

# MATLAB<sup>®</sup>

## The Language of Technical Computing

- Computation
- Visualization
- Programming

## How to Contact The MathWorks:



www.mathworks.com	Web
comp.soft-sys.matlab	Newsgroup



support@mathworks.com	Technical support
suggest@mathworks.com	Product enhancement suggestions
bugs@mathworks.com	Bug reports
doc@mathworks.com	Documentation error reports
service@mathworks.com	Order status, license renewals, passcodes
info@mathworks.com	Sales, pricing, and general information



508-647-7000	Phone
--------------	-------



508-647-7001	Fax
--------------	-----



The MathWorks, Inc. 3 Apple Hill Drive Natick, MA 01760-2098	Mail
--	------

For contact information about worldwide offices, see the MathWorks Web site.

### *MATLAB Programming*

© COPYRIGHT 1984 - 2004 by The MathWorks, Inc.

The software described in this document is furnished under a license agreement. The software may be used or copied only under the terms of the license agreement. No part of this manual may be photocopied or reproduced in any form without prior written consent from The MathWorks, Inc.

**FEDERAL ACQUISITION:** This provision applies to all acquisitions of the Program and Documentation by, for, or through the federal government of the United States. By accepting delivery of the Program or Documentation, the government hereby agrees that this software or documentation qualifies as commercial computer software or commercial computer software documentation as such terms are used or defined in FAR 12.212, DFARS Part 227.72, and DFARS 252.227-7014. Accordingly, the terms and conditions of this Agreement and only those rights specified in this Agreement, shall pertain to and govern the use, modification, reproduction, release, performance, display, and disclosure of the Program and Documentation by the federal government (or other entity acquiring for or through the federal government) and shall supersede any conflicting contractual terms or conditions. If this License fails to meet the government's needs or is inconsistent in any respect with federal procurement law, the government agrees to return the Program and Documentation, unused, to The MathWorks, Inc.

MATLAB, Simulink, Stateflow, Handle Graphics, and Real-Time Workshop are registered trademarks, and TargetBox is a trademark of The MathWorks, Inc.

Other product or brand names are trademarks or registered trademarks of their respective holders.

### Printing History:

June 2004	First printing	New for MATLAB 7.0 (Release 14). Formerly part of <i>Using MATLAB</i> .
October 2004	Online only	Revised for MATLAB 7.0.1 (Release 14SP1)

## Data Structures

1

<b>Creating and Concatenating Matrices</b> .....	<b>1-3</b>
Constructing a Simple Matrix .....	1-4
Specialized Matrix Functions .....	1-4
Concatenating Matrices .....	1-7
Matrix Concatenation Functions .....	1-8
Generating a Numeric Sequence .....	1-10
Combining Unlike Data Types .....	1-11
<b>Matrix Indexing</b> .....	<b>1-17</b>
Accessing Single Elements .....	1-17
Linear Indexing .....	1-18
Functions That Control Indexing Style .....	1-19
Accessing Multiple Elements .....	1-19
Logical Indexing .....	1-21
Indexing on Assignment .....	1-22
<b>Getting Information About a Matrix</b> .....	<b>1-23</b>
Dimensions of the Matrix .....	1-23
Data Types Used in the Matrix .....	1-24
Data Structures Used in the Matrix .....	1-25
<b>Resizing and Reshaping Matrices</b> .....	<b>1-26</b>
Expanding the Size of a Matrix .....	1-26
Diminishing the Size of a Matrix .....	1-30
Reshaping a Matrix .....	1-31
Preallocating Memory .....	1-33
<b>Shifting and Sorting Matrices</b> .....	<b>1-35</b>
Shift and Sort Functions .....	1-35
Shifting the Location of Matrix Elements .....	1-35
Sorting the Data in Each Column .....	1-37
Sorting the Data in Each Row .....	1-37
Sorting Row Vectors .....	1-38

<b>Operating on Diagonal Matrices</b> .....	<b>1-39</b>
Constructing a Matrix from a Diagonal Vector .....	<b>1-39</b>
Returning a Triangular Portion of a Matrix .....	<b>1-40</b>
Concatenating Matrices Diagonally .....	<b>1-40</b>
<b>Empty Matrices, Scalars, and Vectors</b> .....	<b>1-41</b>
The Empty Matrix .....	<b>1-41</b>
Scalars .....	<b>1-44</b>
Vectors .....	<b>1-44</b>
<b>Full and Sparse Matrices</b> .....	<b>1-46</b>
Sparse Matrix Functions .....	<b>1-46</b>
<b>Multidimensional Arrays</b> .....	<b>1-48</b>
Overview .....	<b>1-48</b>
Creating Multidimensional Arrays .....	<b>1-50</b>
Accessing Multidimensional Array Properties .....	<b>1-54</b>
Indexing Multidimensional Arrays .....	<b>1-54</b>
Reshaping Multidimensional Arrays .....	<b>1-58</b>
Permuting Array Dimensions .....	<b>1-60</b>
Computing with Multidimensional Arrays .....	<b>1-62</b>
Organizing Data in Multidimensional Arrays .....	<b>1-64</b>
Multidimensional Cell Arrays .....	<b>1-66</b>
Multidimensional Structure Arrays .....	<b>1-67</b>
<b>Summary of Matrix and Array Functions</b> .....	<b>1-69</b>

## Data Types

# 2

<b>Overview of MATLAB Data Types</b> .....	<b>2-2</b>
<b>Numeric Types</b> .....	<b>2-4</b>
Integers .....	<b>2-4</b>
Floating-Point Numbers .....	<b>2-6</b>
Complex Numbers .....	<b>2-11</b>
Infinity and NaN .....	<b>2-12</b>

Identifying Numeric Types .....	2-14
Display Format for Numeric Values .....	2-14
Function Summary .....	2-16
<b>Logical Types .....</b>	<b>2-20</b>
Creating a Logical Array .....	2-20
How Logical Arrays Are Used .....	2-22
Identifying Logical Arrays .....	2-24
<b>Characters and Strings .....</b>	<b>2-25</b>
Creating Character Arrays .....	2-25
Cell Arrays of Strings .....	2-27
String Comparisons .....	2-30
Searching and Replacing .....	2-33
Converting from Numeric to String .....	2-34
Converting from String to Numeric .....	2-36
Function Summary .....	2-38
<b>Dates and Times .....</b>	<b>2-41</b>
Types of Date Formats .....	2-41
Conversions Between Date Formats .....	2-43
Date String Formats .....	2-44
Output Formats .....	2-44
Current Date and Time .....	2-46
Function Summary .....	2-47
<b>Structures .....</b>	<b>2-49</b>
Building Structure Arrays .....	2-50
Accessing Data in Structure Arrays .....	2-53
Using Dynamic Field Names .....	2-54
Finding the Size of Structure Arrays .....	2-55
Adding Fields to Structures .....	2-56
Deleting Fields from Structures .....	2-56
Applying Functions and Operators .....	2-56
Writing Functions to Operate on Structures .....	2-57
Organizing Data in Structure Arrays .....	2-59
Nesting Structures .....	2-63
Function Summary .....	2-65

<b>Cell Arrays</b> .....	<b>2-66</b>
Creating Cell Arrays .....	<b>2-67</b>
Obtaining Data from Cell Arrays .....	<b>2-70</b>
Deleting Cells .....	<b>2-72</b>
Reshaping Cell Arrays .....	<b>2-72</b>
Replacing Lists of Variables with Cell Arrays .....	<b>2-72</b>
Applying Functions and Operators .....	<b>2-74</b>
Organizing Data in Cell Arrays .....	<b>2-75</b>
Nesting Cell Arrays .....	<b>2-76</b>
Converting Between Cell and Numeric Arrays .....	<b>2-78</b>
Cell Arrays of Structures .....	<b>2-78</b>
Function Summary .....	<b>2-79</b>
<b>Function Handles</b> .....	<b>2-80</b>
Constructing and Invoking a Function Handle .....	<b>2-80</b>
Calling a Function Using Its Handle .....	<b>2-80</b>
Simple Function Handle Example .....	<b>2-81</b>
<b>MATLAB Classes</b> .....	<b>2-82</b>
<b>Java Classes</b> .....	<b>2-83</b>

## Basic Program Components

# 3

<b>Variables</b> .....	<b>3-2</b>
Types of Variables .....	<b>3-2</b>
Naming Variables .....	<b>3-6</b>
Guidelines to Using Variables .....	<b>3-7</b>
Scope of a Variable .....	<b>3-7</b>
Lifetime of a Variable .....	<b>3-8</b>
<b>Keywords</b> .....	<b>3-9</b>
<b>Special Values</b> .....	<b>3-10</b>

<b>Operators</b> .....	<b>3-12</b>
Arithmetic Operators .....	<b>3-12</b>
Relational Operators .....	<b>3-13</b>
Logical Operators .....	<b>3-15</b>
Operator Precedence .....	<b>3-20</b>
<b>MATLAB Expressions</b> .....	<b>3-22</b>
String Evaluation .....	<b>3-22</b>
Shell Escape Functions .....	<b>3-23</b>
<b>Regular Expressions</b> .....	<b>3-25</b>
MATLAB Regular Expression Functions .....	<b>3-25</b>
Elements of an Expression .....	<b>3-26</b>
Character Classes .....	<b>3-26</b>
Character Representation .....	<b>3-29</b>
Logical Operators .....	<b>3-30</b>
Lookaround Operators .....	<b>3-33</b>
Quantifiers .....	<b>3-36</b>
Tokens .....	<b>3-40</b>
Handling Multiple Strings .....	<b>3-46</b>
Operator Summary .....	<b>3-50</b>
<b>Comma-Separated Lists</b> .....	<b>3-54</b>
Generating a List from a Cell Array .....	<b>3-54</b>
Generating a List from a Structure .....	<b>3-55</b>
How to Use the Comma-Separated List .....	<b>3-56</b>
Fast Fourier Transform Example .....	<b>3-58</b>
<b>Program Control Statements</b> .....	<b>3-60</b>
Conditional Control — if, switch .....	<b>3-60</b>
Loop Control — for, while, continue, break .....	<b>3-64</b>
Error Control — try, catch .....	<b>3-67</b>
Program Termination — return .....	<b>3-68</b>
<b>Symbol Reference</b> .....	<b>3-69</b>
Asterisk — * .....	<b>3-70</b>
At — @ .....	<b>3-70</b>
Colon — : .....	<b>3-71</b>
Comma — , .....	<b>3-72</b>

Curly Braces — { }	3-73
Dot — .	3-73
Dot-Dot — ..	3-74
Dot-Dot-Dot (Ellipsis) — ...	3-74
Dot-Parentheses — .()	3-75
Exclamation Point — !	3-75
Parentheses — ()	3-76
Percent — %	3-76
Percent-Brace — %{ %}	3-77
Semicolon — ;	3-77
Single Quotes — ''	3-78
Space Character	3-78
Slash and Backslash — / \	3-79
Square Brackets — []	3-79
<b>MATLAB Functions</b>	<b>3-81</b>
M-File Functions	3-81
Built-In Functions	3-82
Overloaded MATLAB Functions	3-83

## M-File Programming

# 4

<b>Program Development</b>	<b>4-2</b>
Creating a Program	4-2
Getting the Bugs Out	4-3
Cleaning Up the Program	4-4
Improving Performance	4-5
Checking It In	4-6
<b>Working with M-Files</b>	<b>4-7</b>
Types of M-Files	4-7
Basic Parts of an M-File	4-8
Creating a Simple M-File	4-13
Providing Help for Your Program	4-15
Creating P-Code Files	4-15

<b>M-File Scripts and Functions</b> .....	<b>4-16</b>
M-File Scripts .....	<b>4-16</b>
M-File Functions .....	<b>4-17</b>
Types of Functions .....	<b>4-18</b>
Identifying Dependencies .....	<b>4-19</b>
<b>Function Arguments</b> .....	<b>4-21</b>
Checking the Number of Input Arguments .....	<b>4-21</b>
Passing Variable Numbers of Arguments .....	<b>4-23</b>
Returning Output Arguments .....	<b>4-25</b>
<b>Function Handles</b> .....	<b>4-27</b>
Constructing a Function Handle .....	<b>4-27</b>
Calling a Function Using Its Handle .....	<b>4-28</b>
Functions That Operate on Function Handles .....	<b>4-29</b>
Additional Information on Function Handles .....	<b>4-29</b>
<b>Calling Functions</b> .....	<b>4-31</b>
What Happens When You Call a Function .....	<b>4-31</b>
Determining Which Function Is Called .....	<b>4-31</b>
MATLAB Calling Syntax .....	<b>4-34</b>
Passing Certain Argument Types .....	<b>4-37</b>
Passing Arguments in Structures or Cell Arrays .....	<b>4-39</b>
Calling External Functions .....	<b>4-40</b>

## Types of Functions

# 5

<b>Overview of MATLAB Function Types</b> .....	<b>5-2</b>
<b>Anonymous Functions</b> .....	<b>5-3</b>
Constructing an Anonymous Function .....	<b>5-3</b>
Arrays of Anonymous Functions .....	<b>5-5</b>
Outputs from Anonymous Functions .....	<b>5-6</b>
Variables Used in the Expression .....	<b>5-8</b>
Examples of Anonymous Functions .....	<b>5-11</b>

<b>Primary M-File Functions</b> .....	<b>5-14</b>
<b>Nested Functions</b> .....	<b>5-15</b>
Writing Nested Functions .....	<b>5-15</b>
Calling Nested Functions .....	<b>5-17</b>
Variable Scope in Nested Functions .....	<b>5-18</b>
Using Function Handles with Nested Functions .....	<b>5-20</b>
Examples of Nested Functions .....	<b>5-25</b>
<b>Subfunctions</b> .....	<b>5-31</b>
Calling Subfunctions .....	<b>5-32</b>
Accessing Help for a Subfunction .....	<b>5-32</b>
<b>Private Functions</b> .....	<b>5-33</b>
Private Directories .....	<b>5-33</b>
Accessing Help for a Private Function .....	<b>5-33</b>
<b>Overloaded Functions</b> .....	<b>5-34</b>
Class Directories .....	<b>5-34</b>

## Data Import and Export

# 6

<b>Overview</b> .....	<b>6-3</b>
Text Data .....	<b>6-3</b>
Graphics Files .....	<b>6-4</b>
Audio and Audio/Video Data .....	<b>6-4</b>
Spreadsheets .....	<b>6-4</b>
Scientific Formats .....	<b>6-4</b>
The Internet .....	<b>6-5</b>
Low-Level File I/O .....	<b>6-5</b>
Large Data Sets .....	<b>6-5</b>
Toolboxes for Importing Data .....	<b>6-6</b>
<b>Using the Import Wizard</b> .....	<b>6-7</b>
Using the Import Wizard with Text Data .....	<b>6-7</b>
Using the Import Wizard with Binary Data .....	<b>6-13</b>

<b>Supported File Formats</b> .....	<b>6-16</b>
<b>Saving and Loading MAT-Files</b> .....	<b>6-18</b>
Exporting Data to MAT-Files .....	<b>6-18</b>
Importing Data from MAT-Files .....	<b>6-24</b>
<b>Importing Text Data</b> .....	<b>6-28</b>
The MATLAB Import Wizard .....	<b>6-28</b>
Using Import Functions with Text Data .....	<b>6-28</b>
Importing Numeric Text Data .....	<b>6-31</b>
Importing Delimited ASCII Data Files .....	<b>6-32</b>
Importing Numeric Data with Text Headers .....	<b>6-32</b>
Importing Mixed Alphabetic and Numeric Data .....	<b>6-33</b>
Importing from XML Documents .....	<b>6-35</b>
<b>Exporting Text Data</b> .....	<b>6-36</b>
Exporting Delimited ASCII Data Files .....	<b>6-37</b>
Using the diary Function to Export Data .....	<b>6-39</b>
Exporting to XML Documents .....	<b>6-40</b>
<b>Working with Graphics Files</b> .....	<b>6-41</b>
Getting Information About Graphics Files .....	<b>6-41</b>
Importing Graphics Data .....	<b>6-42</b>
Exporting Graphics Data .....	<b>6-42</b>
<b>Working with Audio and Video Data</b> .....	<b>6-44</b>
Getting Information About Audio/Video Files .....	<b>6-44</b>
Importing Audio/Video Data .....	<b>6-45</b>
Exporting Audio/Video Data .....	<b>6-46</b>
<b>Working with Spreadsheets</b> .....	<b>6-48</b>
Microsoft Excel Spreadsheets .....	<b>6-48</b>
Lotus 123 Spreadsheets .....	<b>6-51</b>
<b>Working with Scientific Data Formats</b> .....	<b>6-54</b>
Working with Common Data Format (CDF) Files .....	<b>6-54</b>
Working with Flexible Image Transport System (FITS) Files .....	<b>6-57</b>
Working with Hierarchical Data Format (HDF5) Files .....	<b>6-59</b>

<b>Importing HDF4 and HDF-EOS Data</b> .....	<b>6-69</b>
Using the HDF Import Tool .....	<b>6-69</b>
Using the HDF4 Import Tool Subsetting Options .....	<b>6-74</b>
Using the MATLAB hdfread Function .....	<b>6-84</b>
Using the HDF4 Command-Line Interface .....	<b>6-88</b>
<b>Exporting MATLAB Data to an HDF4 File</b> .....	<b>6-96</b>
Example: Exporting Data to an HDF4 File .....	<b>6-96</b>
Using the MATLAB HDF Utility API .....	<b>6-103</b>
<b>Using Low-Level File I/O Functions</b> .....	<b>6-105</b>
Opening Files .....	<b>6-106</b>
Reading Binary Data .....	<b>6-108</b>
Writing Binary Data .....	<b>6-110</b>
Controlling Position in a File .....	<b>6-110</b>
Reading Strings Line by Line from Text Files .....	<b>6-112</b>
Reading Formatted ASCII Data .....	<b>6-114</b>
Writing Formatted Text Files .....	<b>6-115</b>
Closing a File .....	<b>6-117</b>
<b>Exchanging Files over the Internet</b> .....	<b>6-118</b>
Downloading Web Content and Files .....	<b>6-118</b>
Creating and Uncompressing Zip Archives .....	<b>6-120</b>
Sending E-Mail .....	<b>6-121</b>
Performing FTP File Operations .....	<b>6-123</b>

## Error Handling

# 7

<b>Checking for Errors with try-catch</b> .....	<b>7-2</b>
Nested try-catch Blocks .....	<b>7-3</b>
<b>Handling and Recovering from an Error</b> .....	<b>7-4</b>
Reporting an Error .....	<b>7-4</b>
Identifying the Cause .....	<b>7-5</b>
Regenerating an Error .....	<b>7-7</b>

<b>Message Identifiers</b> .....	7-8
Identifier Format .....	7-8
Using Message Identifiers with lasterr .....	7-9
<b>Warnings</b> .....	7-12
Reporting a Warning .....	7-12
Identifying the Cause .....	7-13
<b>Warning Control</b> .....	7-14
Warning Statements .....	7-14
Warning Control Statements .....	7-15
Output from Control Statements .....	7-17
Saving and Restoring State .....	7-20
Debug, Backtrace, and Verbose Modes .....	7-21
<b>Debugging Errors and Warnings</b> .....	7-23

## Classes and Objects

# 8

<b>Classes and Objects: An Overview</b> .....	8-2
Features of Object-Oriented Programming .....	8-2
MATLAB Data Class Hierarchy .....	8-3
Creating Objects .....	8-4
Invoking Methods on Objects .....	8-4
Private Methods .....	8-5
Helper Functions .....	8-5
Debugging Class Methods .....	8-5
Setting Up Class Directories .....	8-6
Data Structure .....	8-7
Tips for C++ and Java Programmers .....	8-7
<b>Designing User Classes in MATLAB</b> .....	8-9
The MATLAB Canonical Class .....	8-9
The Class Constructor Method .....	8-10
Examples of Constructor Methods .....	8-11
Identifying Objects Outside the Class Directory .....	8-11

The display Method .....	8-12
Accessing Object Data .....	8-13
The set and get Methods .....	8-13
Indexed Reference Using subsref and subsasgn .....	8-14
Handling Subscripted Reference .....	8-15
Handling Subscripted Assignment .....	8-17
Object Indexing Within Methods .....	8-18
Defining end Indexing for an Object .....	8-19
Indexing an Object with Another Object .....	8-19
Converter Methods .....	8-20
<b>Overloading Operators and Functions .....</b>	<b>8-21</b>
Overloading Operators .....	8-21
Overloading Functions .....	8-23
<b>Example — A Polynomial Class .....</b>	<b>8-24</b>
Polynom Data Structure .....	8-24
Polynom Methods .....	8-24
The Polynom Constructor Method .....	8-24
Converter Methods for the Polynom Class .....	8-25
The Polynom display Method .....	8-28
The Polynom subsref Method .....	8-29
Overloading Arithmetic Operators for polynom .....	8-29
Overloading Functions for the Polynom Class .....	8-31
Listing Class Methods .....	8-33
<b>Building on Other Classes .....</b>	<b>8-35</b>
Simple Inheritance .....	8-35
Multiple Inheritance .....	8-37
Aggregation .....	8-37
<b>Example — Assets and Asset Subclasses .....</b>	<b>8-38</b>
Inheritance Model for the Asset Class .....	8-38
Asset Class Design .....	8-39
Other Asset Methods .....	8-39
The Asset Constructor Method .....	8-39
The Asset get Method .....	8-41
The Asset set Method .....	8-41
The Asset subsref Method .....	8-42
The Asset subsasgn Method .....	8-43

The Asset display Method .....	8-44
The Asset fieldcount Method .....	8-45
Designing the Stock Class .....	8-45
The Stock Constructor Method .....	8-46
The Stock get Method .....	8-48
The Stock set Method .....	8-49
The Stock subsref Method .....	8-50
The Stock subsasgn Method .....	8-51
The Stock display Method .....	8-53
<b>Example — The Portfolio Container .....</b>	<b>8-54</b>
Designing the Portfolio Class .....	8-54
The Portfolio Constructor Method .....	8-55
The Portfolio display Method .....	8-56
The Portfolio pie3 Method .....	8-57
Creating a Portfolio .....	8-58
<b>Saving and Loading Objects .....</b>	<b>8-60</b>
Modifying Objects During Save or Load .....	8-60
<b>Example — Defining saveobj and loadobj for Portfolio ..</b>	<b>8-61</b>
Summary of Code Changes .....	8-61
The saveobj Method .....	8-62
The loadobj Method .....	8-62
Changing the Portfolio Constructor .....	8-63
The Portfolio subsref Method .....	8-64
<b>Object Precedence .....</b>	<b>8-65</b>
Specifying Precedence of User-Defined Classes .....	8-66
<b>How MATLAB Determines Which Method to Call .....</b>	<b>8-67</b>
Selecting a Method .....	8-67
Querying Which Method MATLAB Will Call .....	8-70

## Scheduling Program Execution with Timers

# 9

<b>Using a MATLAB Timer Object</b> .....	<b>9-2</b>
Example: Displaying a Message .....	<b>9-3</b>
<b>Creating Timer Objects</b> .....	<b>9-4</b>
Timer Object Naming .....	<b>9-5</b>
<b>Working with Timer Object Properties</b> .....	<b>9-6</b>
Retrieving the Value of Timer Object Properties .....	<b>9-6</b>
Setting the Value of Timer Object Properties .....	<b>9-7</b>
<b>Starting and Stopping Timers</b> .....	<b>9-9</b>
Starting a Timer .....	<b>9-9</b>
Starting a Timer at a Specified Time .....	<b>9-10</b>
Stopping Timer Objects .....	<b>9-10</b>
Blocking the MATLAB Command Line .....	<b>9-11</b>
<b>Creating and Executing Callback Functions</b> .....	<b>9-13</b>
Associating Commands with Timer Object Events .....	<b>9-13</b>
Creating Callback Functions .....	<b>9-14</b>
Specifying the Value of Callback Function Properties .....	<b>9-16</b>
<b>Timer Object Execution Modes</b> .....	<b>9-18</b>
Executing a Timer Callback Function Once .....	<b>9-18</b>
Executing a Timer Callback Function Multiple Times .....	<b>9-19</b>
Handling Callback Function Queuing Conflicts .....	<b>9-20</b>
<b>Deleting Timer Objects from Memory</b> .....	<b>9-22</b>
Testing the Validity of a Timer Object .....	<b>9-22</b>
Deleting All Existing Timer Objects .....	<b>9-22</b>
<b>Finding All Timer Objects in Memory</b> .....	<b>9-23</b>
Finding Invisible Timer Objects .....	<b>9-23</b>

<b>Analyzing Your Program's Performance</b> .....	10-2
The M-File Profiler Utility .....	10-2
Stopwatch Timer Functions .....	10-2
<b>Techniques for Improving Performance</b> .....	10-4
Vectorizing Loops .....	10-4
Preallocating Arrays .....	10-7
Coding Loops in a MEX-File .....	10-9
Assigning to Variables .....	10-9
Operating on Real Data .....	10-10
Using Appropriate Logical Operators .....	10-10
Overloading Built-In Functions .....	10-11
Functions Are Generally Faster Than Scripts .....	10-11
Load and Save Are Faster Than File I/O Functions .....	10-11
Avoid Large Background Processes .....	10-11
<b>Making Efficient Use of Memory</b> .....	10-12
Memory Management Functions .....	10-12
Preallocating Arrays to Reduce Fragmentation .....	10-13
Enlarging Arrays with repmat .....	10-13
Working with Variables .....	10-13
Converting Full Matrices into Sparse .....	10-15
Structure of Arrays vs. Array of Structures .....	10-15
Working with Large Amounts of Data .....	10-15
<b>Resolving "Out of Memory" Errors</b> .....	10-16
General Suggestions For Reclaiming Memory .....	10-16
Compressing Data in Memory .....	10-16
Increasing System Swap Space .....	10-17
Freeing Up System Resources on Windows Systems .....	10-18
Reloading Variables on UNIX Systems .....	10-18

<b>Command and Function Syntax</b> .....	11-3
Syntax Help .....	11-3
Command and Function Syntaxes .....	11-3
Command Line Continuation .....	11-3
Completing Commands Using the Tab Key .....	11-4
Recalling Commands .....	11-4
Clearing Commands .....	11-5
Suppressing Output to the Screen .....	11-5
<b>Help</b> .....	11-6
Using the Help Browser .....	11-6
Help on Functions from the Help Browser .....	11-7
Help on Functions from the Command Window .....	11-7
Topical Help .....	11-7
Paged Output .....	11-8
Writing Your Own Help .....	11-8
Help for Subfunctions and Private Functions .....	11-9
Help for Methods and Overloaded Functions .....	11-9
<b>Development Environment</b> .....	11-10
Workspace Browser .....	11-10
Using the Find and Replace Utility .....	11-10
Commenting Out a Block of Code .....	11-11
Creating M-Files from Command History .....	11-11
Editing M-Files in EMACS .....	11-11
<b>M-File Functions</b> .....	11-12
M-File Structure .....	11-12
Using Lowercase for Function Names .....	11-12
Getting a Function's Name and Path .....	11-13
What M-Files Does a Function Use? .....	11-13
Dependent Functions, Built-Ins, Classes .....	11-13
<b>Function Arguments</b> .....	11-14
Getting the Input and Output Arguments .....	11-14
Variable Numbers of Arguments .....	11-14
String or Numeric Arguments .....	11-15

Passing Arguments in a Structure .....	11-15
Passing Arguments in a Cell Array .....	11-15
<b>Program Development</b> .....	<b>11-16</b>
Planning the Program .....	11-16
Using Pseudo-Code .....	11-16
Selecting the Right Data Structures .....	11-16
General Coding Practices .....	11-17
Naming a Function Uniquely .....	11-17
The Importance of Comments .....	11-17
Coding in Steps .....	11-18
Making Modifications in Steps .....	11-18
Functions with One Calling Function .....	11-18
Testing the Final Program .....	11-18
<b>Debugging</b> .....	<b>11-19</b>
The MATLAB Debug Functions .....	11-19
More Debug Functions .....	11-19
The MATLAB Graphical Debugger .....	11-20
A Quick Way to Examine Variables .....	11-20
Setting Breakpoints from the Command Line .....	11-21
Finding Line Numbers to Set Breakpoints .....	11-21
Stopping Execution on an Error or Warning .....	11-21
Locating an Error from the Error Message .....	11-21
Using Warnings to Help Debug .....	11-22
Making Code Execution Visible .....	11-22
Debugging Scripts .....	11-22
<b>Variables</b> .....	<b>11-23</b>
Rules for Variable Names .....	11-23
Making Sure Variable Names Are Valid .....	11-23
Don't Use Function Names for Variables .....	11-24
Checking for Reserved Keywords .....	11-24
Avoid Using i and j for Variables .....	11-24
Avoid Overwriting Variables in Scripts .....	11-25
Persistent Variables .....	11-25
Protecting Persistent Variables .....	11-25
Global Variables .....	11-26

<b>Strings</b> .....	<b>11-27</b>
Creating Strings with Concatenation .....	11-27
Comparing Methods of Concatenation .....	11-27
Store Arrays of Strings in a Cell Array .....	11-28
Converting Between Strings and Cell Arrays .....	11-28
Search and Replace Using Regular Expressions .....	11-29
<b>Evaluating Expressions</b> .....	<b>11-30</b>
Find Alternatives to Using eval .....	11-30
Assigning to a Series of Variables .....	11-30
Short-Circuit Logical Operators .....	11-31
Changing the Counter Variable within a for Loop .....	11-31
<b>MATLAB Path</b> .....	<b>11-32</b>
Precedence Rules .....	11-32
File Precedence .....	11-33
Adding a Directory to the Search Path .....	11-33
Handles to Functions Not on the Path .....	11-33
Making Toolbox File Changes Visible to MATLAB .....	11-34
Making Nontoolbox File Changes Visible to MATLAB .....	11-35
Change Notification on Windows .....	11-35
<b>Program Control</b> .....	<b>11-36</b>
Using break, continue, and return .....	11-36
Using switch Versus if .....	11-37
MATLAB case Evaluates Strings .....	11-37
Multiple Conditions in a case Statement .....	11-37
Implicit Break in switch-case .....	11-38
Variable Scope in a switch .....	11-38
Catching Errors with try-catch .....	11-38
Nested try-catch Blocks .....	11-39
Forcing an Early Return from a Function .....	11-39
<b>Save and Load</b> .....	<b>11-40</b>
Saving Data from the Workspace .....	11-40
Loading Data into the Workspace .....	11-40
Viewing Variables in a MAT-File .....	11-41
Appending to a MAT-File .....	11-41
Save and Load on Startup or Quit .....	11-42
Saving to an ASCII File .....	11-42

<b>Files and Filenames</b> .....	<b>11-43</b>
Naming M-files .....	<b>11-43</b>
Naming Other Files .....	<b>11-43</b>
Passing Filenames as Arguments .....	<b>11-44</b>
Passing Filenames to ASCII Files .....	<b>11-44</b>
Determining Filenames at Run-Time .....	<b>11-44</b>
Returning the Size of a File .....	<b>11-45</b>
<b>Input/Output</b> .....	<b>11-46</b>
File I/O Function Overview .....	<b>11-46</b>
Common I/O Functions .....	<b>11-46</b>
Readable File Formats .....	<b>11-46</b>
Using the Import Wizard .....	<b>11-47</b>
Loading Mixed Format Data .....	<b>11-47</b>
Reading Files with Different Formats .....	<b>11-47</b>
Reading ASCII Data into a Cell Array .....	<b>11-48</b>
Interactive Input into Your Program .....	<b>11-48</b>
<b>Starting MATLAB</b> .....	<b>11-49</b>
Getting MATLAB to Start Up Faster .....	<b>11-49</b>
<b>Operating System Compatibility</b> .....	<b>11-50</b>
Executing O/S Commands from MATLAB .....	<b>11-50</b>
Searching Text with grep .....	<b>11-50</b>
Constructing Paths and Filenames .....	<b>11-50</b>
Finding the MATLAB Root Directory .....	<b>11-51</b>
Temporary Directories and Filenames .....	<b>11-51</b>
<b>Demos</b> .....	<b>11-52</b>
Demos Available with MATLAB .....	<b>11-52</b>
<b>For More Information</b> .....	<b>11-53</b>

## External Interfaces

---

### A

**Finding the Documentation in Online Help** ..... A-2

**Reference Documentation** ..... A-5

## Index

---

# Data Structures

---

The most basic data structure in MATLAB<sup>®</sup> is the *matrix*: a two-dimensional, rectangularly shaped data structure capable of storing multiple elements of data in an easily accessible format. These data elements can be numbers, characters, logical states of true or false, or even other MATLAB structure types. MATLAB uses these two-dimensional matrices to store single numbers and linear series of numbers as well. In these cases, the dimensions are 1-by-1 and 1-by-n respectively, where n is the length of the numeric series. MATLAB also supports data structures that have more than two dimensions. These data structures are referred to as *arrays* in the MATLAB documentation.

Creating and Concatenating  
Matrices (p. 1-3)

Create a matrix or construct one from other matrices.

Matrix Indexing (p. 1-17)

Access or assign to elements of a matrix using methods of row and column indexing.

Getting Information About a  
Matrix (p. 1-23)

Retrieve information about the structure or contents of a matrix.

Resizing and Reshaping Matrices  
(p. 1-26)

Change the size, shape, or arrangement of elements in an existing matrix.

Shifting and Sorting Matrices  
(p. 1-35)

Shift matrix elements along one or more dimensions, or sort them into an ascending or descending order.

Operating on Diagonal Matrices  
(p. 1-39)

Construct and manipulate matrices along a diagonal of the rectangular shape.

Empty Matrices, Scalars, and  
Vectors (p. 1-41)

Work with matrices that have one or more dimensions equal to zero or one.

Full and Sparse Matrices (p. 1-46)

Conserve memory and get optimal performance with more efficient storage of matrices that contain a large number of zero values.

Multidimensional Arrays (p. 1-48) Create and work with arrays that have more than two dimensions.

Summary of Matrix and Array Functions (p. 1-69) Quick reference to the functions commonly used in working with matrices.

## Creating and Concatenating Matrices

MATLAB is a matrix-based computing environment. All of the data that you enter into MATLAB is stored in the form of a matrix or a multidimensional array. Even a single numeric value like 100 is stored as a matrix (in this case, a matrix having dimensions 1-by-1):

```
A = 100;
```

```
whos A
  Name      Size      Bytes  Class
  ----      -
  A         1x1         8   double array
```

Regardless of the data type being used, whether it is numeric, character, or logical true or false data, MATLAB stores this data in matrix (or array) form. For example, the string 'Hello World' is a 1-by-11 matrix of individual character elements in MATLAB. You can also build matrices composed of more complex data types, such as MATLAB structures and cell arrays.

To create a matrix of basic data elements such as numbers or characters, see

- “Constructing a Simple Matrix” on page 1-4
- “Specialized Matrix Functions” on page 1-4

To build a matrix composed of other matrices, see

- “Concatenating Matrices” on page 1-7
- “Matrix Concatenation Functions” on page 1-8

This section also describes

- “Generating a Numeric Sequence” on page 1-10
- “Combining Unlike Data Types” on page 1-11

## Constructing a Simple Matrix

The simplest way to create a matrix in MATLAB is to use the matrix constructor operator, `[]`. Create a row in the matrix by entering elements (shown as  $E$  below) within the brackets. Separate each element with a comma or space:

$$\text{row} = [E_1, E_2, \dots, E_m] \qquad \text{row} = [E_1 \ E_2 \ \dots \ E_m]$$

For example, to create a one row matrix of five elements, type

$$A = [12 \ 62 \ 93 \ -8 \ 22];$$

To start a new row, terminate the current row with a semicolon:

$$A = [\text{row}_1; \text{row}_2; \dots; \text{row}_n]$$

This example constructs a 3 row, 5 column (or 3-by-5) matrix of numbers. Note that all rows must have the same number of elements:

$$A = [12 \ 62 \ 93 \ -8 \ 22; \ 16 \ 2 \ 87 \ 43 \ 91; \ -4 \ 17 \ -72 \ 95 \ 6]$$

$$A =$$

12	62	93	-8	22
16	2	87	43	91
-4	17	-72	95	6

The square brackets operator constructs two-dimensional matrices only, (including 0-by-0, 1-by-1, and 1-by- $n$  matrices). To construct arrays of more than two dimensions, see “Creating Multidimensional Arrays” on page 1-50.

For instructions on how to read or overwrite any matrix element, see “Matrix Indexing” on page 1-17.

## Specialized Matrix Functions

MATLAB has a number of functions that create different kinds of matrices. Some create specialized matrices like the Hankel or Vandermonde matrix. The functions shown in the table below create a matrices for more general use.

Function	Description
ones	Create a matrix or array of all ones.
zeros	Create a matrix or array of all zeros.

Function	Description
<code>eye</code>	Create a matrix with ones on the diagonal and zeros elsewhere.
<code>accumarray</code>	Distribute elements of an input matrix to specified locations in an output matrix, also allowing for accumulation.
<code>diag</code>	Create a diagonal matrix from a vector.
<code>magic</code>	Create a square matrix with rows, columns, and diagonals that add up to the same number.
<code>rand</code>	Create a matrix or array of uniformly distributed random numbers.
<code>randn</code>	Create a matrix or array of normally distributed random numbers and arrays.
<code>randperm</code>	Create a vector (1-by-n matrix) containing a random permutation of the specified integers.

Most of these functions return matrices of type `double` (double-precision floating point). However, you can easily build basic arrays of any numeric type using the `ones`, `zeros`, and `eye` functions.

To do this, specify the MATLAB class name as the last argument:

```
A = zeros(4, 6, 'uint32')
A =
     0     0     0     0     0     0
     0     0     0     0     0     0
     0     0     0     0     0     0
     0     0     0     0     0     0
```

## Examples

Here are some examples of how you can use these functions.

**Creating a Magic Square Matrix.** A magic square is a matrix in which the sum of the elements in each column, or each row, or each main diagonal is the same. To create a 5-by-5 magic square matrix, use the `magic` function as shown.

```
A = magic(5)
A =
    17    24     1     8    15
    23     5     7    14    16
     4     6    13    20    22
    10    12    19    21     3
    11    18    25     2     9
```

Note that the elements of each row, each column, and each main diagonal add up to the same value: 65.

**Creating a Random Matrix.** The `rand` function creates a matrix or array with elements uniformly distributed between zero and one. This example multiplies each element by 20:

```
A = rand(5) * 20
A =
    3.8686    13.9580    9.9310    13.2046    14.5423
   13.6445     7.5675    17.9954     6.8394     6.1858
    6.0553    17.2002    16.4326     5.7945    16.7699
   10.8335    17.0731    12.8982     6.8239    11.3614
    3.0175    11.8713    16.3595    10.6816     7.4083
```

**Creating a Diagonal Matrix.** Use `diag` to create a diagonal matrix from a vector. You can place the vector along the main diagonal of the matrix, or on a diagonal that is above or below the main one, as shown here. The `-1` input places the vector one row below the main diagonal:

```
A = [12 62 93 -8 22];

B = diag(A, -1)
B =
     0     0     0     0     0     0
    12     0     0     0     0     0
     0    62     0     0     0     0
     0     0    93     0     0     0
     0     0     0    -8     0     0
     0     0     0     0    22     0
```

## Concatenating Matrices

Matrix concatenation is the process of joining one or more matrices to make a new matrix. The brackets `[]` operator discussed earlier in this section serves not only as a matrix constructor, but also as the MATLAB concatenation operator. The expression `C = [A B]` horizontally concatenates matrices A and B. The expression `C = [A; B]` vertically concatenates them.

This example constructs a new matrix C by concatenating matrices A and B in a vertical direction:

```
A = ones(2, 5) * 6;           % 2-by-5 matrix of 6's
B = rand(3, 5);              % 3-by-5 matrix of random values

C = [A; B]                   % Vertically concatenate A and B
C =
    6.0000    6.0000    6.0000    6.0000    6.0000
    6.0000    6.0000    6.0000    6.0000    6.0000
    0.6154    0.7382    0.9355    0.8936    0.8132
    0.7919    0.1763    0.9169    0.0579    0.0099
    0.9218    0.4057    0.4103    0.3529    0.1389
```

## Keeping Matrices Rectangular

You can construct matrices, or even multidimensional arrays, using concatenation as long as the resulting matrix does not have an irregular shape (as in the second illustration shown below). If you are building a matrix horizontally, then each component matrix must have the same number of rows. When building vertically, each component must have the same number of columns.

