RGB image to use as the truecolor image displayed on push buttons or toggle buttons. The values must be between 0 and 1.
- `children` - Not used.
- `createfcn` - Not used.
- `deletefcn` - Not used;

- **enable** - {'on', 'inactive', 'off'} - For **on** (the default) the uicontrol behaves normally. For **inactive**, it is not operational, but looks the same as **on**. For **off**, the control is grayed out.
- **extent** - a read only property that contains the extent of the text for the control.
- **fontangle** - {'normal', 'italic', 'oblique'} - The angle of the fonts used for text labels (e.g., tick labels).
- **fontsize** - scalar - The size of fonts used for text labels (tick labels).
- **fontunits** - Not used.
- **fontname** - string - The name of the font to use for the widget.
- **fontweight** - {'normal', 'bold', 'light', 'demi'} - The weight of the font used
- **foregroundcolor** - colorspec - the foreground color for text.
- **handlevisibility** - Not used.
- **hitest** - Not used.
- **horizontalalignment** - {'left', 'center', 'right'} - determines the justification of text.
- **interruptible** - Not used.
- **keypressfcn** - functionspec - a string or function handle that is called when a key is pressed and a uicontrol object has focus.
- **listboxtop** - a scalar (used only by the listbox style of uicontrols) that specifies which string appears at the top of the list box.
- **max** - a scalar that specifies the largest value allowed for the **value** property. The interpretation varies depending on the type of the control
 - **check boxes** - specifies what **value** is set to when the check box is selected.
 - **edit box** - if **max-min>1** then the text box allows for multiple lines of input. Otherwise, it is a single line only.
 - **list box** - if **max-min>1** then multiple item selections are allowed. Otherwise, only single item selections are allowed.
 - **radio buttons** - specifies what **value** is set to when the radio button is selected.
 - **slider** - the maximum value the slider can take.
 - **toggle button** - specifies what **value** is set to when the toggle button is selected.
- **min** - a scalar that specifies the smallest value for the **value** property. The interpretation of it depends on the type of the control
 - **check boxes** - specifies what **value** is set to when the check box is not selected.

- **edit box** - if **max-min**>1 then the text box allows for multiple lines of input. Otherwise, it is a single line only.
- **list box** - if **max-min**>1 then multiple item selections are allowed. Otherwise, only single item selections are allowed.
- **radio buttons** - specifies what **value** is set to when the radio button is not selected.
- **slider** - the minimum value the slider can take.
- **toggle button** - specifies what **value** is set to when the toggle button is not selected.
- **parent** - the handle of the parent object.
- **position** - size and location of the uicontrol as a four vector [**left**, **bottom**, **width**, **height**]. If **width**>**height** then sliders are horizontal, otherwise the slider is oriented vertically.
- **selected** - {'on', 'off'} - not used.
- **selectionhighlight** - {'on', 'off'} - not used.
- **sliderstep** - a two vector [**min_step** **max_step**] that controls the amount the slider **value** changes when you click the mouse on the control. If you click the arrow for the slider, the value changes by **min_step**, while if you click the trough, the value changes by **max_step**. Each value must be in the range [0,1], and is a percentage of the range **max-min**.
- **string** - **string** - the text for the control.
- **style** - @—'pushbutton', 'toggle', 'radiobutton', 'checkbox', 'edit', 'text', 'slider', 'frame', 'listbox', 'popupmenu'—.
- **tag** - **string** - user specified label.
- **tooltipstring** - **string** the tooltip for the control.
- **type** - **string** - the text is set to 'uicontrol'.
- **uicontextmenu** - **handle** the handle of the **uicontextmenu** that shows up when you right-click over the control.
- **units** - not used.
- **userdata** - **array** - any data you want to associate with the control.
- **value** - The meaning of this property depends on the type of the control:
 - check box - set to **max** when checked, and **min** when off.
 - list box - set to a vector of indices corresponding to selected items, with 1 corresponding to the first item in the list.
 - pop up menu - set to the index of the item selected (starting with 1)
 - radio buttons - set to **max** when selected, and set to **min** when not selected.
 - sliders - set to the value of the slider
 - toggle buttons - set to **max** when selected, and set to **min** when not selected.
 - text controls, push buttons - do not use this property.
- **visible** - {'on', 'off'} - controls whether the control is visible or not

23.70 VIEW Set Graphical View

23.70.1 Usage

The `view` function sets the view into the current plot. The simplest form is

```
view(n)
```

where `n=2` sets a standard view (azimuth 0 and elevation 90), and `n=3` sets a standard 3D view (azimuth 37.5 and elevation 30). With two arguments,

```
view(az,el)
```

you set the viewpoint to azimuth `az` and elevation `el`.

23.70.2 Example


Here is a 3D surface plot shown with a number of viewpoints. First, the default view for a 3D plot.

```
--> x = repmat(linspace(-1,1),[100,1]);
--> y = x';
--> r = x.^2+y.^2;
--> z = exp(-r*3).*cos(5*pi*r);
--> surf(x,y,z);
--> axis equal
--> view(3)
```

Warning: Newly defined variable xlim shadows a function of the same name. Use `clear xlim`

Warning: Newly defined variable ylim shadows a function of the same name. Use `clear ylim`

Warning: Newly defined variable zlim shadows a function of the same name. Use `clear zlim`




Next, we look at it as a 2D plot

```
--> surf(x,y,z);
--> axis equal
--> view(2)
```

Warning: Newly defined variable xlim shadows a function of the same name. Use `clear xlim`

Warning: Newly defined variable ylim shadows a function of the same name. Use `clear ylim`

Warning: Newly defined variable zlim shadows a function of the same name. Use `clear zlim`




Finally, we generate a different view of the same surface.

```
--> surf(x,y,z);
--> axis equal
--> view(25,50);
```

Warning: Newly defined variable xlim shadows a function of the same name. Use clear xlim to recover ac

Warning: Newly defined variable ylim shadows a function of the same name. Use clear ylim to recover ac

Warning: Newly defined variable zlim shadows a function of the same name. Use clear zlim to recover ac



23.71 WINLEV Image Window-Level Function

23.71.1 Usage

Adjusts the data range used to map the current image to the current colormap. The general syntax for its use is

```
winlev(window,level)
```

where `window` is the new window, and `level` is the new level, or

```
winlev
```

in which case it returns a vector containing the current window and level for the active image.

23.71.2 Function Internals

FreeMat deals with scalar images on the range of $[0,1]$, and must therefore map an arbitrary image x to this range before it can be displayed. By default, the `image` command chooses

$$\text{window} = \max x - \min x,$$

and

$$\text{level} = \frac{\text{window}}{2}$$

This ensures that the entire range of image values in x are mapped to the screen. With the `winlev` function, you can change the range of values mapped. In general, before display, a pixel x is mapped to $[0,1]$ via:

$$\max\left(0, \min\left(1, \frac{x - \text{level}}{\text{window}}\right)\right)$$

23.71.3 Examples

The window level function is fairly easy to demonstrate. Consider the following image, which is a Gaussian pulse image that is very narrow:

```
--> t = linspace(-1,1,256);
--> xmat = ones(256,1)*t; ymat = xmat';
--> A = exp(-(xmat.^2 + ymat.^2)*100);
--> image(A);
```

The data range of A is $[0,1]$, as we can verify numerically:

```
--> min(A(:))
```

```
ans =
  1.3839e-87
```

```
--> max(A(:))
```

```
ans =
  0.9969
```

To see the tail behavior, we use the `winlev` command to force FreeMat to map a smaller range of A to the colormap.

```
--> image(A);
--> winlev(1e-4,0.5e-4)
```

The result is a look at more of the tail behavior of A . We can also use the `winlev` function to find out what the window and level are once set, as in the following example.

```
--> image(A);
--> winlev(1e-4,0.5e-4)
--> winlev
```

Warning: Newly defined variable `clim` shadows a function of the same name. Use `clear clim` to recover ac

```
ans =  
1.0000e-04
```

23.72 XLABEL Plot X-axis Label Function

23.72.1 Usage

This command adds a label to the x-axis of the plot. The general syntax for its use is

```
xlabel('label')
```

or in the alternate form

```
xlabel 'label'
```

or simply

```
xlabel label
```

Here `label` is a string variable. You can also specify properties for that label using the syntax


```
xlabel('label',properties...)
```

23.72.2 Example

Here is an example of a simple plot with a label on the x-axis.

```
--> x = linspace(-1,1);  
--> y = cos(2*pi*x);  
--> plot(x,y,'r-');  
--> xlabel('time');
```

which results in the following plot.