This diagram shows two matrices of the same height (i.e., same number of rows) being combined horizontally to form a new matrix.

$$\begin{array}{|c|c|} \hline 7 & 23 \\ \hline 41 & 11 \\ \hline -1 & 90 \\ \hline \end{array}
 \begin{array}{c} + \\ \\ \\ \end{array}
 \begin{array}{|c|c|c|c|} \hline 46 & 0 & 13 & -4 \\ \hline 44 & 62 & 31 & 98 \\ \hline 3 & 51 & -9 & 25 \\ \hline \end{array}
 \begin{array}{c} = \\ \\ \\ \end{array}
 \begin{array}{|c|c|c|c|c|c|} \hline 7 & 23 & 46 & 0 & 13 & -4 \\ \hline 41 & 11 & 44 & 62 & 31 & 98 \\ \hline -1 & 90 & 3 & 51 & -9 & 25 \\ \hline \end{array}$$

3-by-2
3-by-4
3-by-6

The next diagram illustrates an attempt to horizontally combine two matrices of unequal height. MATLAB does not allow this.

$$\begin{array}{|c|c|} \hline 7 & 23 \\ \hline 41 & 11 \\ \hline -1 & 90 \\ \hline \end{array} + \begin{array}{|c|c|c|c|} \hline 46 & 0 & 13 & -4 \\ \hline 44 & 62 & 31 & 98 \\ \hline \end{array} \neq \begin{array}{|c|c|c|c|c|c|} \hline 7 & 23 & 46 & 0 & 13 & -4 \\ \hline 41 & 11 & 44 & 62 & 31 & 98 \\ \hline -1 & 90 & & & & \\ \hline \end{array}$$

3-by-2
2-by-4

## Matrix Concatenation Functions

The following functions combine existing matrices to form a new matrix.

Function	Description
cat	Concatenate matrices along the specified dimension
horzcat	Horizontally concatenate matrices
vertcat	Vertically concatenate matrices
repmat	Create a new matrix by replicating and tiling existing matrices
blkdiag	Create a block diagonal matrix from existing matrices

### Examples

Here are some examples of how you can use these functions.

**Concatenating Matrices and Arrays.** An alternative to using the [ ] operator for concatenation are the three functions `cat`, `horzcat`, and `vertcat`. With these functions, you can construct matrices (or multidimensional arrays) along a specified dimension. Either of the following commands accomplish the same task as the command `C = [A; B]` used in the section on “Concatenating Matrices” on page 1-7:

```

C = cat(1, A, B);           % Concatenate along the first dimension
C = vertcat(A, B);         % Concatenate vertically

```

**Replicating a Matrix.** Use the `repmat` function to create a matrix composed of copies of an existing matrix. When you enter

```
repmat(M, v, h)
```

MATLAB replicates input matrix `M` `v` times vertically and `h` times horizontally. For example, to replicate existing matrix `A` into a new matrix `B`, use

```
A = [8 1 6; 3 5 7; 4 9 2]
```

```
A =
```

```
8   1   6
3   5   7
4   9   2
```

```
B = repmat(A, 2, 4)
```

```
B =
```

```
8   1   6   8   1   6   8   1   6   8   1   6
3   5   7   3   5   7   3   5   7   3   5   7
4   9   2   4   9   2   4   9   2   4   9   2
8   1   6   8   1   6   8   1   6   8   1   6
3   5   7   3   5   7   3   5   7   3   5   7
4   9   2   4   9   2   4   9   2   4   9   2
```

**Creating a Block Diagonal Matrix.** The `blkdiag` function combines matrices in a diagonal direction, creating what is called a block diagonal matrix. All other elements of the newly created matrix are set to zero:

```
A = magic(3);
```

```
B = [-5 -6 -9; -4 -4 -2];
```

```
C = eye(2) * 8;
```

```
D = blkdiag(A, B, C)
```

```
D =
```

```
8   1   6   0   0   0   0   0
3   5   7   0   0   0   0   0
4   9   2   0   0   0   0   0
0   0   0  -5  -6  -9   0   0
0   0   0  -4  -4  -2   0   0
0   0   0   0   0   0   8   0
0   0   0   0   0   0   0   8
```

## Generating a Numeric Sequence

Because numeric sequences can often be useful in constructing and indexing into matrices and arrays, MATLAB provides a special operator to assist in creating them.

This section covers

- “The Colon Operator”
- “Using the Colon Operator with a Step Value”

### The Colon Operator

The colon operator (`first:last`) generates a 1-by- $n$  matrix (or *vector*) of sequential numbers from the first value to the last. The default sequence is made up of incremental values, each 1 greater than the previous one:

```
A = 10:15
A =
    10    11    12    13    14    15
```

The numeric sequence does not have to be made up of positive integers. It can include negative numbers and fractional numbers as well:

```
A = -2.5:2.5
A =
 -2.5000  -1.5000  -0.5000   0.5000   1.5000   2.5000
```

By default, MATLAB always increments by exactly 1 when creating the sequence, even if the ending value is not an integral distance from the start:

```
A = 1:6.3
A =
     1     2     3     4     5     6
```

Also, the default series generated by the colon operator always increments rather than decrementing. The operation shown in this example attempts to increment from 9 to 1 and thus MATLAB returns an empty matrix:

```
A = 9:1
A =
Empty matrix: 1-by-0
```

The next section explains how to generate a nondefault numeric series.

## Using the Colon Operator with a Step Value

To generate a series that does not use the default of incrementing by 1, specify an additional value with the colon operator (`first:step:last`). In between the starting and ending value is a step value that tells MATLAB how much to increment (or decrement, if step is negative) between each number it generates.

To generate a series of numbers from 10 to 50, incrementing by 5, use

```
A = 10:5:50
A =
    10    15    20    25    30    35    40    45    50
```

You can increment by noninteger values. This example increments by 0.2:

```
A = 3:0.2:3.8
A =
    3.0000    3.2000    3.4000    3.6000    3.8000
```

To create a sequence with a decrementing interval, specify a negative step value:

```
A = 9:-1:1
A =
     9     8     7     6     5     4     3     2     1
```

## Combining Unlike Data Types

Matrices and arrays can be composed of elements of most any MATLAB data type as long as all elements in the matrix are of the same type. If you do include elements of unlike data types when constructing a matrix, MATLAB converts some elements so that all elements of the resulting matrix are of the same type. (See “Data Types” on page 2-1 for information on any of the MATLAB data types discussed here.)

Data type conversion is done with respect to a preset precedence of data types. The following table shows the five data types you can concatenate with an unlike type without generating an error.

<b>TYPE</b>	<b>character</b>	<b>integer</b>	<b>single</b>	<b>double</b>	<b>logical</b>
<b>character</b>	character	character	character	character	invalid
<b>integer</b>	character	integer	integer	integer	integer
<b>single</b>	character	integer	single	single	single
<b>double</b>	character	integer	single	double	double
<b>logical</b>	invalid	integer	single	double	logical

For example, concatenating a double and single matrix always yields a matrix of type single. MATLAB converts the double element to single to accomplish this.

### Combining Unlike Integer Types

If you combine different integer types in a matrix (e.g., signed with unsigned, or 8-bit integers with 16-bit integers), MATLAB returns a matrix in which all elements are of one common type. MATLAB sets all elements of the resulting matrix to the data type of the left-most element in the input matrix. For example, the result of the following concatenation is a vector of three 16-bit signed integers:

```
A = [int16(450) uint8(250) int32(1000000)]
```

MATLAB also displays a warning to inform you that the result may not be what you had expected:

```
A = [int16(450) uint8(250) int32(1000000)];
```

```
Warning: Concatenation with dominant (left-most) integer class  
may overflow other operands on conversion to return class.
```

You can disable this warning by entering the following two commands directly after the operation that caused the warning. The first command gets the message identifier associated with the most recent warning issued by MATLAB. The second command uses this identifier to disable any further warnings of that type from being issued:

```
[msg, intcat_msgid] = lastwarn;
warning('off', intcat_msgid);
```

To reenable the warning so that it will now be displayed, use

```
warning('on', intcat_msgid);
```

You can use these commands to disable or enable the display of any MATLAB warning.

**Example of Combining Unlike Integer Sizes.** After disabling the integer concatenation warnings as shown above, concatenate the following two numbers once, and then switch their order. The return value depends on the order in which the integers are concatenated. The left-most type determines the data type for all elements in the vector:

```
A = [int16(5000) int8(50)]
A =
    5000    50
```

```
B = [int8(50) int16(5000)]
B =
    50   127
```

The first operation returns a vector of 16-bit integers. The second returns a vector of 8-bit integers. The element `int16(5000)` is set to 127, the maximum value for an 8-bit signed integer.

The same rules apply to vertical concatenation:

```
C = [int8(50); int16(5000)]
C =
    50
   127
```

---

**Note** You can find the maximum or minimum values for any MATLAB integer type using the `intmax` and `intmin` functions. For floating-point types, use `realmax` and `realmin`.

---

**Example of Combining Signed with Unsigned.** Now do the same exercise with signed and unsigned integers. Again, the left-most element determines the data type for all elements in the resulting matrix:

```
A = [int8(-100) uint8(100)]
A =
    -100    100
```

```
B = [uint8(100) int8(-100)]
B =
    100     0
```

The element `int8(-100)` is set to zero because it is no longer signed:

### Combining Integer and Noninteger Data

If you combine integers with `double`, `single`, or `logical` data types, all elements of the resulting matrix are given the data type of the left-most integer. For example, all elements of the following vector are set to `int32`:

```
A = [true pi int32(1000000) single(17.32) uint8(250)]
```

### Empty Matrices

If you construct a matrix using empty matrix elements, the empty matrices are ignored in the resulting matrix:

```
A = [5.36; 7.01; []; 9.44]
A =
    5.3600
    7.0100
    9.4400
```

## Concatenation Examples

Here are some examples of data type conversion during matrix construction.

**Combining Single and Double Types.** Combining single values with double values yields a single matrix. Note that  $5.73 \times 10^{300}$  is too big to be stored as a single, thus the conversion from double to single sets it to infinity. (The class function used in this example returns the data type for the input value):

```
x = [single(4.5) single(-2.8) pi 5.73*10^300]
x =
    4.5000    -2.8000    3.1416         Inf

class(x)           % Display the data type of x
ans =
    single
```

**Combining Integer and Double Types.** Combining integer values with double values yields an integer matrix. Note that the fractional part of pi is rounded to the nearest integer. (The int8 function used in this example converts its numeric argument to an 8-bit integer):

```
x = [int8(21) int8(-22) int8(23) pi 45/6]
x =
    21    -22    23     3     7

class(x)
ans =
    int8
```

**Combining Character and Double Types.** Combining character values with double values yields a character matrix. MATLAB converts the double elements in this example to their character equivalents:

```
x = ['A' 'B' 'C' 68 69 70]
x =
    ABCDEF

class(x)
ans =
    char
```

**Combining Logical and Double Types.** Combining logical values with double values yields a double matrix. MATLAB converts the logical true and false elements in this example to double:

```
x = [true false false pi sqrt(7)]
x =
    1.0000         0         0    3.1416    2.6458

class(x)
ans =
    double
```

## Matrix Indexing

This section explains how to use subscripting and indexing to access and assign values to the elements of a MATLAB matrix. It covers the following:

- “Accessing Single Elements” on page 1-17
- “Linear Indexing” on page 1-18
- “Functions That Control Indexing Style” on page 1-19
- “Accessing Multiple Elements” on page 1-19
- “Logical Indexing” on page 1-21
- “Indexing on Assignment” on page 1-22

### Accessing Single Elements

To reference a particular element in a matrix, specify its row and column number using the following syntax, where A is the matrix variable. Always specify the row first and column second:

```
A(row, column)
```

For example, for a 4-by-4 magic square A,

```
A = magic(4)
A =
    16     2     3    13
     5    11    10     8
     9     7     6    12
     4    14    15     1
```

you would access the element at row 4, column 2 with

```
A(4, 2)
ans =
    14
```

For arrays with more than two dimensions, specify additional indices following the row and column indices. See the section on “Multidimensional Arrays” on page 1-48.

## Linear Indexing

With MATLAB, you can refer to the elements of a matrix with a single subscript,  $A(k)$ . MATLAB stores matrices and arrays not in the shape that they appear when displayed in the MATLAB Command Window, but as a single column of elements. This single column is composed of all of the columns from the matrix, each appended to the last.

So, matrix A

```
A = [2 6 9; 4 2 8; 3 0 1]
A =
     2     6     9
     4     2     8
     3     5     1
```

is actually stored in memory as the sequence

```
2, 4, 3, 6, 2, 5, 9, 8, 1
```

The element at row 3, column 2 of matrix A (value = 5) can also be identified as element 6 in the actual storage sequence. To access this element, you have a choice of using the standard  $A(3,2)$  syntax, or you can use  $A(6)$ , which is referred to as *linear indexing*.

If you supply more subscripts, MATLAB calculates an index into the storage column based on the dimensions you assigned to the array. For example, assume a two-dimensional array like A has size  $[d1 \ d2]$ , where  $d1$  is the number of rows in the array and  $d2$  is the number of columns. If you supply two subscripts  $(i, j)$  representing row-column indices, the offset is

$$(j-1) * d1 + i$$

Given the expression  $A(3,2)$ , MATLAB calculates the offset into A's storage column as  $(2-1) * 3 + 3$ , or 6. Counting down six elements in the column accesses the value 5.

## Functions That Control Indexing Style

If you have row-column subscripts but want to use linear indexing instead, you can convert to the latter using the `sub2ind` function. In the 3-by-3 matrix `A` used in the previous section, `sub2ind` changes a standard row-column index of (3, 2) to a linear index of 6:

```
A = [2 6 9; 4 2 8; 3 0 1];

linearindex = sub2ind(size(A), 3, 2)
linearindex =
    6
```

To get the row-column equivalent of a linear index, use the `ind2sub` function:

```
[row col] = ind2sub(size(A), 6)
row =
    3
col =
    2
```

## Accessing Multiple Elements

For the 4-by-4 matrix `A` shown below, it is possible to compute the sum of the elements in the fourth column of `A` by typing

```
A = magic(4);
A(1,4) + A(2,4) + A(3,4) + A(4,4)
```

You can reduce the size of this expression using the colon operator. Subscript expressions involving colons refer to portions of a matrix. The expression

```
A(1:m, n)
```

refers to the elements in rows 1 through `m` of column `n` of matrix `A`. Using this notation, you can compute the sum of the fourth column of `A` more succinctly:

```
sum(A(1:4, 4))
```

### Nonconsecutive Elements

To refer to nonconsecutive elements in a matrix, use the colon operator with a step value. The `m:3:n` in this expression means to make the assignment to every third element in the matrix. Note that this example uses linear indexing:

```
B = A;  
  
B(1:3:16) = -10  
B =  
    -10     2     3    -10  
     5    11   -10     8  
     9   -10     6    12  
    -10    14    15   -10
```

### The end Keyword

MATLAB provides a keyword called `end` that designates the last element in the dimension in which it appears. This keyword can be useful in instances where your program doesn't know how many rows or columns there are in a matrix. You can replace the expression in the previous example with

```
B(1:3:end) = -10
```

---

**Note** The keyword `end` has two meanings in MATLAB. It can be used as explained above, or to terminate a certain block of code (e.g., `if` and `for` blocks).

---

### Specifying All Elements of a Row or Column

The colon by itself refers to *all* the elements in a row or column of a matrix. Using the following syntax, you can compute the sum of all elements in the second column of a 4-by-4 magic square `A`:

```
sum(A(:, 2))  
ans =  
    34
```

By using the colon with linear indexing, you can refer to all elements in the entire matrix. This example displays all the elements of matrix A, returning them in a column-wise order:

```
A(:)
ans =
    16
     5
     9
     4
     .
     .
     .
    12
     1
```

### Using a Matrix As an Index

You can repeatedly access an array element using the ones function. To create a new 2-by-6 matrix out of the ninth element of a 4-by-4 magic square A,

```
B = A(9 * ones(2, 6))
B =
     3     3     3     3     3     3
     3     3     3     3     3     3
```

### Logical Indexing

A logical matrix provides a different type of array indexing in MATLAB. While most indices are numeric, indicating a certain row or column number, logical indices are positional. That is, it is the *position* of each 1 in the logical matrix that determines which array element is being referred to.

See “Using Logicals in Array Indexing” on page 2-22 for more information on this subject.

## Indexing on Assignment

When assigning values from one matrix to another matrix, you can use any of the styles of indexing covered in this section. Matrix assignment statements also have the following requirement.

In the assignment  $A(J, K, \dots) = B(M, N, \dots)$ , subscripts  $J, K, M, N$ , etc. may be scalar, vector, or array, provided that all of the following are true:

- The number of subscripts specified for  $B$ , not including trailing subscripts equal to 1, does not exceed `ndims(B)`.
- The number of nonscalar subscripts specified for  $A$  equals the number of nonscalar subscripts specified for  $B$ . For example,  $A(5, 1:4, 1, 2) = B(5:8)$  is valid because both sides of the equation use one nonscalar subscript.
- The order and length of all nonscalar subscripts specified for  $A$  matches the order and length of nonscalar subscripts specified for  $B$ . For example,  $A(1:4, 3, 3:9) = B(5:8, 1:7)$  is valid because both sides of the equation (ignoring the one scalar subscript 3) use a 4-element subscript followed by a 7-element subscript.

## Getting Information About a Matrix

This section explains how to get the following information about an existing matrix:

- “Dimensions of the Matrix” on page 1-23
- “Data Types Used in the Matrix” on page 1-24
- “Data Structures Used in the Matrix” on page 1-25

### Dimensions of the Matrix

These functions return information about the shape and size of a matrix.

Function	Description
<code>length</code>	Return the length of the longest dimension. (The length of a matrix or array with any zero dimension is zero.)
<code>ndims</code>	Return the number of dimensions.
<code>numel</code>	Return the number of elements.
<code>size</code>	Return the length of each dimension.

The following examples show some simple ways to use these functions. Both use the 3-by-5 matrix A shown here:

```
A = rand(5) * 10;
A(4:5, :) = []
A =
    9.5013    7.6210    6.1543    4.0571    0.5789
    2.3114    4.5647    7.9194    9.3547    3.5287
    6.0684    0.1850    9.2181    9.1690    8.1317
```

#### Example Using `numel`

Using the `numel` function, find the average of all values in matrix A:

```
sum(A(:))/numel(A)
ans =
    5.8909
```

### Example Using ndims, numel, and size

Using ndims and size, go through the matrix and find those values that are between 5 and 7, inclusive:

```
if ndims(A) ~= 2
    return
end

[rows cols] = size(A);
for m = 1:rows
    for n = 1:cols
        x = A(m, n);
        if x >= 5 && x <= 7
            disp(sprintf('A(%d, %d) = %5.2f', m, n, A(m,n)))
        end
    end
end
```

The code returns the following:

```
A(1, 3) = 6.15
A(3, 1) = 6.07
```

### Data Types Used in the Matrix

These functions test elements of a matrix for a specific data type.

Function	Description
isa	Detect if input is of a given data type.
iscell	Determine if input is a cell array.
iscellstr	Determine if input is a cell array of strings.
ischar	Determine if input is a character array.
isfloat	Determine if input is a floating-point array.
isinteger	Determine if input is an integer array.
islogical	Determine if input is a logical array.

Function	Description
<code>isnumeric</code>	Determine if input is a numeric array.
<code>isreal</code>	Determine if input is an array of real numbers.
<code>isstruct</code>	Determine if input is a MATLAB structure array.

### Example Using `isnumeric` and `isreal`

Pick out the real numeric elements from this vector:

```
A = [5+7i 8/7 4.23 39j pi 9-2i];

for m = 1:numel(A)
    if isnumeric(A(m)) && isreal(A(m))
        disp(A(m))
    end
end
```

The values returned are

```
1.1429
4.2300
3.1416
```

### Data Structures Used in the Matrix

These functions test elements of a matrix for a specific data structure.

Function	Description
<code>isempty</code>	Determine if input has any dimension with size zero.
<code>isscalar</code>	Determine if input is a 1-by-1 matrix.
<code>issparse</code>	Determine if input is a sparse matrix.
<code>isvector</code>	Determine if input is a 1-by-n or n-by-1 matrix.

## Resizing and Reshaping Matrices

You can easily enlarge or shrink the size of a matrix, modify its shape, or rotate it about various axes. This section covers

- “Expanding the Size of a Matrix” on page 1-26
- “Diminishing the Size of a Matrix” on page 1-30
- “Reshaping a Matrix” on page 1-31
- “Preallocating Memory” on page 1-33

### Expanding the Size of a Matrix

You can expand the size of any existing matrix as long as doing so does not give the resulting matrix an irregular shape. (See “Keeping Matrices Rectangular” on page 1-7). For example, you can vertically combine a 4-by-3 matrix and 7-by-3 matrix because all rows of the resulting matrix have the same number of columns (3).

Two ways of expanding the size of an existing matrix are

- Concatenating new elements onto the matrix
- Storing to a location outside the bounds of the matrix

---

**Note** If you intend to expand the size of a matrix repeatedly over time as it requires more room (usually done in a programming loop), it is advisable to preallocate space for the matrix when you initially create it. See “Preallocating Memory” on page 1-33.

---

### Concatenating Onto the Matrix

Concatenation is most useful when you want to expand a matrix by adding new elements or blocks that are compatible in size with the original matrix. This means that the size of all matrices being joined along a specific dimension must be equal along that dimension. See “Concatenating Matrices” on page 1-7.

This example runs a user-defined function `compareResults` on the data in matrices `stats04` and `stats03`. Each time through the loop, it concatenates the results of this function onto the end of the data stored in `comp04`:

```
col = 10;
comp04 = [];

for k = 1:50
    t = compareResults(stats04(k,1:col), stats03(k,1:col));
    comp04 = [comp04; t];
end
```

**Concatenating to a Structure or Cell Array.** You can add on to arrays of structures or cells in the same way as you do with ordinary matrices. This example creates a 3-by-8 matrix of structures `S`, each having 3 fields: `x`, `y`, and `z`, and then concatenates a second structure matrix `S2` onto the original:

Create a 3-by-8 structure array `S`:

```
for k = 1:24
    S(k) = struct('x', 10*k, 'y', 10*k+1, 'z', 10*k+2);
end
S = reshape(S, 3, 8);
```

Create a second array that is 3-by-2 and uses the same field names:

```
for k = 25:30
    S2(k-24) = struct('x', 10*k, 'y', 10*k+1, 'z', 10*k+2);
end
S2= reshape(S2, 3, 2);
```

Concatenate `S2` onto `S` along the horizontal dimension:

```
S = [S S2]
S =
3x10 struct array with fields:
    x
    y
    z
```

### Adding Smaller Blocks to a Matrix

To add one or more elements to a matrix where the sizes are not compatible, you can often just store the new elements outside the boundaries of the original matrix. MATLAB automatically pads the matrix with zeros to keep it rectangular.

Construct a 3-by-5 matrix, and attempt to add a new element to it using concatenation. The operation fails because you are attempting to join a one-column matrix with one that has five columns:

```
A = [ 10  20  30  40  50; ...  
      60  70  80  90 100; ...  
      110 120 130 140 150];
```

```
A = [A; 160]  
??? Error using ==> vertcat  
All rows in the bracketed expression must have the same  
number of columns.
```

Try this again, but this time do it in such a way that enables MATLAB to make adjustments to the size of the matrix. Store the new element in row 4, a row that does not yet exist in this matrix. MATLAB expands matrix A by an entire new row by padding columns 2 through 5 with zeros:

```
A(4,1) = 160  
A =  
    10    20    30    40    50  
    60    70    80    90   100  
   110   120   130   140   150  
   160     0     0     0     0
```

---

**Note** Attempting to read from nonexistent matrix locations generates an error. You can only write to these locations.

---

You can also expand the matrix by adding a matrix instead of just a single element:

```
A(4:6,1:3) = magic(3)+100
A =
    10    20    30    40    50
    60    70    80    90   100
   110   120   130   140   150
   108   101   106     0     0
   103   105   107     0     0
   104   109   102     0     0
```

You do not have to add new elements sequentially. Wherever you store the new elements, MATLAB pads with zeros to make the resulting matrix rectangular in shape:

```
A(4,8) = 300
A =
    10    20    30    40    50     0     0     0
    60    70    80    90   100     0     0     0
   110   120   130   140   150     0     0     0
     0     0     0     0     0     0     0   300
```

**Expanding a Structure or Cell Array.** You can expand a structure or cell array in the same way that you can a matrix. This example adds an additional cell to a cell array by storing it beyond the bounds of the original array. MATLAB pads the data structure with empty cells ([]) to keep it rectangular.

The original array is 2-by-3:

```
C = {'Madison', 'G', [5 28 1967]; ...
     46, '325 Maple Dr', 3015.28}
```

Add a cell to C{3,1} and MATLAB appends an entire row:

```
C{3, 1} = ...
struct('Fund_A', .45, 'Fund_E', .35, 'Fund_G', 20);
C =
'Madison'          'G'          [1x3 double]
[         46]      '325 Maple Dr' [3.0153e+003]
[1x1 struct]                []                []
```

## Diminishing the Size of a Matrix

You can delete rows and columns from a matrix by assigning the empty array `[]` to those rows or columns. Start with

```
A = magic(4)
A =
    16     2     3    13
     5    11    10     8
     9     7     6    12
     4    14    15     1
```

Then, delete the second column of A using

```
A(:, 2) = []
```

This changes matrix A to

```
A =
    16     3    13
     5    10     8
     9     6    12
     4    15     1
```

If you delete a single element from a matrix, the result isn't a matrix anymore. So expressions like

```
A(1,2) = []
```

result in an error. However, you can use linear indexing to delete a single element, or a sequence of elements. This reshapes the remaining elements into a row vector:

```
A(2:2:10) = []
```

results in

```
A =
    16     9     3     6    13    12     1
```

## Reshaping a Matrix

The following functions change the shape of a matrix.

Function	Description
reshape	Modify the shape of a matrix.
rot90	Rotate the matrix by 90 degrees.
fliplr	Flip the matrix about a vertical axis.
flipud	Flip the matrix about a horizontal axis.
flipdim	Flip the matrix along the specified direction.
transpose	Flip a matrix about its main diagonal, turning row vectors into column vectors and vice versa.
ctranspose	Transpose a matrix and replace each element with its complex conjugate.

### Examples

Here are a few examples to illustrate some of the ways you can reshape matrices.

**Reshaping a Matrix.** Reshape 3-by-4 matrix A to have dimensions 2-by-6:

```
A = [1 4 7 10; 2 5 8 11; 3 6 9 12]
```

```
A =
```

```

1     4     7    10
2     5     8    11
3     6     9    12
```

```
B = reshape(A, 2, 6)
```

```
B =
```

```

1     3     5     7     9    11
2     4     6     8    10    12
```

**Transposing a Matrix.** Transpose A so that the row elements become columns. You can use either the transpose function or the transpose operator (.'') to do this:

```
B = A.'
B =
     1     2     3
     4     5     6
     7     8     9
    10    11    12
```

There is a separate function called ctranspose that performs a complex conjugate transpose of a matrix. The equivalent operator for ctranspose on a matrix A is A':

```
A = [1+9i 2-8i 3+7i; 4-6i 5+5i 6-4i]
A =
 1.0000 + 9.0000i  2.0000 -8.0000i  3.0000 + 7.0000i
 4.0000 -6.0000i  5.0000 + 5.0000i  6.0000 -4.0000i

B = A'
B =
 1.0000 -9.0000i  4.0000 + 6.0000i
 2.0000 + 8.0000i  5.0000 -5.0000i
 3.0000 -7.0000i  6.0000 + 4.0000i
```

**Rotating a Matrix.** Rotate the matrix by 90 degrees:

```
B = rot90(A)
B =
    10    11    12
     7     8     9
     4     5     6
     1     2     3
```

**Flipping a Matrix.** Flip A in a left-to-right direction:

```
B = fliplr(A)
B =
    10     7     4     1
    11     8     5     2
    12     9     6     3
```

## Preallocating Memory

Repeatedly expanding the size of an array over time, (for example, adding more elements to it each time through a programming loop), can adversely affect the performance of your program. This is because

- MATLAB has to spend time allocating more memory each time you increase the size of the array.
- This newly allocated memory is likely to be noncontiguous, thus slowing down any operations that MATLAB needs to perform on the array.

The preferred method for sizing an array that is expected to grow over time is to estimate the maximum possible size for the array, and preallocate this amount of memory for it at the time the array is created. In this way, your program performs one memory allocation that reserves one contiguous block.

The following command preallocates enough space for a 25,000 by 10,000 matrix, and initializes each element to zero:

```
A = zeros(25000, 10000);
```

## Building a Preallocated Array

Once memory has been preallocated for the maximum estimated size of the array, you can store your data in the array as you need it, each time appending to the existing data. This example preallocates a large array, and then reads blocks of data from a file into the array until it gets to the end of the file:

```
blocksize = 5000;
maxrows = 2500000; cols = 20;
rp = 1;      % row pointer

% Preallocate A to its maximum possible size
A = zeros(maxrows, cols);

% Open the data file, saving the file pointer.
fid = fopen('statfile.dat', 'r');

while true
    % Read from file into a cell array. Stop at EOF.
    block = textscan(fid, '%n', blocksize*cols);
    if isempty(block{1}) break, end;
```

```
% Convert cell array to matrix, reshape, place into A.
A(rp:rp+blocksize-1, 1:cols) = ...
    reshape(cell2mat(block), blocksize, cols);

% Process the data in A.
evaluate_stats(A); % User-defined function

% Update row pointer
rp = rp + blocksize;
end
```

---

**Note** If you eventually need more room in a matrix than you had preallocated, you can preallocate additional storage in the same manner, and concatenate this additional storage onto the original array.

---

## Shifting and Sorting Matrices

You can sort matrices, multidimensional arrays, and cell arrays of strings along any dimension and in ascending or descending order of the elements. The sort functions also return an optional array of indices showing the order in which elements were rearranged during the sorting operation.

This section covers

- “Shift and Sort Functions” on page 1-35
- “Shifting the Location of Matrix Elements” on page 1-35
- “Sorting the Data in Each Column” on page 1-37
- “Sorting the Data in Each Row” on page 1-37
- “Sorting Row Vectors” on page 1-38

### Shift and Sort Functions

Use these functions to shift or sort the elements of a matrix.

Function	Description
<code>circshift</code>	Circularly shift matrix contents.
<code>sort</code>	Sort array elements in ascending or descending order.
<code>sortrows</code>	Sort rows in ascending order.
<code>issorted</code>	Determine if matrix elements are in sorted order.

### Shifting the Location of Matrix Elements

The `circshift` function shifts the elements of a matrix in a circular manner along one or more dimensions. Rows or columns that are shifted out of the matrix circulate back into the opposite end. For example, shifting a 4-by-7 matrix one place to the left moves the elements in columns 2 through 7 to columns 1 through 6, and moves column 1 to column 7.

Create a 5-by-8 matrix named A and shift it to the right along the second (horizontal) dimension by three places. (You would use [0, -3] to shift to the left by three places):

```
A = [1:8; 11:18; 21:28; 31:38; 41:48]
A =
     1     2     3     4     5     6     7     8
    11    12    13    14    15    16    17    18
    21    22    23    24    25    26    27    28
    31    32    33    34    35    36    37    38
    41    42    43    44    45    46    47    48
```

```
B = circshift(A, [0, 3])
B =
     6     7     8     1     2     3     4     5
    16    17    18    11    12    13    14    15
    26    27    28    21    22    23    24    25
    36    37    38    31    32    33    34    35
    46    47    48    41    42    43    44    45
```

Now take A and shift it along both dimensions: three columns to the right and two rows up:

```
A = [1:8; 11:18; 21:28; 31:38; 41:48];
B = circshift(A, [-2, 3])
B =
    26    27    28    21    22    23    24    25
    36    37    38    31    32    33    34    35
    46    47    48    41    42    43    44    45
     6     7     8     1     2     3     4     5
    16    17    18    11    12    13    14    15
```

Since `circshift` circulates shifted rows and columns around to the other end of a matrix, shifting by the exact size of A returns all rows and columns to their original location:

```
B = circshift(A, size(A));
all(B(:) == A(:))           % Do all elements of B equal A?
ans =
     1                       % Yes
```

## Sorting the Data in Each Column

The `sort` function sorts matrix elements along a specified dimension. The syntax for the function is

```
sort(matrix, dimension)
```

To sort the columns of a matrix, specify 1 as the dimension argument. To sort along rows, specify dimension as 2.

This example first constructs a 3-by-5 matrix:

```
A = rand(3,5) * 10
A =
    9.5013    7.6210    6.1543    4.0571    0.5789
    2.3114    4.5647    7.9194    9.3547    3.5287
    6.0684    0.1850    9.2181    9.1690    8.1317
```

Sort each column of A in ascending order:

```
c = sort(A, 1)
c =
    2.3114    0.1850    6.1543    4.0571    0.5789
    6.0684    4.5647    7.9194    9.1690    3.5287
    9.5013    7.6210    9.2181    9.3547    8.1317
```

```
issorted(c(:, 1))
ans =
    1
```

## Sorting the Data in Each Row

Sort each row of A in descending order. Note that `issorted` tests for an ascending sequence. You can flip the vector to test for a sorted descending sequence:

```
r = sort(A, 2, 'descend')
r =
    9.5013    7.6210    6.1543    4.0571    0.5789
    9.3547    7.9194    4.5647    3.5287    2.3114
    9.2181    9.1690    8.1317    6.0684    0.1850
```

```

issorted(fliplr(r(1, :)))
ans =
     1

```

When you specify a second output, `sort` returns the indices of the original matrix `A` positioned in the order they appear in the output matrix. In the following example, the second row of `index` contains the sequence 4 3 2 5 1, which means that the sorted elements in output matrix `r` were taken from `A(2,4)`, `A(2,3)`, `A(2,2)`, `A(2,5)`, and `A(2,1)`:

```

[r index] = sort(A, 2, 'descend');
index
index =
     1     2     3     4     5
     4     3     2     5     1
     3     4     5     1     2

```

## Sorting Row Vectors

The `sortrows` function keeps the elements of each row in their original order but sorts the entire row vectors according to the order of the elements in the first column:

```

rowsort = sortrows(A)
rowsort =
     2.3114     4.5647     7.9194     9.3547     3.5287
     6.0684     0.1850     9.2181     9.1690     8.1317
     9.5013     7.6210     6.1543     4.0571     0.5789

```

To run the sort based on a different column, include a second input argument that indicates which column to use. This example sorts the row vectors so that the elements in the third column are in ascending order:

```

rowsort = sortrows(A, 3)
rowsort =
     9.5013     7.6210     6.1543     4.0571     0.5789
     2.3114     4.5647     7.9194     9.3547     3.5287
     6.0684     0.1850     9.2181     9.1690     8.1317

issorted(rowsort(:, 3))
ans =
     1

```

## Operating on Diagonal Matrices

There are several MATLAB functions that work specifically on diagonal matrices.

Function	Description
<code>blkdiag</code>	Construct a block diagonal matrix from input arguments.
<code>diag</code>	Return a diagonal matrix or the diagonals of a matrix.
<code>trace</code>	Compute the sum of the elements on the main diagonal.
<code>tril</code>	Return the lower triangular part of a matrix.
<code>triu</code>	Return the upper triangular part of a matrix.

### Constructing a Matrix from a Diagonal Vector

The `diag` function has two operations that it can perform. You can use it to generate a diagonal matrix:

```
A = diag([12:4:32])
A =
    12     0     0     0     0     0
     0    16     0     0     0     0
     0     0    20     0     0     0
     0     0     0    24     0     0
     0     0     0     0    28     0
     0     0     0     0     0    32
```

You can also use the `diag` function to scan an existing matrix and return the values found along one of the diagonals:

```
A = magic(5)
A =
    17    24     1     8    15
    23     5     7    14    16
     4     6    13    20    22
    10    12    19    21     3
    11    18    25     2     9
```

```
diag(A, 2)      % Return contents of second diagonal of A
ans =
     1
    14
    22
```

## Returning a Triangular Portion of a Matrix

The `tril` and `triu` functions return a triangular portion of a matrix, the former returning the piece from the lower left and the latter from the upper right. By default, the main diagonal of the matrix divides these two segments. You can use an alternate diagonal by specifying an offset from the main diagonal as a second input argument:

```
A = magic(6);

B = tril(A, -1)
B =
     0     0     0     0     0     0
     3     0     0     0     0     0
    31     9     0     0     0     0
     8    28    33     0     0     0
    30     5    34    12     0     0
     4    36    29    13    18     0
```

## Concatenating Matrices Diagonally

You can diagonally concatenate matrices to form a composite matrix using the `blkdiag` function. See “Creating a Block Diagonal Matrix” on page 1-9 for more information on how this works.

## Empty Matrices, Scalars, and Vectors

Although matrices are two dimensional, they do not always appear to have a rectangular shape. A 1-by-8 matrix, for example, has two dimensions yet is linear. These matrices are described in the following sections:

- “The Empty Matrix” on page 1-41

An *empty matrix* has one or more dimensions that are equal to zero. A two-dimensional matrix with both dimensions equal to zero appears in MATLAB as `[]`. The expression `A = []` assigns a 0-by-0 empty matrix to A.

- “Scalars” on page 1-44

A *scalar* is 1-by-1 and appears in MATLAB as a single real or complex number (e.g., 7, 583.62, -3.51, 5.46097e-14, 83+4i).

- “Vectors” on page 1-44

A *vector* is 1-by-n or n-by-1, and appears in MATLAB as a row or column of real or complex numbers:

### Column Vector

```
53.2
87.39
4-12i
43.9
```

### Row Vector

```
53.2 87.39 4-12i 43.9
```

## The Empty Matrix

A matrix having at least one dimension equal to zero is called an empty matrix. The simplest empty matrix is 0-by-0 in size. Examples of more complex matrices are those of dimension 0-by-5 or 10-by-0.

To create a 0-by-0 matrix, use the square bracket operators with no value specified:

```
A = [];
```

```
whos A
```

Name	Size	Bytes	Class
A	0x0	0	double array

You can create empty matrices (and arrays) of other sizes using the `zeros`, `ones`, `rand`, or `eye` functions. To create a 0-by-5 matrix, for example, use

```
A = zeros(0,5)
```

### Operating on an Empty Matrix

The basic model for empty matrices is that any operation that is defined for  $m$ -by- $n$  matrices, and that produces a result whose dimension is some function of  $m$  and  $n$ , should still be allowed when  $m$  or  $n$  is zero. The size of the result of this operation is consistent with the size of the result generated when working with nonempty values, but instead is evaluated at zero.

For example, horizontal concatenation

```
C = [A B]
```

requires that  $A$  and  $B$  have the same number of rows. So if  $A$  is  $m$ -by- $n$  and  $B$  is  $m$ -by- $p$ , then  $C$  is  $m$ -by- $(n+p)$ . This is still true if  $m$  or  $n$  or  $p$  is zero.

As with all matrices in MATLAB, you must follow the rules concerning compatible dimensions. In the following example, an attempt to add a 1-by-3 matrix to a 0-by-3 empty matrix results in an error:

```
[1 2 3] + ones(0,3)
??? Error using ==> +
Matrix dimensions must agree.
```

### Using Empty Matrices in Relational Operations

You can use empty matrices in relational operations such as “equal to” (`==`) or “greater than” (`>`) as long as both operands have the same dimensions, or the nonempty operand is scalar. The result of any relational operation involving an empty matrix is the empty matrix. Even comparing an empty matrix for equality to itself does not return `true`, but instead yields an empty matrix:

```
x = ones(0,3);
y = x;

y == x
ans =
    Empty matrix: 0-by-3
```

## Using Empty Matrices in Logical Operations

MATLAB has two distinct types of logical operators:

- Short-circuit (&&, ||) — Used in testing multiple logical conditions (e.g.,  $x \geq 50$  &&  $x < 100$ ) where each condition evaluates to a scalar true or false.
- Element-wise (&, |) — Performs a logical AND, OR, or NOT on each element of a matrix or array.

**Short-circuit Operations.** The rule for operands used in short-circuit operations is that each operand must be convertible to a logical scalar value. Because of this rule, empty matrices cannot be used in short-circuit logical operations. Such operations return an error.

The only exception is in the case where MATLAB can determine the result of a logical statement without having to evaluate the entire expression. This is true for the following two statements because the result of the entire statements are known by considering just the first term:

```
true || []  
ans =  
    1
```

```
false && []  
ans =  
    0
```

**Element-wise Operations.** Unlike the short-circuit operators, all element-wise operations on empty matrices are considered valid as long as the dimensions of the operands agree, or the nonempty operand is scalar. Element-wise operations on empty matrices always return an empty matrix:

```
true | []  
ans =  
    []
```

---

**Note** This behavior is consistent with the way MATLAB does scalar expansion with binary operators, wherein the nonscalar operand determines the size of the result.