23.73 XLIM Adjust X Axis limits of plot

23.73.1 Usage

There are several ways to use `xlim` to adjust the X axis limits of a plot. The various syntaxes are

```
xlim
xlim([lo,hi])
xlim('auto')
xlim('manual')
xlim('mode')
xlim(handle,...)
```


The first form (without arguments), returns a 2-vector containing the current limits. The second form sets the limits on the plot to `[lo,hi]`. The third and fourth form set the mode for the limit to `auto` and `manual` respectively. In `auto` mode, FreeMat chooses the range for the axis automatically. The `xlim('mode')` form returns the current mode for the axis (either `'auto'` or `'manual'`). Finally, you can specify the handle of an axis to manipulate instead of using the current one.

23.73.2 Example

```
--> x = linspace(-1,1);
--> y = sin(2*pi*x);
--> plot(x,y,'r-');
--> xlim % what are the current limits?
```

```
ans =
-1 1
```


which results in



Next, we zoom in on the plot using the `xlim` function

```
--> plot(x,y,'r-')
--> xlim([-0.2,0.2])
```

which results in



23.74 YLABEL Plot Y-axis Label Function

23.74.1 Usage

This command adds a label to the y-axis of the plot. The general syntax for its use is

```
ylabel('label')
```

or in the alternate form

```
ylabel 'label'
```

or simply

```
ylabel label
```

You can also specify properties for that label using the syntax


```
ylabel('label',properties...)
```

23.74.2 Example

Here is an example of a simple plot with a label on the y-axis.

```
--> x = linspace(-1,1);  
--> y = cos(2*pi*x);  
--> plot(x,y,'r-');  
--> ylabel('cost');
```

which results in the following plot.



23.75 YLIM Adjust Y Axis limits of plot

23.75.1 Usage

There are several ways to use `ylim` to adjust the Y axis limits of a plot. The various syntaxes are

```
ylim
ylim([lo,hi])
ylim('auto')
ylim('manual')
ylim('mode')
ylim(handle,...)
```


The first form (without arguments), returns a 2-vector containing the current limits. The second form sets the limits on the plot to `[lo,hi]`. The third and fourth form set the mode for the limit to `auto` and `manual` respectively. In `auto` mode, FreeMat chooses the range for the axis automatically. The `ylim('mode')` form returns the current mode for the axis (either `'auto'` or `'manual'`). Finally, you can specify the handle of an axis to manipulate instead of using the current one.

23.75.2 Example

```
--> x = linspace(-1,1);
--> y = sin(2*pi*x);
--> plot(x,y,'r-');
--> ylim % what are the current limits?
```

```
ans =
   -0.9999    0.9999
```


which results in



Next, we zoom in on the plot using the `ylim` function

```
--> plot(x,y,'r-')
--> ylim([-0.2,0.2])
```

which results in



23.76 ZLABEL Plot Z-axis Label Function

23.76.1 Usage

This command adds a label to the z-axis of the plot. The general syntax for its use is

```
zlabel('label')
```

or in the alternate form

```
zlabel 'label'
```

or simply

```
zlabel label
```

Here `label` is a string variable. You can also specify properties for that label using the syntax


```
zlabel('label',properties...)
```

23.76.2 Example

Here is an example of a simple plot with a label on the z-axis.

```
--> t = linspace(0,5*pi);
--> x = cos(t);
--> y = sin(t);
--> z = t;
--> plot3(x,y,z,'r-');
Warning: Newly defined variable xlim shadows a function of the same name. Use clear xlim
Warning: Newly defined variable ylim shadows a function of the same name. Use clear ylim
Warning: Newly defined variable zlim shadows a function of the same name. Use clear zlim
--> view(3);
Warning: Newly defined variable xlim shadows a function of the same name. Use clear xlim
Warning: Newly defined variable ylim shadows a function of the same name. Use clear ylim
Warning: Newly defined variable zlim shadows a function of the same name. Use clear zlim
--> zlabel('time');
```

which results in the following plot.



23.77 ZLIM Adjust Z Axis limits of plot

23.77.1 Usage

There are several ways to use `zlim` to adjust the Z axis limits of a plot. The various syntaxes are

```
zlim
zlim([lo,hi])
zlim('auto')
zlim('manual')
zlim('mode')
zlim(handle,...)
```

The first form (without arguments), returns a 2-vector containing the current limits. The second form sets the limits on the plot to `[lo,hi]`. The third and fourth form set the mode for the limit to `auto` and `manual` respectively. In `auto` mode, FreeMat chooses the range for the axis automatically.


The `zlim('mode')` form returns the current mode for the axis (either `'auto'` or `'manual'`). Finally, you can specify the handle of an axis to manipulate instead of using the current one.

23.77.2 Example

```
--> x = linspace(-1,1);  
--> y = sin(2*pi*x);  
--> plot(x,y,'r-');  
--> zlim % what are the current limits?
```

```
ans =  
-0.5000    0.5000
```


which results in



Next, we zoom in on the plot using the `zlim` function

```
--> plot(x,y,'r-')  
--> zlim([-0.2,0.2])
```

which results in



23.78 ZOOM Image Zoom Function

23.78.1 Usage

This function changes the zoom factor associated with the currently active image. It is a legacy support function only, and thus is not quite equivalent to the `zoom` function from previous versions of FreeMat. However, it should achieve roughly the same effect. The generic syntax for its use is

```
zoom(x)
```


where `x` is the zoom factor to be used. The exact behavior of the zoom factor is as follows:

- `x>0` The image is zoomed by a factor `x` in both directions.
- `x=0` The image on display is zoomed to fit the size of the image window, but the aspect ratio of the image is not changed. (see the Examples section for more details). This is the default zoom level for images displayed with the `image` command.
- `x<0` The image on display is zoomed to fit the size of the image window, with the zoom factor in the row and column directions chosen to fill the entire window. The aspect ratio of the image is not preserved. The exact value of `x` is irrelevant.

23.78.2 Example

To demonstrate the use of the `zoom` function, we create a rectangular image of a Gaussian pulse. We start with a display of the image using the `image` command, and a zoom of 1.


```
--> x = linspace(-1,1,300)'*ones(1,600);
--> y = ones(300,1)*linspace(-1,1,600);
--> Z = exp(-(x.^2+y.^2)/0.3);
--> image(Z);
--> zoom(1.0);
```



At this point, resizing the window accomplishes nothing, as with a zoom factor greater than zero, the size of the image is fixed.