---

## Scalars

Any individual real or complex number is represented in MATLAB as a 1-by-1 matrix called a scalar value:

```
A = 5;

ndims(A)      % Check number of dimensions in A
ans =
     2

size(A)       % Check value of row and column dimensions
ans =
     1     1
```

Use the `isscalar` function to tell if a variable holds a scalar value:

```
isscalar(A)
ans =
     1
```

## Vectors

Matrices with one dimension equal to one and the other greater than one are called vectors. Here is an example of a numeric vector:

```
A = [5.73 2-4i 9/7 25e3 .046 sqrt(32) 8j];

size(A)       % Check value of row and column dimensions
ans =
     1     7
```

You can construct a vector out of other vectors, as long as the critical dimensions agree. All components of a row vector must be scalars or other row vectors. Similarly, all components of a column vector must be scalars or other column vectors:

```
A = [29 43 77 9 21];
B = [0 46 11];

C = [A 5 ones(1,3) B]
C =
    29    43    77     9    21     5     1     1     1     0    46    11
```

Concatenating an empty matrix to a vector has no effect on the resulting vector. The empty matrix is ignored in this case:

```
A = [5.36; 7.01; []; 9.44]
A =
    5.3600
    7.0100
    9.4400
```

Use the `isvector` function to tell if a variable holds a vector:

```
isvector(A)
ans =
    1
```

## Full and Sparse Matrices

It is not uncommon to have matrices with a large number of zero-valued elements and, because MATLAB stores zeros in the same way it stores any other numeric value, these elements can use memory space unnecessarily and can sometimes require extra computing time.

Sparse matrices provide a way to store data that has a large percentage of zero elements more efficiently. While *full matrices* internally store every element in memory regardless of value, *sparse matrices* store only the nonzero elements and their row indices. Using sparse matrices can significantly reduce the amount of memory required for data storage.

You can create sparse matrices for the double and logical data types. All MATLAB built-in arithmetic, logical, and indexing operations can be applied to sparse matrices, or to mixtures of sparse and full matrices. Operations on sparse matrices return sparse matrices and operations on full matrices return full matrices.

See the section on “Sparse Matrices” in the MATLAB Mathematics documentation for more information on working with sparse matrices.

### Sparse Matrix Functions

This table shows some of the functions most commonly used when working with sparse matrices.

Function	Description
full	Convert a sparse matrix to a full matrix.
issparse	Determine if a matrix is sparse.
nnz	Return the number of nonzero matrix elements.
nonzeros	Return the nonzero elements of a matrix.
nzmax	Return the amount of storage allocated for nonzero elements.
spalloc	Allocate space for a sparse matrix.
sparse	Create a sparse matrix or convert full to sparse.

---

<b>Function</b>	<b>Description</b>
speye	Create a sparse identity matrix.
sprand	Create a sparse uniformly distributed random matrix.

## Multidimensional Arrays

An array having more than two dimensions is called a multidimensional array in MATLAB. Most of the operations that you can perform on matrices (i.e., two-dimensional arrays) can also be done on multidimensional arrays. This section shows how to create and manipulate these arrays. It covers

- “Overview” on page 1-48
- “Creating Multidimensional Arrays” on page 1-50
- “Accessing Multidimensional Array Properties” on page 1-54
- “Indexing Multidimensional Arrays” on page 1-54
- “Reshaping Multidimensional Arrays” on page 1-58
- “Permuting Array Dimensions” on page 1-60
- “Computing with Multidimensional Arrays” on page 1-62
- “Organizing Data in Multidimensional Arrays” on page 1-64
- “Multidimensional Cell Arrays” on page 1-66
- “Multidimensional Structure Arrays” on page 1-67

### Overview

Multidimensional arrays in MATLAB are an extension of the normal two-dimensional matrix. Matrices have two dimensions: the row dimension and the column dimension.

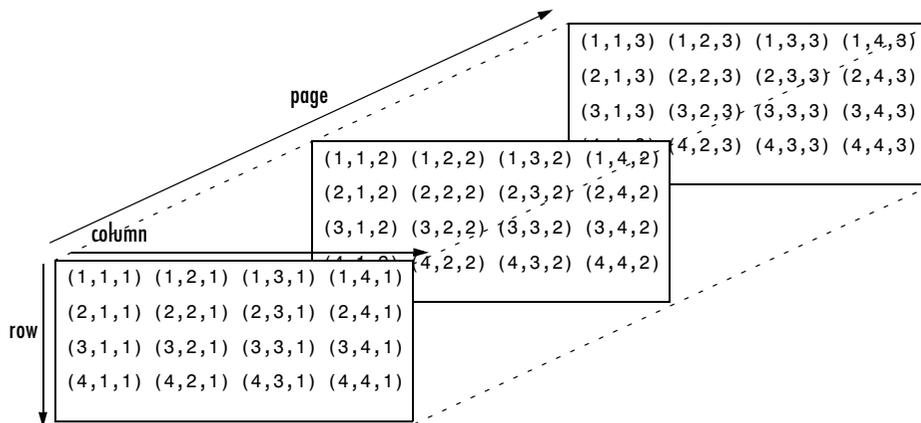
A 4x4 matrix is shown with a grid of cells. Each cell contains a pair of coordinates (row, column). Above the grid, a green arrow points to the right, labeled 'column'. To the left of the grid, a green arrow points downwards, labeled 'row'.

(1,1)	(1,2)	(1,3)	(1,4)
(2,1)	(2,2)	(2,3)	(2,4)
(3,1)	(3,2)	(3,3)	(3,4)
(4,1)	(4,2)	(4,3)	(4,4)

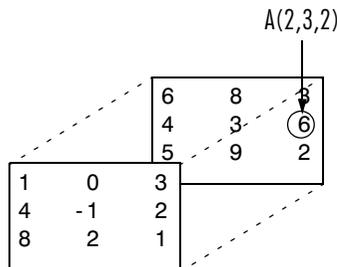
You can access a two-dimensional matrix element with two subscripts: the first representing the row index, and the second representing the column index.

Multidimensional arrays use additional subscripts for indexing. A three-dimensional array, for example, uses three subscripts:

- The first references array dimension 1, the row.
- The second references dimension 2, the column.
- The third references dimension 3. This illustration uses the concept of a *page* to represent dimensions 3 and higher.



To access the element in the second row, third column of page 2, for example, you use the subscripts (2,3,2).



$$A(:, :, 1) =$$

1	0	3
4	-1	2
8	2	1

$$A(:, :, 2) =$$

6	8	3
4	3	6
5	9	2

As you add dimensions to an array, you also add subscripts. A four-dimensional array, for example, has four subscripts. The first two reference a row-column pair; the second two access the third and fourth dimensions of data.

---

**Note** The general multidimensional array functions reside in the `datatypes` directory.

---

## Creating Multidimensional Arrays

You can use the same techniques to create multidimensional arrays that you use for two-dimensional matrices. In addition, MATLAB provides a special concatenation function that is useful for building multidimensional arrays.

This section discusses

- “Generating Arrays Using Indexing” on page 1-50
- “Extending Multidimensional Arrays” on page 1-51
- “Generating Arrays Using MATLAB Functions” on page 1-52
- “Building Multidimensional Arrays with the `cat` Function” on page 1-52

### Generating Arrays Using Indexing

One way to create a multidimensional array is to create a two-dimensional array and extend it. For example, begin with a simple two-dimensional array `A`.

```
A = [5 7 8; 0 1 9; 4 3 6];
```

`A` is a 3-by-3 array, that is, its row dimension is 3 and its column dimension is 3. To add a third dimension to `A`,

```
A(:,:,2) = [1 0 4; 3 5 6; 9 8 7]
```

MATLAB responds with

```
A(:,:,1) =  
    5     7     8  
    0     1     9  
    4     3     6
```

```
A(:,:,2) =
     1     0     4
     3     5     6
     9     8     7
```

You can continue to add rows, columns, or pages to the array using similar assignment statements.

### Extending Multidimensional Arrays

To extend A in any dimension:

- Increment or add the appropriate subscript and assign the desired values.
- Assign the same number of elements to corresponding array dimensions. For numeric arrays, all rows must have the same number of elements, all pages must have the same number of rows and columns, and so on.

You can take advantage of the MATLAB scalar expansion capabilities, together with the colon operator, to fill an entire dimension with a single value:

```
A(:,:,3) = 5;
```

```
A(:,:,3)
ans =
     5     5     5
     5     5     5
     5     5     5
```

To turn A into a 3-by-3-by-3-by-2, four-dimensional array, enter

```
A(:,:,1,2) = [1 2 3; 4 5 6; 7 8 9];
A(:,:,2,2) = [9 8 7; 6 5 4; 3 2 1];
A(:,:,3,2) = [1 0 1; 1 1 0; 0 1 1];
```

Note that after the first two assignments MATLAB pads A with zeros, as needed, to maintain the corresponding sizes of dimensions.

## Generating Arrays Using MATLAB Functions

You can use MATLAB functions such as `randn`, `ones`, and `zeros` to generate multidimensional arrays in the same way you use them for two-dimensional arrays. Each argument you supply represents the size of the corresponding dimension in the resulting array. For example, to create a 4-by-3-by-2 array of normally distributed random numbers:

```
B = randn(4,3,2)
```

To generate an array filled with a single constant value, use the `repmat` function. `repmat` replicates an array (in this case, a 1-by-1 array) through a vector of array dimensions.

```
B = repmat(5, [3 4 2])
```

```
B(:,:,1) =  
    5    5    5    5  
    5    5    5    5  
    5    5    5    5
```

```
B(:,:,2) =  
    5    5    5    5  
    5    5    5    5  
    5    5    5    5
```

---

**Note** Any dimension of an array can have size zero, making it a form of empty array. For example, 10-by-0-by-20 is a valid size for a multidimensional array.

---

## Building Multidimensional Arrays with the `cat` Function

The `cat` function is a simple way to build multidimensional arrays; it concatenates a list of arrays along a specified dimension:

```
B = cat(dim, A1, A2...)
```

where `A1`, `A2`, and so on are the arrays to concatenate, and `dim` is the dimension along which to concatenate the arrays.

For example, to create a new array with `cat`:

```
B = cat(3, [2 8; 0 5], [1 3; 7 9])
```

```
B(:,:,1) =
     2     8
     0     5
```

```
B(:,:,2) =
     1     3
     7     9
```

The `cat` function accepts any combination of existing and new data. In addition, you can nest calls to `cat`. The lines below, for example, create a four-dimensional array.

```
A = cat(3, [9 2; 6 5], [7 1; 8 4])
B = cat(3, [3 5; 0 1], [5 6; 2 1])
D = cat(4, A, B, cat(3, [1 2; 3 4], [4 3; 2 1]))
```

`cat` automatically adds subscripts of 1 between dimensions, if necessary. For example, to create a 2-by-2-by-1-by-2 array, enter

```
C = cat(4, [1 2; 4 5], [7 8; 3 2])
```

In the previous case, `cat` inserts as many singleton dimensions as needed to create a four-dimensional array whose last dimension is not a singleton dimension. If the `dim` argument had been 5, the previous statement would have produced a 2-by-2-by-1-by-1-by-2 array. This adds additional 1s to indexing expressions for the array. To access the value 8 in the four-dimensional case, use

```
C(1,2,1,2)
```

↑  
Singleton dimension  
index

## Accessing Multidimensional Array Properties

You can use the following MATLAB functions to get information about multidimensional arrays you have created.

Array Property	Function	Example
Size	size	<pre>size(C) ans =      2     2     1     2   rows columns dim3 dim4</pre>
Dimensions	ndims	<pre>ndims(C) ans =      4</pre>
Storage and format	whos	<pre>whos Name      Size      Bytes  Class A         2x2x2      64    double array B         2x2x2      64    double array C         4-D        64    double array D         4-D       192    double array  Grand total is 48 elements using 384 bytes</pre>

## Indexing Multidimensional Arrays

Many of the concepts that apply to two-dimensional matrices extend to multidimensional arrays as well.

To access a single element of a multidimensional array, use integer subscripts. Each subscript indexes a dimension—the first indexes the row dimension, the second indexes the column dimension, the third indexes the first page dimension, and so on.

Consider a 10-by-5-by-3 array `nndata` of random integers:

```
nndata = fix(8 * randn(10,5,3));
```

To access element (3,2) on page 2 of `nndata`, for example, use `nndata(3,2,2)`.

You can use vectors as array subscripts. In this case, each vector element must be a valid subscript, that is, within the bounds defined by the dimensions of the array. To access elements (2,1), (2,3), and (2,4) on page 3 of `nndata`, use

```
nndata(2,[1 3 4],3);
```

### The Colon and Multidimensional Array Indexing

The MATLAB colon indexing extends to multidimensional arrays. For example, to access the entire third column on page 2 of `nndata`, use `nndata(:,3,2)`.

The colon operator is also useful for accessing other subsets of data. For example, `nndata(2:3,2:3,1)` results in a 2-by-2 array, a subset of the data on page 1 of `nndata`. This matrix consists of the data in rows 2 and 3, columns 2 and 3, on the first page of the array.

The colon operator can appear as an array subscript on both sides of an assignment statement. For example, to create a 4-by-4 array of zeros:

```
C = zeros(4, 4)
```

Now assign a 2-by-2 subset of array `nndata` to the four elements in the center of `C`.

```
C(2:3,2:3) = nndata(2:3,1:2,2)
```

### Linear Indexing with Multidimensional Arrays

MATLAB linear indexing also extends to multidimensional arrays. In this case, MATLAB operates on a page-by-page basis to create the storage column, again appending elements columnwise. See “Linear Indexing” on page 1-18 for an introduction to this topic.

For example, consider a 5-by-4-by-3-by-2 array C.

MATLAB displays C as

MATLAB stores C as

page(1,1) =

1	4	3	5
2	1	7	9
5	6	3	2
0	1	5	9
3	2	7	5

1
2
5
0
3

page(2,1) =

6	2	4	2
7	1	4	9
0	0	1	5
9	4	4	2
1	8	2	5

4
1
6
1
2
3
7
3
5
7
5

page(3,1) =

2	2	8	3
2	5	1	8
5	1	5	2
0	9	0	9
9	4	5	3

9
2
9
5
6
7
0
9
1
2
1

page(1,2) =

9	8	2	3
0	0	3	3
6	4	9	6
1	9	2	3
0	2	8	7

4
4
1
4
2

page(2,2) =

7	0	1	3
2	4	8	1
7	5	8	6
6	8	8	4
9	4	1	2

2
9
5
2
5
2
2
2
5
0

page(3,2) =

1	6	6	5
2	9	1	3
7	1	1	1
8	0	1	5
3	2	7	6

2
5
1
9
4
:

Again, a single subscript indexes directly into this column. For example, `C(4)` produces the result

```
ans =
     0
```

If you specify two subscripts  $(i, j)$  indicating row-column indices, MATLAB calculates the offset as described above. Two subscripts always access the first page of a multidimensional array, provided they are within the range of the original array dimensions.

If more than one subscript is present, all subscripts must conform to the original array dimensions. For example, `C(6, 2)` is invalid because all pages of `C` have only five rows.

If you specify more than two subscripts, MATLAB extends its indexing scheme accordingly. For example, consider four subscripts  $(i, j, k, l)$  into a four-dimensional array with size  $[d_1 \ d_2 \ d_3 \ d_4]$ . MATLAB calculates the offset into the storage column by

$$(l-1)(d_3)(d_2)(d_1) + (k-1)(d_2)(d_1) + (j-1)(d_1) + i$$

For example, if you index the array `C` using subscripts  $(3, 4, 2, 1)$ , MATLAB returns the value 5 (index 38 in the storage column).

In general, the offset formula for an array with dimensions  $[d_1 \ d_2 \ d_3 \ \dots \ d_n]$  using any subscripts  $(s_1 \ s_2 \ s_3 \ \dots \ s_n)$  is

$$(s_n-1)(d_{n-1})(d_{n-2})\dots(d_1) + (s_{n-1}-1)(d_{n-2})\dots(d_1) + \dots + (s_2-1)(d_1) + s_1$$

Because of this scheme, you can index an array using any number of subscripts. You can append any number of 1s to the subscript list because these terms become zero. For example,

```
C(3,2,1,1,1,1,1,1)
```

is equivalent to

```
C(3,2)
```

## Avoiding Ambiguity in Multidimensional Indexing

Some assignment statements, such as

```
A(:, :, 2) = 1:10
```

are ambiguous because they do not provide enough information about the shape of the dimension to receive the data. In the case above, the statement tries to assign a one-dimensional vector to a two-dimensional destination. MATLAB produces an error for such cases. To resolve the ambiguity, be sure you provide enough information about the destination for the assigned data, and that both data and destination have the same shape. For example:

```
A(1, :, 2) = 1:10;
```

## Reshaping Multidimensional Arrays

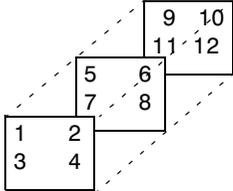
Unless you change its shape or size, a MATLAB array retains the dimensions specified at its creation. You change array size by adding or deleting elements. You change array shape by respecifying the array's row, column, or page dimensions while retaining the same elements. The reshape function performs the latter operation. For multidimensional arrays, its form is

```
B = reshape(A, [s1 s2 s3 ...])
```

s1, s2, and so on represent the desired size for each dimension of the reshaped matrix. Note that a reshaped array must have the same number of elements as the original array (that is, the product of the dimension sizes is constant).

<b>M</b>	<b>reshape(M, [6 5])</b>
<pre> 1  2  3  4  5 9  0  6  3  7 8  1  5  0  2 </pre>	<pre> 1  3  5  7  5 9  6  7  5  5 8  5  2  9  3 2  4  9  8  2 0  3  3  8  1 1  0  6  4  3 </pre>

The reshape function operates in a columnwise manner. It creates the reshaped matrix by taking consecutive elements down each column of the original data construct.

<b>C</b>	<b>reshape(C, [6 2])</b>												
	<table border="1" data-bbox="1061 472 1154 675"> <tr><td>1</td><td>6</td></tr> <tr><td>3</td><td>8</td></tr> <tr><td>2</td><td>9</td></tr> <tr><td>4</td><td>11</td></tr> <tr><td>5</td><td>10</td></tr> <tr><td>7</td><td>12</td></tr> </table>	1	6	3	8	2	9	4	11	5	10	7	12
1	6												
3	8												
2	9												
4	11												
5	10												
7	12												

Here are several new arrays from reshaping nddata:

```
B = reshape(nddata, [6 25])
C = reshape(nddata, [5 3 10])
D = reshape(nddata, [5 3 2 5])
```

### Removing Singleton Dimensions

MATLAB creates singleton dimensions if you explicitly specify them when you create or reshape an array, or if you perform a calculation that results in an array dimension of one:

```
B = repmat(5, [2 3 1 4]);
```

```
size(B)
ans =
     2     3     1     4
```

The squeeze function removes singleton dimensions from an array:

```
C = squeeze(B);
```

```
size(C)
ans =
     2     3     4
```

The squeeze function does not affect two-dimensional arrays; row vectors remain rows.

## Permuting Array Dimensions

The permute function reorders the dimensions of an array:

```
B = permute(A, dims);
```

dims is a vector specifying the new order for the dimensions of A, where 1 corresponds to the first dimension (rows), 2 corresponds to the second dimension (columns), 3 corresponds to pages, and so on.

<b>A</b>	<b>B = permute(A, [2 1 3])</b>	<b>C = permute(A, [3 2 1])</b>																								
$A(:,:,1) =$ <table style="margin-left: 40px;"> <tr><td>1</td><td>2</td><td>3</td></tr> <tr><td>4</td><td>5</td><td>6</td></tr> <tr><td>7</td><td>8</td><td>9</td></tr> </table>	1	2	3	4	5	6	7	8	9	$B(:,:,1) =$ <table style="margin-left: 40px;"> <tr><td>1</td><td>4</td><td>7</td></tr> <tr><td>2</td><td>5</td><td>8</td></tr> <tr><td>3</td><td>6</td><td>9</td></tr> </table> <p style="margin-left: 40px;">Row and column subscripts are reversed (page-by-page transposition).</p>	1	4	7	2	5	8	3	6	9	$C(:,:,1) =$ <table style="margin-left: 40px;"> <tr><td>1</td><td>2</td><td>3</td></tr> <tr><td>0</td><td>5</td><td>4</td></tr> </table> <p style="margin-left: 40px;">Row and page subscripts are reversed.</p>	1	2	3	0	5	4
1	2	3																								
4	5	6																								
7	8	9																								
1	4	7																								
2	5	8																								
3	6	9																								
1	2	3																								
0	5	4																								
$A(:,:,2) =$ <table style="margin-left: 40px;"> <tr><td>0</td><td>5</td><td>4</td></tr> <tr><td>2</td><td>7</td><td>6</td></tr> <tr><td>9</td><td>3</td><td>1</td></tr> </table>	0	5	4	2	7	6	9	3	1	$B(:,:,2) =$ <table style="margin-left: 40px;"> <tr><td>0</td><td>2</td><td>9</td></tr> <tr><td>5</td><td>7</td><td>3</td></tr> <tr><td>4</td><td>6</td><td>1</td></tr> </table>	0	2	9	5	7	3	4	6	1	$C(:,:,2) =$ <table style="margin-left: 40px;"> <tr><td>4</td><td>5</td><td>6</td></tr> <tr><td>2</td><td>7</td><td>6</td></tr> </table>	4	5	6	2	7	6
0	5	4																								
2	7	6																								
9	3	1																								
0	2	9																								
5	7	3																								
4	6	1																								
4	5	6																								
2	7	6																								
		$C(:,:,3) =$ <table style="margin-left: 40px;"> <tr><td>7</td><td>8</td><td>9</td></tr> <tr><td>9</td><td>3</td><td>1</td></tr> </table>	7	8	9	9	3	1																		
7	8	9																								
9	3	1																								

For a more detailed look at the permute function, consider a four-dimensional array A of size 5-by-4-by-3-by-2. Rearrange the dimensions, placing the column dimension first, followed by the second page dimension, the first page dimension, then the row dimension. The result is a 4-by-2-by-3-by-5 array.

```
B = permute(A, [2 4 3 1])
```

Move dimension 2 of A to first subscript position of B, dimension 4 to second subscript position, and so on.

Input array A	Dimension	1	2	3	4
	Size	5	4	3	2

Output array B	Dimension	1	2	3	4
	Size	4	2	3	5

The order of dimensions in `permute`'s argument list determines the size and shape of the output array. In this example, the second dimension of the original array had size 4, the output array's first dimension also has size 4.

You can think of `permute`'s operation as an extension of the `transpose` function, which switches the row and column dimensions of a matrix. For `permute`, the order of the input dimension list determines the reordering of the subscripts. In the example above, element (4, 2, 1, 2) of A becomes element (2, 2, 1, 4) of B, element (5, 4, 3, 2) of A becomes element (4, 2, 3, 5) of B, and so on.

### Inverse Permutation

The `ipermute` function is the inverse of `permute`. Given an input array A and a vector of dimensions v, `ipermute` produces an array B such that `permute(B, v)` returns A.

For example, these statements create an array E that is equal to the input array C:

```
D = ipermute(C, [1 4 2 3]);
E = permute(D, [1 4 2 3])
```

You can obtain the original array after permuting it by calling `ipermute` with the same vector of dimensions.

## Computing with Multidimensional Arrays

Many of the MATLAB computational and mathematical functions accept multidimensional arrays as arguments. These functions operate on specific dimensions of multidimensional arrays; that is, they operate on individual elements, on vectors, or on matrices.

### Operating on Vectors

Functions that operate on vectors, like `sum`, `mean`, and so on, by default typically work on the first nonsingleton dimension of a multidimensional array. Most of these functions optionally let you specify a particular dimension on which to operate. There are exceptions, however. For example, the `cross` function, which finds the cross product of two vectors, works on the first nonsingleton dimension having length 3.

---

**Note** In many cases, these functions have other restrictions on the input arguments—for example, some functions that accept multiple arrays require that the arrays be the same size. Refer to the online help for details on function arguments.

---

### Operating Element-by-Element

MATLAB functions that operate element-by-element on two-dimensional arrays, like the trigonometric and exponential functions in the `elfun` directory, work in exactly the same way for multidimensional cases. For example, the `sin` function returns an array the same size as the function's input argument. Each element of the output array is the sine of the corresponding element of the input array.

Similarly, the arithmetic, logical, and relational operators all work with corresponding elements of multidimensional arrays that are the same size in every dimension. If one operand is a scalar and one an array, the operator applies the scalar to each element of the array.

## Operating on Planes and Matrices

Functions that operate on planes or matrices, such as the linear algebra and matrix functions in the `matfun` directory, do not accept multidimensional arrays as arguments. That is, you cannot use the functions in the `matfun` directory, or the array operators `*`, `^`, `\`, or `/`, with multidimensional arguments. Supplying multidimensional arguments or operands in these cases results in an error.

You can use indexing to apply a matrix function or operator to matrices within a multidimensional array. For example, create a three-dimensional array `A`:

```
A = cat(3, [1 2 3; 9 8 7; 4 6 5], [0 3 2; 8 8 4; 5 3 5], ...
          [6 4 7; 6 8 5; 5 4 3]);
```

Applying the `eig` function to the entire multidimensional array results in an error:

```
eig(A)
??? Error using ==> eig
Input arguments must be 2-D.
```

You can, however, apply `eig` to planes within the array. For example, use colon notation to index just one page (in this case, the second) of the array:

```
eig(A(:, :, 2))
ans =
    12.9129
   -2.6260
    2.7131
```

---

**Note** In the first case, subscripts are not colons; you must use `squeeze` to avoid an error. For example, `eig(A(2, :, :))` results in an error because the size of the input is `[1 3 3]`. The expression `eig(squeeze(A(2, :, :)))`, however, passes a valid two-dimensional matrix to `eig`.

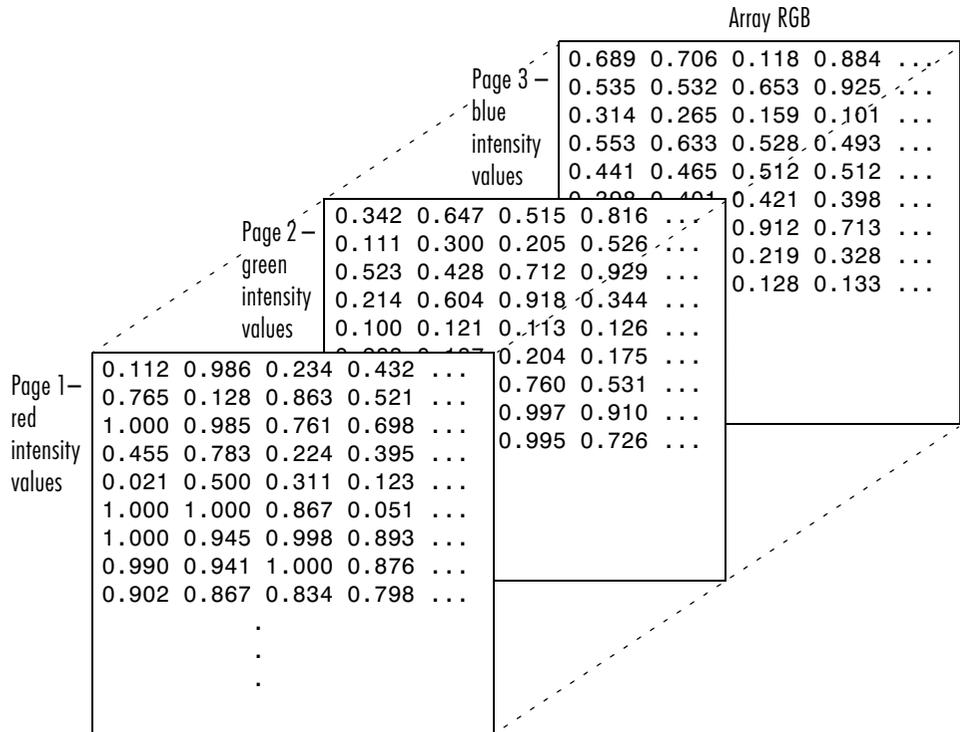
---

## Organizing Data in Multidimensional Arrays

You can use multidimensional arrays to represent data in two ways:

- As planes or pages of two-dimensional data. You can then treat these pages as matrices.
- As multivariate or multidimensional data. For example, you might have a four-dimensional array where each element corresponds to either a temperature or air pressure measurement taken at one of a set of equally spaced points in a room.

For example, consider an RGB image. For a single image, a multidimensional array is probably the easiest way to store and access data.



To access an entire plane of the image, use

```
redPlane = RGB(:,:,1);
```

To access a subimage, use

```
subimage = RGB(20:40,50:85,:);
```

The RGB image is a good example of data that needs to be accessed in planes for operations like display or filtering. In other instances, however, the data itself might be multidimensional. For example, consider a set of temperature measurements taken at equally spaced points in a room. Here the location of each value is an integral part of the data set—the physical placement in three-space of each element is an aspect of the information. Such data also lends itself to representation as a multidimensional array.

Array TEMP

			67.9°	68.0°	67.9°
			67.8°	67.8°	67.9°
		67.9°	68.0°	68.0°	67.7°
		67.7°	67.8°	67.7°	
68.0°	68.0°	67.8°	67.5°		
67.9°	67.8°	67.6°			
67.8°	67.6°	67.6°			

Now to find the average of all the measurements, use

```
mean(mean(mean(TEMP)));
```

To obtain a vector of the “middle” values (element (2,2)) in the room on each page, use

```
B = TEMP(2,2,:);
```

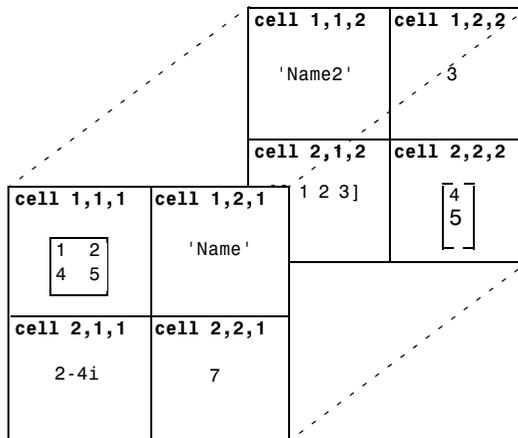
## Multidimensional Cell Arrays

Like numeric arrays, the framework for multidimensional cell arrays in MATLAB is an extension of the two-dimensional cell array model. You can use the `cat` function to build multidimensional cell arrays, just as you use it for numeric arrays.

For example, create a simple three-dimensional cell array `C`:

```
A{1,1} = [1 2;4 5];
A{1,2} = 'Name';
A{2,1} = 2-4i;
A{2,2} = 7;
B{1,1} = 'Name2';
B{1,2} = 3;
B{2,1} = 0:1:3;
B{2,2} = [4 5]';
C = cat(3, A, B);
```

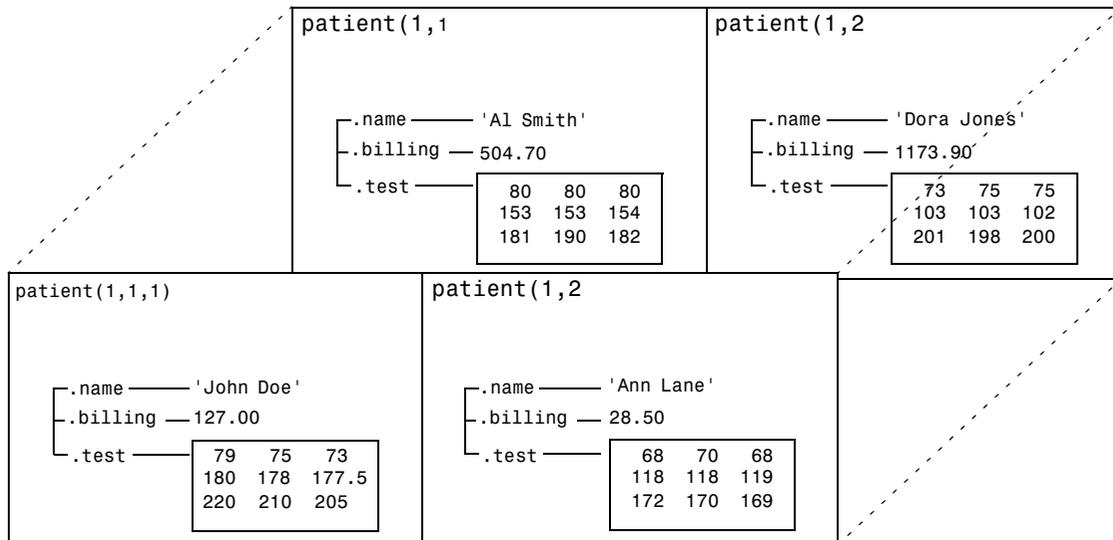
The subscripts for the cells of `C` look like



## Multidimensional Structure Arrays

Multidimensional structure arrays are extensions of rectangular structure arrays. Like other types of multidimensional arrays, you can build them using direct assignment or the `cat` function:

```
patient(1,1,1).name = 'John Doe';patient(1,1,1).billing = 127.00;
patient(1,1,1).test = [79 75 73; 180 178 177.5; 220 210 205];
patient(1,2,1).name = 'Ann Lane';patient(1,2,1).billing = 28.50;
patient(1,2,1).test = [68 70 68; 118 118 119; 172 170 169];
patient(1,1,2).name = 'Al Smith';patient(1,1,2).billing = 504.70;
patient(1,1,2).test = [80 80 80; 153 153 154; 181 190 182];
patient(1,2,2).name = 'Dora Jones';patient(1,2,2).billing =
1173.90;
patient(1,2,2).test = [73 73 75; 103 103 102; 201 198 200];
```



### **Applying Functions to Multidimensional Structure Arrays**

To apply functions to multidimensional structure arrays, operate on fields and field elements using indexing. For example, find the sum of the columns of the test array in patient(1,1,2):

```
sum(patient(1,1,2).test);
```

Similarly, add all the billing fields in the patient array:

```
total = sum([patient.billing]);
```

## Summary of Matrix and Array Functions

This section summarizes the principal functions used in creating and handling matrices. Most of these functions work on multidimensional arrays as well.

### Functions to Create a Matrix

Function	Description
[a,b] or [a;b]	Create a matrix from specified elements, or concatenate matrices together.
accumarray	Construct a matrix using accumulation.
blkdiag	Construct a block diagonal matrix.
cat	Concatenate matrices along the specified dimension.
diag	Create a diagonal matrix from a vector.
horzcat	Concatenate matrices horizontally.
magic	Create a square matrix with rows, columns, and diagonals that add up to the same number.
ones	Create a matrix of all ones.
rand	Create a matrix of uniformly distributed random numbers.
repmat	Create a new matrix by replicating or tiling another.
vertcat	Concatenate two or more matrices vertically.
zeros	Create a matrix of all zeros.

### Functions to Modify the Shape of a Matrix

Function	Description
ctranspose	Flip a matrix about the main diagonal and replace each element with its complex conjugate.
flipdim	Flip a matrix along the specified dimension.

**Functions to Modify the Shape of a Matrix**

<b>Function</b>	<b>Description</b>
flipplr	Flip a matrix about a vertical axis.
flipud	Flip a matrix about a horizontal axis.
reshape	Change the dimensions of a matrix.
rot90	Rotate a matrix by 90 degrees.
transpose	Flip a matrix about the main diagonal.

**Functions to Find the Structure or Shape of a Matrix**

<b>Function</b>	<b>Description</b>
isempty	Return true for 0-by-0 or 0-by-n matrices.
isscalar	Return true for 1-by-1 matrices.
issparse	Return true for sparse matrices.
isvector	Return true for 1-by-n matrices.
length	Return the length of a vector.
ndims	Return the number of dimensions in a matrix.
numel	Return the number of elements in a matrix.
size	Return the size of each dimension.

**Functions to Determine Data Type**

<b>Function</b>	<b>Description</b>
iscell	Return true if the matrix is a cell array.
ischar	Return true if matrix elements are characters or strings.
isfloat	Determine if input is a floating point array.
isinteger	Determine if input is an integer array.

**Functions to Determine Data Type**

<b>Function</b>	<b>Description</b>
islogical	Return true if matrix elements are logicals.
isnumeric	Return true if matrix elements are numeric.
isreal	Return true if matrix elements are real numbers.
isstruct	Return true if matrix elements are MATLAB structures.

**Functions to Sort and Shift Matrix Elements**

<b>Function</b>	<b>Description</b>
circshift	Circularly shift matrix contents.
issorted	Return true if the matrix elements are sorted.
sort	Sort elements in ascending or descending order.
sortrows	Sort rows in ascending order.

**Functions That Work on Diagonals of a Matrix**

<b>Function</b>	<b>Description</b>
blkdiag	Construct a block diagonal matrix.
diag	Return the diagonals of a matrix.
trace	Compute the sum of the elements on the main diagonal.
tril	Return the lower triangular part of a matrix.
triu	Return the upper triangular part of a matrix.

**Functions to Change the Indexing Style**

<b>Function</b>	<b>Description</b>
ind2sub	Convert a linear index to a row-column index.
sub2ind	Convert a row-column index to a linear index.

**Functions for Working with Multidimensional Arrays**

<b>Function</b>	<b>Description</b>
cat	Concatenate arrays.
circshift	Shift array circularly.
ipermute	Inverse permute array dimensions.
ndgrid	Generate arrays for n-dimensional functions and interpolation.
ndims	Return the number of array dimensions.
permute	Permute array dimensions.
shiftdim	Shift array dimensions.
squeeze	Remove singleton dimensions.

# Data Types

---

There are many different types of data that you can work with in MATLAB. You can build matrices and arrays of floating-point and integer data, characters and strings, logical true and false states, etc. You can also develop your own data types using MATLAB classes. Two of the MATLAB data types, structures and cell arrays, provide a way to store dissimilar types of data in the same array. This chapter describes each of these data types and how to use them in your MATLAB programming.

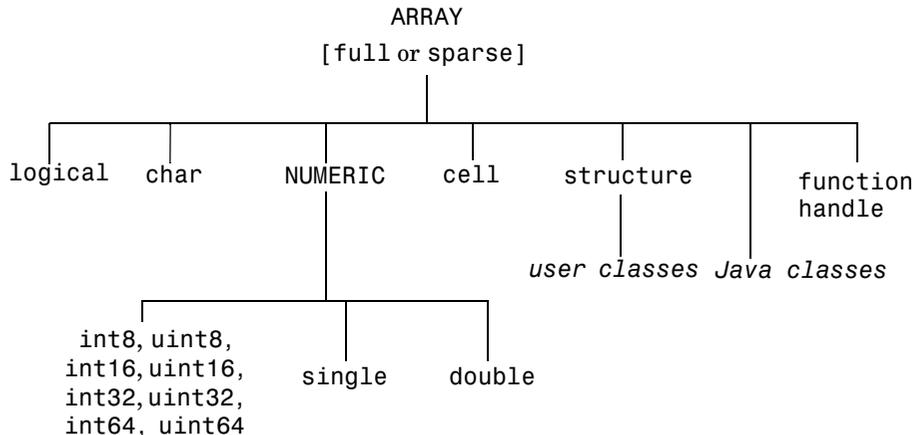
Overview of MATLAB Data Types (p. 2-2)	Brief description of all MATLAB data types
Numeric Types (p. 2-4)	Integer and floating-point data types, complex numbers, NaN, infinity, and numeric display format
Logical Types (p. 2-20)	States of true and false, use of logicals in conditional statements and logical indexing, logical/numeric conversion
Characters and Strings (p. 2-25)	Characters, strings, cell arrays of strings, string comparison, search and replace, character/numeric conversion
Dates and Times (p. 2-41)	Date strings, serial date numbers, date vectors, date type conversion, output display format
Structures (p. 2-49)	C-like structures with named fields, dynamic field names, adding and removing fields
Cell Arrays (p. 2-66)	Arrays of cells containing different data types and shapes, using cell arrays in argument lists, numeric/cell conversion
Function Handles (p. 2-80)	Passing function access data to other functions, extending function scope, extending the lifetime of variables
MATLAB Classes (p. 2-82)	Object-oriented classes and methods using MATLAB classes, creating your own MATLAB data types
Java Classes (p. 2-83)	Working with Java classes within MATLAB using the MATLAB interface to the Java programming language

## Overview of MATLAB Data Types

There are 15 fundamental data types in MATLAB. Each of these data types is in the form of a matrix or array. This matrix or array is a minimum of 0-by-0 in size and can grow to an n-dimensional array of any size.

All of the fundamental data types are shown in lowercase text in the diagram below. Additional data types are user-defined, object-oriented *user classes* and *Java classes*. You can use the latter with the MATLAB interface to Java (see “Calling Java from MATLAB” in the MATLAB External Interfaces documentation).

You can create two-dimensional double and logical matrices using one of two storage formats: full or sparse. For matrices with mostly zero-valued elements, a sparse matrix requires a fraction of the storage space required for an equivalent full matrix. Sparse matrices invoke methods especially tailored to solve sparse problems.



The following table describes these data types in more detail.