If we change the zoom to another factor larger than 1, we enlarge the image by the specified factor (or shrink it, for zoom factors $0 < x < 1$). Here is the same image zoomed out to 60

```
--> image(Z);
--> zoom(0.6);
```




Similarly, we can enlarge it to 130

```
--> image(Z)
--> zoom(1.3);
```




The “free” zoom of $x = 0$ results in the image being zoomed to fit the window without changing the aspect ratio. The image is zoomed as much as possible in one direction.

```
--> image(Z);
--> zoom(0);
--> sizefig(200,400);
```



The case of a negative zoom $x < 0$ results in the image being scaled arbitrarily. This allows the image aspect ratio to be changed, as in the following example.

```
--> image(Z);  
--> zoom(-1);  
--> sizefig(200,400);
```



23.79 ZPLANE Zero-pole plot

23.79.1 Usage

This function makes a zero-pole plot of a discrete-time system defined by its zeros and poles. The various syntaxes are

`zplane(z,p)`

where `z` and `p` are the zeros and the poles of the system stored as column vectors, or

`zplane(b,a)`

where `a` and `b` are the polynomial coefficients of the numerator and denominator stored as line vectors (`roots` is used to find the zeros and poles). The symbol 'o' represents a zero and the symbol 'x' represents a pole. The plot includes the unit circle for reference. Contributed by Paulo Xavier Candeias under GPL

Chapter 24

OpenGL Models

24.1 GLASSEMBLY Create a GL Assembly

24.1.1 Usage

Define a GL Assembly. A GL Assembly consists of one or more GL Nodes or GL Assemblies that are placed relative to the coordinate system of the assembly. For example, if we have `glnode` definitions for 'bread' and 'cheese', then a `glassembly` of sandwich would consist of placements of two 'bread' nodes with a 'cheese' node in between. Furthermore, a 'lunch' assembly could consist of a 'sandwich' a 'chips' and 'soda'. Hopefully, you get the idea. The syntax for the `glassembly` command is

```
glassembly(name,part1,transform1,part2,transform2,...)
```

where `part1` is the name of the first part, and could be either a `glnode` or itself be another `glassembly`. Here `transform1` is the 4 x 4 matrix that transforms the part into the local reference coordinate system.

WARNING!! Currently FreeMat does not detect or gracefully handle self-referential assemblies (i.e, if you try to make a `sandwich` contain a `sandwich`, which you can do by devious methods that I refuse to explain). Do not do this! You have been warned.

24.2 GLCLUMP Create a GL Clump

24.2.1 Usage

Defines an aggregate clump of objects that can be treated as a node. A GL Clump is defined by a vector consisting of the following elements:

```
[r1 g1 b1 n1 p1 p2 p3 ... r2 g2 b2 n2 p1 p2 p3 ... ]
```

i.e., an RGB color spec, followed by a point count `ni`, followed by a length `ni` vector of coordinates that are `x,y,z` triplets. The usage of this function is

```
glclump(name,vector)
```

where `name` is the name of the clump and `vector` is the aforementioned vector of points.

24.3 GLDEFMATERIAL Defines a GL Material

24.3.1 Usage

Define a material. The syntax for its use is

```
gldefmaterial(name,ambient,diffuse,specular,shininess)
```

where **name** is the name of the material, and **ambient** is a 4 x 1 vector containing the ambient component of the material property, and **diffuse** is a 4 x 1 vector and **specular** is a 4 x 1 vector containing the specular component of the material properties and **shininess** is the exponent that governs the shininess of the material.

24.4 GLLINES Create a GL Lineset

24.4.1 Usage

Defines a set of lines that can be treated as a node. A GL Lines is defined by a vector consisting of the following elements:

```
[m1 x1 y1 z1 ... xn yn zn m2 x1 y1 z1 .... ]
```

i.e., a point count followed by that number of triplets. The usage of this function is

```
gllines(name,vector,color)
```

where **name** is the name of the lineset and **vector** is the aforementioned vector of points.

24.5 GLNODE Create a GL Node

24.5.1 Usage

Define a GL Node. A GL Node is an object that can be displayed in a GL Window. It is defined by a triangular mesh of vertices. It must also have a material that defines its appearance (i.e. color, shininess, etc.). The syntax for the **glnode** command is

```
glnode(name,material,pointset)
```

where **material** is the name of a material that has already been defined with **gldefmaterial**, **pointset** is a 3 x N matrix of points that define the geometry of the object. Note that the points are assumed to be connected in triangular facets, with the points defined counter clock-wise as seen from the outside of the facet. **FreeMat** will compute the normals. The **name** argument must be unique. If you want multiple instances of a given **glnode** in your GLWindow, that is fine, as instances of a **glnode** are created through a **glassembly**.

Chapter 25

Object Oriented Programming

25.1 AND Overloaded Logical And Operator

25.1.1 Usage

This is a method that is invoked to combine two variables using a logical and operator, and is invoked when you call

```
c = and(a,b)
```

or for

```
c = a & b
```

25.2 CAT Concatenation of Arrays

25.2.1 Usage

This function concatenates arrays in a given dimension. The syntax for its use is

```
cat (DIM, A, B)  
cat (DIM, A, B, C ...)
```

to return the concatenation along the dimension DIM of all arguments. `cat(1, A, B, C)` is the same as `[A; B; C]` or `vertcat(A, B, C)`. `cat(2, A, B, C)` is the same as `[A, B, C]` or `horzcat(A, B, C)`.

25.3 CLASS Class Support Function

25.3.1 Usage

There are several uses for the `class` function. The first version takes a single argument, and returns the class of that variable. The syntax for this form is

```
classname = class(variable)
```

and it returns a string containing the name of the class for `variable`. The second form of the `class` function is used to construct an object of a specific type based on a structure which contains data elements for the class. The syntax for this version is

```
classvar = class(template, classname, parent1, parent2,...)
```

This should be called inside the constructor for the class. The resulting class will be of the type `classname`, and will be derived from `parent1`, `parent2`, etc. The `template` argument should be a structure array that contains the members of the class. See the `constructors` help for some details on how to use the `class` function. Note that if the `template` argument is an empty structure matrix, then the resulting variable has no fields beyond those inherited from the parent classes.

25.4 COLON Overloaded Colon Operator

25.4.1 Usage

This is a method that is invoked in one of two forms, either the two argument version

```
c = colon(a,b)
```

which is also called using the notation

```
c = a:b
```

and the three argument version

```
d = colon(a,b,c)
```

which is also called using the notation

```
d = a:b:c
```

25.5 CONSTRUCTORS Class Constructors

25.5.1 Usage

When designing a constructor for a FreeMat class, you should design the constructor to take a certain form. The following is the code for the sample `mat` object

```
function p = mat(a)
    if (nargin == 0)
        p.c = [];
        p = class(p,'mat');
    elseif isa(a,'mat')
        p = a;
    else
        p.c = a;
        p = class(p,'mat');
    end
```

Generally speaking when it is provided with zero arguments, the constructor returns a default version of the class using a template structure with the right fields populated with default values. If the constructor is given a single argument that matches the class we are trying to construct, the constructor passes through the argument. This form of the constructor is used for type conversion. In particular,

```
p = mat(a)
```

guarantees that `p` is an array of class `mat`. The last form of the constructor builds a class object given the input. The meaning of this form depends on what makes sense for your class. For example, for a polynomial class, you may want to pass in the coefficients of the polynomial.

25.6 CTRANSPOSE Overloaded Conjugate Transpose Operator

25.6.1 Usage

This is a method that is invoked when a variable has the conjugate transpose operator method applied, and is invoked when you call

```
c = ctranspose(a)
```

or

```
/ c = a'
```

25.7 EQ Overloaded Equals Comparison Operator

25.7.1 Usage

This is a method that is invoked to combine two variables using an equals comparison operator, and is invoked when you call

```
c = eq(a,b)
```

or for

```
c = a == b
```

25.8 GE Overloaded Greater-Than-Equals Comparison Operator

25.8.1 Usage

This is a method that is invoked to combine two variables using a greater than or equals comparison operator, and is invoked when you call

```
c = ge(a,b)
```

or for

```
c = a >= b
```

25.9 GT Overloaded Greater Than Comparison Operator

25.9.1 Usage

This is a method that is invoked to combine two variables using a greater than comparison operator, and is invoked when you call

```
c = gt(a,b)
```

or for

```
c = a > b
```

25.10 HORZCAT Horizontal Array Concatenation

25.10.1 Usage

This function concatenates arrays horizontally (along the column dimension). The syntax for its use is

```
d = horzcat(a,b,c)
```

which is equivalent to the statement `d = [a,b,c]`.