<b>Data Type</b>	<b>Example</b>	<b>Description</b>
int8, uint8, int16, uint16, int32, uint32, int64, uint64	uint16(65000)	Array of signed and unsigned integers. Requires less storage space than single or double. All integer types except for int64 and uint64 can be used in mathematical operations.
single	3 * 10 <sup>38</sup>	Array of single-precision numbers. Requires less storage space than double, but has less precision and a smaller range.
double	3 * 10 <sup>300</sup> 5 + 6i	Array of double-precision numbers. Two-dimensional arrays can be sparse. The default numeric type in MATLAB.
logical	magic(4) > 10	Array of logical values of 1 or 0 to represent true and false respectively. Two-dimensional arrays can be sparse.
char	'Hello'	Array of characters. Strings are represented as vectors of characters. For arrays containing more than one string, it is best to use cell arrays.
cell array	a{1,1} = 12; a{1,2} = 'Red'; a{1,3} = magic(4);	Array of indexed cells, each capable of storing an array of a different dimension and data type.
structure	a.day = 12; a.color = 'Red'; a.mat = magic(3);	Array of C-like structures, each structure having named fields capable of storing an array of a different dimension and data type.
function handle	@sin	Pointer to a function. You can pass function handles to other functions.
user class	polynom([0 -2 -5])	Objects constructed from a user-defined class. See “MATLAB Classes” on page 2-82
Java class	java.awt.Frame	Objects constructed from a Java class. See “Java Classes” on page 2-83.

## Numeric Types

Numeric data types in MATLAB include signed and unsigned integers, and single- and double-precision floating-point numbers. Integer and single-precision arrays offer more memory efficient storage than double precision.

All numeric types support basic array operations, such as subscripting and reshaping. All numeric types except for `int64` and `uint64` can be used in mathematical operations.

This section covers the following topics:

- “Integers” on page 2-4
- “Floating-Point Numbers” on page 2-6
- “Complex Numbers” on page 2-11
- “Infinity and NaN” on page 2-12
- “Identifying Numeric Types” on page 2-14
- “Display Format for Numeric Values” on page 2-14
- “Function Summary” on page 2-16

### Integers

MATLAB has four signed and four unsigned integer data types. Signed types enable you to work with negative integers as well as positive, but cannot represent as wide a range of numbers as the unsigned types because one bit is used to designate a positive or negative sign for the number. Unsigned types give you a wider range of numbers, but these numbers can only be zero or positive.

MATLAB supports 1-, 2-, 4-, and 8-byte storage for integer data. You can save memory and execution time for your programs if you use the smallest integer type that accommodates your data. For example, you don’t need a 32-bit integer to store the value 100.

Here are the eight integer data types, the range of values you can store with each type, and the MATLAB conversion function required to create that type:

<b>Data Type</b>	<b>Range of Values</b>	<b>Conversion Function</b>
Signed 8-bit integer	$-2^7$ to $2^7-1$	int8
Signed 16-bit integer	$-2^{15}$ to $2^{15}-1$	int16
Signed 32-bit integer	$-2^{31}$ to $2^{31}-1$	int32
Signed 64-bit integer	$-2^{63}$ to $2^{63}-1$	int64
Unsigned 8-bit integer	0 to $2^8-1$	uint8
Unsigned 16-bit integer	0 to $2^{16}-1$	uint16
Unsigned 32-bit integer	0 to $2^{32}-1$	uint32
Unsigned 64-bit integer	0 to $2^{64}-1$	uint64

### Creating Integer Data

MATLAB stores numeric data as double-precision floating point by default. To store data as an integer, use one of the conversion functions shown in the table above:

```
x = int16(32501);
```

You can use the whos function to show the dimensions, byte count, and data type of an array represented by a variable:

```
whos x
  Name      Size      Bytes  Class
  x         1x1         2     int16 array
```

Or you can use the `class` function if you want to assign the output as shown here:

```
xType = class(x)
xType =
    int16
```

Use the `isinteger` function if you just want to verify that `x` is an integer:

```
isinteger(x)
ans =
    1
```

The conversion functions are also useful when converting other data types, such as strings, to integers:

```
str = 'Hello World';

int8(str)
ans =
    72  101  108  108  111  32  87  111  114  108  100
```

### Integer Functions

See “Integer Functions” on page 2-16 for a list of functions most commonly used with integers in MATLAB.

### Floating-Point Numbers

MATLAB represents floating-point numbers in either double-precision or single-precision format. The default is double precision, but you can make any number single precision with a simple conversion function.

This section covers

- “Double-Precision Floating Point” on page 2-7
- “Single-Precision Floating Point” on page 2-9
- “Floating-Point Functions” on page 2-10

## Double-Precision Floating Point

MATLAB constructs the double data type according to IEEE Standard 754 for double precision. Any value stored as a double requires 64 bits, formatted as shown in the table below:

Bits	Usage
63	Sign (0 = positive, 1 = negative)
62 to 52	Exponent, biased by 1023
51 to 0	Fraction $f$ of the number $1.f$

**Maximum and Minimum Double-Precision Values.** The MATLAB functions `realmax` and `realmin` return the maximum and minimum values that you can represent with the double data type:

```
str = 'The range for double is:\n\t%g to %g and\n\t %g to %g';
sprintf(str, -realmax, -realmin, realmin, realmax)
```

```
ans =
The range for double is:
-1.79769e+308 to -2.22507e-308 and
2.22507e-308 to 1.79769e+308
```

Numbers larger than `realmax` or smaller than `-realmax` are assigned the values of positive and negative infinity respectively:

```
realmax + .0001e+308
ans =
    Inf

-realmax - .0001e+308
ans =
   -Inf
```

**Creating Double-Precision Data.** Since the default numeric type for MATLAB is double, you can create a double with a simple assignment statement:

```
x = 25.783;
```

The `whos` function shows that MATLAB has created a 1-by-1 array of type double for the value you just stored in `x`:

```
whos x
  Name      Size      Bytes  Class
  x         1x1         8      double array
```

Use the `isfloat` function if you just want to verify that `x` is a floating-point number:

```
isfloat(x)
ans =
     1
```

The next statement creates a much larger value. This value requires a double; you could not store this using the single-precision data type:

```
x = 5.73 * 10^300;
```

**Converting to Double Precision.** You can convert other numeric data, characters or strings, and logical data to double precision using the MATLAB function, `double`. This example converts a signed integer to double-precision floating point:

```
y = int64(-589324077574);           % Create a 64-bit integer
x = double(y)                       % Convert to double
x =
    -5.8932e+011
```

```
whos x
  Name      Size      Bytes  Class
  x         1x1         8      double array
```

## Single-Precision Floating Point

MATLAB constructs the single data type according to IEEE Standard 754 for single precision. Any value stored as a single requires 32 bits, formatted as shown in the table below:

Bits	Usage
31	Sign (0 = positive, 1 = negative)
30 to 23	Exponent, biased by 127
22 to 0	Fraction $f$ of the number $1.f$

**Maximum and Minimum Single-Precision Values.** The MATLAB functions `realmax` and `realmin`, when called with the argument `'single'`, return the maximum and minimum values that you can represent with the single data type:

```
str = 'The range for single is:\n\t%g to %g and\n\t %g to %g';
sprintf(str, -realmax('single'), -realmin('single'), ...
        realmin('single'), realmax('single'))
```

```
ans =
The range for single is:
-3.40282e+038 to -1.17549e-038 and
1.17549e-038 to 3.40282e+038
```

Numbers larger than `realmax('single')` or smaller than `-realmax('single')` are assigned the values of positive and negative infinity respectively:

```
realmax('single') + .0001e+038
ans =
    Inf

-realmax('single') - .0001e+038
ans =
   -Inf
```

**Creating Single-Precision Data.** Because MATLAB stores numeric data as a double by default, you need to use the `single` conversion function to create a single-precision number:

```
x = single(25.783);
```

The `whos` function shows that MATLAB has created a 1-by-1 array of type `single`. This value requires only 4 bytes compared with the 8 bytes used to store a double:

```
whos x
  Name      Size      Bytes  Class
  x         1x1         4      single array
```

Use the `isfloat` function if you just want to verify that `x` is a floating-point number:

```
isfloat(x)
ans =
     1
```

**Converting to Single Precision.** You can convert other numeric data, characters or strings, and logical data to single precision using the MATLAB function, `single`. This example converts a signed integer to single-precision floating point:

```
y = int64(-589324077574);           % Create a 64-bit integer
x = single(y)                       % Convert to single
x =
    -5.8932e+011
```

```
whos x
  Name      Size      Bytes  Class
  x         1x1         4      single array
```

### Floating-Point Functions

See “Floating-Point Functions” on page 2-17 for a list of functions most commonly used with floating-point numbers in MATLAB.

## Complex Numbers

Complex numbers consist of two separate parts: a real part and an imaginary part. The basic imaginary unit is equal to the square root of -1. This is represented in MATLAB by either of two letters: *i* or *j*.

### Creating Complex Numbers

The following statement shows one way of creating a complex value in MATLAB. The variable *x* is assigned a complex number with a real part of 2 and an imaginary part of 3:

```
x = 2 + 3i;
```

Another way to create a complex number is using the `complex` function. This function combines two numeric inputs into a complex output, making the first input real and the second imaginary:

```
x = rand(3) * 5;
y = rand(3) * -8;

z = complex(x, y)
z =
    4.7842 -1.0921i    0.8648 -1.5931i    1.2616 -2.2753i
    2.6130 -0.0941i    4.8987 -2.3898i    4.3787 -3.7538i
    4.4007 -7.1512i    1.3572 -5.2915i    3.6865 -0.5182i
```

You can separate a complex number into its real and imaginary parts using the `real` and `imag` functions:

```
zr = real(z)
zr =
    4.7842    0.8648    1.2616
    2.6130    4.8987    4.3787
    4.4007    1.3572    3.6865

zi = imag(z)
zi =
   -1.0921   -1.5931   -2.2753
   -0.0941   -2.3898   -3.7538
   -7.1512   -5.2915   -0.5182
```

## Complex Number Functions

See “Complex Number Functions” on page 2-18 for a list of functions most commonly used with MATLAB complex numbers in MATLAB.

## Infinity and NaN

MATLAB uses the special values `inf`, `-inf`, and `NaN` to represent values that are positive and negative infinity, and not a number respectively.

### Infinity

MATLAB represents infinity by the special value `inf`. Infinity results from operations like division by zero and overflow, which lead to results too large to represent as conventional floating-point values. MATLAB also provides a function called `inf` that returns the IEEE arithmetic representation for positive infinity as a double scalar value.

Several examples of statements that return positive or negative infinity in MATLAB are shown here.

---

<pre>x = 1/0 x =     Inf</pre>	<pre>x = 1.e1000 x =     Inf</pre>
<pre>x = exp(1000) x =     Inf</pre>	<pre>x = log(0) x =    -Inf</pre>

---

Use the `isinf` function to verify that `x` is positive or negative infinity:

```
x = log(0);

isinf(x)
ans =
     1
```

### NaN

MATLAB represents values that are not real or complex numbers with a special value called `NaN`, which stands for Not a Number. Expressions like `0/0` and `inf/inf` result in `NaN`, as do any arithmetic operations involving a `NaN`.

For example, the statement `n/0`, where `n` is complex, returns NaN:

```
x = 7i/0
x =
     NaN +      Inf i
```

Use the `isnan` function to verify that `x` is NaN:

```
x = log(0);

isnan(x)
ans =
     1
```

MATLAB also provides a function called `NaN` that returns the IEEE arithmetic representation for NaN as a double scalar value:

```
x = NaN;

whos x
      Name      Size      Bytes  Class
      x         1x1         8      double array
```

**Logical Operations on NaN.** Because two NaNs are not equal to each other, logical operations involving NaN always return false, except for a test for inequality, (`NaN ~= NaN`):

```
NaN > NaN
ans =
     0

NaN ~= NaN
ans =
     1
```

## Infinity and NaN Functions

See “Infinity and NaN Functions” on page 2-18 for a list of functions most commonly used with `inf` and `NaN` in MATLAB.

## Identifying Numeric Types

You can check the data type of a variable `x` using any of these commands.

Command	Operation
<code>whos x</code>	Display the data type of <code>x</code> .
<code>xType = class(x);</code>	Assign the data type of <code>x</code> to a variable.
<code>isnumeric(x)</code>	Determine if <code>x</code> is a numeric type.
<code>isa(x, 'integer')</code> <code>isa(x, 'uint64')</code> <code>isa(x, 'float')</code> <code>isa(x, 'double')</code> <code>isa(x, 'single')</code>	Determine if <code>x</code> is the specified numeric type. (Examples for any integer, unsigned 64-bit integer, any floating point, double precision, and single precision are shown here).
<code>isreal(x)</code>	Determine if <code>x</code> is real or complex.
<code>isnan(x)</code>	Determine if <code>x</code> is Not a Number (NaN).
<code>isinf(x)</code>	Determine if <code>x</code> is infinite.
<code>isfinite(x)</code>	Determine if <code>x</code> is finite.

## Display Format for Numeric Values

By default, MATLAB displays numeric output as 5-digit scaled, fixed-point values. You can change the way numeric values are displayed to any of the following:

- 5-digit scaled fixed point, floating point, or the best of the two
- 15-digit scaled fixed point, floating point, or the best of the two
- A ratio of small integers
- Hexadecimal (base 16)
- Bank notation

All available formats are listed on the format reference page.

To change the numeric display setting, use either the `format` function or the **Preferences** dialog box (accessible from the **MATLAB File** menu). The `format` function changes the display of numeric values for the duration of a single

MATLAB session, while your Preferences settings remain active from one session to the next. These settings affect only how numbers are displayed, not how MATLAB computes or saves them.

### Display Format Examples

Here are a few examples of the various formats and the output produced from the following two-element vector  $x$ , with components of different magnitudes.

Check the current format setting:

```
get(0, 'format')
ans =
    short
```

Set the value for  $x$  and display in 5-digit scaled fixed point:

```
x = [4/3 1.2345e-6]
x =
    1.3333    0.0000
```

Set the format to 5-digit floating point:

```
format short e
x
x =
    1.3333e+000    1.2345e-006
```

Set the format to 15-digit scaled fixed point:

```
format long
x
x =
    1.333333333333333    0.00000123450000
```

Set the format to 'rational' for small integer ratio output:

```
format rational
x
x =
    4/3    1/810045
```

Set an integer value for `x` and display it in hexadecimal (base 16) format:

```
format hex
x = uint32(876543210)
x =
    343efcea
```

### Setting Numeric Format in a Program

To temporarily change the numeric format inside a program, get the original format using the `get` function and save it in a variable. When you finish working with the new format, you can restore the original format setting using the `set` function as shown here:

```
origFormat = get(0, 'format');
format('rational');

-- Work in rational format --

set(0,'format', origFormat);
```

### Function Summary

MATLAB provides these functions for working with numeric data types:

- “Integer Functions” on page 2-16
- “Floating-Point Functions” on page 2-17
- “Complex Number Functions” on page 2-18
- “Infinity and NaN Functions” on page 2-18
- “Type Identification Functions” on page 2-18
- “Output Formatting Functions” on page 2-19

#### Integer Functions

Function	Description
<code>int8</code> , <code>int16</code> , <code>int32</code> , <code>int64</code>	Convert to signed 1-, 2-, 4-, or 8-byte integer.
<code>uint8</code> , <code>uint16</code> , <code>uint32</code> , <code>uint64</code>	Convert to unsigned 1-, 2-, 4-, or 8-byte integer.

**Integer Functions (Continued)**

<b>Function</b>	<b>Description</b>
<code>class</code>	Return the data type of an object.
<code>isa</code>	Determine if input value has the specified data type.
<code>isinteger</code>	Determine if input value is an integer array.
<code>isnumeric</code>	Determine if input value is a numeric array.

**Floating-Point Functions**

<b>Function</b>	<b>Description</b>
<code>double</code>	Convert to double precision.
<code>single</code>	Convert to single precision.
<code>class</code>	Return the data type of an object.
<code>isa</code>	Determine if input value has the specified data type.
<code>isfloat</code>	Determine if input value is a floating-point array.
<code>isnumeric</code>	Determine if input value is a numeric array.
<code>eps</code>	Return the floating-point relative accuracy. This value is the tolerance MATLAB uses in its calculations.
<code>realmax</code>	Return the largest floating-point number your computer can represent.
<code>realmin</code>	Return the smallest floating-point number your computer can represent.

**Complex Number Functions**

<b>Function</b>	<b>Description</b>
complex	Construct complex data from real and imaginary components.
i or j	Return the imaginary unit used in constructing complex data.
real	Return the real part of a complex number.
imag	Return the imaginary part of a complex number.
isreal	Determine if a number is real or imaginary.

**Infinity and NaN Functions**

<b>Function</b>	<b>Description</b>
inf	Return the IEEE value for infinity.
isnan	Detect NaN elements of an array.
isinf	Detect infinite elements of an array.
isfinite	Detect finite elements of an array.
nan	Return the IEEE value for Not a Number.

**Type Identification Functions**

<b>Function</b>	<b>Description</b>
class	Return data type (or class).
isa	Determine if input value is of the specified data type.
isfloat	Determine if input value is a floating-point array.
isinteger	Determine if input value is an integer array.
isnumeric	Determine if input value is a numeric array.

**Type Identification Functions (Continued)**

<b>Function</b>	<b>Description</b>
isreal	Determine if input value is real.
whos	Display the data type of input.

**Output Formatting Functions**

<b>Function</b>	<b>Description</b>
format	Control display format for output.

## Logical Types

The `logical` data type represents a logical true or false state using the numbers 1 and 0, respectively. Certain MATLAB functions and operators return logical true or false to indicate whether a certain condition was found to be true or not. For example, the statement `(5 * 10) > 40` returns a logical true value.

Logical data does not have to be scalar; MATLAB supports arrays of logical values as well. For example, the following statement returns a vector of logicals indicating false for the first two elements and true for the last three:

```
[30 40 50 60 70] > 40
ans =
     0     0     1     1     1
```

This section covers the following topics:

- “Creating a Logical Array” on page 2-20
- “How Logical Arrays Are Used” on page 2-22
- “Identifying Logical Arrays” on page 2-24

### Creating a Logical Array

One way of creating an array of logicals is to just enter a true or false value for each element. The `true` function returns logical one; the `false` function returns logical zero:

```
x = [true, true, false, true, false];
```

### Logical Operations on an Array

You can also perform some logical operation on an array that yields an array of logicals:

```
x = magic(4) >= 9
x =
     1     0     0     1
     0     1     1     0
     1     0     0     1
     0     1     1     0
```

The MATLAB functions that have names beginning with `is` (e.g., `ischar`, `issparse`) also return a logical value or array:

```
a = [2.5 6.7 9.2 inf 4.8];

isfinite(a)
ans =
     1     1     1     0     1
```

This table shows some of the MATLAB operations that return a logical true or false.

Function	Operation
<code>true</code> , <code>false</code>	Setting value to true or false
<code>logical</code>	Numeric to logical conversion
<code>&amp;</code> (and), <code> </code> (or), <code>~</code> (not), <code>xor</code> , <code>any</code> , <code>all</code>	Logical operations
<code>&amp;&amp;</code> , <code>  </code>	Short-circuit AND and OR
<code>==</code> (eq), <code>~=</code> (ne), <code>&lt;</code> (lt), <code>&gt;</code> (gt), <code>&lt;=</code> (le), <code>&gt;=</code> (ge)	Relational operations
All <code>is*</code> functions, <code>cellfun</code>	Test operations
<code>strcmp</code> , <code>strncmp</code> , <code>strcmpi</code> , <code>strncmpi</code>	String comparisons

### Sparse Logical Arrays

Logical arrays can also be sparse as long as they have no more than two dimensions:

```
x = sparse(magic(20) > 395)
x =
    (1,1)      1
    (1,4)      1
    (1,5)      1
    (20,18)    1
    (20,19)    1
```

## How Logical Arrays Are Used

MATLAB has two primary uses for logical arrays:

- “Using Logicals in Conditional Statements” on page 2-22
- “Using Logicals in Array Indexing” on page 2-22

Most mathematics operations are not supported on logical values.

### Using Logicals in Conditional Statements

Conditional statements are useful when you want to execute a block of code only when a certain condition is met. For example, the `sprintf` command shown below is valid only if `str` is a nonempty string. The statement

```
if ~isempty(str) && ischar(str)
```

checks for this condition and allows the `sprintf` to execute only if it is true:

```
str = 'Hello';

if ~isempty(str) && ischar(str)
    sprintf('Input string is ''%s''', str)
end

ans =
    Input string is 'Hello'
```

### Using Logicals in Array Indexing

MATLAB supports a type of array indexing that uses one array as the index into another array. For example, array `B` below indexes into elements 1, 3, 6, 7, and 10 of array `A`:

```
A = 5:5:50
A =
     5     10     15     20     25     30     35     40     45     50
B = [1 3 6 7 10];

A(B)
ans =
     5     15     30     35     50
```

In this case, the *numeric values* of array `B` designate the intended elements of `A`.

Another type of array index, a logical index, designates the elements of A based on their *position* in the indexing array, B. In this *masking* type of operation, every true element in the indexing array is treated as a positional index into the array being accessed.

**Logical Indexing Example 1.** This next example creates logical array B that satisfies the condition  $A > 0.5$ , and uses the positions of ones in B to index into A. This is called *logical indexing*:

```
A = rand(5);
B = A > 0.5;

A(B) = 0
A =
    0.2920    0.3567    0.1133         0    0.0595
         0    0.4983         0    0.2009    0.0890
    0.3358    0.4344         0    0.2731    0.2713
         0         0         0         0    0.4091
    0.0534         0         0         0    0.4740
```

A simpler way to express this is

```
A(A > 0.5) = 0
```

**Logical Indexing Example 2.** The next example highlights the location of the prime numbers in a magic square using logical indexing to set the nonprimes to 0:

```
A = magic(4)
A =
    16     2     3    13
     5    11    10     8
     9     7     6    12
     4    14    15     1

B = isprime(A)
B =
     0     1     1     1
     1     1     0     0
     0     1     0     0
     0     0     0     0
```

```
A(-B) = 0; % Logical indexing
```

```
A
A =
     0     2     3    13
     5    11     0     0
     0     7     0     0
     0     0     0     0
```

### Identifying Logical Arrays

This table shows the commands you can use to determine whether or not an array  $x$  is logical. The last function listed, `cellfun`, operates on cell arrays, which you can read about in the section “Cell Arrays” on page 2-66.

Command	Operation
<code>whos(x)</code>	Display value and data type for $x$ .
<code>islogical(x)</code>	Return true if array is logical.
<code>isa(x, 'logical')</code>	Return true if array is logical.
<code>class(x)</code>	Return string with data type name.
<code>cellfun('islogical', x)</code>	Check each cell array element for logical.

## Characters and Strings

In MATLAB, the term *string* refers to an array of Unicode characters. MATLAB represents each character internally as its corresponding numeric value. Unless you want to access these values, however, you can simply work with the characters as they display on screen.

You can use `char` to hold an  $m$ -by- $n$  array of strings as long as each string in the array has the same length. (This is because MATLAB arrays must be rectangular.) To hold an array of strings of unequal length, use a cell array.

The string is actually a vector whose components are the numeric codes for the characters. The actual characters displayed depend on the character set encoding for a given font.

This section covers

- “Creating Character Arrays” on page 2-25
- “Cell Arrays of Strings” on page 2-27
- “String Comparisons” on page 2-30
- “Searching and Replacing” on page 2-33
- “Converting from Numeric to String” on page 2-34
- “Converting from String to Numeric” on page 2-36
- “Function Summary” on page 2-38

### Creating Character Arrays

Specify character data by placing characters inside a pair of single quotes. For example, this line creates a 1-by-13 character array called `name`:

```
name = 'Thomas R. Lee';
```

In the workspace, the output of `whos` shows

Name	Size	Bytes	Class
name	1x13	26	char array

You can see that each character uses two bytes of storage internally.

The `class` and `ischar` functions show `name`'s identity as a character array:

```
class(name)
ans =
    char

ischar(name)
ans =
    1
```

You can also join two or more character arrays together to create a new character array. Use either the string concatenation function, `strcat`, or the MATLAB concatenation operator, `[]`, to do this. The latter preserves any trailing spaces found in the input arrays:

```
name = 'Thomas R. Lee';
title = ' Sr. Developer';

strcat(name, ', ', title)
ans =
    Thomas R. Lee, Sr. Developer
```

To concatenate strings vertically, use `strvcat`.

### Creating Two-Dimensional Character Arrays

When creating a two-dimensional character array, be sure that each row has the same length. For example, this line is legal because both input rows have exactly 13 characters:

```
name = ['Thomas R. Lee' ; 'Sr. Developer']
name =
    Thomas R. Lee
    Sr. Developer
```

When creating character arrays from strings of different lengths, you can pad the shorter strings with blanks to force rows of equal length:

```
name = ['Thomas R. Lee   ' ; 'Senior Developer'];
```

A simpler way to create string arrays is to use the `char` function. `char` automatically pads all strings to the length of the longest input string. In the following example, `char` pads the 13-character input string 'Thomas R. Lee' with three trailing blanks so that it will be as long as the second string:

```
name = char('Thomas R. Lee','Senior Developer')
name =
    Thomas R. Lee
    Senior Developer
```

When extracting strings from an array, use the `deblank` function to remove any trailing blanks:

```
trimname = deblank(name(1,:))
trimname =
    Thomas R. Lee

size(trimname)
ans =
     1     13
```

## Cell Arrays of Strings

Creating strings in a regular MATLAB array requires that all strings in the array be of the same length. This often means that you have to pad blanks at the end of strings to equalize their length. However, another type of MATLAB array, the cell array, can hold different sizes and types of data in an array without padding. Cell arrays provide a more flexible way to store strings of varying length.

For details on cell arrays, see “Cell Arrays” on page 2-66.

### Converting to a Cell Array of Strings

The `cellstr` function converts a character array into a cell array of strings. Consider the character array

```
data = ['Allison Jones'; 'Development   '; 'Phoenix       '];
```

Each row of the matrix is padded so that all have equal length (in this case, 13 characters).

Now use `cellstr` to create a column vector of cells, each cell containing one of the strings from the data array:

```
celldata = cellstr(data)
celldata =
    'Allison Jones'
    'Development'
    'Phoenix'
```

Note that the `cellstr` function strips off the blanks that pad the rows of the input string matrix:

```
length(celldata{3})
ans =
    7
```

The `iscellstr` function determines if the input argument is a cell array of strings. It returns a logical 1 (true) in the case of `celldata`:

```
iscellstr(celldata)
ans =
    1
```

Use `char` to convert back to a standard padded character array:

```
strings = char(celldata)
strings =
    Allison Jones
    Development
    Phoenix

length(strings(3,:))
ans =
    13
```

## Functions for Cell Arrays of Strings

This table describes the MATLAB functions for working with cell arrays.

<b>Function</b>	<b>Description</b>
cellstr	Convert a character array to a cell array of strings.
char	Convert a cell array of strings to a character array.
deblank	Remove trailing blanks from a string.
iscellstr	Return true for a cell array of strings.
sort	Sort elements in ascending or descending order.
strcat	Concatenate strings.
strcmp	Compare strings.
strmatch	Find possible matches for a string.

You can also use the following set functions with cell arrays of strings.

<b>Function</b>	<b>Description</b>
intersect	Set the intersection of two vectors.
ismember	Detect members of a set.
setdiff	Return the set difference of two vectors.
setxor	Set the exclusive OR of two vectors.
union	Set the union of two vectors.
unique	Set the unique elements of a vector.

### String Comparisons

There are several ways to compare strings and substrings:

- You can compare two strings, or parts of two strings, for equality.
- You can compare individual characters in two strings for equality.
- You can categorize every element within a string, determining whether each element is a character or white space.

These functions work for both character arrays and cell arrays of strings.

### Comparing Strings for Equality

You can use any of four functions to determine if two input strings are identical:

- `strcmp` determines if two strings are identical.
- `strncmp` determines if the first `n` characters of two strings are identical.
- `strcmpi` and `strncmpi` are the same as `strcmp` and `strncmp`, except that they ignore case.

Consider the two strings

```
str1 = 'hello';  
str2 = 'help';
```

Strings `str1` and `str2` are not identical, so invoking `strcmp` returns logical 0 (`false`). For example,

```
C = strcmp(str1, str2)  
C =  
    0
```

---

**Note** For C programmers, this is an important difference between the MATLAB `strcmp` and C's `strcmp()`, which returns 0 if the two strings are the same.

---

The first three characters of `str1` and `str2` are identical, so invoking `strncmp` with any value up to 3 returns 1:

```
C = strncmp(str1, str2, 2)
C =
    1
```

These functions work cell-by-cell on a cell array of strings. Consider the two cell arrays of strings

```
A = {'pizza'; 'chips'; 'candy'};
B = {'pizza'; 'chocolate'; 'pretzels'};
```

Now apply the string comparison functions:

```
strcmp(A,B)
ans =
    1
    0
    0

strncmp(A,B,1)
ans =
    1
    1
    0
```

### Comparing for Equality Using Operators

You can use MATLAB relational operators on character arrays, as long as the arrays you are comparing have equal dimensions, or one is a scalar. For example, you can use the equality operator (`==`) to determine which characters in two strings match:

```
A = 'fate';
B = 'cake';

A == B
ans =
    0    1    0    1
```

All of the relational operators (`>`, `>=`, `<`, `<=`, `==`, `~=`) compare the values of corresponding characters.

### Categorizing Characters Within a String

There are three functions for categorizing characters inside a string:

- `isletter` determines if a character is a letter
- `isspace` determines if a character is white space (blank, tab, or new line)
- `isstrprop` checks characters in a string to see if they match a category you specify, such as
  - Alphabetic
  - Alphanumeric
  - Lowercase or uppercase
  - Decimal digits
  - Hexadecimal digits
  - Control characters
  - Graphic characters
  - Punctuation characters
  - White-space characters

For example, create a string named `mystring`:

```
mystring = 'Room 401';
```

`isletter` examines each character in the string, producing an output vector of the same length as `mystring`:

```
A = isletter(mystring)
A =
     1     1     1     1     0     0     0     0
```

The first four elements in `A` are logical 1 (true) because the first four characters of `mystring` are letters.

## Searching and Replacing

MATLAB provides several functions for searching and replacing characters in a string. (MATLAB also supports search and replace operations using regular expressions. See “Regular Expressions” on page 3-25.)

Consider a string named `label`:

```
label = 'Sample 1, 10/28/95';
```

The `strrep` function performs the standard search-and-replace operation. Use `strrep` to change the date from '10/28' to '10/30':

```
newlabel = strrep(label, '28', '30')
newlabel =
    Sample 1, 10/30/95
```

`findstr` returns the starting position of a substring within a longer string. To find all occurrences of the string 'amp' inside `label`, use

```
position = findstr('amp', label)
position =
     2
```

The position within `label` where the only occurrence of 'amp' begins is the second character.

The `strtok` function returns the characters before the first occurrence of a delimiting character in an input string. The default delimiting characters are the set of white-space characters. You can use the `strtok` function to parse a sentence into words. For example,

```
function allWords = words(inputString)
remainder = inputString;
allWords = '';

while (any(remainder))
    [chopped,remainder] = strtok(remainder);
    allWords = strvcat(allWords, chopped);
end
```

The `strmatch` function looks through the rows of a character array or cell array of strings to find strings that begin with a given series of characters. It returns the indices of the rows that begin with these characters:

```
maxstrings = strvcat('max', 'minimax', 'maximum')
maxstrings =
    max
    minimax
    maximum

strmatch('max', maxstrings)
ans =
     1
     3
```

## Converting from Numeric to String

The functions listed in this table provide a number of ways to convert numeric data to character strings.

Function	Description	Example
<code>char</code>	Convert a positive integer to an equivalent character. (Truncates any fractional parts.)	<code>[72 105] → 'Hi'</code>
<code>int2str</code>	Convert a positive or negative integer to a character type. (Rounds any fractional parts.)	<code>[72 105] → '72 105'</code>
<code>num2str</code>	Convert a numeric type to a character type of the specified precision and format.	<code>[72 105] → '72/105/'</code> (format set to <code>%1d/</code> )
<code>mat2str</code>	Convert a numeric type to a character type of the specified precision, returning a string MATLAB can evaluate.	<code>[72 105] → '[72 105]'</code>
<code>dec2hex</code>	Convert a positive integer to a character type of hexadecimal base.	<code>[72 105] → '48 69'</code>

Function	Description	Example
dec2bin	Convert a positive integer to a character type of binary base.	[72 105] → '1001000 1101001'
dec2base	Convert a positive integer to a character type of any base from 2 through 36.	[72 105] → '110 151' (base set to 8)

### Converting to a Character Equivalent

The `char` function converts integers to Unicode character codes and returns a string composed of the equivalent characters:

```
x = [77 65 84 76 65 66];
char(x)
ans =
    MATLAB
```

### Converting to a String of Numbers

The `int2str`, `num2str`, and `mat2str` functions convert numeric values to strings where each character represents a separate digit of the input value. The `int2str` and `num2str` functions are often useful for labeling plots. For example, the following lines use `num2str` to prepare automated labels for the  $x$ -axis of a plot:

```
function plotlabel(x, y)
plot(x, y)
str1 = num2str(min(x));
str2 = num2str(max(x));
out = ['Value of f from ' str1 ' to ' str2];
xlabel(out);
```

### Converting to a Specific Radix

Another class of conversion functions changes numeric values into strings representing a decimal value in another base, such as binary or hexadecimal representation. This includes `dec2hex`, `dec2bin`, and `dec2base`.

### Converting from String to Numeric

The functions listed in this table provide a number of ways to convert character strings to numeric data.

Function	Description	Example
uintN (e.g., uint8)	Convert a character to an integer code that represents that character.	'Hi' → 72 105
str2num	Convert a character type to a numeric type.	'72 105' → [72 105]
str2double	Similar to str2num, but offers better performance and works with cell arrays of strings.	{'72' '105'} → [72 105]
hex2num	Convert a numeric type to a character type of specified precision, returning a string that MATLAB can evaluate.	'A' → '-1.4917e-154'
hex2dec	Convert a character type of hexadecimal base to a positive integer.	'A' → 10
bin2dec	Convert a positive integer to a character type of binary base.	'1010' → 10
base2dec	Convert a positive integer to a character type of any base from 2 through 36.	'12' → 10 (if base == 8)

### Converting from a Character Equivalent

Character arrays store each character as a 16-bit numeric value. Use one of the integer conversion functions (e.g., uint8) or the double function to convert strings to their numeric values, and char to revert to character representation:

```
name = 'Thomas R. Lee';

name = double(name)
name =
    84  104  111  109  97  115  32  82  46  32  76  101  101
```

```
name = char(name)
name =
    Thomas R. Lee
```

### Converting from a Numeric String

Use `str2num` to convert a character array to the numeric value represented by that string:

```
str = '37.294e-1';

val = str2num(str)
val =
    3.7294
```

The `str2double` function converts a cell array of strings to the double-precision values represented by the strings:

```
c = {'37.294e-1'; '-58.375'; '13.796'};

d = str2double(c)
d =
    3.7294
   -58.3750
    13.7960
```

```
whos
  Name      Size      Bytes  Class

  c         3x1         224   cell array
  d         3x1         24    double array
```

### Converting from a Specific Radix

To convert from a character representation of a nondecimal number to the value of that number, use one of these functions: `hex2num`, `hex2dec`, `bin2dec`, or `base2dec`.

The `hex2num` and `hex2dec` functions both take hexadecimal (base 16) inputs, but `hex2num` returns the IEEE double-precision floating-point number it represents while `hex2dec` converts to a decimal integer.

## Function Summary

MATLAB provides these functions for working with character arrays:

- “Functions to Create Character Arrays” on page 2-38
- “Functions to Modify Character Arrays” on page 2-38
- “Functions to Read and Operate on Character Arrays” on page 2-39
- “Functions to Search or Compare Character Arrays” on page 2-39
- “Functions to Determine Data Type or Content” on page 2-39
- “Functions to Convert Between Numeric and String Data Types” on page 2-40
- “Functions to Work with Cell Arrays of Strings as Sets” on page 2-40

### Functions to Create Character Arrays

Function	Description
<code>'str'</code>	Create the string specified between quotes.
<code>blanks</code>	Create a string of blanks.
<code>sprintf</code>	Write formatted data to a string.
<code>strcat</code>	Concatenate strings.
<code>strvcat</code>	Concatenate strings vertically.

### Functions to Modify Character Arrays

Function	Description
<code>deblank</code>	Remove trailing blanks.
<code>lower</code>	Make all letters lowercase.
<code>sort</code>	Sort elements in ascending or descending order.
<code>strjust</code>	Justify a string.
<code>strrep</code>	Replace one string with another.

**Functions to Modify Character Arrays**

<b>Function</b>	<b>Description</b>
strtrim	Remove leading and trailing white space.
upper	Make all letters uppercase.

**Functions to Read and Operate on Character Arrays**

<b>Function</b>	<b>Description</b>
eval	Execute a string with MATLAB expression.
sscanf	Read a string under format control.

**Functions to Search or Compare Character Arrays**

<b>Function</b>	<b>Description</b>
findstr	Find one string within another.
strcmp	Compare strings.
strcmpi	Compare strings, ignoring case.
strmatch	Find matches for a string.
strncmp	Compare the first N characters of strings.
strncmpi	Compare the first N characters, ignoring case.
strtok	Find a token in a string.

**Functions to Determine Data Type or Content**

<b>Function</b>	<b>Description</b>
iscellstr	Return true for a cell array of strings.
ischar	Return true for a character array.
isletter	Return true for letters of the alphabet.

**Functions to Determine Data Type or Content**

<b>Function</b>	<b>Description</b>
isstrprop	Determine if a string is of the specified category.
isspace	Return true for white-space characters.

**Functions to Convert Between Numeric and String Data Types**

<b>Function</b>	<b>Description</b>
char	Convert to a character or string.
cellstr	Convert a character array to a cell array of strings.
double	Convert a string to numeric codes.
int2str	Convert an integer to a string.
mat2str	Convert a matrix to an eval'able string.
num2str	Convert a number to a string.
str2num	Convert a string to a number.
str2double	Convert a string to a double-precision value.

**Functions to Work with Cell Arrays of Strings as Sets**

<b>Function</b>	<b>Description</b>
intersect	Set the intersection of two vectors.
ismember	Detect members of a set.
setdiff	Return the set difference of two vectors.
setxor	Set the exclusive OR of two vectors.
union	Set the union of two vectors.
unique	Set the unique elements of a vector.

## Dates and Times

MATLAB represents date and time information in either of three formats: date strings, serial date numbers, or date vectors. You have the choice of using any of these formats. If you work with more than one date and time format, MATLAB provides functions to help you easily convert from one format to another, (e.g., from a string to a serial date number).

When using date strings, you have an additional option of choosing from 19 different string styles to express date and/or time information.

This section covers the following topics:

- “Types of Date Formats” on page 2-41
- “Conversions Between Date Formats” on page 2-43
- “Date String Formats” on page 2-44
- “Output Formats” on page 2-44
- “Current Date and Time” on page 2-46
- “Function Summary” on page 2-47

### Types of Date Formats

The three MATLAB date and time formats are

- “Date Strings”
- “Serial Date Numbers”
- “Date Vectors”

This table shows examples of the three formats.

<b>Date Format</b>	<b>Example</b>
Date string	02-Oct-1996
Serial date number	729300
Date vector	1996 10 2 0 0 0

### Date Strings

There are a number of different styles in which to express date and time information as a date string. For example, several possibilities for October 31, 2003 at 3:45:17 in the afternoon are

```
31-Oct-2003 15:45:17
```

```
10/31/03
```

```
15:45:17
```

```
03:45:17 PM
```

If you are working with a small number of dates at the MATLAB command line, then date strings are often the most convenient format to use.

---

**Note** The MATLAB date function returns the current date and time as a date string.

---

### Serial Date Numbers

A serial date number represents a calendar date as the number of days that has passed since a fixed base date. In MATLAB, serial date number 1 is January 1, 0000. MATLAB also uses serial time to represent fractions of days beginning at midnight; for example, 6 p.m. equals 0.75 serial days. So the string '31-Oct-2003, 6:00 pm' in MATLAB is date number 731885.75.

MATLAB works internally with serial date numbers. If you are using functions that handle large numbers of dates or doing extensive calculations with dates, you get better performance if you use date numbers.

---

**Note** The MATLAB now function returns the current date and time as a serial date number.

---

## Date Vectors

Date vectors are an internal format for some MATLAB functions; you do not typically use them in calculations. A date vector contains the elements [year month day hour minute second].

---

**Note** The MATLAB `clock` function returns the current date and time as a serial vector.

---

## Conversions Between Date Formats

Functions that convert between date formats are shown below.

Function	Description
<code>datenum</code>	Convert a date string to a serial date number.
<code>datestr</code>	Convert a serial date number to a date string.
<code>datevec</code>	Split a date number or date string into individual date elements.

Here are some examples of conversions from one date format to another:

```
d1 = datenum('02-Oct-1996')
d1 =
    729300

d2 = datestr(d1 + 10)
d2 =
    12-Oct-1996

dv1 = datevec(d1)
dv1 =
    1996     10     2     0     0     0

dv2 = datevec(d2)
dv2 =
    1996     10    12     0     0     0
```

## Date String Formats

The `datetime` function is important for doing date calculations efficiently. `datetime` takes an input string in any of several formats, with `'dd-mmm-yyyy'`, `'mm/dd/yyyy'`, or `'dd-mmm-yyyy, hh:mm:ss.ss'` most common. You can form up to six fields from letters and digits separated by any other characters:

- The day field is an integer from 1 to 31.
- The month field is either an integer from 1 to 12 or an alphabetic string with at least three characters.
- The year field is a nonnegative integer: if only two digits are specified, then a year 19yy is assumed; if the year is omitted, then the current year is used as a default.
- The hours, minutes, and seconds fields are optional. They are integers separated by colons or followed by `'AM'` or `'PM'`.

For example, if the current year is 1996, then these are all equivalent:

```
'17-May-1996'  
'17-May-96'  
'17-May'  
'May 17, 1996'  
'5/17/96'  
'5/17'
```

and both of these represent the same time:

```
'17-May-1996, 18:30'  
'5/17/96/6:30 pm'
```

Note that the default format for numbers-only input follows the American convention. Thus `3/6` is March 6, not June 3.

If you create a vector of input date strings, use a column vector and be sure all strings are the same length. Fill in with spaces or zeros.

## Output Formats

The command `datestr(D, dateform)` converts a serial date `D` to one of 19 different date string output formats showing date, time, or both. The default output for dates is a day-month-year string: `01-Mar-1996`. You select an alternative output format by using the optional integer argument `dateform`.

This table shows the date string formats that correspond to each dateform value.

<b>dateform</b>	<b>Format</b>	<b>Description</b>
0	01-Mar-1996 15:45:17	day-month-year hour:minute:second
1	01-Mar-1996	day-month-year
2	03/01/96	month/day/year
3	Mar	month, three letters
4	M	month, single letter
5	3	month
6	03/01	month/day
7	1	day of month
8	Wed	day of week, three letters
9	W	day of week, single letter
10	1996	year, four digits
11	96	year, two digits
12	Mar96	month year
13	15:45:17	hour:minute:second
14	03:45:17 PM	hour:minute:second AM or PM
15	15:45	hour:minute
16	03:45 PM	hour:minute AM or PM
17	Q1-96	calendar quarter-year
18	Q1	calendar quarter

### Converting Output Format with `datestr`

Here are some examples of converting the date March 1, 1996 to various forms using the `datestr` function:

```
d = '01-Mar-1999'  
d =  
    01-Mar-1999
```

```
datestr(d)  
ans =  
    01-Mar-1999
```

```
datestr(d, 2)  
ans =  
    03/01/99
```

```
datestr(d, 17)  
ans =  
    Q1-99
```

### Current Date and Time

The function `date` returns a string for today's date:

```
date  
ans =  
    02-Oct-1996
```

The function `now` returns the serial date number for the current date and time:

```
now  
ans =  
    729300.71  
  
datestr(now)  
ans =  
    02-Oct-1996 16:56:16  
  
datestr(floor(now))  
ans =  
    02-Oct-1996
```

## Function Summary

MATLAB provides functions for time and date handling. These functions are in a directory called `timefun` in the MATLAB Toolbox:

- “Current Date and Time Functions”
- “Conversion Functions”
- “Utility Functions”
- “Timing Measurement Functions”

### Current Date and Time Functions

Function	Description
<code>clock</code>	Return the current date and time as a date vector.
<code>date</code>	Return the current date as date string.
<code>now</code>	Return the current date and time as serial date number.

### Conversion Functions

Function	Description
<code>datenum</code>	Convert to a serial date number.
<code>datestr</code>	Convert to a string representation of the date.
<code>datevec</code>	Convert to a date vector.

### Utility Functions

Function	Description
<code>addtodate</code>	Modify a date number by field.
<code>calendar</code>	Return a matrix representing a calendar.
<code>datetick</code>	Label axis tick lines with dates.
<code>eomday</code>	Return the last day of a year and month.
<code>weekday</code>	Return the current day of the week.

### Timing Measurement Functions

Function	Description
<code>cputime</code>	Return the total CPU time used by MATLAB since it was started.
<code>etime</code>	Return the time elapsed between two date vectors.
<code>tic, toc</code>	Measure the time elapsed between invoking <code>tic</code> and <code>toc</code> .



- “Writing Functions to Operate on Structures” on page 2-57
- “Organizing Data in Structure Arrays” on page 2-59
- “Nesting Structures” on page 2-63
- “Function Summary” on page 2-65

## Building Structure Arrays

You can build structures in two ways:

- Using assignment statements
- Using the struct function

### Building Structure Arrays Using Assignment Statements

You can build a simple 1-by-1 structure array by assigning data to individual fields. MATLAB automatically builds the structure as you go along. For example, create the 1-by-1 patient structure array shown at the beginning of this section:

```
patient.name = 'John Doe';  
patient.billing = 127.00;  
patient.test = [79 75 73; 180 178 177.5; 220 210 205];
```

Now entering

```
patient
```

at the command line results in

```
name: 'John Doe'  
billing: 127  
test: [3x3 double]
```

patient is an array containing a structure with three fields. To expand the structure array, add subscripts after the structure name:

```
patient(2).name = 'Ann Lane';  
patient(2).billing = 28.50;  
patient(2).test = [68 70 68; 118 118 119; 172 170 169];
```

The `patient` structure array now has size `[1 2]`. Note that once a structure array contains more than a single element, MATLAB does not display individual field contents when you type the array name. Instead, it shows a summary of the kind of information the structure contains:

```
patient
patient =
1x2 struct array with fields:
    name
    billing
    test
```

You can also use the `fieldnames` function to obtain this information. `fieldnames` returns a cell array of strings containing field names.

As you expand the structure, MATLAB fills in unspecified fields with empty matrices so that

- All structures in the array have the same number of fields.
- All fields have the same field names.

For example, entering `patient(3).name = 'Alan Johnson'` expands the `patient` array to size `[1 3]`. Now both `patient(3).billing` and `patient(3).test` contain empty matrices.

---

**Note** Field sizes do not have to conform for every element in an array. In the `patient` example, the `name` fields can have different lengths, the `test` fields can be arrays of different sizes, and so on.

---

## Building Structure Arrays Using the `struct` Function

You can preallocate an array of structures with the `struct` function. Its basic form is

```
strArray = struct('field1',val1,'field2',val2, ...)
```

where the arguments are field names and their corresponding values. A field value can be a single value, represented by any MATLAB data construct, or a cell array of values. All field values in the argument list must be of the same scale (single value or cell array).

You can use different methods for preallocating structure arrays. These methods differ in the way in which the structure fields are initialized. As an example, consider the allocation of a 1-by-3 structure array, `weather`, with the structure fields `temp` and `rainfall`. Three different methods for allocating such an array are shown in this table.

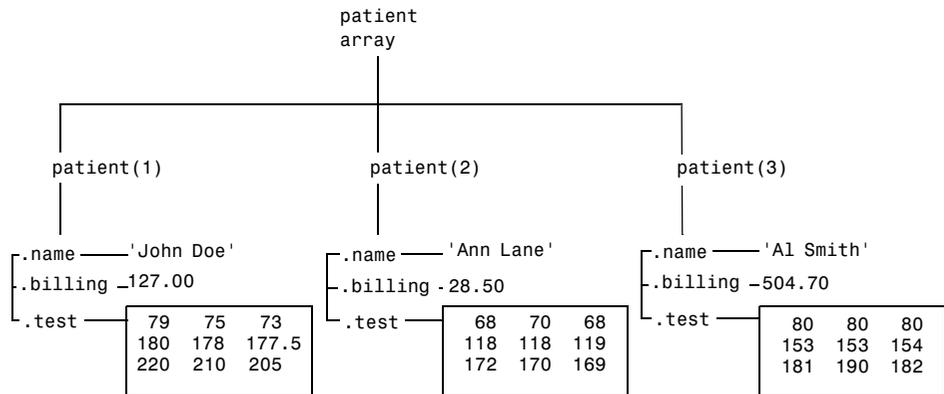
<b>Method</b>	<b>Syntax</b>	<b>Initialization</b>
struct	<pre>weather(3) = struct('temp', 72, ...     'rainfall', 0.0);</pre>	<code>weather(3)</code> is initialized with the field values shown. The fields for the other structures in the array, <code>weather(1)</code> and <code>weather(2)</code> , are initialized to the empty matrix.
struct with repmat	<pre>weather = repmat(struct('temp', ...     72, 'rainfall', 0.0), 1, 3);</pre>	All structures in the <code>weather</code> array are initialized using one set of field values.
struct with cell array syntax	<pre>weather = ...     struct('temp', {68, 80, 72}, ...     'rainfall', {0.2, 0.4, 0.0});</pre>	The structures in the <code>weather</code> array are initialized with distinct field values specified with cell arrays.

### Memory Requirements for Structures

You do not necessarily need a contiguous block of memory to store a structure. The memory for each field in the structure needs to be contiguous, but not the entire structure itself.

## Accessing Data in Structure Arrays

Using structure array indexing, you can access the value of any field or field element in a structure array. Likewise, you can assign a value to any field or field element. For the examples in this section, consider this structure array.



To access a field of a particular structure, include a period (.) after the structure name followed by the field name:

```
str = patient(2).name
str =
    Ann Lane
```

To access elements within fields, append the appropriate indexing mechanism to the field name. That is, if the field contains an array, use array subscripting; if the field contains a cell array, use cell array subscripting, and so on:

```
test2b = patient(3).test(2,2)
test2b =
    153
```

Use the same notations to assign values to structure fields, for example,

```
patient(3).test(2,2) = 7;
```

You can extract field values for multiple structures at a time. For example, the line below creates a 1-by-3 vector containing all of the billing fields:

```
bills = [patient.billing]
bills =
    127.0000    28.5000    504.7000
```

Similarly, you can create a cell array containing the test data for the first two structures:

```
tests = {patient(1:2).test}
tests =
    [3x3 double]    [3x3 double]
```

## Using Dynamic Field Names

The most common way to access the data in a structure is by specifying the name of the field that you want to reference. Another means of accessing structure data is to use dynamic field names. These names express the field as a variable expression that MATLAB evaluates at run-time. The dot-parentheses syntax shown here makes expression a dynamic field name:

```
structName.(expression)
```

Index into this field using the standard MATLAB indexing syntax. For example, to evaluate *expression* into a field name and obtain the values of that field at columns 1 through 25 of row 7, use

```
structName.(expression)(7,1:25)
```

### Dynamic Field Names Example

The `avgscore` function shown below computes an average test score, retrieving information from the `testscores` structure using dynamic field names:

```
function avg = avgscore(testscores, student, first, last)
for k = first:last
    scores(k) = testscores.(student).week(k);
end
avg = sum(scores)/(last - first + 1);
```

You can run this function using different values for the dynamic field, `student`:

```
avgscore(testscores, 'Ann Lane', 1, 20)
ans =
    83.5000
```

```
avgscore(testscores, 'William King', 1, 20)
ans =
    92.1000
```

### Finding the Size of Structure Arrays

Use the `size` function to obtain the size of a structure array, or of any structure field. Given a structure array name as an argument, `size` returns a vector of array dimensions. Given an argument in the form `array(n).field`, the `size` function returns a vector containing the size of the field contents.

For example, for the 1-by-3 structure array `patient`, `size(patient)` returns the vector `[1 3]`. The statement `size(patient(1,2).name)` returns the length of the name string for element `(1,2)` of `patient`.

## Adding Fields to Structures

You can add a field to every structure in an array by adding the field to a single structure. For example, to add a social security number field to the patient array, use an assignment like

```
patient(2).ssn = '000-00-0000';
```

Now `patient(2).ssn` has the assigned value. Every other structure in the array also has the `ssn` field, but these fields contain the empty matrix until you explicitly assign a value to them.

To add new fields to a structure, specifying the names for these fields at run-time, see the section on “Using Dynamic Field Names” on page 2-54.

## Deleting Fields from Structures

You can remove a given field from every structure within a structure array using the `rmfield` function. Its most basic form is

```
struc2 = rmfield(array, 'field')
```

where `array` is a structure array and `'field'` is the name of a field to remove from it. To remove the `name` field from the patient array, for example, enter

```
patient = rmfield(patient, 'name');
```

## Applying Functions and Operators

Operate on fields and field elements the same way you operate on any other MATLAB array. Use indexing to access the data on which to operate. For example, this statement finds the mean across the rows of the `test` array in `patient(2)`:

```
mean((patient(2).test)');
```

There are sometimes multiple ways to apply functions or operators across fields in a structure array. One way to add all the billing fields in the patient array is

```
total = 0;
for k = 1:length(patient)
    total = total + patient(k).billing;
end
```

To simplify operations like this, MATLAB enables you to operate on all like-named fields in a structure array. Simply enclose the array.field expression in square brackets within the function call. For example, you can sum all the billing fields in the patient array using

```
total = sum ([patient.billing]);
```

This is equivalent to using the comma-separated list:

```
total = sum ([patient(1).billing, patient(2).billing, ...]);
```

This syntax is most useful in cases where the operand field is a scalar field:

## Writing Functions to Operate on Structures

You can write functions that work on structures with specific field architectures. Such functions can access structure fields and elements for processing.

---

**Note** When writing M-file functions to operate on structures, you must perform your own error checking. That is, you must ensure that the code checks for the expected fields.

---

As an example, consider a collection of data that describes measurements, at different times, of the levels of various toxins in a water source. The data consists of fifteen separate observations, where each observation contains three separate measurements.

You can organize this data into an array of 15 structures, where each structure has three fields, one for each of the three measurements taken.

The function `concen`, shown below, operates on an array of structures with specific characteristics. Its arguments must contain the fields `lead`, `mercury`, and `chromium`:

```
function [r1, r2] = concen(toxtest);  
% Create two vectors. r1 contains the ratio of mercury to lead  
% at each observation. r2 contains the ratio of lead to chromium.  
r1 = [toxtest.mercury] ./ [toxtest.lead];  
r2 = [toxtest.lead] ./ [toxtest.chromium];
```

```
% Plot the concentrations of lead, mercury, and chromium
% on the same plot, using different colors for each.
lead = [toxtest.lead];
mercury = [toxtest.mercury];
chromium = [toxtest.chromium];

plot(lead, 'r'); hold on
plot(mercury, 'b')
plot(chromium, 'y'); hold off
```

Try this function with a sample structure array like test:

```
test(1).lead = .007;
test(2).lead = .031;
test(3).lead = .019;

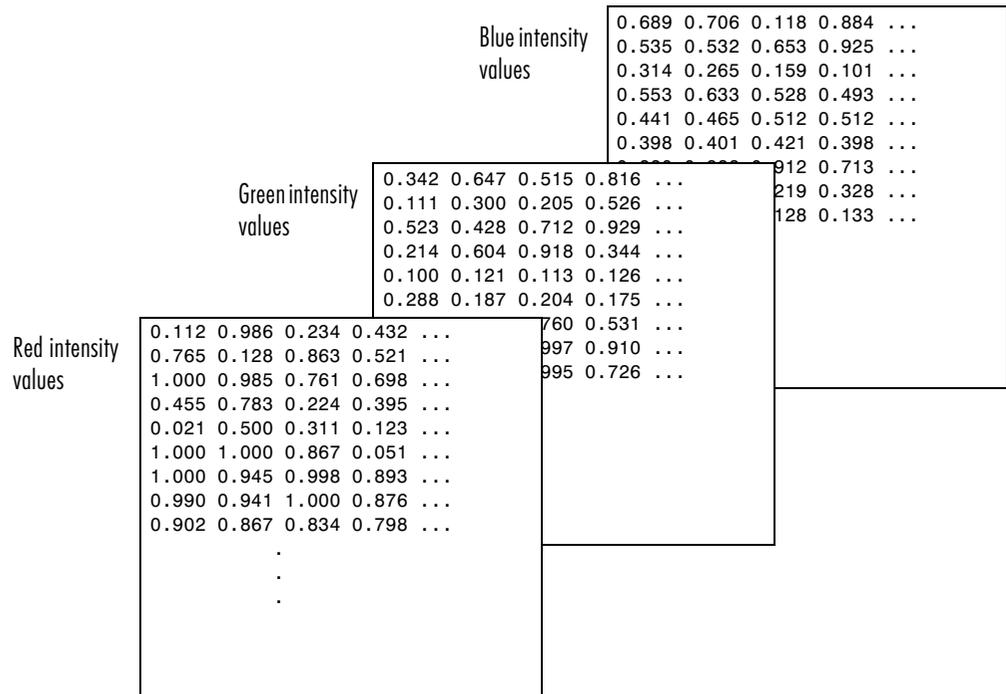
test(1).mercury = .0021;
test(2).mercury = .0009;
test(3).mercury = .0013;

test(1).chromium = .025;
test(2).chromium = .017;
test(3).chromium = .10;
```

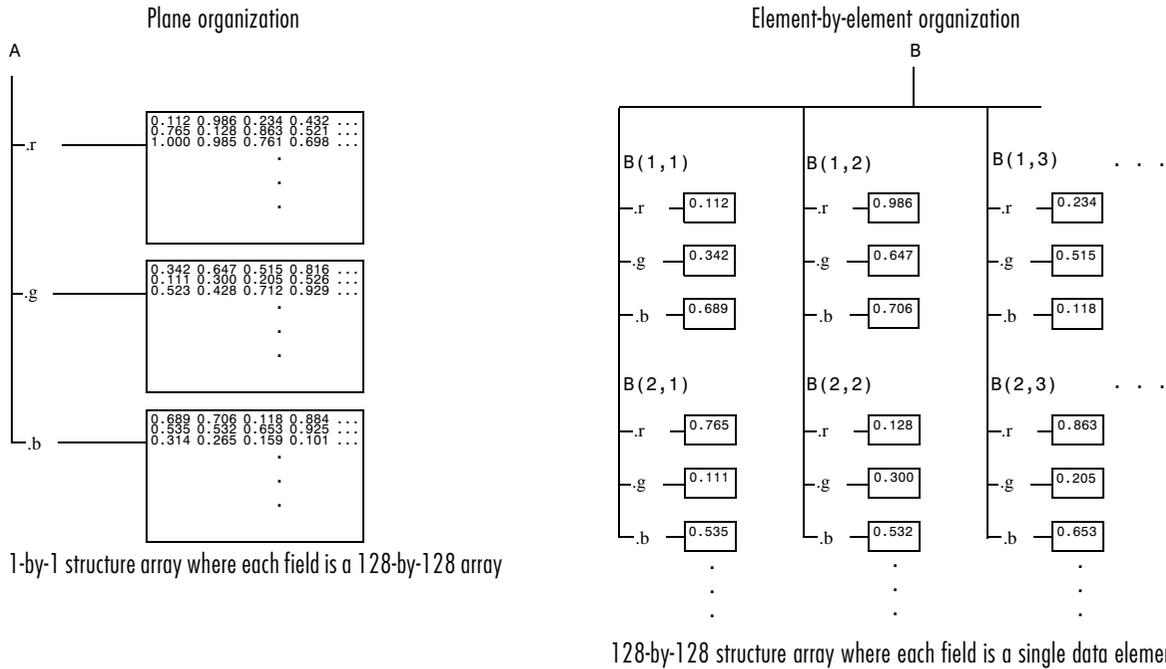
## Organizing Data in Structure Arrays

The key to organizing structure arrays is to decide how you want to access subsets of the information. This, in turn, determines how you build the array that holds the structures, and how you break up the structure fields.

For example, consider a 128-by-128 RGB image stored in three separate arrays; RED, GREEN, and BLUE.



There are at least two ways you can organize such data into a structure array.



## Plane Organization

In the plane organization, shown to the left in the figure above, each field of the structure is an entire plane of the image. You can create this structure using

```

A.r = RED;
A.g = GREEN;
A.b = BLUE;

```

This approach allows you to easily extract entire image planes for display, filtering, or other tasks that work on the entire image at once. To access the entire red plane, for example, use

```
redPlane = A.r;
```

Plane organization has the additional advantage of being extensible to multiple images in this case. If you have a number of images, you can store them as A(2), A(3), and so on, each containing an entire image.

The disadvantage of plane organization is evident when you need to access subsets of the planes. To access a subimage, for example, you need to access each field separately:

```
redSub = A.r(2:12,13:30);
greenSub = A.g(2:12,13:30);
blueSub = A.b(2:12,13:30);
```

### Element-by-Element Organization

The element-by-element organization, shown to the right in the figure above, has the advantage of allowing easy access to subsets of data. To set up the data in this organization, use

```
for m = 1:size(RED,1)
    for n = 1:size(RED,2)
        B(m,n).r = RED(m,n);
        B(m,n).g = GREEN(m,n);
        B(m,n).b = BLUE(m,n);
    end
end
```

With element-by-element organization, you can access a subset of data with a single statement:

```
Bsub = B(1:10,1:10);
```

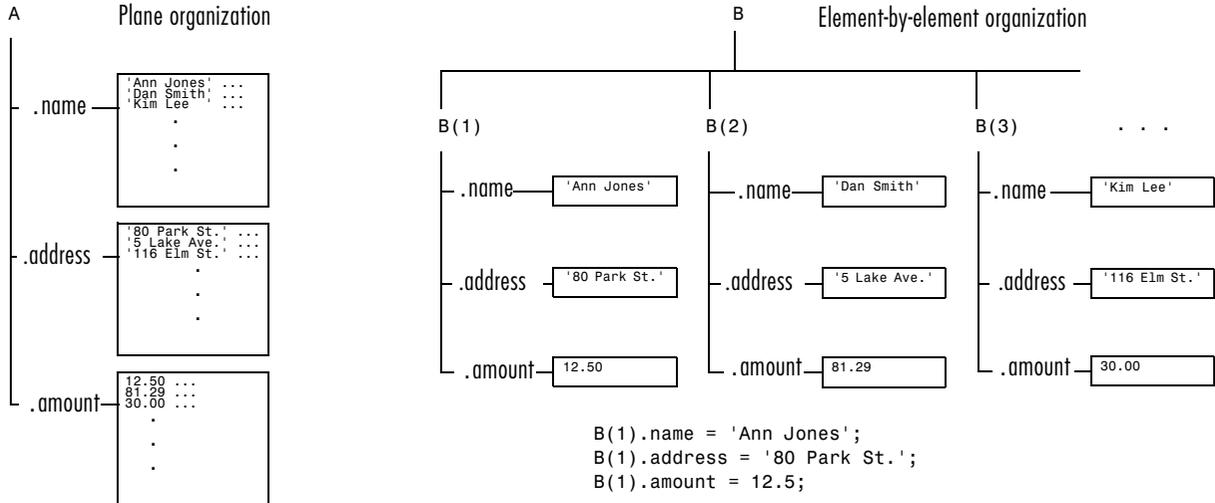
To access an entire plane of the image using the element-by-element method, however, requires a loop:

```
redPlane = zeros(128, 128);
for k = 1:(128 * 128)
    redPlane(k) = B(k).r;
end
```

Element-by-element organization is not the best structure array choice for most image processing applications; however, it can be the best for other applications wherein you will routinely need to access corresponding subsets of structure fields. The example in the following section demonstrates this type of application.

## Example — A Simple Database

Consider organizing a simple database.



```
A.name = strcat('Ann Jones', ...
'Dan Smith',...);
A.address = strcat('80 Park St.', ...
'5 Lake Ave.',...);
A.amount = [12.5; 81.29; 30; ...];
```

```
B(1).name = 'Ann Jones';
B(1).address = '80 Park St.';
B(1).amount = 12.5;
```

```
B(2).name = 'Dan Smith';
B(2).address = '5 Lake Ave.';
B(2).amount = 81.29;
```

⋮  
⋮  
⋮

Each of the possible organizations has advantages depending on how you want to access the data:

- Plane organization makes it easier to operate on all field values at once. For example, to find the average of all the values in the amount field,
  - Using plane organization
 

```
avg = mean(A.amount);
```
  - Using element-by-element organization
 

```
avg = mean([B.amount]);
```

- Element-by-element organization makes it easier to access all the information related to a single client. Consider an M-file, `client.m`, which displays the name and address of a given client on screen.

Using plane organization, pass individual fields.

```
function client(name,address)
disp(name)
disp(address)
```

To call the `client` function,

```
client(A.name(2,:),A.address(2,:))
```

Using element-by-element organization, pass an entire structure.

```
function client(B)
disp(B)
```

To call the `client` function,

```
client(B(2))
```

- Element-by-element organization makes it easier to expand the string array fields. If you do not know the maximum string length ahead of time for plane organization, you may need to frequently recreate the name or address field to accommodate longer strings.

Typically, your data does not dictate the organization scheme you choose. Rather, you must consider how you want to access and operate on the data.

## Nesting Structures

A structure field can contain another structure, or even an array of structures. Once you have created a structure, you can use the `struct` function or direct assignment statements to nest structures within existing structure fields.

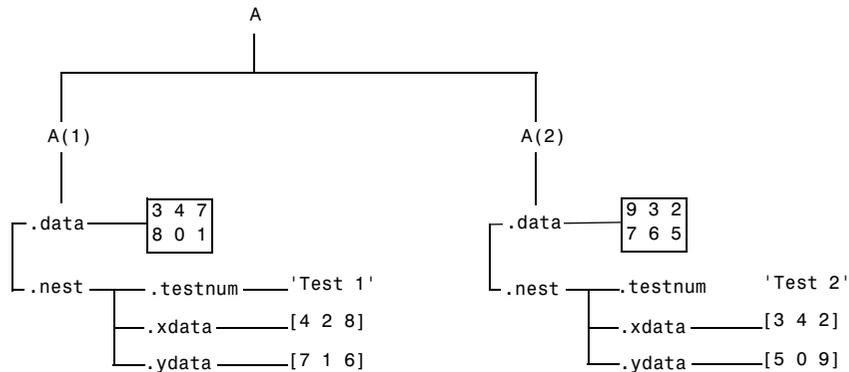
### Building Nested Structures with the `struct` Function

To build nested structures, you can nest calls to the `struct` function. For example, create a 1-by-1 structure array:

```
A = struct('data', [3 4 7; 8 0 1], 'nest',...
    struct('testnum', 'Test 1', 'xdata', [4 2 8],...
    'ydata', [7 1 6]));
```

You can build nested structure arrays using direct assignment statements. These statements add a second element to the array:

```
A(2).data = [9 3 2; 7 6 5];  
A(2).nest.testnum = 'Test 2';  
A(2).nest.xdata = [3 4 2];  
A(2).nest.ydata = [5 0 9];
```



### Indexing Nested Structures

To index nested structures, append nested field names using dot notation. The first text string in the indexing expression identifies the structure array, and subsequent expressions access field names that contain other structures.

For example, the array `A` created earlier has three levels of nesting:

- To access the nested structure inside `A(1)`, use `A(1).nest`.
- To access the `xdata` field in the nested structure in `A(2)`, use `A(2).nest.xdata`.
- To access element 2 of the `ydata` field in `A(1)`, use `A(1).nest.ydata(2)`.

## Function Summary

This table describes the MATLAB functions for working with structures.

<b>Function</b>	<b>Description</b>
<code>deal</code>	Deal inputs to outputs.
<code>fieldnames</code>	Get structure field names.
<code>isfield</code>	Return true if the field is in a structure array.
<code>isstruct</code>	Return true for structures.
<code>rmfield</code>	Remove a structure field.
<code>struct</code>	Create or convert to a structure array.
<code>struct2cell</code>	Convert a structure array into a cell array.

## Cell Arrays

A cell array provides a storage mechanism for dissimilar kinds of data. You can store arrays of different types and/or sizes within the cells of a cell array. For example, you can store a 1-by-50 char array, a 7-by-13 double array, and a 1-by-1 uint32 in cells of the same cell array.

This illustration shows a cell array that contains arrays of: unsigned integers (cell 1,1), strings (cell 1,2), complex numbers (cell 1,3), floating-point numbers (cell 2,1), signed integers (cell 2,2), and another cell array (cell 2,3).

<p><b>cell 1,1</b></p> <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"> <table style="border-collapse: collapse; text-align: center;"> <tr><td>3</td><td>4</td><td>2</td></tr> <tr><td>9</td><td>7</td><td>6</td></tr> <tr><td>8</td><td>5</td><td>1</td></tr> </table> </div>	3	4	2	9	7	6	8	5	1	<p><b>cell 1,2</b></p> <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"> <table style="border-collapse: collapse; text-align: left;"> <tr><td>'Anne Smith'</td></tr> <tr><td>'9/12/94'</td></tr> <tr><td>'Class II'</td></tr> <tr><td>'Obs. 1'</td></tr> <tr><td>'Obs. 2'</td></tr> </table> </div>	'Anne Smith'	'9/12/94'	'Class II'	'Obs. 1'	'Obs. 2'	<p><b>cell 1,3</b></p> <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"> <table style="border-collapse: collapse; text-align: left;"> <tr><td>.25+3i</td><td>8-16i</td></tr> <tr><td>34+5i</td><td>7+.92i</td></tr> </table> </div>	.25+3i	8-16i	34+5i	7+.92i											
3	4	2																													
9	7	6																													
8	5	1																													
'Anne Smith'																															
'9/12/94'																															
'Class II'																															
'Obs. 1'																															
'Obs. 2'																															
.25+3i	8-16i																														
34+5i	7+.92i																														
<p><b>cell 2,1</b></p> <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"> <table style="border-collapse: collapse; text-align: left;"> <tr><td>1.43</td><td>2.98</td><td>7.83</td><td>5.67</td></tr> <tr><td>4.21</td><td></td><td></td><td></td></tr> </table> </div>	1.43	2.98	7.83	5.67	4.21				<p><b>cell 2,2</b></p> <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"> <table style="border-collapse: collapse; text-align: left;"> <tr><td>-7</td><td>2</td><td>-14</td></tr> <tr><td>8</td><td>3</td><td>-45</td></tr> <tr><td>52</td><td>-16</td><td>3</td></tr> </table> </div>	-7	2	-14	8	3	-45	52	-16	3	<p><b>cell 2,3</b></p> <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"> <table style="border-collapse: collapse; text-align: left;"> <tr> <td style="border: 1px solid black; padding: 2px;">'text'</td> <td style="border: 1px solid black; padding: 2px;"> <table style="border-collapse: collapse; text-align: center;"> <tr><td>4</td><td>2</td></tr> <tr><td>1</td><td>5</td></tr> </table> </td> </tr> <tr> <td style="border: 1px solid black; padding: 2px;"> <table style="border-collapse: collapse; text-align: left;"> <tr><td>7.3</td><td>2.5</td></tr> <tr><td>1.4</td><td>0</td></tr> </table> </td> <td style="border: 1px solid black; padding: 2px;">.02 + 8i</td> </tr> </table> </div>	'text'	<table style="border-collapse: collapse; text-align: center;"> <tr><td>4</td><td>2</td></tr> <tr><td>1</td><td>5</td></tr> </table>	4	2	1	5	<table style="border-collapse: collapse; text-align: left;"> <tr><td>7.3</td><td>2.5</td></tr> <tr><td>1.4</td><td>0</td></tr> </table>	7.3	2.5	1.4	0	.02 + 8i
1.43	2.98	7.83	5.67																												
4.21																															
-7	2	-14																													
8	3	-45																													
52	-16	3																													
'text'	<table style="border-collapse: collapse; text-align: center;"> <tr><td>4</td><td>2</td></tr> <tr><td>1</td><td>5</td></tr> </table>	4	2	1	5																										
4	2																														
1	5																														
<table style="border-collapse: collapse; text-align: left;"> <tr><td>7.3</td><td>2.5</td></tr> <tr><td>1.4</td><td>0</td></tr> </table>	7.3	2.5	1.4	0	.02 + 8i																										
7.3	2.5																														
1.4	0																														

To access data in a cell array, you use the same matrix indexing as with other MATLAB matrices and arrays. However, with cell array indexing, you use curly braces, {}, instead of square brackets and parentheses around the array indices. For example, A{2,5} accesses the cell in row 2 and column 5 of cell array A.

---

**Note** The examples in this section focus on two-dimensional cell arrays. For examples of higher-dimension cell arrays, see “Multidimensional Arrays” on page 1-48.

---

The following list summarizes the contents of this section:

- “Creating Cell Arrays” on page 2-67
- “Obtaining Data from Cell Arrays” on page 2-70
- “Deleting Cells” on page 2-72
- “Reshaping Cell Arrays” on page 2-72
- “Replacing Lists of Variables with Cell Arrays” on page 2-72
- “Applying Functions and Operators” on page 2-74
- “Organizing Data in Cell Arrays” on page 2-75
- “Nesting Cell Arrays” on page 2-76
- “Converting Between Cell and Numeric Arrays” on page 2-78
- “Cell Arrays of Structures” on page 2-78
- “Function Summary” on page 2-79

## Creating Cell Arrays

You can create cell arrays by

- Using assignment statements
- Preallocating the array using the `cell` function, then assigning data to cells

### Creating Cell Arrays with Assignment Statements

You can build a cell array by assigning data to individual cells, one cell at a time. MATLAB automatically builds the array as you go along. There are two ways to assign data to cells:

- Cell indexing

Enclose the cell subscripts in parentheses using standard array notation. Enclose the cell contents on the right side of the assignment statement in curly braces `{}`. For example, create a 2-by-2 cell array `A`:

```
A(1,1) = {[1 4 3; 0 5 8; 7 2 9]};  
A(1,2) = {'Anne Smith'};  
A(2,1) = {3+7i};  
A(2,2) = {-pi:pi/10:pi};
```

---

**Note** The notation {} denotes the empty cell array, just as [] denotes the empty matrix for numeric arrays. You can use the empty cell array in any cell array assignments.

---

- Content indexing

Enclose the cell subscripts in curly braces using standard array notation. Specify the cell contents on the right side of the assignment statement:

```
A{1,1} = [1 4 3; 0 5 8; 7 2 9];  
A{1,2} = 'Anne Smith';  
A{2,1} = 3+7i;  
A{2,2} = -pi:pi/10:pi;
```

The various examples in this guide do not use one syntax throughout, but attempt to show representative usage of cell and content addressing. You can use the two forms interchangeably.

---

**Note** If you already have a numeric array of a given name, don't try to create a cell array of the same name by assignment without first clearing the numeric array. If you do not clear the numeric array, MATLAB assumes that you are trying to "mix" cell and numeric syntaxes, and generates an error. Similarly, MATLAB does not clear a cell array when you make a single assignment to it. If any of the examples in this section give unexpected results, clear the cell array from the workspace and try again.

---

MATLAB displays the cell array A in a condensed form:

```
A =  
      [3x3 double]      'Anne Smith'  
      [3.0000+ 7.0000i]  [1x21 double]
```

To display the full cell contents, use the `celldisp` function. For a high-level graphical display of cell architecture, use `cellplot`.

If you assign data to a cell that is outside the dimensions of the current array, MATLAB automatically expands the array to include the subscripts you

specify. It fills any intervening cells with empty matrices. For example, the assignment below turns the 2-by-2 cell array A into a 3-by-3 cell array.

```
A(3,3) = {5};
```

<b>cell 1,1</b> <table border="1"> <tbody> <tr><td>1</td><td>4</td><td>3</td></tr> <tr><td>0</td><td>5</td><td>8</td></tr> <tr><td>7</td><td>2</td><td>9</td></tr> </tbody> </table>	1	4	3	0	5	8	7	2	9	<b>cell 1,2</b> 'Anne Smith'	<b>cell 1,3</b> [ ]
1	4	3									
0	5	8									
7	2	9									
<b>cell 2,1</b> 3+7i	<b>cell 2,2</b> <table border="1"> <tbody> <tr><td>[-3.14...3.14]</td></tr> </tbody> </table>	[-3.14...3.14]	<b>cell 2,3</b> [ ]								
[-3.14...3.14]											
<b>cell 3,1</b> [ ]	<b>cell 3,2</b> [ ]	<b>cell 3,3</b> 5									

### Cell Array Syntax: Using Braces

The curly braces {} are cell array constructors, just as square brackets are numeric array constructors. Curly braces behave similarly to square brackets, except that you can nest curly braces to denote nesting of cells (see “Nesting Cell Arrays” on page 2-76 for details).

Curly braces use commas or spaces to indicate column breaks and semicolons to indicate row breaks between cells. For example,

```
C = {[1 2], [3 4]; [5 6], [7 8]};
```

results in

<b>cell 1,1</b> [1 2]	<b>cell 1,2</b> [3 4]
<b>cell 2,1</b> [5 6]	<b>cell 2,2</b> [7 8]

Use square brackets to concatenate cell arrays, just as you do for numeric arrays.

### Preallocating Cell Arrays with the cell Function

The `cell` function allows you to preallocate empty cell arrays of the specified size. For example, this statement creates an empty 2-by-3 cell array:

```
B = cell(2, 3);
```

Use assignment statements to fill the cells of B:

```
B(1,3) = {1:3};
```

The `cell` function offers the most memory-efficient way of preallocating a cell array.

### Memory Requirements for Cell Arrays

You do not necessarily need a contiguous block of memory to store a cell array. The memory for each cell needs to be contiguous, but not the entire array of cells.

### Obtaining Data from Cell Arrays

You can obtain data from cell arrays and store the result as either a standard array or a new cell array. This section covers

- Accessing cell contents using content indexing
- Accessing a subset of cells using cell indexing

### Accessing Cell Contents Using Content Indexing

You can use content indexing on the right side of an assignment to access some or all of the data in a single cell. Specify the variable to receive the cell contents on the left side of the assignment. Enclose the cell index expression on the right side of the assignment in curly braces. This indicates that you are assigning cell contents, not the cells themselves.

Consider the 2-by-2 cell array N:

```
N{1,1} = [1 2; 4 5];  
N{1,2} = 'Name';  
N{2,1} = 2-4i;  
N{2,2} = 7;
```

You can obtain the string in  $N\{1,2\}$  using

```
c = N{1,2}
c =
    Name
```

---

**Note** In assignments, you can use content indexing to access only a single cell, not a subset of cells. For example, the statements  $A\{1, : \} = value$  and  $B = A\{1, : \}$  are both invalid. However, you can use a subset of cells any place you would normally use a comma-separated list of variables (for example, as function inputs or when building an array). See “Replacing Lists of Variables with Cell Arrays” on page 2-72 for details.

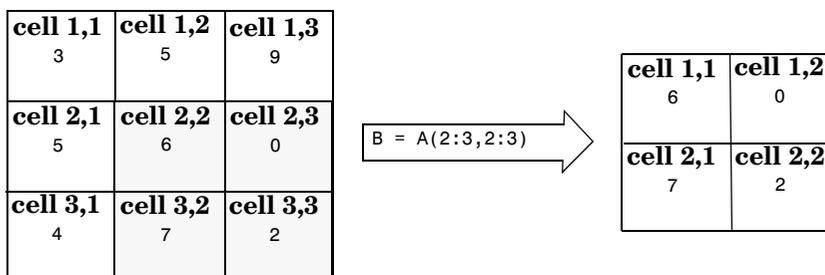
---

To obtain subsets of a cell’s contents, concatenate indexing expressions. For example, to obtain element (2,2) of the array in cell  $N\{1,1\}$ , use

```
d = N{1,1}(2,2)
d =
    5
```

### Accessing a Subset of Cells Using Cell Indexing

Use cell indexing to assign any set of cells to another variable, creating a new cell array. Use the colon operator to access subsets of cells within a cell array.



## Deleting Cells

You can delete an entire dimension of cells using a single statement. Like standard array deletion, use vector subscripting when deleting a row or column of cells and assign the empty matrix to the dimension:

```
A(cell_subscripts) = []
```

When deleting cells, curly braces do not appear in the assignment statement at all.

## Reshaping Cell Arrays

Like other arrays, you can reshape cell arrays using the reshape function. The number of cells must remain the same after reshaping; you cannot use reshape to add or remove cells:

```
A = cell(3, 4);
```

```
size(A)
ans =
     3     4
```

```
B = reshape(A, 6, 2);
```

```
size(B)
ans =
     6     2
```

## Replacing Lists of Variables with Cell Arrays

Cell arrays can replace comma-separated lists of MATLAB variables in

- Function input lists
- Function output lists
- Display operations
- Array constructions (square brackets and curly braces)

If you use the colon to index multiple cells in conjunction with the curly brace notation, MATLAB treats the contents of each cell as a separate variable. For example, assume you have a cell array `T` where each cell contains a separate vector. The expression `T{1:5}` is equivalent to a comma-separated list of the vectors in the first five cells of `T`.