25.11 HORZCAT Horizontal Array Concatenation

25.11.1 Usage

This function concatenates arrays horizontally (along the column dimension). The syntax for its use is

```
d = horzcat(a,b,c)
```

which is equivalent to the statement `d = [a,b,c]`.

25.12 LDIVIDE Overloaded Left Divide Operator

25.12.1 Usage

This is a method that is invoked when two variables are divided and is invoked when you call

```
c = ldivide(a,b)
```

or for

```
c = a .\ b
```

25.13 LE Overloaded Less-Than-Equals Comparison Operator

25.13.1 Usage

This is a method that is invoked to compare two variables using a less than or equals comparison operator, and is invoked when you call

```
c = le(a,b)
```

or for

```
c = a <= b
```

25.14 LT Overloaded Less Than Comparison Operator

25.14.1 Usage

This is a method that is invoked to compare two variables using a less than comparison operator, and is invoked when you call

```
c = lt(a,b)
```

or for

```
c = a < b
```

25.15 MINUS Overloaded Addition Operator

25.15.1 Usage

This is a method that is invoked when two variables are subtracted and is invoked when you call

```
c = minus(a,b)
```

or for

```
c = a - b
```

25.16 MLDIVIDE Overloaded Matrix Left Divide Operator

25.16.1 Usage

This is a method that is invoked when two variables are divided using the matrix (left) divide operator, and is invoked when you call

```
c = mldivide(a,b)
```

or for

```
c = a \ b
```

25.17 MPOWER Overloaded Matrix Power Operator

25.17.1 Usage

This is a method that is invoked when one variable is raised to another variable using the matrix power operator, and is invoked when you call

```
c = mpower(a,b)
```

or

```
c = a^b
```

25.18 MRDIVIDE Overloaded Matrix Right Divide Operator

25.18.1 Usage

This is a method that is invoked when two variables are divided using the matrix divide operator, and is invoked when you call

```
c = mrdivide(a,b)
```

or for

```
c = a / b
```

25.19 MTIMES Overloaded Matrix Multiplication Operator

25.19.1 Usage

This is a method that is invoked when two variables are multiplied using the matrix operator and is invoked when you call

```
c = mtimes(a,b)
```

or for

```
c = a * b
```

25.20 NE Overloaded Not-Equals Comparison Operator

25.20.1 Usage

This is a method that is invoked to combine two variables using a not-equals comparison operator, and is invoked when you call

```
c = ne(a,b)
```

or for

```
c = a != b
```

25.21 NOT Overloaded Logical Not Operator

25.21.1 Usage

This is a method that is invoked when a variable is logically inverted, and is invoked when you call

```
c = not(a)
```

or for

```
c = ~a
```

25.22 OR Overloaded Logical Or Operator

25.22.1 Usage

This is a method that is invoked to combine two variables using a logical or operator, and is invoked when you call

```
c = or(a,b)
```

or for

```
c = a | b
```

25.23 PLUS Overloaded Addition Operator

25.23.1 Usage

This is a method that is invoked when two variables are added and is invoked when you call

```
c = plus(a,b)
```

or for

```
c = a + b
```

25.24 POWER Overloaded Power Operator

25.24.1 Usage

This is a method that is invoked when one variable is raised to another variable using the dot-power operator, and is invoked when you call

```
c = power(a,b)
```

or

```
c = a.^b
```

25.25 RDIVIDE Overloaded Right Divide Operator

25.25.1 Usage

This is a method that is invoked when two variables are divided and is invoked when you call

```
c = rdivide(a,b)
```

or for

```
c = a ./ b
```

25.26 SUBSASGN Overloaded Class Assignment

25.26.1 Usage

This method is called for expressions of the form

```
a(b) = c, a{b} = c, a.b = c
```

and overloading the `subsasgn` method can allow you to define the meaning of these expressions for objects of class `a`. These expressions are mapped to a call of the form

```
a = subsasgn(a,s,b)
```

where `s` is a structure array with two fields. The first field is

- `type` is a string containing either `'()'` or `'{}'` or `'.'` depending on the form of the call.
- `subs` is a cell array or string containing the the subscript information.

When multiple indexing expressions are combined together such as `a(5).foo{:} = b`, the `s` array contains the following entries

```
s(1).type = '()'   s(1).subs = {5}
s(2).type = '.'   s(2).subs = 'foo'
s(3).type = '{}'  s(3).subs = ':'
```

25.27 SUBSINDEX Overloaded Class Indexing

25.27.1 Usage

This method is called for classes in the expressions of the form

```
c = subsindex(a)
```

where `a` is an object, and `c` is an index vector. It is also called for

```
c = b(a)
```

in which case `subsindex(a)` must return a vector containing integers between 0 and `N-1` where `N` is the number of elements in the vector `b`.

25.28 SUBSREF Overloaded Class Indexing

25.28.1 Usage

This method is called for expressions of the form

```
c = a(b), c = a{b}, c = a.b
```

and overloading the `subsref` method allows you to define the meaning of these expressions for objects of class `a`. These expressions are mapped to a call of the form

```
b = subsref(a,s)
```

where `s` is a structure array with two fields. The first field is

- `type` is a string containing either `'()'` or `'{'` or `','` depending on the form of the call.
- `subs` is a cell array or string containing the the subscript information.

When multiple indexing expressions are combined together such as `b = a(5).foo{:}`, the `s` array contains the following entries

```
s(1).type = '()'   s(1).subs = {5}
s(2).type = ','   s(2).subs = 'foo'
s(3).type = '{'   s(3).subs = ':'
```

25.29 TIMES Overloaded Multiplication Operator

25.29.1 Usage

This is a method that is invoked when two variables are multiplied and is invoked when you call

```
c = times(a,b)
```

or for

```
c = a .* b
```

25.30 TRANSPOSE Overloaded Transpose Operator

25.30.1 Usage

This is a method that is invoked when a variable has the transpose operator method applied, and is invoked when you call

```
c = transpose(a)
```

or

```
/ c = a.'
```

25.31 UMINUS Overloaded Unary Minus Operator

25.31.1 Usage

This is a method that is invoked when a variable is negated, and is invoked when you call

```
c = uminus(a)
```

or for

```
c = -a
```

25.32 UPLUS Overloaded Unary Plus Operator

25.32.1 Usage

This is a method that is invoked when a variable is preceded by a "+", and is invoked when you call

```
c = uplus(a)
```

or for

```
c = +a
```

25.33 VERTCAT Horizontal Array Concatenation

25.33.1 Usage

This function concatenates arrays vertically (along the row dimension). The syntax for its use is

```
d = vertcat(a,b,c)
```

which is equivalent to the statement `d = [a;b;c]`.

25.34 VERTCAT Horizontal Array Concatenation

25.34.1 Usage

This function concatenates arrays vertically (along the row dimension). The syntax for its use is

```
d = vertcat(a,b,c)
```

which is equivalent to the statement `d = [a;b;c]`.

Chapter 26

Bitwise Operations

26.1 BITAND Bitwise Boolean And Operation

26.1.1 Usage

Performs a bitwise binary and operation on the two arguments and returns the result. The syntax for its use is

```
y = bitand(a,b)
```

where **a** and **b** are multi-dimensional unsigned integer arrays. The and operation is performed using 32 bit unsigned intermediates. Note that if **a** or **b** is a scalar, then each element of the other array is and'ed with that scalar. Otherwise the two arrays must match in size.