Consider the cell array `C`:

```
C(1) = {[1 2 3]};
C(2) = {[1 0 1]};
C(3) = {1:10};
C(4) = {[9 8 7]};
C(5) = {3};
```

To convolve the vectors in `C(1)` and `C(2)` using `conv`,

```
d = conv(C{1:2})
d =
     1     2     4     2     3
```

Display vectors two, three, and four with

```
C{2:4}
ans =
     1     0     1

ans =
     1     2     3     4     5     6     7     8     9    10

ans =
     9     8     7
```

Similarly, you can create a new numeric array using the statement

```
B = [C{1}; C{2}; C{4}]
B =
     1     2     3
     1     0     1
     9     8     7
```

You can also use content indexing on the left side of an assignment to create a new cell array where each cell represents a separate output argument:

```
[D{1:2}] = eig(B)
D =
    [3x3 double]    [3x3 double]
```

You can display the actual eigenvectors and eigenvalues using `D{1}` and `D{2}`.

---

**Note** The `varargin` and `varargout` arguments allow you to specify variable numbers of input and output arguments for MATLAB functions that you create. Both `varargin` and `varargout` are cell arrays, allowing them to hold various sizes and kinds of MATLAB data. See “Passing Variable Numbers of Arguments” on page 4-23 for details.

---

## Applying Functions and Operators

Use indexing to apply functions and operators to the contents of cells. For example, use content indexing to call a function with the contents of a single cell as an argument:

```
A{1,1} = [1 2; 3 4];
A{1,2} = randn(3, 3);
A{1,3} = 1:5;

B = sum(A{1,1})
B =
     4     6
```

To apply a function to several cells of a non-nested cell array, use a loop:

```
for k = 1:length(A)
    M{k} = sum(A{1,k});
end
```

## Organizing Data in Cell Arrays

Cell arrays are useful for organizing data that consists of different sizes or kinds of data. Cell arrays are better than structures for applications where

- You need to access multiple fields of data with one statement.
- You want to access subsets of the data as comma-separated variable lists.
- You don't have a fixed set of field names.
- You routinely remove fields from the structure.

As an example of accessing multiple fields with one statement, assume that your data consists of

- A 3-by-4 array consisting of measurements taken for an experiment.
- A 15-character string containing a technician's name.
- A 3-by-4-by-5 array containing a record of measurements taken for the past five experiments.

For many applications, the best data construct for this data is a structure. However, if you routinely access only the first two fields of information, then a cell array might be more convenient for indexing purposes.

This example shows how to access the first and second elements of the cell array TEST:

```
[newdata,name] = deal(TEST{1:2})
```

This example shows how to access the first and second elements of the structure TEST:

```
newdata = TEST.measure  
name = TEST.name
```

The `varargin` and `varargout` arguments are examples of the utility of cell arrays as substitutes for comma-separated lists. Create a 3-by-3 numeric array A:

```
A = [0 1 2; 4 0 7; 3 1 2];
```

Now apply the `normest` (2-norm estimate) function to `A`, and assign the function output to individual cells of `B`:

```
[B{1:2}] = normest(A)
B =
    [8.8826]    [4]
```

All of the output values from the function are stored in separate cells of `B`. `B(1)` contains the norm estimate; `B(2)` contains the iteration count.

## Nesting Cell Arrays

A cell can contain another cell array, or even an array of cell arrays. (Cells that contain noncell data are called *leaf cells*.) You can use nested curly braces, the `cell` function, or direct assignment statements to create nested cell arrays. You can then access and manipulate individual cells, subarrays of cells, or cell elements.

### Building Nested Arrays with Nested Curly Braces

You can nest pairs of curly braces to create a nested cell array. For example,

```
clear A
A(1,1) = {magic(5)};

A(1,2) = {[5 2 8; 7 3 0; 6 7 3] 'Test 1'; [2-4i 5+7i] {17 []}}
A =
    [5x5 double]    {2x2 cell}
```

Note that the right side of the assignment is enclosed in two sets of curly braces. The first set represents cell (1,2) of cell array `A`. The second “packages” the 2-by-2 cell array inside the outer cell.

### Building Nested Arrays with the cell Function

To nest cell arrays with the `cell` function, assign the output of `cell` to an existing cell:

- 1 Create an empty 1-by-2 cell array.

```
A = cell(1,2);
```

**2** Create a 2-by-2 cell array inside  $A(1,2)$ .

```
A(1,2) = {cell(2,2)};
```

**3** Fill A, including the nested array, using assignments.

```
A(1,1) = {magic(5)};
A{1,2}(1,1) = {[5 2 8; 7 3 0; 6 7 3]};
A{1,2}(1,2) = {'Test 1'};
A{1,2}(2,1) = {[2-4i 5+7i]};
A{1,2}(2,2) = {cell(1, 2)}
A{1,2}{2,2}(1) = {17};
```

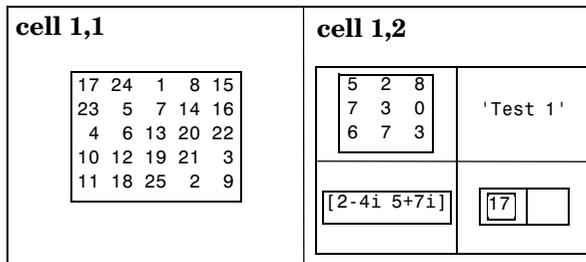
Note the use of curly braces until the final level of nested subscripts. This is required because you need to access cell contents to access cells within cells.

You can also build nested cell arrays with direct assignments using the statements shown in step 3 above.

## Indexing Nested Cell Arrays

To index nested cells, concatenate indexing expressions. The first set of subscripts accesses the top layer of cells, and subsequent sets of parentheses access successively deeper layers.

For example, array A has three levels of nesting:



- To access the 5-by-5 array in cell (1,1), use  $A\{1,1\}$ .
- To access the 3-by-3 array in position (1,1) of cell (1,2), use  $A\{1,2\}\{1,1\}$ .
- To access the 2-by-2 cell array in cell (1,2), use  $A\{1,2\}$ .
- To access the empty cell in position (2,2) of cell (1,2), use  $A\{1,2\}\{2,2\}\{1,2\}$ .

## Converting Between Cell and Numeric Arrays

Use for loops to convert between cell and numeric formats. For example, create a cell array F:

```
F{1,1} = [1 2; 3 4];  
F{1,2} = [-1 0; 0 1];  
F{2,1} = [7 8; 4 1];  
F{2,2} = [4i 3+2i; 1-8i 5];
```

Now use three for loops to copy the contents of F into a numeric array NUM:

```
for k = 1:4  
    for m = 1:2  
        for n = 1:2  
            NUM(m,n,k) = F{k}(m,n);  
        end  
    end  
end
```

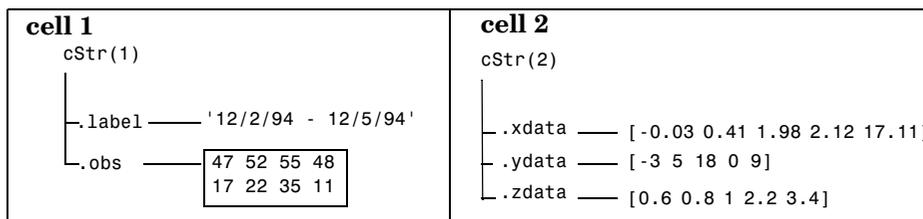
Similarly, you must use for loops to assign each value of a numeric array to a single cell of a cell array:

```
G = cell(1,16);  
for m = 1:16  
    G{m} = NUM(m);  
end
```

## Cell Arrays of Structures

Use cell arrays to store groups of structures with different field architectures:

```
cStr = cell(1,2);  
cStr{1}.label = '12/2/94 - 12/5/94';  
cStr{1}.obs = [47 52 55 48; 17 22 35 11];  
cStr{2}.xdata = [-0.03 0.41 1.98 2.12 17.11];  
cStr{2}.ydata = [-3 5 18 0 9];  
cStr{2}.zdata = [0.6 0.8 1 2.2 3.4];
```



Cell 1 of the `cStr` array contains a structure with two fields, one a string and the other a vector. Cell 2 contains a structure with three vector fields.

When building cell arrays of structures, you must use content indexing. Similarly, you must use content indexing to obtain the contents of structures within cells. The syntax for content indexing is

```
cellArray{index}.field
```

For example, to access the `label` field of the structure in cell 1, use `cStr{1}.label`.

## Function Summary

This table describes the MATLAB functions for working with cell arrays.

Function	Description
<code>cell</code>	Create a cell array.
<code>cell2struct</code>	Convert a cell array into a structure array.
<code>celldisp</code>	Display cell array contents.
<code>cellfun</code>	Apply a cell function to a cell array.
<code>cellplot</code>	Display a graphical depiction of a cell array.
<code>deal</code>	Copy input to separate outputs.
<code>iscell</code>	Return true for a cell array.
<code>num2cell</code>	Convert a numeric array into a cell array.

## Function Handles

A *function handle* is a MATLAB value and data type that provides a means of calling a function indirectly. You can pass function handles in calls to other functions (often called *function functions*). You can also store function handles in data structures for later use (for example, as Handle Graphics® callbacks).

Read more about function handles in the section, “Function Handles” on page 4-27.

### Constructing and Invoking a Function Handle

You construct a handle for a specific function by preceding the function name with an @ sign. Use only the function *name* (with no path information) after the @ sign:

```
fhandle = @functionname
```

### Handles to Anonymous Functions

Another way to construct a function handle is to create an anonymous function. For example,

```
sqr = @(x) x.^2;
```

creates an anonymous function that computes the square of its input argument *x*. The variable *sqr* contains a handle to the anonymous function. See “Anonymous Functions” on page 5-3 for more information.

### Calling a Function Using Its Handle

To execute a function associated with a function handle, use the syntax shown here, treating the function handle *fhandle* as if it were a function name:

```
fhandle(arg1, arg2, ..., argn)
```

If the function being called takes no input arguments, then use empty parentheses after the function handle name:

```
fhandle()
```

## Simple Function Handle Example

The following example calls a function `plotFHandle`, passing it a handle for the MATLAB `sin` function. `plotFHandle` then calls the `plot` function, passing it some data and the function handle to `sin`. The `plot` function calls the function associated with the handle to compute its y-axis values:

```
function x = plotFHandle(fhandle, data)
    plot(data, fhandle(data))
```

Call `plotFHandle` with a handle to the `sin` function and the value shown below:

```
plotFHandle(@sin, -pi:0.01:pi)
```

## MATLAB Classes

All MATLAB data types are implemented as object-oriented classes. You can add data types of your own to your MATLAB environment by creating additional classes. These user-defined classes define the structure of your new data type, and the M-file functions, or *methods*, that you write for each class define the behavior for that data type.

These methods can also define the way various MATLAB operators, including arithmetic operations, subscript referencing, and concatenation, apply to the new data types. For example, a class called `polynomial` might redefine the addition operator (+) so that it correctly performs the operation of addition on polynomials.

With MATLAB classes you can

- Create methods that override existing MATLAB functionality
- Restrict the operations that are allowed on an object of a class
- Enforce common behavior among related classes by inheriting from the same parent class
- Significantly increase the reuse of your code

Read more about MATLAB classes in “Classes and Objects” on page 8-1.

## Java Classes

MATLAB provides an interface to the Java programming language that enables you to create objects from Java classes and call Java methods on these objects. A Java class is a MATLAB data type. Native and third-party classes are already available through the MATLAB interface. You can also create your own Java class definitions and bring them into MATLAB.

The MATLAB Java interface enables you to

- Access Java API (application programming interface) class packages that support essential activities such as I/O and networking
- Access third-party Java classes
- Easily construct Java objects in MATLAB
- Call Java object methods, using either Java or MATLAB syntax
- Pass data between MATLAB variables and Java objects

Read more about Java classes in MATLAB in “Calling Java from MATLAB” in the MATLAB External Interfaces documentation.



# Basic Program Components

---

This chapter introduces some of the principal building blocks used in writing MATLAB programs.

Variables (p. 3-2)	Guidelines for creating variables; global and persistent variables; variable scope and lifetime
Keywords (p. 3-9)	Reserved words that you should avoid using
Special Values (p. 3-10)	Functions that return constant values, like <code>pi</code> or <code>inf</code>
Operators (p. 3-12)	Arithmetic, relational, and logical operators
MATLAB Expressions (p. 3-22)	Executing user-supplied strings; constructing executable strings, shell escape functions
Regular Expressions (p. 3-25)	A versatile way to search and replace character strings
Comma-Separated Lists (p. 3-54)	Using lists with structures and cell arrays to simplify your code
Program Control Statements (p. 3-60)	Using statements such as <code>if</code> , <code>for</code> , and <code>try-catch</code> to control the code path your program follows
Symbol Reference (p. 3-69)	Quick reference to the symbols used by MATLAB
MATLAB Functions (p. 3-81)	Description of the M-file, built-in, and overloaded function types supplied with MATLAB

# Variables

A MATLAB variable is essentially a tag that you assign to a value while that value remains in memory. The tag gives you a way to reference the value in memory so that your programs can read it, operate on it with other data, and save it back to memory. This section covers the following topics on using variables in MATLAB:

- “Types of Variables” on page 3-2
- “Naming Variables” on page 3-6
- “Guidelines to Using Variables” on page 3-7
- “Scope of a Variable” on page 3-7
- “Lifetime of a Variable” on page 3-8

## Types of Variables

MATLAB provides three basic types of variables:

- “Local Variables” on page 3-2
- “Global Variables” on page 3-3
- “Persistent Variables” on page 3-5

### Local Variables

Each MATLAB function has its own local variables. These are separate from those of other functions (except for nested functions), and from those of the base workspace. Variables defined in a function do not remain in memory from one function call to the next, unless they are defined as `global` or `persistent`.

Scripts, on the other hand, do not have a separate workspace. They store their variables in a workspace that is shared with the caller of the script. When called from the command line, they share the base workspace. When called from a function, they share that function’s workspace.

---

**Note** If you run a script that alters a variable that already exists in the caller’s workspace, that variable is overwritten by the script.

---

## Global Variables

If several functions, and possibly the base workspace, all declare a particular name as `global`, then they all share a single copy of that variable. Any assignment to that variable, in any function, is available to all the other functions declaring it `global`.

Suppose, for example, you want to study the effect of the interaction coefficients,  $\alpha$  and  $\beta$ , in the Lotka-Volterra predator-prey model.

$$\begin{aligned} \dot{y}_1 &= y_1 - \alpha y_1 y_2 \\ \dot{y}_2 &= -y_2 + \beta y_1 y_2 \end{aligned}$$

Create an M-file, `lotka.m`.

```
function yp = lotka(t,y)
%LOTKA Lotka-Volterra predator-prey model.
global ALPHA BETA
yp = [y(1) - ALPHA*y(1)*y(2); -y(2) + BETA*y(1)*y(2)];
```

Then interactively enter the statements

```
global ALPHA BETA
ALPHA = 0.01
BETA = 0.02
[t,y] = ode23('lotka',0,10,[1; 1]);
plot(t,y)
```

The two `global` statements make the values assigned to `ALPHA` and `BETA` at the command prompt available inside the function defined by `lotka.m`. They can be modified interactively and new solutions obtained without editing any files.

**Creating Global Variables.** Each function that uses a global variable must first declare the variable as `global`. It is usually best to put global declarations toward the beginning of the function. You would declare global variable `MAXLEN` as follows:

```
global MAXLEN
```

If the M-file contains subfunctions as well, then each subfunction requiring access to the global variable must declare it as `global`. To access the variable from the MATLAB command line, you must declare it as `global` at the command line.

MATLAB global variable names are typically longer and more descriptive than local variable names, and often consist of all uppercase characters. These are not requirements, but guidelines to increase the readability of MATLAB code, and to reduce the chance of accidentally redefining a global variable.

**Displaying Global Variables.** To see only those variables you have declared as global, use the `who` or `whos` functions with the literal, `global`.

```
global MAXLEN MAXWID
MAXLEN = 36; MAXWID = 78;
len = 5; wid = 21;
```

```
whos global
  Name          Size          Bytes  Class
  -----
  MAXLEN        1x1          8  double array (global)
  MAXWID        1x1          8  double array (global)
```

Grand total is 2 elements using 16 bytes

**Suggestions for Using Global Variables.** A certain amount of risk is associated with using global variables and, because of this, it is recommended that you use them sparingly. You might, for example, unintentionally give a global variable in one function a name that is already used for a global variable in another function. When you run your application, one function may overwrite the variable used by the other. This error can be difficult to track down.

Another problem comes when you want to change the variable name. To make a change without introducing an error into the application, you must find every occurrence of that name in your code (and other people's code, if you share functions).

**Alternatives to Using Global Variables.** Instead of using a global variable, you may be able to

- Pass the variable to other functions as an additional argument. In this way, you make sure that any shared access to the variable is intentional.

If this means that you have to pass a number of additional variables, you can put them into a structure or cell array and just pass it as one additional argument.

- Use a persistent variable (described in the next section), if you only need to make the variable persist in memory from one function call to the next.

### Persistent Variables

Characteristics of persistent variables are

- You can declare and use them within M-file functions only.
- Only the function in which the variables are declared is allowed access to it.
- MATLAB does not clear them from memory when the function exits, so their value is retained from one function call to the next.

You must declare persistent variables before you can use them in a function. It is usually best to put your persistent declarations toward the beginning of the function. You would declare persistent variable `SUM_X` as follows:

```
persistent SUM_X
```

If you clear a function that defines a persistent variable (i.e., using `clear functionname` or `clear all`), or if you edit the M-file for that function, MATLAB clears all persistent variables used in that function.

You can use the `mlock` function to keep an M-file from being cleared from memory, thus keeping persistent variables in the M-file from being cleared as well.

**Initializing Persistent Variables.** When you declare a persistent variable, MATLAB initializes its value to an empty matrix, `[]`. After the declaration statement, you can assign your own value to it. This is often done using an `isempty` statement, as shown here:

```
function findSum(inputvalue)
persistent SUM_X

if isempty(SUM_X)
    SUM_X = 0;
end

SUM_X = SUM_X + inputvalue
```

This initializes the variable to 0 the first time you execute the function, and then accumulates the value on each iteration.

## Naming Variables

MATLAB variable names must begin with a letter, which may be followed by any combination of letters, digits, and underscores. MATLAB distinguishes between uppercase and lowercase characters, so A and a are not the same variable.

Although variable names can be of any length, MATLAB uses only the first N characters of the name, (where N is the number returned by the function `namelengthmax`), and ignores the rest. Hence, it is important to make each variable name unique in the first N characters to enable MATLAB to distinguish variables.

```
N = namelengthmax
N =
    63
```

The `genvarname` function can be useful in creating variable names that are both valid and unique.

## Verifying a Variable Name

You can use the `isvarname` function to make sure a name is valid before you use it. `isvarname` returns 1 if the name is valid, and 0 otherwise.

```
isvarname 8th_column
ans =
    0                                % Not valid - begins with a number
```

## Avoid Using Function Names for Variables

When naming a variable, make sure you are not using a name that is already used as a function name. If you define a variable with a function name, you won't be able to call that function until you either remove the variable from memory with the `clear` function, or invoke the function using `builtin`.

For example, if you enter the following command, you will not be able to use the MATLAB `disp` function until you clear the variable with `clear disp`.

```
disp = 50;
```

To test whether a proposed variable name is already used as a function name, use

```
which -all name
```

## Guidelines to Using Variables

The same guidelines that apply to MATLAB variables at the command line also apply to variables in M-files:

- You do not need to type or declare variables used in M-files (with the possible exception of designating them as `global` or `persistent`).
- Before assigning one variable to another, you must be sure that the variable on the right-hand side of the assignment has a value.
- Any operation that assigns a value to a variable creates the variable, if needed, or overwrites its current value, if it already exists.

## Scope of a Variable

MATLAB stores variables in a part of memory called a workspace. The *base workspace* holds variables created during your interactive MATLAB session and also any variables created by running M-file scripts. Variables created at the MATLAB command prompt can also be used by scripts without having to declare them as `global`.

Functions do not use the base workspace. Every function has its own *function workspace*. Each function workspace is kept separate from the base workspace and all other workspaces to protect the integrity of the data used by that function. Even subfunctions that are defined in the same M-file have a separate function workspace.

## Extending Variable Scope

In most cases, variables created within a function are known only within that function. These variables are not available at the MATLAB command prompt or to any other function or subfunction. The most secure way to extend the scope of a function variable is to pass it to other functions as an argument in the function call. Since MATLAB passes data only by value, you also need to add the variable to the return values of any function that modifies its value.

Another way to extend the variable scope is to declare the variable as `global` within every function that needs access to it. If you do this, you need make sure that no functions with access to the variable overwrite its value unintentionally.

### **Scope in Nested Functions**

Variables within nested functions are accessible to more than just their immediate function. As a general rule, the scope of a local variable is the largest containing function body in which the variable appears, and all functions nested within that function. For more information on nested functions, see “Variable Scope in Nested Functions” on page 5-18.

### **Lifetime of a Variable**

Variables created at the MATLAB command prompt or in an M-file script exist until you clear them or end your MATLAB session. Variables in functions exist only until the function completes unless they have been declared as `global` or `persistent`.

## Keywords

MATLAB reserves certain words for its own use as keywords of the language. To list the keywords, type

```
iskeyword
ans =
    'break'
    'case'
    'catch'
    'continue'
    'else'
    'elseif'
    'end'
    'for'
    'function'
    'global'
    'if'
    'otherwise'
    'persistent'
    'return'
    'switch'
    'try'
    'while'
```

See the online function reference pages to learn how to use these keywords.

You should not use MATLAB keywords other than for their intended purpose. For example, a keyword should not be used as follows:

```
while = 5;
??? while = 5;
    |
Error: Expected a variable, function, or constant, found "=".
```

## Special Values

Several functions return important special values that you can use in your M-files.

Function	Return Value
ans	Most recent answer (variable). If you do not assign an output variable to an expression, MATLAB automatically stores the result in ans.
eps	Floating-point relative accuracy. This is the tolerance MATLAB uses in its calculations.
intmax	Largest 8-, 16-, 32-, or 64-bit integer your computer can represent.
intmin	Smallest 8-, 16-, 32-, or 64-bit integer your computer can represent.
realmax	Largest floating-point number your computer can represent.
realmin	Smallest positive floating-point number your computer can represent.
pi	3.1415926535897...
i, j	Imaginary unit.
inf	Infinity. Calculations like $n/0$ , where $n$ is any nonzero real value, result in inf.
NaN	Not a Number, an invalid numeric value. Expressions like $0/0$ and $inf/inf$ result in a NaN, as do arithmetic operations involving a NaN. Also, if $n$ is complex with a zero real part, then $n/0$ returns a value with a NaN real part.
computer	Computer type.
version	MATLAB version string.

Here are some examples that use these values in MATLAB expressions.

```
x = 2 * pi
```

```
x =
```

```
6.2832
```

```
A = [3+2i 7-8i]
```

```
A =
```

```
3.0000 + 2.0000i    7.0000 - 8.0000i
```

```
tol = 3 * eps
```

```
tol =
```

```
6.6613e-016
```

```
intmax('uint64')
```

```
ans =
```

```
18446744073709551615
```

## Operators

The MATLAB operators fall into three categories:

- “Arithmetic Operators” on page 3-12 perform numeric computations, for example, adding two numbers or raising the elements of an array to a given power.
- “Relational Operators” on page 3-13 compare operands quantitatively, using operators like “less than” and “not equal to.”
- “Logical Operators” on page 3-15 use the logical operators AND, OR, and NOT.

This section also discusses “Operator Precedence” on page 3-20.

### Arithmetic Operators

MATLAB provides these arithmetic operators.

<b>Operator</b>	<b>Description</b>
+	Addition
-	Subtraction
.*	Multiplication
./	Right division
.\	Left division
+	Unary plus
-	Unary minus
:	Colon operator
.^	Power
.'	Transpose
'	Complex conjugate transpose
*	Matrix multiplication

Operator	Description
/	Matrix right division
\	Matrix left division
^	Matrix power

### Arithmetic Operators and Arrays

Except for some matrix operators, MATLAB arithmetic operators work on corresponding elements of arrays with equal dimensions. For vectors and rectangular arrays, both operands must be the same size unless one is a scalar. If one operand is a scalar and the other is not, MATLAB applies the scalar to every element of the other operand — this property is known as *scalar expansion*.

This example uses scalar expansion to compute the product of a scalar operand and a matrix.

```
A = magic(3)
A =
     8     1     6
     3     5     7
     4     9     2

3 * A
ans =
    24     3    18
     9    15    21
    12    27     6
```

### Relational Operators

MATLAB provides these relational operators.

Operator	Description
<	Less than
<=	Less than or equal to

<b>Operator</b>	<b>Description</b>
>	Greater than
>=	Greater than or equal to
==	Equal to
~=	Not equal to

### Relational Operators and Arrays

The MATLAB relational operators compare corresponding elements of arrays with equal dimensions. Relational operators always operate element-by-element. In this example, the resulting matrix shows where an element of A is equal to the corresponding element of B.

```
A = [2 7 6;9 0 5;3 0.5 6];  
B = [8 7 0;3 2 5;4 -1 7];
```

```
A == B  
ans =  
     0     1     0  
     0     0     1  
     0     0     0
```

For vectors and rectangular arrays, both operands must be the same size unless one is a scalar. For the case where one operand is a scalar and the other is not, MATLAB tests the scalar against every element of the other operand. Locations where the specified relation is true receive logical 1. Locations where the relation is false receive logical 0.

### Relational Operators and Empty Arrays

The relational operators work with arrays for which any dimension has size zero, as long as both arrays are the same size or one is a scalar. However, expressions such as

```
A == []
```

return an error if A is not 0-by-0 or 1-by-1. This behavior is consistent with that of all other binary operators, such as +, -, >, <, &, |, etc.

To test for empty arrays, use the function

```
isempty(A)
```

## Logical Operators

MATLAB offers three types of logical operators and functions:

- Element-wise — operate on corresponding elements of logical arrays.
- Bit-wise — operate on corresponding bits of integer values or arrays.
- Short-circuit — operate on scalar, logical expressions.

The values returned by MATLAB logical operators and functions, with the exception of bit-wise functions, are of type `logical` and are suitable for use with logical indexing.

### Element-Wise Operators and Functions

The following logical operators and functions perform element-wise logical operations on their inputs to produce a like-sized output array. The examples shown in the following table use vector inputs A and B, where

```
A = [0 1 1 0 1];
B = [1 1 0 0 1];
```

Operator	Description	Example
&	Returns 1 for every element location that is true (nonzero) in both arrays, and 0 for all other elements.	A & B = 01001
	Returns 1 for every element location that is true (nonzero) in either one or the other, or both arrays, and 0 for all other elements.	A   B = 11101
~	Complements each element of the input array, A.	~A = 10010
xor	Returns 1 for every element location that is true (nonzero) in only one array, and 0 for all other elements.	xor(A,B)=10100

For operators and functions that take two array operands, (&, |, and xor), both arrays must have equal dimensions, with each dimension being the same size. The one exception to this is where one operand is a scalar and the other is not. In this case, MATLAB tests the scalar against every element of the other operand.

---

**Note** MATLAB converts any finite nonzero, numeric values used as inputs to logical expressions to logical 1, or true.

---

**Operator Overloading.** You can overload the &, |, and ~ operators to make their behavior dependent upon the data type on which they are being used. Each of these operators has a representative function that is called whenever that operator is used. These are shown in the table below.

Logical Operation	Equivalent Function
A & B	and(A, B)
A   B	or(A, B)
~A	not(A)

**Other Array Functions.** Two other MATLAB functions that operate logically on arrays, but not in an element-wise fashion, are any and all. These functions show whether *any* or *all* elements of a vector, or a vector within a matrix or an array, are nonzero.

When used on a matrix, any and all operate on the columns of the matrix. When used on an N-dimensional array, they operate on the first nonsingleton dimension of the array. Or, you can specify an additional dimension input to operate on a specific dimension of the array.

The examples shown in the following table use array input A, where

```
A = [0  1  2;  
     0 -3  8;  
     0  5  0];
```

Function	Description	Example
<code>any(A)</code>	Returns 1 for a vector where <i>any</i> element of the vector is true (nonzero), and 0 if no elements are true.	<code>any(A)</code> <code>ans =</code> 0 1 1
<code>all(A)</code>	Returns 1 for a vector where <i>all</i> elements of the vector are true (nonzero), and 0 if all elements are not true.	<code>all(A)</code> <code>ans =</code> 0 1 0

---

**Note** The `all` and `any` functions ignore any NaN values in the input arrays.

---

**Logical Expressions Using the `find` Function.** The `find` function determines the indices of array elements that meet a given logical condition. The function is useful for creating masks and index matrices. In its most general form, `find` returns a single vector of indices. This vector can be used to index into arrays of any size or shape.

For example,

```
A = magic(4)
A =
    16     2     3    13
     5    11    10     8
     9     7     6    12
     4    14    15     1

i = find(A > 8);
A(i) = 100
A =
    100     2     3    100
     5    100    100     8
    100     7     6    100
     4    100    100     1
```

---

**Note** An alternative to using `find` in this context is to index into the matrix using the logical expression itself. See the example below.

---

The last two statements of the previous example can be replaced with this one statement:

```
A(A > 8) = 100;
```

You can also use `find` to obtain both the row and column indices of a rectangular matrix for the array values that meet the logical condition:

```
A = magic(4)
A =
    16     2     3    13
     5    11    10     8
     9     7     6    12
     4    14    15     1

[row, col] = find(A > 12)
row =
     1
     4
     4
     1
col =
     1
     2
     3
     4
```

#### Bit-Wise Functions

The following functions perform bit-wise logical operations on nonnegative integer inputs. Inputs may be scalar or in arrays. If in arrays, these functions produce a like-sized output array.

The examples shown in the following table use scalar inputs A and B, where

```
A = 28;           % binary 11100
B = 21;           % binary 10101
```

Function	Description	Example
<code>bitand</code>	Returns the bit-wise AND of two nonnegative integer arguments.	<code>bitand(A,B) = 20</code> (binary 10100)
<code>bitor</code>	Returns the bit-wise OR of two nonnegative integer arguments.	<code>bitor(A,B) = 29</code> (binary 11101)
<code>bitcmp</code>	Returns the bit-wise complement as an n-bit number, where n is the second input argument to <code>bitcmp</code> .	<code>bitcmp(A,5) = 3</code> (binary 00011)
<code>bitxor</code>	Returns the bit-wise exclusive OR of two nonnegative integer arguments.	<code>bitxor(A,B) = 9</code> (binary 01001)

### Short-Circuit Operators

The following operators perform AND and OR operations on logical expressions containing scalar values. They are *short-circuit* operators in that they evaluate their second operand only when the result is not fully determined by the first operand.

Operator	Description
<code>&amp;&amp;</code>	Returns logical 1 (true) if both inputs evaluate to true, and logical 0 (false) if they do not.
<code>  </code>	Returns logical 1 (true) if either input, or both, evaluate to true, and logical 0 (false) if they do not.

The statement shown here performs an AND of two logical terms, A and B:

```
A && B
```

If A equals zero, then the entire expression will evaluate to logical 0 (false), regardless of the value of B. Under these circumstances, there is no need to evaluate B because the result is already known. In this case, MATLAB short-circuits the statement by evaluating only the first term.

A similar case is when you OR two terms and the first term is true. Again, regardless of the value of B, the statement will evaluate to true. There is no need to evaluate the second term, and MATLAB does not do so.

**Advantage of Short-Circuiting.** You can use the short-circuit operators to evaluate an expression only when certain conditions are satisfied. For example, you want to execute an M-file function only if the M-file resides on the current MATLAB path.

Short-circuiting keeps the following code from generating an error when the file, `myfun.m`, cannot be found:

```
comp = (exist('myfun.m') == 2) && (myfun(x) >= y)
```

Similarly, this statement avoids divide-by-zero errors when `b` equals zero:

```
x = (b ~= 0) && (a/b > 18.5)
```

You can also use the `&&` and `||` operators in `if` and `while` statements to take advantage of their short-circuiting behavior:

```
if (nargin >= 3) && (ischar(varargin{3}))
```

## Operator Precedence

You can build expressions that use any combination of arithmetic, relational, and logical operators. Precedence levels determine the order in which MATLAB evaluates an expression. Within each precedence level, operators have equal precedence and are evaluated from left to right. The precedence rules for MATLAB operators are shown in this list, ordered from highest precedence level to lowest precedence level:

- 1 Parentheses (`()`)
- 2 Transpose (`'`), power (`^`), complex conjugate transpose (`'`), matrix power (`^`)
- 3 Unary plus (`+`), unary minus (`-`), logical negation (`~`)
- 4 Multiplication (`*`), right division (`/`), left division (`\`), matrix multiplication (`*`), matrix right division (`/`), matrix left division (`\`)
- 5 Addition (`+`), subtraction (`-`)
- 6 Colon operator (`:`)

- 7 Less than (<), less than or equal to (<=), greater than (>), greater than or equal to (>=), equal to (==), not equal to (~=)
- 8 Element-wise AND (&)
- 9 Element-wise OR (|)
- 10 Short-circuit AND (&&)
- 11 Short-circuit OR (||)

### Precedence of AND and OR Operators

MATLAB always gives the & operator precedence over the | operator. Although MATLAB typically evaluates expressions from left to right, the expression  $a|b\&c$  is evaluated as  $a|(b\&c)$ . It is a good idea to use parentheses to explicitly specify the intended precedence of statements containing combinations of & and |.

The same precedence rule holds true for the && and || operators.

### Overriding Default Precedence

The default precedence can be overridden using parentheses, as shown in this example:

```
A = [3 9 5];
B = [2 1 5];
C = A./B.^2
C =
    0.7500    9.0000    0.2000

C = (A./B).^2
C =
    2.2500   81.0000    1.0000
```

## MATLAB Expressions

Two specific types of MATLAB expressions covered in this section are

- “String Evaluation” on page 3-22
- “Shell Escape Functions” on page 3-23

### String Evaluation

String evaluation adds power and flexibility to the MATLAB language, letting you perform operations like executing user-supplied strings and constructing executable strings through concatenation of strings stored in variables.

#### eval

The `eval` function evaluates a string that contains a MATLAB expression, statement, or function call. In its simplest form, the `eval` syntax is

```
eval('string')
```

For example, this code uses `eval` on an expression to generate a Hilbert matrix of order `n`.

```
t = '1/(m + n - 1)';  
for m = 1:k  
    for n = 1:k  
        a(m,n) = eval(t);  
    end  
end
```

Here is an example that uses `eval` on a statement.

```
eval('t = clock');
```

**Constructing Strings for Evaluation.** You can concatenate strings to create a complete expression for input to `eval`. This code shows how `eval` can create 10 variables named `P1`, `P2`, ..., `P10`, and set each of them to a different value.

```
for n = 1:10  
    eval(['P', int2str(n), '= n .^ 2'])  
end
```

## feval

The `feval` function differs from `eval` in that it executes a function rather than a MATLAB expression. The function to be executed is specified in the first argument by either a function handle or a string containing the function name.

You can use `feval` and the `input` function to choose one of several tasks defined by M-files. This example uses function handles for the `sin`, `cos`, and `log` functions.

```
fun = [@sin; @cos; @log];
k = input('Choose function number: ');
x = input('Enter value: ');
feval(fun(k), x)
```

## Shell Escape Functions

It is sometimes useful to access your own C or Fortran programs using *shell escape functions*. Shell escape functions use the shell escape command `!` to make external stand-alone programs act like new MATLAB functions. A shell escape M-function is an M-file that

- 1 Saves the appropriate variables on disk.
- 2 Runs an external program (which reads the data file, processes the data, and writes the results back out to disk).
- 3 Loads the processed file back into the workspace.

For example, look at the code for `garfield.m`, below. This function uses an external function, `gareqn`, to find the solution to Garfield's equation.

```
function y = garfield(a,b,q,r)
save gardata a b q r
!gareqn
load gardata
```

### This M-file

- 1** Saves the input arguments `a`, `b`, `q`, and `r` to a MAT-file in the workspace using the `save` command.
- 2** Uses the shell escape operator to access a C or Fortran program called `gareqn` that uses the workspace variables to perform its computation. `gareqn` writes its results to the `gardata` MAT-file.
- 3** Loads the `gardata` MAT-file to obtain the results.

## Regular Expressions

A regular expression is a string of characters that defines a certain pattern. You would normally use a regular expression in searching through text for a group of words that matches this pattern, perhaps while parsing program input, or while processing a block of text.

The string `'Joh?n\w*'` is an example of a regular expression. It defines a pattern that starts with the letters `Jo`, is optionally followed by the letter `h` (indicated by `'h?'`), is then followed by the letter `n`, and ends with any number of non-whitespace characters (indicated by `'\w*'`). This pattern matches any of the following:

Jon, John, Jonathan, Johnny

MATLAB supports most of the special characters, or *metacharacters*, commonly used with regular expressions and provides several functions to use in searching and replacing text with these expressions.

This section discusses the following topics:

- “MATLAB Regular Expression Functions” on page 3-25
- “Elements of an Expression” on page 3-26
- “Tokens” on page 3-40
- “Handling Multiple Strings” on page 3-46
- “Operator Summary” on page 3-50

### MATLAB Regular Expression Functions

Several MATLAB functions support searching and replacing characters using regular expressions:

Function	Description
<code>regexp</code>	Match regular expression.
<code>regexpi</code>	Match regular expression, ignoring case.
<code>regexprep</code>	Replace string using regular expression.

See the function reference pages to obtain more information on these functions. For more information on how to use regular expressions in general, consult a reference on that subject.

## Elements of an Expression

The following tables show the syntax supported by the `regexp`, `regextpi`, and `regextprep` functions. Expressions shown in the left column have special meaning and match one or more characters according to the usage described in the right column. Any character not having a special meaning, (e.g., any alphabetic character) matches that same character literally.

These elements are presented under these categories:

- “Character Classes” on page 3-26
- “Character Representation” on page 3-29
- “Logical Operators” on page 3-30
- “Lookaround Operators” on page 3-33
- “Quantifiers” on page 3-36
- “Tokens” on page 3-40

Each table is followed by a set of examples that show how to use the syntax presented in that table.

## Character Classes

Character classes represent either a specific set of characters (e.g., uppercase) or a certain type of character (e.g., non-white-space).

Operator	Usage
<code>.</code>	Any single character, including white space
<code>[c<sub>1</sub>c<sub>2</sub>c<sub>3</sub>]</code>	Any character contained within the brackets: <code>c<sub>1</sub></code> or <code>c<sub>2</sub></code> or <code>c<sub>3</sub></code>
<code>[^c<sub>1</sub>c<sub>2</sub>c<sub>3</sub>]</code>	Any character not contained within the brackets: anything but <code>c<sub>1</sub></code> or <code>c<sub>2</sub></code> or <code>c<sub>3</sub></code>
<code>[c<sub>1</sub>-c<sub>2</sub>]</code>	Any character in the range of <code>c<sub>1</sub></code> through <code>c<sub>2</sub></code>
<code>\s</code>	Any white-space character; equivalent to <code>[ \f\n\r\t\v]</code>

Operator	Usage
<code>\S</code>	Any non-white-space character; equivalent to <code>[^\f\n\r\t\v]</code>
<code>\w</code>	Any alphabetic, numeric, or underscore character; equivalent to <code>[a-zA-Z_0-9]</code>
<code>\W</code>	Any character that is not alphabetic, numeric, or underscore; equivalent to <code>[^a-zA-Z_0-9]</code>
<code>\d</code>	Any numeric digit; equivalent to <code>[0-9]</code>
<code>\D</code>	Any nondigit character; equivalent to <code>[^0-9]</code>

These examples demonstrate how to use the character classes listed above. See the `regex` reference page for help with syntax.

Most of these examples use the following string:

```
str = 'The rain in Spain falls mainly on the plain.'
```

### Any Character — `.`

Use `..ain` in an expression to match a sequence of five characters ending in `ain`. Note that `.` matches white-space characters as well:

```
regexp(str, '..ain')
ans =
     4     13     24     39
```

Matches `' rain'`, `'Spain'`, `' main'`, and `'plain'`.

**Returning Strings Rather than Indices.** Here is the same example, this time specifying the command qualifier `'match'`. In this case, `regexp` returns the *text* of the matching strings rather than the starting index:

```
regexp(str, '..ain', 'match')
ans =
    ' rain'    'Spain'    ' main'    'plain'
```

**Selected Characters — [c<sub>1</sub>c<sub>2</sub>c<sub>3</sub>]**

Use [c<sub>1</sub>c<sub>2</sub>c<sub>3</sub>] in an expression to match selected characters r, p, or m followed by 'ain'. Specify two qualifiers this time, 'match' and 'start', along with an output argument for each, mat and idx. This returns the matching strings and the starting indices of those strings:

```
[mat idx] = regexp(str, '[rpm]ain', 'match', 'start')
mat =
    'rain'    'pain'    'main'
idx =
     5     14     25
```

**Range of Characters — [c<sub>1</sub> - c<sub>2</sub>]**

Use [c<sub>1</sub>-c<sub>2</sub>] in an expression to find words that begin with a letter in the range of A through Z:

```
[mat idx] = regexp(str, '[A-Z]\w*', 'match', 'start')
mat =
    'The'    'Spain'
idx =
     1     13
```

**Word and White-Space Characters — \w, \s**

Use \w and \s in an expression to find words that end with the letter n followed by a white-space character. Add a new qualifier, 'end', to return the str index that marks the end of each match:

```
[mat ix1 ix2] = regexp(str, '\w*n\s', 'match', 'start', 'end')
mat =
    'rain '    'in '    'Spain '    'on '
ix1 =
     5     10     13     32
ix2 =
     9     12     18     34
```

## Numeric Digits — \d

Use \d to find numeric digits in the following string:

```
numstr = 'Easy as 1, 2, 3';

[mat idx] = regexp(numstr, '\d', 'match', 'start')
mat =
    '1'    '2'    '3'
idx =
    9     12     15
```

## Character Representation

The following character combinations represent specific character and numeric values.