26.1.2 Example

```
--> bitand(uint16([1,16,255]),uint16([3,17,128]))
```

```
ans =  
    1  16 128
```

```
--> bitand(uint16([1,16,255]),uint16(3))
```

```
ans =  
    1  0  3
```

26.2 BITCMP Bitwise Boolean Complement Operation

26.2.1 Usage

Usage

Performs a bitwise binary complement operation on the argument and returns the result. The syntax for its use is

```
y = bitcmp(a)
```

where `a` is an unsigned integer arrays. This version of the command uses as many bits as required by the type of `a`. For example, if `a` is an `uint8` type, then the complement is formed using 8 bits. The second form of `bitcmp` allows you to specify the number of bits to use,

```
y = bitcmp(a,n)
```

in which case the complement is taken with respect to `n` bits, where `n` must be less than the length of the integer type.

26.2.2 Example

```
--> bitcmp(uint16(2^14-2))
```

```
ans =
    49153
```

```
--> bitcmp(uint16(2^14-2),14)
```

```
ans =
     1
```

26.3 BITOR Bitwise Boolean Or Operation

26.3.1 Usage

Performs a bitwise binary or operation on the two arguments and returns the result. The syntax for its use is

```
y = bitor(a,b)
```

where `a` and `b` are multi-dimensional unsigned integer arrays. The `and` operation is performed using 32 bit unsigned intermediates. Note that if `a` or `b` is a scalar, then each element of the other array is or'ed with that scalar. Otherwise the two arrays must match in size.

26.3.2 Example

```
--> bitand(uint16([1,16,255]),uint16([3,17,128]))
```

```
ans =
     1  16 128
```

```
--> bitand(uint16([1,16,255]),uint16(3))
```

```
ans =
     1  0  3
```

26.4 BITXOR Bitwise Boolean Exclusive-Or (XOR) Operation

26.4.1 Usage

Performs a bitwise binary xor operation on the two arguments and returns the result. The syntax for its use is

```
y = bitxor(a,b)
```

where **a** and **b** are multi-dimensional unsigned integer arrays. The and operation is performed using 32 bit unsigned intermediates. Note that if **a** or **b** is a scalar, then each element of the other array is xor'ed with that scalar. Otherwise the two arrays must match in size.

26.4.2 Example

```
--> bitand(uint16([1,16,255]),uint16([3,17,128]))
```

```
ans =
```

```
1 16 128
```

```
--> bitand(uint16([1,16,255]),uint16(3))
```

```
ans =
```

```
1 0 3
```


Chapter 27

FreeMat Threads

27.1 THREADCALL Call Function In A Thread

27.1.1 Usage

The `threadcall` function is a convenience function for executing a function call in a thread. The syntax for its use is

```
[val1,...,valn] = threadcall(threadid,timeout,funcname,arg1,arg2,...)
```

where `threadid` is the ID of the thread (as returned by the `threadnew` function), `funcname` is the name of the function to call, and `argi` are the arguments to the function, and `timeout` is the amount of time (in milliseconds) that the function is allowed to take.

27.1.2 Example

Here is an example of executing a simple function in a different thread.

```
--> id = threadnew

id =
  3

--> d = threadcall(id,1000,'cos',1.02343)

d =
  0.5204

--> threadfree(id)
```

27.2 THREADFREE Free thread resources

27.2.1 Usage

The `threadfree` is a function to free the resources claimed by a thread that has finished. The syntax for its use is

```
threadfree(handle)
```

where `handle` is the handle returned by the call to `threadnew`. The `threadfree` function requires that the thread be completed. Otherwise it will wait for the thread to complete, potentially for an arbitrarily long period of time. To fix this, you can either call `threadfree` only on threads that are known to have completed, or you can call it using the syntax

```
threadfree(handle,timeout)
```

where `timeout` is a time to wait in milliseconds. If the thread fails to complete before the timeout expires, an error occurs.

27.3 THREADID Get Current Thread Handle

27.3.1 Usage

The `threadid` function in FreeMat tells you which thread is executing the context you are in. Normally, this is thread 1, the main thread. However, if you start a new thread using `threadnew`, you will be operating in a new thread, and functions that call `threadid` from the new thread will return their handles.

27.3.2 Example

From the main thread, we have

```
--> threadid
```

```
ans =  
  2
```

But from a launched auxilliary thread, we have

```
--> t_id = threadnew
```

```
t_id =  
  3
```

```
--> id = threadcall(t_id,1000,'threadid')
```

```
id =  
  3
```

```
--> threadfree(t_id);
```

27.4 THREADKILL Halt execution of a thread

27.4.1 Usage

The `threadkill` function stops (or attempts to stop) execution of the given thread. It works only for functions defined in M-files (i.e., not for built in or imported functions), and it works by setting a flag that causes the thread to stop execution at the next available statement. The syntax for this function is

```
threadkill(handle)
```

where `handle` is the value returned by a `threadnew` call. Note that the `threadkill` function returns immediately. It is still your responsibility to call `threadfree` to free the thread you have halted.

You cannot kill the main thread (thread id 1).

27.4.2 Example

Here is an example of stopping a runaway thread using `threadkill`. Note that the thread function in this case is an M-file function. We start by setting up a free running counter, where we can access the counter from the global variables.

```
freecount.m
function freecount
global count
if (~exist('count')) count = 0; end % Initialize the counter
while (1)
    count = count + 1;           % Update the counter
end
```

We now launch this function in a thread, and use `threadkill` to stop it:

```
--> a = threadnew;
--> global count           % register the global variable count
--> count = 0;
--> threadstart(a,'freecount',0) % start the thread
--> count                 % it is counting

ans =
    7145

--> sleep(1)             % Wait a bit
--> count                 % it is still counting

ans =
    485030

--> threadkill(a)       % kill the counter
--> threadwait(a,1000) % wait for it to finish
```

```

ans =
  1

--> count                                % The count will no longer increase

ans =
  485031

--> sleep(1)
--> count

ans =
  485031

--> threadfree(a)

```

27.5 THREADNEW Create a New Thread

27.5.1 Usage

The `threadnew` function creates a new FreeMat thread, and returns a handle to the resulting thread. The `threadnew` function takes no arguments. The general syntax for the `threadnew` function is

```
handle = threadnew
```

Once you have a handle to a thread, you can start the thread on a computation using the `threadstart` function. The threads returned by `threadnew` are in a dormant state (i.e., not running). Once you are finished with the thread you must call `threadfree` to free the resources associated with that thread.

Some additional important information. Thread functions operate in their own context or workspace, which means that data cannot be shared between threads. The exception is `global` variables, which provide a thread-safe way for multiple threads to share data. Accesses to global variables are serialized so that they can be used to share data. Threads and FreeMat are a new feature, so there is room for improvement in the API and behavior. The best way to improve threads is to experiment with them, and send feedback.

27.6 THREADSTART Start a New Thread Computation

27.6.1 Usage

The `threadstart` function starts a new computation on a FreeMat thread, and you must provide a function (no scripts are allowed) to run inside the thread, pass any parameters that the thread function requires, as well as the number of output arguments expected. The general syntax for the `threadstart` function is

```
threadstart(threadid,function,nargout,arg1,arg2,...)
```

where `threadid` is a thread handle (returned by `threadnew`), where `function` is a valid function name (it can be a built-in imported or M-function), `nargout` is the number of output arguments expected from the function, and `arg1` is the first argument that is passed to the function. Because the function runs in its own thread, the return values of the function are not available immediately. Instead, execution of that function will continue in parallel with the current thread. To retrieve the output of the thread function, you must wait for the thread to complete using the `threadwait` function, and then call `threadvalue` to retrieve the result. You can also stop the running thread prematurely by using the `threadkill` function. It is important to call `threadfree` on the handle you get from `threadnew` when you are finished with the thread to ensure that the resources are properly freed.

It is also perfectly reasonable to use a single thread multiple times, calling `threadstart` and `threadreturn` multiple times on a single thread. The context is preserved between threads. When calling `threadstart` on a pre-existing thread, FreeMat will attempt to wait on the thread. If the wait fails, then an error will occur.