Operator	Usage
\a	Alarm (beep)
\b	Backspace
\e	Escape
\f	Form feed
\n	New line
\r	Carriage return
\t	Horizontal tab
\v	Vertical tab
\oN or \o{N}	Character of octal value N
\xN or \x{N}	Character of hexadecimal value N
\char	If a character has special meaning in a regular expression, precede it with backslash (\) to match it literally.

### Octal and Hexadecimal — \o, \x

Use \x and \o in an expression to find a comma (hex 2C) followed by a space (octal 40) followed by the character 2:

```
numstr = 'Easy as 1, 2, 3';

[mat idx] = regexp(numstr, '\x2C{o{40}2', 'match', 'start')
mat =
    ', 2'
idx =
    10
```

### Special Characters — \char

Use \ before a character that has a special meaning to the regular expression functions if you want that character to be interpreted literally. The intention in this example is to have the string '(ab[XY|Z]c)' interpreted literally. The first expression does not do that because regexp interprets the parentheses and | sign as the special characters for grouping and logical OR:

```
regexp('(ab[XY|Z]c)', '(ab[XY|Z]c)', 'match')
ans =
    'ab[XY'    'Z]c'
```

This next expression uses a \ before any special characters. As a result the entire string is matched:

```
regexp('(ab[XY|Z]c)', '\\(ab\\[XY\\|Z\\]c\\)', 'match')
ans =
    '(ab[XY|Z]c)'
```

### Logical Operators

Logical operators do not match any specific characters. They are used to specify the context for matching an accompanying regular expression.

Operator	Usage
(expr)	Group regular expressions and capture tokens.
(?:expr)	Group regular expressions, but do not capture tokens.

Operator	Usage
(?>expr)	Group atomically.
(?#expr)	Insert a comment into the expression. Comments are ignored in matching.
expr <sub>1</sub>  expr <sub>2</sub>	Match expression expr <sub>1</sub> or expression expr <sub>2</sub> .
^expr	Match the expression only at the beginning of the string.
expr\$	Match the expression only at the end of the string.
\<expr	Match the characters when they start a word.
expr\>	Match the characters when they end a word.
\<expr\>	Match an exact word.

### Grouping and Capture — (expr)

You can group elements together using either (expr) to group and capture or (?:expr) for grouping alone. For an example of the former, see “Using Tokens — Example 1” on page 3-43. For the latter, see the “Grouping-Only” example below.

### Grouping Only — (?:expr)

Use (?:expr) to group a consonant followed by a vowel in the palindrome pstr. Specify at least two consecutive occurrences ({2,}) of this group. Return the starting and ending indices of the matched substrings:

```
pstr = 'Marge lets Norah see Sharon''s telegram';
expr = '(?:[^\aeiou][aeiou]){2,}';

[mat ix1 ix2] = regexp(pstr, expr, 'match', 'start', 'end')
mat =
    'Nora'    'haro'    'tele'
ix1 =
    12     23     31
ix2 =
    15     26     34
```

Remove the grouping, and the {2,} now applies only to [aeiou]. The command is entirely different now as it looks for a consonant followed by at least two consecutive vowels:

```
expr = '[^aeiou][aeiou]{2,}';

[mat ix1 ix2] = regexp(pstr, expr, 'match', 'start', 'end')
mat =
    'see'
ix1 =
    18
ix2 =
    20
```

### Including Comments — (?#expr)

Use (?#expr) to add a comment to this expression that matches capitalized words in pstr. Comments are ignored in the process of finding a match:

```
regexp(pstr, '(?# Match words in caps)[A-Z]\w+', 'match')
ans =
    'Marge'    'Norah'    'Sharon'
```

### Alternative Match — expr<sub>1</sub>|expr<sub>2</sub>

Use p<sub>1</sub>|p<sub>2</sub> to pick out words in the string that start with let or tel:

```
regexpi(pstr, '(let|tel)\w+', 'match')
ans =
    'lets'    'telegram'
```

---

**Note** Be careful about using the '|' operator within square brackets in MATLAB regular expressions. Because the '|' operator has a higher precedence than '[ ]', MATLAB interprets the expression '[A|B]' not as 'A' or 'B', but as '[A' or 'B]'. The recommended way to match "the letter A or the letter B" in a MATLAB regexp expression is to use '[AB]'.

---

### Start and End of String Match — `^expr, expr$`

Use `^expr` to match words starting with the letter `m` or `M` only when it begins the string, and `expr$` to match words ending with `m` or `M` only when it ends the string:

```
regexpi(pstr, '^m\w*|\w*m$', 'match')
ans =
  'Marge'      'telegram'
```

### Start and End of Word Match — `\<expr, expr\>`

Use `\<expr` to match any words starting with `n` or `N`, or ending with `e` or `E`:

```
regexpi(pstr, '\<n\w*|\w*e\>', 'match')
ans =
  'Marge'      'Norah'      'see'
```

### Exact Word Match — `\<expr\>`

Use `\<expr\>` to match a word starting with an `n` or `N` and ending with an `h` or `H`:

```
regexpi(pstr, '\<n\w*h\>', 'match')
ans =
  'Norah'
```

## Lookaround Operators

Lookaround operators have two components: a match pattern and a test pattern. If you call the match pattern  $p_1$  and the test pattern  $p_2$ , then the simplest form of lookaround operator looks like this:

$$p_1(?:p_2)$$

The match pattern  $p_1$  is just like any other element in an expression. For example, it can be `\<[A-Za-z]+\>` to make `regexp` find any word.

The test pattern  $p_2$  places a condition on this match. There can be a match for  $p_1$  only if there is also a match for  $p_2$ , and the  $p_2$  match must immediately precede (for *lookbehind* operators) or follow (for *lookahead* operators) the match for  $p_1$ .

In the following expression, the match pattern is '`\<[A-Za-z]+\>`' and the test pattern is '`\S`'. The entire expression can be read as “Find those words that are followed by a non-white-space character”:

```
'\<[A-Za-z]+\>(?\S)'
```

When used on the following string, this lookahead expression matches the letters of the words Raven and Nevermore:

```
str = 'Quoth the Raven, "Nevermore"';  
  
regexp(str, '\<[A-Za-z]+\>(?\S)', 'match')  
ans =  
    'Raven'    'Nevermore'
```

One important characteristic of lookaround operators is how they affect the parsing of the input string. The parser can be said to “consume” pieces of the string as it looks for matching phrases. With lookaround operators, only the match pattern  $p_1$  affects the current parsing location. Finding a match for the test pattern  $p_2$  does not move the parser location.

---

**Note** You can also use lookaround operators to perform a logical AND of two elements. See “Using Lookaround as a Logical Operator” on page 3-36.

---

This table shows the four lookaround expressions: lookahead, negative lookahead, lookbehind, and negative lookbehind.

Operator	Usage
<code>expr<sub>1</sub>(?=expr<sub>2</sub>)</code>	Match expression <code>expr<sub>1</sub></code> if followed by expression <code>expr<sub>2</sub></code> .
<code>expr<sub>1</sub>(?!expr<sub>2</sub>)</code>	Match expression <code>expr<sub>1</sub></code> if not followed by expression <code>expr<sub>2</sub></code> .
<code>(?&lt;=expr<sub>1</sub>)expr<sub>2</sub></code>	Match expression <code>expr<sub>2</sub></code> if preceded by expression <code>expr<sub>1</sub></code> .
<code>(?&lt;!expr<sub>1</sub>)expr<sub>2</sub></code>	Match expression <code>expr<sub>2</sub></code> if not preceded by expression <code>expr<sub>1</sub></code> .

**Lookahead —  $\text{expr}_1(=\text{expr}_2)$** 

Use  $p_1(=p_2)$  to find all words of this string that precede a comma:

```
poestr = ['While I nodded, nearly napping, ' ...
          'suddenly there came a tapping,'];

[mat idx] = regexp(poestr, '\w*(?=,)', 'match', 'start')
mat =
    'nodded'      'napping'      'tapping'
idx =
    9    24    55
```

**Negative Lookahead —  $\text{expr}_1(!\text{expr}_2)$** 

Use  $p_1(!p_2)$  to find all words that do not precede a comma:

```
[mat idx] = regexp(poestr, '\w+(?!\w*,)', 'match', 'start')
mat =
    'While' 'I' 'nearly' 'suddenly' 'there' 'came' 'a'
idx =
    1    7    17    33    42    48    53
```

**Lookbehind —  $(?<=\text{expr}_1)\text{expr}_2$** 

Use  $(?<=p_1)p_2$  to find all words that follow a comma and zero or more spaces:

```
[mat idx] = regexp(poestr, '(?<=,\s*)\w*', 'match', 'start')
mat =
    'nearly'      'suddenly'
idx =
    17    33
```

**Negative Lookbehind —  $(?<!\text{expr}_1)\text{expr}_2$** 

Use  $(?<!p_1)p_2$  to find all words that do not follow a comma and zero or more spaces:

```
[mat idx] = regexp(poestr, '(?<!\s*\w*)\w*', 'match', 'start')
mat =
    'While' 'I' 'nodded' 'napping' 'there' 'came' 'a' 'tapping'
idx =
    1    7    9    24    42    48    53    55
```

### Using Lookaround as a Logical Operator

You can use lookaround operators to perform a logical AND, as shown in this example. The expression used here finds all words that contain a sequence of two letters under the condition that the two letters are identical *and* are in the range a through m. (The expression '(?=[a-m])' is a lookahead test for the range a through m, and the expression '(.)\1' tests for identical characters using a token):

```
[mat idx] = regexp(poestr, '\<\w*(?=[a-m])\1\w*\>', ...
    'match', 'start')
mat =
    'noddod'      'suddenly'
idx =
    9      33
```

Note that when using a lookahead operator to perform an AND, you need to place the match expression `expr1` *after* the test expression `expr2`:

`(?=expr2)expr1`            or            `(?!expr2)expr1`

### Quantifiers

You can use quantifiers to specify how many instances of an element are to be matched. The first six rows of this table show the basic quantifiers. When used alone, they match as much of the string as possible. Thus these are sometimes called *greedy* quantifiers.

When one of these quantifiers is followed by a plus sign (e.g., '`\w*+`'), it is known as a *possessive* quantifier. Possessive quantifiers also match as much of the string as possible, but they do not rescan any portions of the string should the initial match fail.

When you follow a quantifier with a question mark (e.g., '`\w*?`'), it is known as a *lazy* quantifier. Lazy quantifiers match as little of the string as possible.

See the examples for each quantifier and quantifier type following the table.

Operator	Usage
<code>expr?</code>	Match the preceding element 0 times or 1 time. Equivalent to <code>{0, 1}</code> .
<code>expr*</code>	Match the preceding element 0 or more times. Equivalent to <code>{0, }</code> .
<code>expr+</code>	Match the preceding element 1 or more times. Equivalent to <code>{1, }</code> .
<code>expr{n}</code>	Must match exactly <code>n</code> times. Equivalent to <code>{n, n}</code> .
<code>expr{n, }</code>	Must occur at least <code>n</code> times.
<code>expr{n, m}</code>	Must occur at least <code>n</code> times but no more than <code>m</code> times.
<code>qu_expr?</code>	Match the quantified expression according to the guidelines stated above for lazy quantifiers, where <code>qu_expr</code> represents any one of the expressions shown in the top six rows of this table.
<code>qu_expr+</code>	Match the quantified expression according to the guidelines stated above for possessive quantifiers, where <code>qu_expr</code> represents any one of the expressions shown in the top six rows of this table.

### Zero or One — `expr?`

Use `?` to make the HTML `<code>` and `</code>` tags optional in the string. The first string, `hstr1`, contains one occurrence of each tag. Since the expression uses `()?` around the tags, one occurrence is a match:

```
hstr1 = '<td><a name="18854"></a><code>%%</code><br></td>';
expr = '</a>( <code>)?..(</code>)?<br>';
```

```
regex(hstr1, expr, 'match')
ans =
    '</a><code>%%</code><br>'
```

The second string, `hstr2`, does not contain the code tags at all. Just the same, the expression matches because `()?` allows for zero occurrences of the tags:

```
hstr2 = '<td><a name="18854"></a>%<br></td>';
expr = '</a>(<code>)?..(</code>)?<br>';

regex(hstr2, expr, 'match')
ans =
    '</a>%<br>'
```

### Zero or More — `expr*`

Use `*` to match strings having any number of line breaks, including no line breaks at all.

```
hstr1 = '<p>This string has <br><br>line breaks</p>';
expr = '<p>.*(<br>)*.*</p>';

regex(hstr1, expr, 'match')
ans =
    '<p>This string has <br><br>line breaks</p>'

hstr2 = '<p>This string has no line breaks</p>';
regex(hstr2, expr, 'match')
ans =
    '<p>This string has no line breaks</p>'
```

### One or More — `expr+`

Use `+` to verify that the HTML image source is not empty. This looks for one or more characters in the gif filename:

```
hstr = '<a href="s12.html">';
expr = '</a><a href="s13.html#18760">';
expr = '<a href="\w{1,}(\.html){1}(\#\d{5,8}){0,1}"';

regexp(hstr, expr, 'match')
ans =
    '<a href="s13.html#18760"'
```

### Greedy Quantifiers — expr\*

Use \* to match as many characters as possible between any < and > signs in the string. Because of the .\* in the expression, regexp reads all characters in the string up to the end. Finding no closing > at the end, regexp then backs up to the </a> and ends the phrase there:

```
hstr = '<tr valign=top><td><a name="19184"></a>xyz';

regexp(hstr, '<.*>', 'match')
ans =
    '<tr valign=top><td><a name="19184"></a>'
```

### Possessive Quantifiers — expr\*+

Except for the possessive \*+ quantifier, this expression is the same as that used in the last example. Unlike the greedy quantifier, possessive quantifiers do not reevaluate parts of the string that have already been evaluated. This command scans the entire string because of the .\* quantifier, but then cannot back up to locate the </a> sequence that would satisfy the expression. As a result, no match is found and regexp returns an empty cell array:

```
regexp(hstr, '<.*+>', 'match')
ans =
    {}
```

### Lazy Quantifiers — `expr*?`

This example shows the difference between lazy and greedy quantifiers. The first expression uses lazy `. *?` to match the minimum number of characters between `<tr`, `<td`, or `</td` tags:

```
hstr = '<tr valign=top><td><a name="19184"></a><br></td>';  
  
regexp(hstr, '</?t.*?>', 'match')  
ans =  
    '<tr valign=top>'    '<td>'    '</td>'
```

The second expression uses greedy `. *` to match all characters from the opening `<tr` to the ending `</td`:

```
regexp(hstr, '</?t.*>', 'match')  
ans =  
    '<tr valign=top><td><a name="19184"></a><br></td>'
```

### Tokens

Parentheses used in a regular expression not only group elements of that expression together, but also designate any matches found for that group as *tokens*. You can use tokens to match other parts of the same string. One advantage of using tokens is that they remember what they matched, so you can recall and reuse matched text in the process of searching or replacing.

This section covers

- “Introduction to Using Tokens” on page 3-41
- “Using the token Parameter” on page 3-41
- “Operators Used with Tokens” on page 3-42
- “Using Tokens — Example 1” on page 3-43
- “Using Tokens — Example 2” on page 3-43
- “Using Tokens in a Replacement String” on page 3-44
- “Named Capture — `(?<name>expr)`” on page 3-44
- “Conditional Expressions — `(?(token)expr1 | expr2)`” on page 3-45

## Introduction to Using Tokens

You can turn any pattern being matched into a token by enclosing the pattern in parentheses within the expression. For example, to create a token for a dollar amount, you could use `'(\$d+)`'. Each token in the expression is assigned a number from 1 to 255 going from left to right. To make a reference to a token later in the expression, refer to it using a backslash followed by the token number. For example, when referencing a token generated by the third set of parentheses in the expression, use `\3`.

As a simple example, if you wanted to search for identical sequential letters in a string, you could capture the first letter as a token and then search for a matching character immediately afterwards. In the expression shown below, the `(\S)` phrase creates a token whenever `regex` matches any non-white-space character in the string. The second part of the expression, `'\1'`, looks for a second instance of the same character immediately following the first:

```
poestr = ['While I nodded, nearly napping, ' ...
         'suddenly there came a tapping,'];

[mat tok ext] = regex(poestr, '(\S)\1', 'match', ...
    'tokens', 'tokenExtents');
mat
mat =
    'dd'    'pp'    'dd'    'pp'
```

The tokens returned in cell array `tok` are:

```
'd', 'p', 'd', 'p'
```

Starting and ending indices for each token in the input string `poestr` are:

```
11 11, 26 26, 35 35, 57 57
```

## Using the token Parameter

You can have `regex` and `regexpi` return the actual tokens rather than token indices by specifying the optional `'token'` parameter in the command. The following example is the same as the one above, except that it returns the text of the tokens found by the pattern `\S`.

```

tok = regexp(poestr, '(\S)\1', 'tokens')
tok =
    {1x1 cell}    {1x1 cell}    {1x1 cell}    {1x1 cell}

tok{:}
ans =
    'd'
ans =
    'p'
ans =
    'd'
ans =
    'p'

```

### Operators Used with Tokens

Here are the operators you can use with tokens in MATLAB.

Operator	Usage
(expr)	Capture in a token all characters matched by the expression within the parentheses.
\N	Match the N <sup>th</sup> token generated by this command. That is, use \1 to match the first token, \2 to match the second, and so on.
\$N	Insert the match for the N <sup>th</sup> token in a replacement string. Used only by the regexp function.
(?<name>expr)	Capture in a token all characters matched by the expression within the parentheses. Assign a name to the token.
\k<name>	Match the token referred to by name.
(?(tok)expr)	If token tok is generated, then match expression expr.
(?(tok)expr <sub>1</sub>   expr <sub>2</sub> )	If token tok is generated, then match expression expr <sub>1</sub> . Otherwise, match expression expr <sub>2</sub> .

## Using Tokens — Example 1

Here is an example of how tokens are assigned values. Suppose that you are going to search the following text:

```
andy ted bob jim andrew andy ted mark
```

You choose to search the above text with the following search pattern:

```
and(y|rew)|(t)e(d)
```

This pattern has three parenthetical expressions that generate tokens. When you finally perform the search, the following tokens are generated for each match.

Match	Token 1	Token 2
andy	y	
ted	t	d
andrew	rew	
andy	y	
ted	t	d

Only the highest level parentheses are used. For example, if the search pattern `and(y|rew)` finds the text `andrew`, token 1 is assigned the value `rew`. However, if the search pattern `(and(y|rew))` is used, token 1 is assigned the value `andrew`.

## Using Tokens — Example 2

Use `(expr)` and `\N` to capture pairs of matching HTML tags (e.g., `<a>` and `<\a>`) and the text between them. The expression used for this example is

```
expr = '<(\w+).*?>.*?</\1>';
```

The first part of the expression, `'<(\w+)'`, matches an opening bracket (`<`) followed by one or more alphabetic, numeric, or underscore characters. The enclosing parentheses capture token characters following the opening bracket.

The second part of the expression, `'.*?>.*?'`, matches the remainder of this HTML tag (characters up to the `>`), and any characters that may precede the next opening bracket.

The last part, `'</\1>'`, matches all characters in the ending HTML tag. This tag is composed of the sequence `</tag>`, where `tag` is whatever characters were captured as a token.

```
hstr = '<!comment><a name="752507"></a><b>Default</b><br>';
expr = '<(\w+).*?>.*?</\1>';

[mat tok] = regexp(hstr, expr, 'match', 'tokens');
mat{:}
ans =
    <a name="752507"></a>
ans =
    <b>Default</b>

tok{:}
ans =
    'a'
ans =
    'b'
```

### Using Tokens in a Replacement String

When using tokens in a replacement string, reference them using `$1`, `$2`, etc. instead of `\1`, `\2`, etc. This example captures two tokens and reverses their order. The first, `$1`, is `'Norma Jean'` and the second, `$2`, is `'Baker'`. Note that `regexp` returns the modified string, not a vector of starting indices, by default:

```
regexp('Norma Jean Baker', '(\w+\s\w+)\s(\w+)', '$2, $1')
ans =
    Baker, Norma Jean
```

### Named Capture — (?<name>expr)

If you use a lot of tokens in your expressions, it may be helpful to assign them names rather than having to keep track of which token number is assigned to which token. Use the operator `(?<name>expr)` to assign name to the token matching expression `expr`.

When referencing a named token within the expression, use the syntax `\k<name>` instead of the numeric `\1`, `\2`, etc.:

```
poestr = ['While I nodded, nearly napping, ' ...
         'suddenly there came a tapping,'];

regexp(poestr, '(?<anychar>.)\k<anychar>', 'match')
ans =
    'dd'    'pp'    'dd'    'pp'
```

### Conditional Expressions — `(?(token)expr1|expr2)`

With conditional regular expressions, you can select which pattern to match, depending on whether a token elsewhere in the string is found. The expression appears as

```
(?(token)expr1|expr2)
```

This expression can be translated as an if-then-else statement, as follows:

```
if the specified token is found
    then match expression expr1
    else match expression expr2
```

The next example uses the conditional expression `expr` to match the string regardless of the gender used. The expression creates a token if `Mr` is followed by the letter `s`. It later matches either `her` or `his`, depending on whether this token was found. The phrase `(?(1)her|his)` means that **if** token 1 is found, **then** match `her`, **else** match `his`:

```
expr = 'Mr(s?)\..*?(?(1)her|his) son';

[mat tok] = regexp('Mr. Clark went to see his son', ...
    expr, 'match', 'tokens')
mat =
    'Mr. Clark went to see his son'
tok =
    {1x2 cell}

tok{:}
ans =
    ''    'his'
```

In the second part of the example, the token `s` is found and MATLAB matches the word `her`:

```
[mat tok] = regexp('Mrs. Clark went to see her son', ...
    expr, 'match', 'tokens')
mat =
    'Mrs. Clark went to see her son'
tok =
    {1x2 cell}

tok{:}
ans =
    's'    'her'
```

---

**Note** The MATLAB regular expression functions support both `if-then` and `if-then-else` statements.

---

## Handling Multiple Strings

You can use any of the MATLAB regular expression functions with cell arrays of strings as well as with single strings. Any or all of the input parameters (the string, expression, or replacement string) can be a cell array of strings. The `regexp` function requires that the string and expression arrays have the same number of elements if both are vectorized (i.e., if they have dimensions greater than 1-by-N). The `regexprep` function requires that the expression and replacement arrays have the same number of elements if the replacement array is vectorized. (The cell arrays do not have to have the same shape.)

Whenever the first input parameter to a regular expression function is a cell array, all output values are cell arrays of the same size.

This section covers the following topics:

- “Finding a Single Pattern in Multiple Strings” on page 3-47
- “Finding Multiple Patterns in Multiple Strings” on page 3-48
- “Replacing Multiple Strings” on page 3-49

## Finding a Single Pattern in Multiple Strings

The example shown here uses the `regexp` function on a cell array of strings `cstr`. It searches each string of the cell array for consecutive matching letters (e.g., 'oo'). The function returns a cell array of the same size as the input array. Each row of the return array contains the indices for which there was a match against the input cell array.

Here is the input cell array:

```
cstr = {
    'Whose woods these are I think I know.' ; ...
    'His house is in the village though;' ; ...
    'He will not see me stopping here' ; ...
    'To watch his woods fill up with snow.'};
```

Find consecutive matching letters by capturing a letter as a token (.) and then repeating that letter as a token reference, \1:

```
idx = regexp(cstr, '(.)\1');
```

```
whos idx
  Name      Size      Bytes  Class
  idx      4x1      296    cell array
```

```
idx{:}
ans =          % 'Whose woods these are I think I know.'
      8          %          |8

ans =          % 'His house is in the village though;'
      23          %          |23

ans =          % 'He will not see me stopping here'
      6      14      23  %          |6          |14          |23

ans =          % 'To watch his woods fill up with snow.'
      15      22          %          |15          |22
```



## Replacing Multiple Strings

When replacing multiple strings with `regexprep`, use a single replacement string if the expression consists of a single string. This example uses a common replacement value ('--') for all matches found in the multiple string input `cstr`. The function returns a cell array of strings having the same dimensions as the input cell array:

```
s = regexprep(cstr, '(.)\1', '--', 'ignorecase')
s =
    'Whose w--ds these are I think I know.'
    'His house is in the vi--age though;'
    'He wi-- not s-- me sto--ing here'
    'To watch his w--ds fi-- up with snow.'
```

You can use multiple replacement strings if the expression consists of multiple strings. In this example, the input string and replacement string are both 4-by-1 cell arrays, and the expression is a 1-by-4 cell array. As long as the expression and replacement arrays contain the same number of elements, the statement is valid. The dimensions of the return value match the dimensions of the input string:

```
expr = {'i\s', 'hou', '(.)\1', '\<w[aeiou]'};
repl = {'-1-', '-2-', '-3-', '-4-'};

s = regexprep(cstr, expr, repl, 'ignorecase')
s =
    'Whose w-3-ds these are -1-think -1-know.'
    'His -2-se is in the vi-3-age t-2-gh;'
    'He -4--3- not s-3- me sto-3-ing here'
    'To -4-tch his w-3-ds fi-3- up -4-th snow.'
```

## Operator Summary

MATLAB provides these operators for working with regular expressions:

- “Character Classes” on page 3-50
- “Character Representation” on page 3-51
- “Logical Operators” on page 3-51
- “Lookaround Operators” on page 3-52
- “Quantifiers” on page 3-52
- “Token Operators” on page 3-53

### Character Classes

Operator	Usage
.	Any single character, including white space
[c <sub>1</sub> c <sub>2</sub> c <sub>3</sub> ]	Any character contained within the brackets: c <sub>1</sub> or c <sub>2</sub> or c <sub>3</sub>
[^c <sub>1</sub> c <sub>2</sub> c <sub>3</sub> ]	Any character not contained within the brackets: anything but c <sub>1</sub> or c <sub>2</sub> or c <sub>3</sub>
[c <sub>1</sub> -c <sub>2</sub> ]	Any character in the range of c <sub>1</sub> through c <sub>2</sub>
\s	Any white-space character; equivalent to [ \f\n\r\t]
\S	Any non-white-space character; equivalent to [^ \f\n\r\t]
\w	Any alphabetic, numeric, or underscore character; equivalent to [a-zA-Z_0-9]
\W	Any character that is not alphabetic, numeric, or underscore; equivalent to [^a-zA-Z_0-9]
\d	Any numeric digit; equivalent to [0-9]
\D	Any nondigit character; equivalent to [^0-9]

**Character Representation**

<b>Operator</b>	<b>Usage</b>
<code>\a</code>	Alarm (beep)
<code>\b</code>	Backspace
<code>\e</code>	Escape
<code>\f</code>	Form feed
<code>\n</code>	New line
<code>\r</code>	Carriage return
<code>\t</code>	Horizontal tab
<code>\v</code>	Vertical tab
<code>\oN</code> or <code>\o{N}</code>	Character of octal value N
<code>\xN</code> or <code>\x{N}</code>	Character of hexadecimal value N
<code>\char</code>	If a character has special meaning in a regular expression, precede it with backslash ( <code>\</code> ) to match it literally.

**Logical Operators**

<b>Operator</b>	<b>Usage</b>
<code>(expr)</code>	Group regular expressions and capture tokens.
<code>(?:expr)</code>	Group regular expressions, but do not capture tokens.
<code>(?&gt;expr)</code>	Group atomically.
<code>(?#expr)</code>	Insert a comment into the expression. Comments are ignored in matching.
<code>expr<sub>1</sub> expr<sub>2</sub></code>	Match expression <code>expr<sub>1</sub></code> or expression <code>expr<sub>2</sub></code> .
<code>^expr</code>	Match the expression only at the beginning of the string.
<code>expr\$</code>	Match the expression only at the end of the string.

**Logical Operators**

<b>Operator</b>	<b>Usage</b>
<code>\&lt;expr</code>	Match the characters when they start a word.
<code>expr\&gt;</code>	Match the characters when they end a word.
<code>\&lt;expr\&gt;</code>	Match an exact word.

**Lookaround Operators**

<b>Operator</b>	<b>Usage</b>
<code>expr<sub>1</sub>(?=expr<sub>2</sub>)</code>	Match expression <code>expr<sub>1</sub></code> if followed by expression <code>expr<sub>2</sub></code> .
<code>expr<sub>1</sub>(?!expr<sub>2</sub>)</code>	Match expression <code>expr<sub>1</sub></code> if not followed by expression <code>expr<sub>2</sub></code> .
<code>(?&lt;=expr<sub>1</sub>)expr<sub>2</sub></code>	Match expression <code>expr<sub>2</sub></code> if preceded by expression <code>expr<sub>1</sub></code> .
<code>(?&lt;!expr<sub>1</sub>)expr<sub>2</sub></code>	Match expression <code>expr<sub>2</sub></code> if not preceded by expression <code>expr<sub>1</sub></code> .

**Quantifiers**

<b>Operator</b>	<b>Usage</b>
<code>?</code>	Match the preceding element 0 times or 1 time. Equivalent to <code>{0, 1}</code> .
<code>*</code>	Match the preceding element 0 or more times. Equivalent to <code>{0, }</code> .
<code>+</code>	Match the preceding element 1 or more times. Equivalent to <code>{1, }</code> .
<code>{n}</code>	Must match exactly <code>n</code> times. Equivalent to <code>{n, n}</code> .
<code>{n, }</code>	Must occur at least <code>n</code> times.

**Quantifiers (Continued)**

<b>Operator</b>	<b>Usage</b>
<code>{n,m}</code>	Must occur at least <i>n</i> times but no more than <i>m</i> times.
<code>q?</code>	Match the preceding element <i>q</i> times according to the guidelines stated above for lazy quantifiers, where <i>q</i> represents any one of the expressions shown in the top six rows of this table (e.g., <code>{2,5}?</code> ).
<code>q+</code>	Match the preceding element <i>q</i> times according to the guidelines stated above for possessive quantifiers, where <i>q</i> represents any one of the expressions shown in the top six rows of this table (e.g., <code>{2,5}+</code> ).

**Token Operators**

<b>Operator</b>	<b>Usage</b>
<code>(expr)</code>	Capture in a token all characters matched by the expression within the parentheses.
<code>\N</code>	Match the <i>N</i> <sup>th</sup> token generated by this command. That is, use <code>\1</code> to match the first token, <code>\2</code> to match the second, and so on.
<code>\$N</code>	Insert the match for the <i>N</i> <sup>th</sup> token in a replacement string. Used only by the <code>regexprep</code> function.
<code>(?&lt;name&gt;expr)</code>	Capture in a token all characters matched by the expression within the parentheses. Assign a name to the token.
<code>\k&lt;name&gt;</code>	Match the token referred to by name.
<code>(?(tok)expr)</code>	If token <code>tok</code> is generated, then match expression <code>expr</code> .
<code>(?(tok)expr<sub>1</sub> expr<sub>2</sub>)</code>	If token <code>tok</code> is generated, then match expression <code>expr<sub>1</sub></code> . Otherwise, match expression <code>expr<sub>2</sub></code> .

## Comma-Separated Lists

Typing in a series of numbers separated by commas gives you what is called a *comma-separated list*. MATLAB returns each value individually:

```
1, 2, 3
ans =
     1
ans =
     2
ans =
     3
```

Such a list, by itself, is not very useful. But when used with large and more complex data structures like MATLAB structures and cell arrays, the comma-separated list can enable you to simplify your MATLAB code.

This section covers

- “Generating a List from a Cell Array” on page 3-54
- “Generating a List from a Structure” on page 3-55
- “How to Use the Comma-Separated List” on page 3-56
- “Fast Fourier Transform Example” on page 3-58

### Generating a List from a Cell Array

Extracting multiple elements from a cell array yields a comma-separated list. Given a 4-by-6 cell array as shown here

```
C = cell(4, 6);
for k = 1:24, C{k} = k * 2; end

C
C =
     [2]     [10]     [18]     [26]     [34]     [42]
     [4]     [12]     [20]     [28]     [36]     [44]
     [6]     [14]     [22]     [30]     [38]     [46]
     [8]     [16]     [24]     [32]     [40]     [48]
```

extracting the fifth column generates the following comma-separated list:

```
C(:, 5)
ans =
    34
ans =
    36
ans =
    38
ans =
    40
```

This is the same as explicitly typing

```
C{1, 5}, C{2, 5}, C{3, 5}, C{4, 5}
```

## Generating a List from a Structure

For structures, extracting a field of the structure that exists across one of its dimensions yields a comma-separated list.

Start by converting the cell array used above into a 4-by-1 MATLAB structure with six fields: f1 through f6. Read field f5 for all rows and MATLAB returns a comma-separated list:

```
S = cell2struct(C, {'f1', 'f2', 'f3', 'f4', 'f5', 'f6'}, 2);

S.f5
ans =
    34
ans =
    36
ans =
    38
ans =
    40
```

This is the same as explicitly typing

```
S(1).f5, S(2).f5, S(3).f5, S(4).f5
```

## How to Use the Comma-Separated List

Common uses for comma-separated lists are

- “Constructing Arrays” on page 3-56
- “Displaying Arrays” on page 3-56
- “Concatenation” on page 3-57
- “Function Call Arguments” on page 3-57
- “Function Return Values” on page 3-58

The following sections provide examples of using comma-separated lists with cell arrays. Each of these examples applies to MATLAB structures as well.

### Constructing Arrays

You can use a comma-separated list to enter a series of elements when constructing a matrix or array. Note what happens when you insert a *list* of elements as opposed to adding the cell itself.

When you specify a list of elements with `C{: , 5}`, MATLAB inserts the four individual elements:

```
A = {'Hello', C{: , 5}, magic(4)}
A =
    'Hello'    [34]    [36]    [38]    [40]    [4x4 double]
```

When you specify the `C` cell itself, MATLAB inserts the entire cell array:

```
A = {'Hello', C, magic(4)}
A =
    'Hello'    {4x6 cell}    [4x4 double]
```

### Displaying Arrays

Use a list to display all or part of a structure or cell array:

```
A{:}
ans =
    Hello
ans =
    34
ans =
    36
```

```
ans =
    38
    .
    .
    .
```

### Concatenation

Putting a comma-separated list inside square brackets extracts the specified elements from the list and concatenates them:

```
A = [C{:, 5:6}]
A =
    34    36    38    40    42    44    46    48

whos A
  Name      Size           Bytes  Class

  A         1x8             64    double array
```

### Function Call Arguments

When writing the code for a function call, you enter the input arguments as a list with each argument separated by a comma. If you have these arguments stored in a structure or cell array, then you can generate all or part of the argument list from the structure or cell array instead. This can be especially useful when passing in variable numbers of arguments.

This example passes several attribute-value arguments to the plot function:

```
X = -pi:pi/10:pi;
Y = tan(sin(X)) - sin(tan(X));

C{1,1} = 'LineWidth';           C{2,1} = 2;
C{1,2} = 'MarkerEdgeColor';    C{2,2} = 'k';
C{1,3} = 'MarkerFaceColor';    C{2,3} = 'g';

plot(X, Y, '--rs', C{:})
```

### Function Return Values

MATLAB functions can also return more than one value to the caller. These values are returned in a list with each value separated by a comma. Instead of listing each return value, you can use a comma-separated list with a structure or cell array. This becomes more useful for those functions that have variable numbers of return values.

This example returns four values to a cell array:

```
C = cell(1, 4);
[C{:}] = fileparts('work/mytests/strArrays.mat')
C =
    'work/mytests'    'strArrays'    '.mat'    ''
```

### Fast Fourier Transform Example

The `fftshift` function swaps the left and right halves of each dimension of an array. For a simple vector such as `[0 2 4 6 8 10]` the output would be `[6 8 10 0 2 4]`. For a multidimensional array, `fftshift` performs this swap along each dimension.

`fftshift` uses vectors of indices to perform the swap. For the vector shown above, the index `[1 2 3 4 5 6]` is rearranged to form a new index `[4 5 6 1 2 3]`. The function then uses this index vector to reposition the elements. For a multidimensional array, `fftshift` must construct an index vector for each dimension. A comma-separated list makes this task much simpler.

Here is the `fftshift` function:

```
function y = fftshift(x)

numDims = ndims(x);
idx = cell(1, numDims);

for k = 1:numDims
    m = size(x, k);
    p = ceil(m/2);
    idx{k} = [p+1:m 1:p];
end

y = x(idx{:});
```

The function stores the index vectors in cell array `idx`. Building this cell array is relatively simple. For each of the `N` dimensions, determine the size of that dimension and find the integer index nearest the midpoint. Then, construct a vector that swaps the two halves of that dimension.

By using a cell array to store the index vectors and a comma-separated list for the indexing operation, `fftshift` shifts arrays of any dimension using just a single operation: `y = x(idx{:})`. If you were to use explicit indexing, you would need to write one `if` statement for each dimension you want the function to handle:

```
if ndims(x) == 1
    y = x(index1);
else if ndims(x) == 2
    y = x(index1, index2);
end
```

Another way to handle this without a comma-separated list would be to loop over each dimension, converting one dimension at a time and moving data each time. With a comma-separated list, you move the data just once. A comma-separated list makes it very easy to generalize the swapping operation to an arbitrary number of dimensions.

## Program Control Statements

Program control is divided into these four categories:

- “Conditional Control — if, switch” on page 3-60
- “Loop Control — for, while, continue, break” on page 3-64
- “Error Control — try, catch” on page 3-67
- “Program Termination — return” on page 3-68

---

**Note** You can often speed up the execution of MATLAB code by replacing for and while loops with vectorized code. See “Techniques for Improving Performance” on page 10-4 for more information on this.

---

### Conditional Control — if, switch

This group of control statements enables you to select at run-time which block of code is executed. To make this selection based on whether a condition is true or false, use the `if` statement (which may include `else` or `elseif`). To select from a number of possible options depending on the value of an expression, use the `switch` and `case` statements (which may include `otherwise`).

#### if, else, and elseif

`if` evaluates a logical expression and executes a group of statements based on the value of the expression. In its simplest form, its syntax is

```
if logical_expression
    statements
end
```

If the logical expression is true (that is, if it evaluates to logical 1), MATLAB executes all the statements between the `if` and `end` lines. It resumes execution at the line following the `end` statement. If the condition is false (evaluates to logical 0), MATLAB skips all the statements between the `if` and `end` lines, and resumes execution at the line following the `end` statement.

For example,

```
if rem(a, 2) == 0
    disp('a is even')
    b = a/2;
end
```

You can nest any number of `if` statements.

If the logical expression evaluates to a nonscalar value, all the elements of the argument must be nonzero. For example, assume `X` is a matrix. Then the statement

```
if X
    statements
end
```

is equivalent to

```
if all(X(:))
    statements
end
```

The `else` and `elseif` statements further conditionalize the `if` statement:

- The `else` statement has no logical condition. The statements associated with it execute if the preceding `if` (and possibly `elseif` condition) evaluates to logical 0 (false).
- The `elseif` statement has a logical condition that it evaluates if the preceding `if` (and possibly `elseif` condition) is false. The statements associated with it execute if its logical condition evaluates to logical 1 (true). You can have multiple `elseif` statements within an `if` block.

```
if n < 0 % If n negative, display error message.
    disp('Input must be positive');
elseif rem(n,2) == 0 % If n positive and even, divide by 2.
    A = n/2;
else
    A = (n+1)/2; % If n positive and odd, increment and divide.
end
```

**if Statements and Empty Arrays.** An if condition that reduces to an empty array represents a false condition. That is,

```
if A
    S1
else
    S0
end
```

executes statement S0 when A is an empty array.

### **switch, case, and otherwise**

switch executes certain statements based on the value of a variable or expression. Its basic form is

```
switch expression (scalar or string)
    case value1
        statements           % Executes if expression is value1
    case value2
        statements           % Executes if expression is value2
    .
    .
    .
    otherwise
        statements           % Executes if expression does not
                             % match any case
end
```

This block consists of

- The word `switch` followed by an expression to evaluate.
- Any number of case groups. These groups consist of the word `case` followed by a possible value for the expression, all on a single line. Subsequent lines contain the statements to execute for the given value of the expression. These can be any valid MATLAB statement including another `switch` block. Execution of a case group ends when MATLAB encounters the next case statement or the `otherwise` statement. Only the first matching case is executed.
- An optional `otherwise` group. This consists of the word `otherwise`, followed by the statements to execute if the expression's value is not handled by any

of the preceding case groups. Execution of the otherwise group ends at the end statement.