Some additional important information. Thread functions operate in their own context or workspace, which means that data cannot be shared between threads. The exception is `global` variables, which provide a thread-safe way for multiple threads to share data. Accesses to global variables are serialized so that they can be used to share data. Threads and FreeMat are a new feature, so there is room for improvement in the API and behavior. The best way to improve threads is to experiment with them, and send feedback.

27.6.2 Example

Here we do something very simple. We want to obtain a listing of all files on the system, but do not want the results to stop our computation. So we run the `system` call in a thread.

```
--> a = threadnew;                % Create the thread
--> threadstart(a,'system',1,'ls -lrt /'); % Start the thread
--> b = rand(100)\rand(100,1);    % Solve some equations simultaneously
--> c = threadvalue(a);          % Retrieve the file list
--> size(c)                      % It is large!
```

```
ans =
     1 20
```

```
--> threadfree(a);
```

The possibilities for threads are significant. For example, we can solve equations in parallel, or take Fast Fourier Transforms on multiple threads. On multi-processor machines or multicore CPUs, these threaded calculations will execute in parallel. Neat.

The reason for the `nargout` argument is best illustrated with an example. Suppose we want to compute the Singular Value Decomposition `svd` of a matrix `A` in a thread. The documentation for the `svd` function tells us that the behavior depends on the number of output arguments we request. For example, if we want a full decomposition, including the left and right singular vectors, and a diagonal singular matrix, we need to use the three-output syntax, instead of the single output syntax (which returns only the singular values in a column vector):

```

--> A = float(rand(4))

A =
    0.3256    0.8138    0.8819    0.2759
    0.4954    0.7569    0.9476    0.9368
    0.6082    0.1744    0.0073    0.5028
    0.8448    0.1609    0.8433    0.9130

--> [u,s,v] = svd(A)    % Compute the full decomposition
u =
   -0.4551   -0.6585   -0.1575   -0.5783
   -0.6358   -0.1877   -0.1019    0.7417
   -0.2455    0.5238   -0.7940   -0.1870
   -0.5731    0.5067    0.5782   -0.2836

s =
    2.5027         0         0         0
         0    0.8313         0         0
         0         0    0.3646         0
         0         0         0    0.2569

v =
   -0.4382    0.5283   -0.2637   -0.6777
   -0.3942   -0.6076   -0.6878    0.0489
   -0.5949   -0.3938    0.6757   -0.1854
   -0.5465    0.4434   -0.0280    0.7099

--> sigmas = svd(A)    % Only want the singular values

sigmas =
    2.5027
    0.8313
    0.3646
    0.2569

```

Normally, FreeMat uses the left hand side of an assignment to calculate the number of outputs for the function. When running a function in a thread, we separate the assignment of the output from the invocation of the function. Hence, we have to provide the number of arguments at the time we invoke the function. For example, to compute a full decomposition in a thread, we specify that we want 3 output arguments:

```

--> a = threadnew;           % Create the thread
--> threadstart(a,'svd',3,A); % Start a full decomposition
--> [u1,s1,v1] = threadvalue(a); % Retrieve the function values
--> threadfree(a);

```

If we want to compute just the singular values, we start the thread function with only one output argument:

```
--> a = threadnew;
--> threadstart(a,'svd',1,A);
--> sigmas = threadvalue(a);
--> threadfree(a)
```

27.7 THREADVALUE Retrieve the return values from a thread

27.7.1 Usage

The `threadvalue` function retrieves the values returned by the function specified in the `threadnew` call. The syntax for its use is

```
[arg1,arg2,...,argN] = threadvalue(handle)
```

where `handle` is the value returned by a `threadnew` call. Note that there are issues with `nargout`. See the examples section of `threadnew` for details on how to work around this limitation. Because the function you have spawned with `threadnew` may still be executing, `threadvalue` must first `threadwait` for the function to complete before retrieving the output values. This wait may take an arbitrarily long time if the thread function is caught in an infinite loop. Hence, you can also specify a timeout parameter to `threadvalue` as

```
[arg1,arg2,...,argN] = threadvalue(handle,timeout)
```

where the `timeout` is specified in milliseconds. If the wait times out, an error is raised (that can be caught with a `try` and `catch` block).

In either case, if the thread function itself caused an error and ceased execution abruptly, then calling `threadvalue` will cause that function to raise an error, allowing you to retrieve the error that was caused and correct it. See the examples section for more information.

27.7.2 Example

Here we do something very simple. We want to obtain a listing of all files on the system, but do not want the results to stop our computation. So we run the `system` call in a thread.

```
--> a = threadnew; % Create the thread
--> threadstart(a,'system',1,'ls -lrt /'); % Start the thread
--> b = rand(100)\rand(100,1); % Solve some equations simultaneously
--> c = threadvalue(a); % Retrieve the file list
--> size(c) % It is large!

ans =
     1 20

--> threadfree(a);
```

In this example, we force the threaded function to cause an exception (by calling the `error` function as the thread function). When we call `threadvalue`, we get an error, instead of the return value of the function

```
--> a = threadnew

a =
    3

--> threadstart(a,'error',0,'Hello world!'); % Will immediately stop due to error
--> c = threadvalue(a)                    % The error comes to us
Error: Thread: Hello world!
--> threadfree(a)
```

Note that the error has the text `Thread:` prepended to the message to help you identify that this was an error in a different thread.

27.8 THREADWAIT Wait on a thread to complete execution

27.8.1 Usage

The `threadwait` function waits for the given thread to complete execution, and stops execution of the current thread (the one calling `threadwait`) until the given thread completes. The syntax for its use is

```
success = threadwait(handle)
```

where `handle` is the value returned by `threadnew` and `success` is a logical variable that will be 1 if the wait was successful or 0 if the wait times out. By default, the wait is indefinite. It is better to use the following form of the function

```
success = threadwait(handle,timeout)
```

where `timeout` is the amount of time (in milliseconds) for the `threadwait` function to wait before a timeout occurs. If the `threadwait` function succeeds, then the return value is a logical 1, and if it fails, the return value is a logical 0. Note that you can call `threadwait` multiple times on a thread, and if the thread is completed, each one will succeed.

27.8.2 Example

Here we launch the `sleep` function in a thread with a time delay of 10 seconds. This means that the thread function will not complete until 10 seconds have elapsed. When we call `threadwait` on this thread with a short timeout, it fails, but not when the timeout is long enough to capture the end of the function call.

```
--> a = threadnew;
--> threadstart(a,'sleep',0,10); % start a thread that will sleep for 10
--> threadwait(a,2000)          % 2 second wait is not long enough

ans =
    0

--> threadwait(a,10000)        % 10 second wait is long enough

ans =
    1

--> threadfree(a)
```


Chapter 28

Function Related Functions

28.1 INLINE Construct Inline Function

28.1.1 Usage

Constructs an inline function object. The syntax for its use is either

```
y = inline(expr)
```

which uses the `symvar` function to identify the variables in the expression, or the explicit form

```
y = inline(expr,var1,var2,...,varn)
```

where the variables are explicitly given. Note that inline functions are only partially supported in FreeMat. If you need features of the inline function that are not currently implemented, please file a feature request at the FreeMat website.

28.1.2 Example

Here we construct an inline expression using the autodetection of `symvar`

```
--> a = inline('x^2')
```

```
a =  
  inline function object  
  f(x) = x^2  
--> a(3)
```

```
ans =  
  9
```

```
--> a(i)
```

```
ans =  
 -1.0000 + 0.0000i
```

In this case, we have multiple arguments (again, autodetected)

```
--> a = inline('x+y-cos(x+y)')

a =
  inline function object
  f(x,y) = x+y-cos(x+y)
--> a(pi,-pi)

ans =
  -1
```

In this form, we specify which arguments we want to use (thereby also specifying the order of the arguments)

```
--> a = inline('x+t-sin(x)', 'x', 't')

a =
  inline function object
  f(x,t) = x+t-sin(x)
--> a(0.5,1)

ans =
  1.0206
```

Inline objects can also be used with `feval`

```
--> a = inline('cos(t)')

a =
  inline function object
  f(t) = cos(t)
--> feval(a,pi/2)

ans =
  6.1230e-17
```

28.2 SYMVAR Find Symbolic Variables in an Expression

28.2.1 Usage

Finds the symbolic variables in an expression. The syntax for its use is

```
syms = symvar(expr)
```

where `expr` is a string containing an expression, such as `'x^2 + cos(t+alpha)'`. The result is a cell array of strings containing the non-function identifiers in the expression. Because they are usually not used as identifiers in expressions, the strings `'pi'`, `'inf'`, `'nan'`, `'eps'`, `'i'`, `'j'` are ignored.

28.2.2 Example

Here are some simple examples:

```
--> symvar('x^2+sqrt(x)') % sqrt is eliminated as a function
```

```
ans =  
[x]
```

```
--> symvar('pi+3') % No identifiers here
```

```
ans =  
Empty array 0 1  
--> symvar('x + t*alpha') % x, t and alpha
```

```
ans =  
[alpha]  
[t]  
[x]
```


Chapter 29

FreeMat External Interface

29.1 CENUM Lookup Enumerated C Type

29.1.1 Usage

The `cenum` function allows you to use the textual strings of C enumerated types (that have been defined using `ctypedefine`) in your FreeMat scripts instead of the hardcoded numerical values. The general syntax for its use is

```
enum_int = cenum(enum_type,enum_string)
```

which looks up the integer value of the enumerated type based on the string. You can also supply an integer argument, in which case `cenum` will find the matching string

```
enum_string = cenum(enum_type,enum_int)
```

29.2 CTYPECAST Cast FreeMat Structure to C Structure

29.2.1 Usage

The `ctypecast` function is a convenience function for ensuring that a FreeMat structure fits the definition of a C struct (as defined via `ctypedefine`). It does so by encoding the structure to a byte array using `ctyfefreeze` and then recovering it using the `ctypethaw` function. The usage is simply

```
s = ctypecast(s,typename)
```

where `s` is the structure and `typename` is the name of the C structure that describes the desired layout and types for elements of `s`. This function is equivalent to calling `ctyfefreeze` and `ctypethaw` in succession on a FreeMat structure.

29.3 CTYPEDEFINE Define C Type

29.3.1 Usage

The `ctypedefine` function allows you to define C types for use with FreeMat. Three variants of C types can be used. You can use structures, enumerations, and aliases (typedefs). All three are defined through a single function `ctypedefine`. The general syntax for its use is

```
ctypedefine(typeclass,typename,...)
```

where `typeclass` is the variant of the type (legal values are `'struct'`, `'alias'`, `'enum'`). The second argument is the name of the C type. The remaining arguments depend on what the class of the typedef is.

To define a C structure, use the `'struct'` type class. The usage in this case is:

```
ctypedefine('struct',typename,field1,type1,field2,type2,...)
```

The argument `typename` must be a valid identifier string. Each of the `field` arguments is also a valid identifier string that describe in order, the elements of the C structure. The `type` arguments are `typespecs`. They can be of three types:

- Built in types, e.g. `'uint8'` or `'double'` to name a couple of examples.
- C types that have previously been defined with a call to `ctypedefine`, e.g. `'mytype'` where `'mytype'` has already been defined through a call to `ctypedefine`.
- Arrays of either built in types or previously defined C types with the length of the array coded as an integer in square brackets, for example: `'uint8[10]'` or `'double[1000]'`.

To define a C enumeration, use the `'enum'` type class. The usage in this case is: `ctypedefine('enum',typename,name1,value1,name2,value2,...) @]` The argument `typename` must be a valid identifier string. Each of the `name` arguments must also be valid identifier strings that describe the possible values that the enumeration can take an, and their corresponding integer values. Note that the names should be unique. The behavior of the various `enum` functions is undefined if the names are not unique.

To define a C alias (or typedef), use the following form of `ctypedefine`:

```
ctypedefine('alias',typename,aliased_typename)
```

where `aliased_typename` is the type that is being aliased to.

29.4 CTYPEFREEZE Convert FreeMat Structure to C Type

29.4.1 Usage

The `ctypedefreeze` function is used to convert a FreeMat structure into a C struct as defined by a C structure typedef. To use the `cstructfreeeze` function, you must first define the type of the C structure using the `ctypedefine` function. The `ctypedefreeze` function then serializes a FreeMat structure to a set of bytes, and returns it as an array. The usage for `ctypedefreeze` is

```
byte_array = ctypesfreeze(mystruct, 'typename')
```

where `mystruct` is the array we want to 'freeze' to a memory array, and `typename` is the name of the C type that we want the resulting byte array to conform to.

29.5 CTYPENEW Create New Instance of C Structure

29.5.1 Usage

The `ctypenew` function is a convenience function for creating a FreeMat structure that corresponds to a C structure. The entire structure is initialized with zeros. This has some negative implications, because if the structure definition uses `enums`, they may come out as `'unknown'` values if there are no enumerations corresponding to zero. The use of the function is

```
a = ctypenew('typename')
```

which creates a single structure of C structure type `'typename'`. To create an array of structures, we can provide a second argument

```
a = ctypenew('typename',count)
```

where `count` is the number of elements in the structure array.

29.6 CTYPEPRINT Print C Type

29.6.1 Usage

The `ctypeprint` function prints a C type on the console. The usage is

```
ctypeprint(typename)
```

where `typename` is a string containing the name of the C type to print. Depending on the class of the C type (e.g., structure, alias or enumeration) the `ctypeprint` function will dump information about the type definition.

29.7 CTYPEREAD Read a C Structure From File

29.7.1 Usage

The `ctyperead` function is a convenience function for reading a C structure from a file. This is generally a very bad idea, as direct writing of C structures to files is notoriously unportable. Consider yourself warned. The syntax for this function is

```
a = ctyperead(fid,'typename')
```

where `'typename'` is a string containing the name of the C structure as defined using `ctypedefine`, and `fid` is the file handle returned by the `fopen` command. Note that this form will read a single structure from the file. If you want to read multiple structures into an array, use the following form

```
a = ctperead(fid,'typename',count)
```

Note that the way this function works is by using `ctypesize` to compute the size of the structure, reading that many bytes from the file, and then calling `ctypethaw` on the resulting buffer. A consequence of this behavior is that the byte-endian corrective behavior of FreeMat does not work.

29.8 CTYPESIZE Compute Size of C Struct

29.8.1 Usage

The `ctypesize` function is used to compute the size of a C structure that is defined using the `ctypedefine` function. The usage of `ctypesize` is

```
size = ctypesize('typename')
```

where `typename` is the name of the C structure you want to compute the size of. The returned count is measured in bytes. Note that as indicated in the help for `ctypedefine` that FreeMat does not automatically pad the entries of the structure to match the particulars of your C compiler. Instead, the assumption is that you have adequate padding entries in your structure to align the FreeMat members with the C ones. See `ctypedefine` for more details. You can also specify an optional count parameter if you want to determine how large multiple structures are

```
size = ctypesize('typename',count)
```

29.9 CTYPETHAW Convert C Struct to FreeMat Structure

29.9.1 Usage

The `ctypethaw` function is used to convert a C structure that is encoded in a byte array into a FreeMat structure using a C structure typedef. To use the `ctypethaw` function, you must first define the type of the C structure using the `ctypedefine` function. The usage of `ctypethaw` is

```
mystruct = ctypethaw(byte_array, 'typename')
```

where `byte_array` is a `uint8` array containing the bytes that encode the C structure, and `typename` is a string that contains the type description as registered with `ctypedefine`. If you want to retrieve multiple structures from a single byte array, you can specify a count as

```
mystruct = ctypethaw(byte_array, 'typename', count)
```

where `count` is an integer containing the number of structures to retrieve. Sometimes it is also useful to retrieve only part of the structure from a byte array, and then (based on the contents of the structure) retrieve more data. In this case, you can retrieve the residual byte array using the optional second output argument of `ctypethaw`:

```
[mystruct,byte_array_remaining] = ctypethaw(byte_array, 'typename',...)
```

29.10 CTYPEWRITE Write a C Typedef To File

29.10.1 Usage

The `ctypewrite` function is a convenience function for writing a C typedef to a file. This is generally a very bad idea, as writing of C typedefs to files is notoriously unportable. Consider yourself warned. The syntax for this function is

```
ctypewrite(fid,a,'typename')
```

where `a` is the FreeMat typedef to write, `'typename'` is a string containing the name of the C typedef to use when writing the typedef to the file (previously defined using `ctypedefine`), and `fid` is the file handle returned by `fopen`.

29.11 IMPORT Foreign Function Import

29.11.1 Usage

The `import` function allows you to call functions that are compiled into shared libraries, as if they were FreeMat functions. The usage is

```
import(libraryname,symbol,function,return,arguments)
```

The argument `libraryname` is the name of the library (as a string) to import the function from. The second argument `symbol` (also a string), is the name of the symbol to import from the library. The third argument `function` is the the name of the function after its been imported into Freemat. The fourth argument is a string that specifies the return type of the function. It can take on one of the following types:

- `'uint8'` for an unsigned, 8-bit integer.
- `'int8'` for a signed, 8-bit integer.
- `'uint16'` an unsigned, 16-bit integer.
- `'int16'` a signed, 16-bit integer.
- `'uint32'` for an unsigned, 32-bit integer.
- `'int32'` for a signed, 32-bit integer.
- `'single'` for a 32-bit floating point.
- `'double'` for a 64-bit floating point.
- `'void'` for no return type.

The fourth argument is more complicated. It encodes the arguments of the imported function using a special syntax. In general, the argument list is a string consisting of entries of the form:

```
type[optional bounds check] {optional &}name
```

Here is a list of various scenarios (expressed in 'C'), and the corresponding entries, along with snippets of code.

Scalar variable passed by value: Suppose a function is defined in the library as

```
int fooFunction(float t),
```

i.e., it takes a scalar value (or a string) that is passed by value. Then the corresponding argument string would be

```
'float t'
```

For a C-string, which corresponds to a function prototype of

```
int fooFunction(const char *t),
```

the corresponding argument string would be

```
'string t'
```

Other types are as listed above. Note that FreeMat will automatically promote the type of scalar variables to the type expected by the C function. For example, if we call a function expecting a float with a double or int16 argument, then FreeMat will automatically apply type promotion rules prior to calling the function.

Scalar variable passed by reference: Suppose a function is defined in the library as

```
int fooFunction(float *t),
```

i.e., it takes a scalar value (or a string) that is passed as a pointer. Then the corresponding argument string would be

```
'float &t'
```

If the function `fooFunction` modifies `t`, then the argument passed in FreeMat will also be modified.

Array variable passed by value: In C, it is impossible to distinguish an array being passed from a simple pointer being passed. More often than not, another argument indicates the length of the array. FreeMat has the ability to perform bounds-checking on array values. For example, suppose we have a function of the form

```
int sum_onehundred_ints(int *t),
```

where `sum_onehundred_ints` assumes that `t` is a length 100 vector. Then the corresponding FreeMat argument is

```
'float32[100] t'.
```

Note that because the argument is marked as not being passed by reference, that if `sub_onehundred_ints` modifies the array `t`, this will not affect the FreeMat argument. Note that the bounds-check expression can be any legal scalar expression that evaluates to an integer, and can be a function of the arguments. For example to pass a square $N \times N$ matrix to the following function:

```
float determinantmatrix(int N, float *A),
```

we can use the following argument to `import`:

```
'int32 N, float[N*N] t'.
```

Array variable passed by reference: If the function in C modifies an array, and we wish this to be reflected in the FreeMat side, we must pass that argument by reference. Hence, consider the following hypothetical function that squares the elements of an array (functionally equivalent to $x.^2$):

```
void squarearray(int N, float *A)
```

we can use the following argument to `import`:

```
'int32 N, float[N] &A'.
```

Note that to avoid problems with memory allocation, external functions are not allowed to return pointers. As a result, as a general operating mechanism, the FreeMat code must allocate the proper arrays, and then pass them by reference to the external function.

Sending text to the FreeMat console: Starting with FreeMat 4.0, it is possible for external code that is called using the `import` statement to send text to the FreeMat console. To do so, you must define in each library that wants to send text a function with the name `freemat_io_handler` that captures its argument (a function pointer), and stores it for use by functions in the library. That function pointer takes a standard C string argument. Here is a snippet of code to demonstrate how this works:

```
/* just to make it readable */
typedef void (*strfunc)(const char*);

/* The name we want to use for the function */
strfunc FreeMatText;

/* this name is case sensitive and must not be mangled (i.e., use extern "C") */
void freemat_io_handler(strfunc printFunc) {
    FreeMatText = printFunc;
}

double SomeImportedFunction(double t) {
    FreeMatText("I am going to double that argument!\n");
    return (t*2);
}
```

In this case, once `SomeImportedFunction` is called from within FreeMat, the text "I am going to double that argument" will appear in the FreeMat console.

Your `freemat_io_handler` function is automatically called when your library is loaded by FreeMat, which happens with the first `import` statement that references it.

29.11.2 Example

Here is a complete example. We have a C function that adds two float vectors of the same length, and stores the result in a third array that is modified by the function. First, the C code:

```

    addArrays.c
void addArrays(int N, float *a, float *b, float *c) {
    int i;

    for (i=0;i<N;i++)
        c[i] = a[i] + b[i];
}

```

We then compile this into a dynamic library, say, `add.so`. The import command would then be:

```

import('add.so', 'addArrays', 'addArrays', 'void', ...
       'int32 N, float[N] a, float[N] b, float[N] &c');

```

We could then exercise the function exactly as if it had been written in FreeMat. The following only works on systems using the GNU C Compiler:

```

--> if (strcmp(computer,'MAC')) system('gcc -bundle -flat_namespace -undefined suppress -o
--> if (~strcmp(computer,'MAC')) system('gcc -shared -fPIC -o add.so addArrays.c'); end;
--> import('add.so', 'addArrays', 'addArrays', 'void', 'int32 N, float[N] a, float[N] b, float[N] c');
--> a = [3,2,3,1];
--> b = [5,6,0,2];
--> c = [0,0,0,0];
--> addArrays(length(a),a,b,c)

ans =
 []
--> c

ans =
 8 8 3 3

```

29.12 LOADLIB Load Library Function

29.12.1 Usage

The `loadlib` function allows a function in an external library to be added to FreeMat dynamically. This interface is generally to be used as last resort, as the form of the function being called is assumed to match the internal implementation. In short, this is not the interface mechanism of choice. For all but very complicated functions, the `import` function is the preferred approach. Thus, only a very brief summary of it is presented here. The syntax for `loadlib` is

```

loadlib(libfile, symbolname, functionname, nargin, nargout)

```

where `libfile` is the complete path to the library to use, `symbolname` is the name of the symbol in the library, `functionname` is the name of the function after it is imported into FreeMat (this is optional, it defaults to the `symbolname`), `nargin` is the number of input arguments (defaults to 0), and `nargout` is the number of output arguments (defaults to 0). If the number of (input or output) arguments is variable then set the corresponding argument to `-1`.