- An end statement.

switch works by comparing the input expression to each case value. For numeric expressions, a case statement is true if (value==expression). For string expressions, a case statement is true if strcmp(value,expression).

The code below shows a simple example of the switch statement. It checks the variable input\_num for certain values. If input\_num is -1, 0, or 1, the case statements display the value as text. If input\_num is none of these values, execution drops to the otherwise statement and the code displays the text 'other value'.

```
switch input_num
    case -1
        disp('negative one');
    case 0
        disp('zero');
    case 1
        disp('positive one');
    otherwise
        disp('other value');
end
```

---

**Note** For C programmers, unlike the C language switch construct, the MATLAB switch does not “fall through.” That is, if the first case statement is true, other case statements do not execute. Therefore, break statements are not used.

---

switch can handle multiple conditions in a single case statement by enclosing the case expression in a cell array.

```
switch var
    case 1
        disp('1')
    case {2,3,4}
        disp('2 or 3 or 4')
```

```
        case 5
            disp('5')
        otherwise
            disp('something else')
    end
```

### Loop Control – for, while, continue, break

With loop control statements, you can repeatedly execute a block of code, looping back through the block while keeping track of each iteration with an incrementing index variable. Use the `for` statement to loop a specific number of times. The `while` statement is more suitable for basing the loop execution on how long a condition continues to be true or false. The `continue` and `break` statements give you more control on exiting the loop.

#### for

The `for` loop executes a statement or group of statements a predetermined number of times. Its syntax is

```
    for index = start:increment:end
        statements
    end
```

The default increment is 1. You can specify any increment, including a negative one. For positive indices, execution terminates when the value of the index exceeds the *end* value; for negative increments, it terminates when the index is less than the end value.

For example, this loop executes five times.

```
    for n = 2:6
        x(n) = 2 * x(n - 1);
    end
```

You can nest multiple `for` loops.

```
    for m = 1:5
        for n = 1:100
            A(m, n) = 1/(m + n - 1);
        end
    end
```

---

**Note** You can often speed up the execution of MATLAB code by replacing `for` and `while` loops with vectorized code. See “Vectorizing Loops” on page 10-4 for details.

---

**Using Arrays as Indices.** The index of a `for` loop can be an array. For example, consider an  $m$ -by- $n$  array  $A$ . The statement

```
for k = A
    statements
end
```

sets  $k$  equal to the vector  $A(:, i)$ , where  $i$  is the iteration number of the loop. For the first loop iteration,  $k$  is equal to  $A(:, 1)$ ; for the second,  $k$  is equal to  $A(:, 2)$ ; and so on until  $k$  equals  $A(:, n)$ . That is, the loop iterates for a number of times equal to the number of columns in  $A$ . For each iteration,  $k$  is a vector containing one of the columns of  $A$ .

## **while**

The `while` loop executes a statement or group of statements repeatedly as long as the controlling expression is true (1). Its syntax is

```
while expression
    statements
end
```

If the expression evaluates to a matrix, all its elements must be 1 for execution to continue. To reduce a matrix to a scalar value, use the `all` and `any` functions.

For example, this `while` loop finds the first integer  $n$  for which  $n!$  ( $n$  factorial) is a 100-digit number.

```
n = 1;
while prod(1:n) < 1e100
    n = n + 1;
end
```

Exit a `while` loop at any time using the `break` statement.

**while Statements and Empty Arrays.** A while condition that reduces to an empty array represents a false condition. That is,

```
while A, S1, end
```

never executes statement S1 when A is an empty array.

#### **continue**

The continue statement passes control to the next iteration of the for or while loop in which it appears, skipping any remaining statements in the body of the loop. In nested loops, continue passes control to the next iteration of the for or while loop enclosing it.

The example below shows a continue loop that counts the lines of code in the file, `magic.m`, skipping all blank lines and comments. A continue statement is used to advance to the next line in `magic.m` without incrementing the count whenever a blank line or comment line is encountered.

```
fid = fopen('magic.m', 'r');
count = 0;
while ~feof(fid)
    line = fgetl(fid);
    if isempty(line) | strncmp(line, '%', 1)
        continue
    end
    count = count + 1;
end
disp(sprintf('%d lines', count));
```

#### **break**

The break statement terminates the execution of a for loop or while loop. When a break statement is encountered, execution continues with the next statement outside of the loop. In nested loops, break exits from the innermost loop only.

The example below shows a while loop that reads the contents of the file `fft.m` into a MATLAB character array. A break statement is used to exit the while loop when the first empty line is encountered. The resulting character array contains the M-file help for the `fft` program.

```
fid = fopen('fft.m', 'r');
s = '';
while ~feof(fid)
    line = fgetl(fid);
    if isempty(line)
        break
    end
    s = strvcat(s, line);
end
disp(s)
```

## Error Control – try, catch

Error control statements provide a way for you to take certain actions in the event of an error. Use the try statement to test whether a certain command in your code generates an error. If an error does occur within the try block, MATLAB immediately jumps to the corresponding catch block. The catch part of the statement needs to respond in some way to the error.

### try and catch

The general form of a try-catch statement sequence is

```
try
    statement
    ...
    statement
catch
    statement
    ...
    statement
end
```

In this sequence, the statements between try and catch are executed until an error occurs. The statements between catch and end are then executed. Use lasterr to see the cause of the error. If an error occurs between catch and end, MATLAB terminates execution unless another try-catch sequence has been established.

### **Program Termination – return**

Program termination control enables you to exit from your program at some point prior to its normal termination point.

#### **return**

After a MATLAB function runs to completion, it terminates and returns control either to the function that called it, or to the keyboard. If you need to exit a function prior to the point of normal completion, you can force an early termination using the `return` function. `return` immediately terminates the current sequence of commands and exits the currently running function.

`return` is also used to terminate keyboard mode.

## Symbol Reference

This section provides a quick reference to the following symbols used in MATLAB programming:

- “Asterisk — `*`” on page 3-70
- “At — `@`” on page 3-70
- “Colon — `:`” on page 3-71
- “Comma — `,`” on page 3-72
- “Curly Braces — `{ }`” on page 3-73
- “Dot — `.`” on page 3-73
- “Dot-Dot — `..`” on page 3-74
- “Dot-Dot-Dot (Ellipsis) — `...`” on page 3-74
- “Dot-Parentheses — `.()`” on page 3-75
- “Exclamation Point — `!`” on page 3-75
- “Parentheses — `()`” on page 3-76
- “Percent — `%`” on page 3-76
- “Percent-Brace — `%{ %}`” on page 3-77
- “Semicolon — `;`” on page 3-77
- “Single Quotes — `' '`” on page 3-78
- “Slash and Backslash — `/ \`” on page 3-79
- “Space Character” on page 3-78
- “Square Brackets — `[ ]`” on page 3-79

This section does not include symbols used in arithmetic, relational, and logical operations. For a description of these symbols, see the top of the list “Functions — Alphabetical List” in the MATLAB Help browser.

### Asterisk — \*

An asterisk in a filename specification is used as a wildcard specifier, as described below.

#### Filename Wildcard

Wildcards are generally used in file operations that act on multiple files or directories. They usually appear in the string containing the file or directory specification. MATLAB matches all characters in the name exactly except for the wildcard character \*, which can match any one or more characters.

To locate all files with names that start with 'january\_' and have a mat file extension, use

```
dir('january_*.mat')
```

You can also use wildcards with the `who` and `whos` functions. To get information on all variables with names starting with 'image' and ending with 'Offset', use

```
whos image*Offset
```

### At — @

The @ sign signifies either a function handle constructor or a directory that supports a MATLAB class.

#### Function Handle Constructor

The @ operator forms a handle to either the named function that follows the @ sign, or to the anonymous function that follows the @ sign.

**Function Handles in General.** Function handles are commonly used in passing functions as arguments to other functions. Construct a function handle by preceding the function name with an @ sign:

```
fhandle = @myfun
```

You can read more about function handles in “Function Handles” on page 4-27.

**Handles to Anonymous Functions.** Anonymous functions give you a quick means of creating simple functions without having to create M-files each time. You can construct an anonymous function and a handle to that function using the syntax

```
fhandle = @(arglist) body
```

where `body` defines the body of the function and `arglist` is the list of arguments you can pass to the function.

See “Anonymous Functions” on page 5-3 for more information.

### **Class Directory Designator**

A MATLAB class directory contains source files that define the methods and properties of a class. All MATLAB class directory names must begin with an @ sign:

```
\@myclass\get.m
```

See “MATLAB Classes” on page 2-82 for more information.

### **Colon — :**

The colon operator generates a sequence of numbers that you can use in creating or indexing into arrays. See “Generating a Numeric Sequence” on page 1-10 for more information on using the colon operator.

### **Numeric Sequence Range**

Generate a sequential series of regularly spaced numbers from `first` to `last` using the syntax `first:last`. For an incremental sequence from 6 to 17, use

```
N = 6:17
```

### **Numeric Sequence Step**

Generate a sequential series of numbers, each number separated by a step value, using the syntax `first:step:last`. For a sequence from 2 through 38, stepping by 4 between each entry, use

```
N = 2:4:38
```

### Indexing Range Specifier

Index into multiple rows or columns of a matrix using the colon operator to specify a range of indices:

```
B = A(7, 1:5);           % Read columns 1-5 of row 7.  
B = A(4:2:8, 10:15);    % Read columns 1-5 of rows 4, 6, and 8.  
B = A(:, 1:5);          % Read columns 1-5 of all rows.
```

### Conversion to Column Vector

Convert a matrix or array to a column vector using the colon operator as a single index:

```
A = rand(3,4);  
B = A(:);
```

### Preserving Array Shape on Assignment

Using the colon operator on the left side of an assignment statement, you can assign new values to array elements without changing the shape of the array:

```
A = rand(3,4);  
A(:) = 1:12;
```

### Comma — ,

A comma is used to separate the following types of elements.

#### Row Element Separator

When constructing an array, use a comma to separate elements that belong in the same row:

```
A = [5.92, 8.13, 3.53]
```

#### Array Index Separator

When indexing into an array, use a comma to separate the indices into each dimension:

```
X = A(2, 7, 4)
```

### Function Input and Output Separator

When calling a function, use a comma to separate output and input arguments:

```
function [data, text] = xlsread(file, sheet, range, mode)
```

### Command or Statement Separator

To enter more than one MATLAB command or statement on the same line, separate each command or statement with a comma:

```
for k = 1:10,    sum(A(k)),    end
```

### Curly Braces — { }

Use curly braces to construct or get the contents of cell arrays.

### Cell Array Constructor

To construct a cell array, enclose all elements of the array in curly braces:

```
C = {[2.6 4.7 3.9], rand(8)*6, 'C. Coolidge'}
```

### Cell Array Indexing

Index to a specific cell array element by enclosing all indices in curly braces:

```
A = C{4,7,2}
```

See “Cell Arrays” on page 2-66 for more information.

### Dot — .

The single dot operator has the following different uses in MATLAB.

### Structure Field Definition

Add fields to a MATLAB structure by following the structure name with a dot and then a field name:

```
funds(5,2).bondtype = 'Corporate';
```

See “Structures” on page 2-49 for more information.

### Object Method Specifier

Specify the properties of an instance of a MATLAB class using the object name followed by a dot, and then the property name:

```
val = asset.current_value
```

See “MATLAB Classes” on page 2-82 for more information.

### Dot-Dot – ..

Two dots in sequence refer to the parent of the current directory.

### Parent Directory

Specify the directory immediately above your current directory using two dots. For example, to go up two levels in the directory tree and down into the `testdir` directory, use

```
cd ..\..\testdir
```

### Dot-Dot-Dot (Ellipsis) – ...

The ellipsis is the line continuation operator in MATLAB.

### Line Continuation

Continue any MATLAB command or expression by placing an ellipsis at the end of the line to be continued:

```
fprintf('The current value of %s is %d', ...  
       vname, value)
```

**Entering Long Strings.** You cannot use an ellipsis within single quotes to continue a string to the next line:

```
string = 'This is not allowed and will generate an ...  
         error in MATLAB.'
```

To enter a string that extends beyond a single line, piece together shorter strings using either the concatenation operator (`[]`) or the `fprintf` function.

Here are two examples:

```
quote1 = [  
    'Tiger, tiger, burning bright in the forests of the night, ' ...  
    'what immortal hand or eye could frame thy fearful symmetry?'];  
  
quote2 = sprintf('%s%s%s', ...  
    'In Xanadu did Kubla Khan a stately pleasure-dome decree, ', ...  
    'where Alph, the sacred river, ran ', ...  
    'through caverns measureless to man down to a sunless sea.');
```

## Dot-Parentheses – .( )

Use dot-parentheses to specify the name of a dynamic structure field.

### Dynamic Structure Fields

Sometimes it is useful to reference structures with field names that can vary. For example, the referenced field might be passed as an argument to a function. Dynamic field names specify a variable name for a structure field.

The variable `fundtype` shown here is a dynamic field name:

```
type = funds(5,2).(fundtype);
```

See “Using Dynamic Field Names” on page 2-54 for more information.

## Exclamation Point – !

The exclamation point precedes operating system commands that you want to execute from within MATLAB.

### Shell Escape

The exclamation point initiates a shell escape function. Such a function is to be performed directly by the operating system:

```
!rmdir oldtests
```

See “Shell Escape Functions” on page 3-23 for more information.

### Parentheses – ( )

Parentheses are used mostly for indexing into elements of an array or for specifying arguments passed to a called function.

#### Array Indexing

When parentheses appear to the right of a variable name, they are indices into the array stored in that variable:

```
A(2, 7, 4)
```

#### Function Input Arguments

When parentheses follow a function name in a function declaration or call, the enclosed list contains input arguments used by the function:

```
function sendmail(to, subject, message, attachments)
```

### Percent – %

The percent sign is most commonly used to indicate nonexecutable text within the body of a program. This text is normally used to include comments in your code. Some functions also interpret the percent sign as a conversion specifier.

See “Help Text” on page 4-11 for more information.

#### Single Line Comments

Precede any one-line comments in your code with a percent sign. MATLAB does not execute anything that follows a percent sign (that is, unless the sign is quoted, '%'):

```
% The purpose of this routine is to compute  
% the value of ...
```

#### Conversion Specifiers

Some functions, like `sscanf` and `sprintf`, precede conversion specifiers with the percent sign:

```
sprintf('%s = %d', name, value)
```

## Percent-Brace — %{        %}

The %{ and %} symbols enclose a block of comments that extend beyond one line.

### Block Comments

Enclose any multiline comments with percent followed by an opening or closing brace.

```
%{  
  The purpose of this routine is to compute  
  the value of ...  
%}
```

---

**Note** The %{ and %} operators must appear alone on the lines that immediately precede and follow the block of help text. Do not include any other text on these lines.

---

## Semicolon — ;

The semicolon can be used to construct arrays, suppress output from a MATLAB command, or to separate commands entered on the same line.

### Array Row Separator

When used within square brackets to create a new array or concatenate existing arrays, the semicolon creates a new row in the array:

```
A = [5, 8; 3, 4]  
A =  
     5     8  
     3     4
```

### Output Suppression

When placed at the end of a command, the semicolon tells MATLAB not to display any output from that command. In this example, MATLAB does not display the resulting 100-by-100 matrix:

```
A = ones(100, 100);
```

### Command or Statement Separator

Like the comma operator, you can enter more than one MATLAB command on a line by separating each command with a semicolon. MATLAB suppresses output for those commands terminated with a semicolon, and displays the output for commands terminated with a comma.

In this example, assignments to variables A and C are terminated with a semicolon, and thus do not display. Because the assignment to B is comma-terminated, the output of this one command is displayed:

```
A = 12.5; B = 42.7, C = 1.25;  
B =  
42.7000
```

### Single Quotes – ' '

Single quotes are the constructor symbol for MATLAB character arrays.

### Character and String Constructor

MATLAB constructs a character array from all characters enclosed in single quotes. If only one character is in quotes, then MATLAB constructs a 1-by-1 array:

```
S = 'Hello World'
```

See “Characters and Strings” on page 2-25 for more information.

### Space Character

The space character serves a purpose similar to the comma in that it can be used to separate row elements in an array constructor, or the values returned by a function.

### Row Element Separator

You have the option of using either commas or spaces to delimit the row elements of an array when constructing the array. To create a 1-by-3 array, use

```
A = [5.92 8.13 3.53]  
A =  
5.9200 8.1300 3.5300
```

When indexing into an array, you must always use commas to reference each dimension of the array.

### Function Output Separator

Spaces are allowed when specifying a list of values to be returned by a function. You can use spaces to separate return values in both function declarations and function calls:

```
function [data text] = xlsread(file, sheet, range, mode)
```

### Slash and Backslash – / \

The slash (/) and backslash (\) characters separate the elements of a path or directory string. On Windows-based systems, both slash and backslash have the same effect. On UNIX-based systems, you must use slash only.

On a Windows system, you can use either backslash or slash:

```
dir([matlabroot '\toolbox\matlab\elmat\shiftdim.m'])  
dir([matlabroot '/toolbox/matlab/elmat/shiftdim.m'])
```

On a UNIX system, use only the forward slash:

```
dir([matlabroot '/toolbox/matlab/elmat/shiftdim.m'])
```

### Square Brackets – [ ]

Square brackets are used in array construction and concatenation, and also in declaring and capturing values returned by a function.

#### Array Constructor

To construct a matrix or array, enclose all elements of the array in square brackets:

```
A = [5.7, 9.8, 7.3; 9.2, 4.5, 6.4]
```

#### Concatenation

To combine two or more arrays into a new array through concatenation, enclose all array elements in square brackets:

```
A = [B, eye(6), diag([0:2:10])]
```

### Function Declarations and Calls

When declaring or calling a function that returns more than one output, enclose each return value that you need in square brackets:

```
[data, text] = xlsread(file, sheet, range, mode)
```

## MATLAB Functions

Many of the functions provided with MATLAB are implemented as M-files just like the M-files that you will create with MATLAB. Other MATLAB functions are precompiled executable programs called built-ins that run much more efficiently.

This section discusses both types of functions and also functions that are overloaded to handle different data types appropriately:

- “M-File Functions” on page 3-81
- “Built-In Functions” on page 3-82
- “Overloaded MATLAB Functions” on page 3-83

### M-File Functions

If you look in the subdirectories of the `toolbox\matlab` directory, you can find the M-file sources to many of the functions supplied with MATLAB. You can locate your `toolbox\matlab` directory by typing

```
dir([matlabroot '\toolbox\matlab\'])
```

MATLAB functions with an M-file source are just like any other functions coded with MATLAB. When one of these M-file functions is called, MATLAB parses and executes each line of code in the M-file. It saves the parsed version of the function in memory, eliminating parsing time on any further calls to this function.

### Identifying M-File Functions

To find out if a function is implemented with an M-file, use the `exist` function. The `exist` function searches for the name you enter on the MATLAB path and returns a number identifying the source. If the source is an M-file, then `exist` returns the number 2. This example identifies the source for the `repmat` function as an M-file:

```
exist repmat
ans =
     2
```

The `exist` function also returns 2 for files that have a file type unknown to MATLAB. However, if you invoke `exist` on a MATLAB function name, the file type will be known to MATLAB and will return 2 only on M-files.

### Viewing the Source Code

One advantage of functions implemented as M-files is that you can look at the source code. This may help when you need to understand why the function returns a value you didn't expect, if you need to figure out how to code something in MATLAB that is already coded in a function, or perhaps to help you create a function that overloads one of the MATLAB functions.

To find the source code for any MATLAB M-file function, use which:

```
which repmat
D:\matlabR14\toolbox\matlab\elmat\repmat.m
```

### Built-In Functions

Functions that are frequently used or that can take more time to execute are often implemented as executable files. These functions are called *built-ins*.

Unlike M-file functions, you cannot see the source code for built-ins. Although most built-in functions do have an M-file associated with them, this file is there mainly to supply the help documentation for the function. Take the `reshape` function, for example, and find it on the MATLAB path:

```
which reshape
D:\matlabR14\toolbox\matlab\elmat\reshape.m
```

If you type this M-file out, you will see that it consists almost entirely of help text. At the bottom is a call to the built-in executable image.

### Identifying Built-In Functions

As with M-file functions, you can identify which functions are built-ins using the `exist` function. This function identifies built-ins by returning the number 5:

```
exist reshape
ans =
    5
```

## Forcing a Built-In Call

If you overload any of the MATLAB built-in functions to handle a specific data type, then MATLAB will always call the overloaded function on that type. If, for some reason, you need to call the built-in version, you can override the usual calling mechanism using a function called `builtin`. The expression

```
builtin('reshape', arg1, arg2, ..., argN);
```

forces a call to MATLAB built-in `reshape`, passing the arguments shown even though an overload exists for the data types in this argument list.

## Overloaded MATLAB Functions

An overloaded function is an additional implementation of an existing function that has been designed specifically to handle a certain data type. When you pass an argument of this type in a call to the function, MATLAB looks for the function implementation that handles that type and executes that function code.

Each overloaded MATLAB function has an M-file on the MATLAB path. The M-files for a certain data type (or class) are placed in a directory named with an @ sign followed by the class name. For example, to overload the MATLAB `plot` function to plot expressions of a class named `polynom` differently than other data types, you would create a directory called `@polynom` and store your own version of `plot.m` in that directory.

You can add your own overloads to any function by creating a class directory for the data type you wish to support for that function, and creating an M-file that handles that type in a manner different from the default. See “Setting Up Class Directories” on page 8-6 and “Designing User Classes in MATLAB” on page 8-9.

When you use the `which` command with the `-all` option, MATLAB returns all occurrences of the file you are looking for. This is an easy way to find functions that are overloaded:

```
which -all set           % Show all implementations for 'set'
```



# M-File Programming

---

When you write a program in MATLAB, you save it to a file called an M-file (named after its `.m` file extension). There are two types of M-files that you can write: scripts and functions. Most of this chapter is about functions, as they are the more complex and useful of the two M-file types.

This chapter covers basic program development, describes how to write and call scripts and functions, and shows how to pass different types of data in a function call.

Program Development (p. 4-2)	Procedures and tools used in creating, debugging, optimizing, and checking in a program
Working with M-Files (p. 4-7)	Introduction to the basic MATLAB program file
M-File Scripts and Functions (p. 4-16)	Overview of scripts, simple programs that require no input or output, and functions, more complex programs that exchange input and output data with the caller
Function Arguments (p. 4-21)	Handling the data passed into and out of an M-file function, checking input data, passing variable numbers of arguments
Function Handles (p. 4-27)	Packaging the access to a function into a function handle, and passing that handle to other functions
Calling Functions (p. 4-31)	Calling syntax, determining which function will be called, passing different types of arguments, passing arguments in structures and cell arrays, identifying function dependencies

## Program Development

This section explains the basic steps in developing an M-file program in MATLAB. Associated with each step in this process are certain MATLAB tools and utilities that are fully documented in the Desktop Tools and Development Environment documentation.

This section covers

- “Creating a Program” on page 4-2
- “Getting the Bugs Out” on page 4-3
- “Cleaning Up the Program” on page 4-4
- “Improving Performance” on page 4-5
- “Checking It In” on page 4-6

For more ideas on good programming style, see “Program Development” on page 11-16. The MATLAB Programming Tips is a compilation of useful pieces of information that can show you alternate and often more efficient ways to accomplish common programming tasks while also expanding your knowledge of MATLAB.

### Creating a Program

You can type in your program code using any text editor. This section focuses on using the MATLAB Editor/Debugger for this purpose. The Editor/Debugger is fully documented in “Ways to Edit and Debug Files” in the Desktop Tools and Development Environment documentation.

The first step in creating a program is to open an editing window. To create a new M-file, type the word `edit` at the MATLAB command prompt. To edit an existing M-file, type `edit` followed by the filename:

```
edit drawPlot.m
```

MATLAB opens a new window for entering your program code. As you type in your program, MATLAB keeps track of the line numbers in the left column.

## Saving the Program

It is usually a good idea to save your program periodically while you are in the development process. To do this, click **File -> Save** in the Editor/Debugger. Enter a filename with a `.m` extension in the **Save file as** dialog box that appears and click **OK**. It is customary and less confusing if you give the M-file the same name as the first function in the M-file.

## Running the Program

Before trying to run your program, make sure that its M-file is on the MATLAB path. The MATLAB path defines those directories that you want MATLAB to know about when executing M-files. The path includes all the directories that contain functions provided with MATLAB. It should also include any directories that you use for your own functions.

Use the `which` function to see if your program is on the path:

```
which drawPlot
D:\matlabR14\work\drawPlot.m
```

If not, add its directory to the path using the `addpath` function:

```
addpath('D:\matlabR14\work')
```

Now you can run the program just by typing the name of the M-file at the MATLAB command prompt:

```
drawPlot(xdata, ydata)
```

## Getting the Bugs Out

In all but the simplest programs, you are likely to encounter some type of unexpected behavior when you run the program for the first time. Program defects can show up in the form of warning or error messages displayed in the command window, programs that hang (never terminate), inaccurate results, or some number of other symptoms. This is where the second functionality of the MATLAB Editor/Debugger becomes useful.

The MATLAB Debugger enables you to examine the inner workings of your program while you run it. You can stop the execution of your program at any point and then continue from that point, stepping through the code line by line and examining the results of each operation performed. You have the choice of

operating the debugger from the Editor window that displays your program, from the MATLAB command line, or both.

### The Debugging Process

You can step through the program right from the start if you want. For longer programs, you will probably save time by stopping the program somewhere in the middle and stepping through from there. You can do this by approximating where the program code breaks and setting a stopping point (or *breakpoint*) at that line. Once a breakpoint has been set, start your program from the MATLAB command prompt. MATLAB opens an Editor/Debugger window (if it is not already open) showing a green arrow pointing to the next line to execute.

From this point, you can examine any values passed into the program, or the results of each operation performed. You can step through the program line by line to see which path is taken and why. You can step into any functions that your program calls, or choose to step over them and just see the end results. You can also modify the values assigned to a variable and see how that affects the outcome.

To learn about using the MATLAB Debugger, see “Debugging and Improving M-Files” in the Desktop Tools and Development Environment documentation. Type `help debug` for a listing of all MATLAB debug functions.

For programming tips on how to debug, see “Debugging” on page 11-19.

### Cleaning Up the Program

Even after your program is bug-free, there are still some steps you can take to improve its performance and readability. The MATLAB M-Lint utility generates a report that can highlight potential problems in your code. For example, you may be using the element-wise AND operator (`&`) where the short-circuit AND (`&&`) is more appropriate. You may be using the `find` function in a context where logical subscripting would be faster.

MATLAB offers M-Lint and several other reporting utilities to help you make the finishing touches to your program code. These tools are described under “Tuning and Refining M-Files” in the Desktop Tools and Development Environment documentation.

## Improving Performance

The MATLAB Profiler generates a report that shows how your program spends its processing time. For details about using the MATLAB Profiler, see “Profiling for Improving Performance” in the MATLAB Desktop Tools and Development Environment documentation. For tips on other ways to improve the performance of your programs, see Chapter 10, “Improving Performance and Memory Usage.”.

Three types of reports are available:

- “Summary Report” on page 4-5
- “Detail Report” on page 4-5
- “File Listing” on page 4-5

### Summary Report

The summary report provides performance information on your main program and on every function it calls. This includes how many times each function is called, the total time spent in that function, along with a bar graph showing the relative time spent by each function.

### Detail Report

When you click a function name in the summary report, MATLAB displays a detailed report on that function. This report shows the lines of that function that take up the most time, the time spent executing that line, the percentage of total time for that function that is spent on that line, and a bar graph showing the relative time spent on the line.

### File Listing

The detail report for a function also displays the entire M-file code for that function. This listing enables you to view the time-consuming code in the context of the entire function body. For every line of code that takes any significant time, additional performance information is provided by the statistics and by the color and degree of highlighting of the program code.

### Checking It In

Source control systems offer a way to manage large numbers of files while they are under development. They keep track of the work done on these files as your project progresses, and also ensure that changes are made in a secure and orderly fashion.

If you have a source control system available to you, you will probably want to check your M-files into the system once they are complete. If further work is required on one of those files, you just check it back out, make the necessary modifications, and then check it back in again.

MATLAB provides an interface to external source control systems so that you can check files in and out directly from your MATLAB session. See “Revision Control” in the Desktop Tools and Development Environment documentation for instructions on how to use this interface.

## Working with M-Files

MATLAB provides a full programming language that enables you to write a series of MATLAB statements into a file and then execute them with a single command. You write your program in an ordinary text file, giving the file a name of `filename.m`. The term you use for `filename` becomes the new command that MATLAB associates with the program. The file extension of `.m` makes this a MATLAB M-file.

This section covers

- “Types of M-Files” on page 4-7
- “Basic Parts of an M-File” on page 4-8
- “Creating a Simple M-File” on page 4-13
- “Providing Help for Your Program” on page 4-15
- “Creating P-Code Files” on page 4-15

### Types of M-Files

M-files can be *scripts* that simply execute a series of MATLAB statements, or they can be *functions* that also accept input arguments and produce output.

MATLAB scripts:

- Are useful for automating a series of steps you need to perform many times.
- Do not accept input arguments or return output arguments.
- Store variables in a workspace that is shared with other scripts and with the MATLAB command line interface.

MATLAB functions:

- Are useful for extending the MATLAB language for your application.
- Can accept input arguments and return output arguments.
- Store variables in a workspace internal to the function.

## Basic Parts of an M-File

This simple function shows the basic parts of an M-file. Note that any line that begins with % is not executable:

```

function f = fact(n)           Function definition line
% Compute a factorial value.  H1 line
% FACT(N) returns the factorial of N, Help text
% usually denoted by N!

% Put simply, FACT(N) is PROD(1:N). Comment
f = prod(1:n);                Function body

```

The table below briefly describes each of these M-file parts. Both functions and scripts can have all of these parts, except for the function definition line which applies to functions only. These parts are described in greater detail following the table.

<b>M-File Element</b>	<b>Description</b>
Function definition line (functions only)	Defines the function name, and the number and order of input and output arguments
H1 line	A one line summary description of the program, displayed when you request help on an entire directory, or when you use lookfor
Help text	A more detailed description of the program, displayed together with the H1 line when you request help on a specific function
Function or script body	Program code that performs the actual computations and assigns values to any output arguments
Comments	Text in the body of the program that explains the internal workings of the program

## Function Definition Line

The function definition line informs MATLAB that the M-file contains a function, and specifies the argument calling sequence of the function. The function definition line for the `fact` function is

```
function y = fact(x)
```

The diagram illustrates the components of the function definition line `function y = fact(x)`. Arrows point from labels to the corresponding parts of the line:
 

- `function`: keyword
- `y`: output argument
- `fact`: function name
- `x`: input argument

All MATLAB functions have a function definition line that follows this pattern.

**Function Name.** Function names must begin with a letter, may contain any alphanumeric characters or underscores, and must be no longer than the maximum allowed length (returned by the function `namelengthmax`). Because variables must obey similar rules, you can use the `isvarname` function to check whether a function name is valid:

```
isvarname myfun
```

Although function names can be of any length, MATLAB uses only the first `N` characters of the name (where `N` is the number returned by the function `namelengthmax`) and ignores the rest. Hence, it is important to make each function name unique in the first `N` characters:

```
N = namelengthmax
N =
    63
```

---

**Note** Some operating systems may restrict file names to shorter lengths.

---

The name of the text file that contains a MATLAB function consists of the function name with the extension `.m` appended. For example,

```
average.m
```

If the filename and the function definition line name are different, the internal (function) name is ignored. Thus, if `average.m` is the file that defines a function named `computeAverage`, you would invoke the function by typing

```
average
```

---

**Note** While the function name specified on the function definition line does not have to be the same as the filename, it is best to use the same name for both to avoid confusion.

---

**Function Arguments.** If the function has multiple output values, enclose the output argument list in square brackets. Input arguments, if present, are enclosed in parentheses following the function name. Use commas to separate multiple input or output arguments. Here is the declaration for a function named `sphere` that has three inputs and three outputs:

```
function [x, y, z] = sphere(theta, phi, rho)
```

If there is no output, leave the output blank

```
function printresults(x)
```

or use empty square brackets:

```
function [] = printresults(x)
```

The variables that you pass to the function do not need to have the same name as those in the function definition line.

### The H1 Line

The H1 line, so named because it is the first help text line, is a comment line immediately following the function definition line. Because it consists of comment text, the H1 line begins with a percent sign, `%`. For the `average` function, the H1 line is

```
% AVERAGE Mean of vector elements.
```

This is the first line of text that appears when a user types `help functionname` at the MATLAB prompt. Further, the `lookfor` function searches on and displays only the H1 line. Because this line provides important summary information about the M-file, it is important to make it as descriptive as possible.

### Help Text

You can create online help for your M-files by entering help text on one or more consecutive comment lines at the start of your M-file program. MATLAB considers the first group of consecutive lines immediately following the H1 line that begin with `%` to be the online help text for the function. The first line without `%` as the left-most character ends the help.

The help text for the `average` function is

```
% AVERAGE(X), where X is a vector, is the mean of vector elements.  
% Nonvector input results in an error.
```

When you type `help functionname` at the command prompt, MATLAB displays the H1 line followed by the online help text for that function. The help system ignores any comment lines that appear after this help block.

---

**Note** Help text in an M-file can be viewed at the MATLAB command prompt only (using `help functionname`). You cannot display this text using the MATLAB Help browser. You can, however, use the Help browser to get help on MATLAB functions and also to read the documentation on any MathWorks products.

---

### The Function or Script Body

The function body contains all the MATLAB code that performs computations and assigns values to output arguments. The statements in the function body can consist of function calls, programming constructs like flow control and interactive input/output, calculations, assignments, comments, and blank lines.

For example, the body of the average function contains a number of simple programming statements:

```
[m,n] = size(x);
if (~(m == 1) | (n == 1)) | (m == 1 & n == 1) % Flow control
    error('Input must be a vector') % Error message display
end
y = sum(x)/length(x); % Computation and assignment
```

### Comments

As mentioned earlier, comment lines begin with a percent sign (%). Comment lines can appear anywhere in an M-file, and you can append comments to the end of a line of code. For example,

```
% Add up all the vector elements.
y = sum(x) % Use the sum function.
```

In addition to comment lines, you can insert blank lines anywhere in an M-file. Blank lines are ignored. However, a blank line can indicate the end of the help text entry for an M-file.

**Block Comments.** To write comments that require more than one line, use the block comment operators, %{ and %}:

```
{
This next block of code checks the number of inputs
passed in, makes sure that each input is a valid data
type, and then branches to start processing the data.
}
```

---

**Note** The %{ and %} operators must appear alone on the lines that immediately precede and follow the block of help text. Do not include any other text on these lines.

---

## Creating a Simple M-File

You create M-files using a text editor. MATLAB provides a built-in editor, but you can use any text editor you like. Once you have written and saved the M-file, you can run the program as you would any other MATLAB function or command.

The process looks like this:

- 1 Create an M-file using a text editor.

```
function c = myfile(a,b)
c = sqrt((a.^2)+(b.^2))
```

- 2 Call the M-file from the command line, or from within another M-file.

```
a = 7.5
b = 3.342
c = myfile(a,b)

c =

    8.2109
```

## Using Text Editors

M-files are ordinary text files that you create using a text editor. If you use the MATLAB Editor/Debugger, open a new file by selecting **New -> M-File** from the **File** menu at the top of the MATLAB Command Window.

Another way to edit an M-file is from the MATLAB command line using the `edit` function. For example,

```
edit foo
```

opens the editor on the file `foo.m`. Omitting a filename opens the editor on an untitled file.

You can create the `fact` function shown in “Basic Parts of an M-File” on page 4-8 by opening your text editor, entering the lines shown, and saving the text in a file called `fact.m` in your current directory.

Once you have created this file, here are some things you can do:

- List the names of the files in your current directory:

```
what
```

- List the contents of M-file `fact.m`:

```
type fact
```

- Call the `fact` function:

```
fact(5)
ans =
    120
```

### A Word of Caution on Saving M-Files

Save any M-files you create and any MathWorks supplied M-files that you edit in directories outside of the directory tree in which the MATLAB software is installed. If you keep your files in any of the installed directories, your files may be overwritten when you install a new version of MATLAB.

MATLAB installs its software into directories under `$matlabroot/toolbox`. To see what `$matlabroot` is on your system, type `matlabroot` at the MATLAB command prompt.

Also note that locations of files in the `$matlabroot/toolbox` directory tree are loaded and cached in memory at the beginning of each MATLAB session to improve performance. If you save files to `$matlabroot/toolbox` directories using an external editor, or if you add or remove files from these directories using file system operations, enter the commands `clear functionname` and `rehash toolbox` before you use the files in the current session.

For more information, see the `rehash` function reference page or the section “Toolbox Path Caching” in the Desktop Tools and Development Environment documentation.

## Providing Help for Your Program

You can provide user information for the programs you write by including a help text section at the beginning of your M-file. (See “Help Text” on page 4-11).

You can also make help entries for an entire directory by creating a file with the special name `Contents.m` that resides in the directory. This file must contain only comment lines; that is, every line must begin with a percent sign. MATLAB displays the lines in a `Contents.m` file whenever you type

```
help directoryname
```

If a directory does not contain a `Contents.m` file, typing `help directoryname` displays the first help line (the H1 line) for each M-file in the directory.

## Creating P-Code Files

You can save a parsed version of a function or script, called P-code files, for later MATLAB sessions using the `pcode` function. For example,

```
pcode average
```

parses `average.m` and saves the resulting pseudocode to the file named `average.p`. This saves MATLAB from reparsing `average.m` the first time you call it in each session.

MATLAB is very fast at parsing so the `pcode` function rarely makes much of a speed difference.

One situation where `pcode` does provide a speed benefit is for large GUI applications. In this case, many M-files must be parsed before the application becomes visible.

You can also use `pcode` to hide algorithms you have created in your M-file, if you need to do this for proprietary reasons.

## M-File Scripts and Functions

This section covers the following topics regarding functions:

- “M-File Scripts” on page 4-16
- “M-File Functions” on page 4-17
- “Types of Functions” on page 4-18
- “Identifying Dependencies” on page 4-19

### M-File Scripts

Scripts are the simplest kind of M-file because they have no input or output arguments. They are useful for automating series of MATLAB commands, such as computations that you have to perform repeatedly from the command line.

### The Base Workspace

Scripts share the base workspace with your interactive MATLAB session and with other scripts. They operate on existing data in the workspace, or they can create new data on which to operate. Any variables that scripts create remain in the workspace after the script finishes so you can use them for further computations. You should be aware, though, that running a script can unintentionally overwrite data stored in the base workspace by commands entered at the MATLAB command prompt.

### Simple Script Example

These statements calculate  $\rho$  for several trigonometric functions of  $\theta$ , then create a series of polar plots:

```
% An M-file script to produce          % Comment lines
% "flower petal" plots
theta = -pi:0.01:pi;                  % Computations
rho(1,:) = 2 * sin(5 * theta) .^ 2;
rho(2,:) = cos(10 * theta) .^ 3;
rho(3,:) = sin(theta) .^ 2;
rho(4,:) = 5 * cos(3.5 * theta) .^ 3;
for k = 1:4
    polar(theta, rho(k,:))            % Graphics output
    pause
end
```

Try entering these commands in an M-file called `petals.m`. This file is now a MATLAB script. Typing `petals` at the MATLAB command line executes the statements in the script.

After the script displays a plot, press **Enter** or **Return** to move to the next plot. There are no input or output arguments; `petals` creates the variables it needs in the MATLAB workspace. When execution completes, the variables (`i`, `theta`, and `rho`) remain in the workspace. To see a listing of them, enter `whos` at the command prompt.

## M-File Functions

Functions are program routines, usually implemented in M-files, that accept input arguments and return output arguments. They operate on variables within their own workspace. This workspace is separate from the workspace you access at the MATLAB command prompt.

### The Function Workspace

Each M-file function has an area of memory, separate from the MATLAB base workspace, in which it operates. This area, called the function workspace, gives each function its own workspace context.

While using MATLAB, the only variables you can access are those in the calling context, be it the base workspace or that of another function. The variables that you pass to a function must be in the calling context, and the function returns its output arguments to the calling workspace context. You can, however, define variables as global variables explicitly, allowing more than one workspace context to access them.

### Simple Function Example

The average function is a simple M-file that calculates the average of the elements in a vector:

```
function y = average(x)
% AVERAGE Mean of vector elements.
% AVERAGE(X), where X is a vector, is the mean of vector elements.
% Nonvector input results in an error.
[m,n] = size(x);
```

```
if (~(m == 1) | (n == 1)) | (m == 1 & n == 1)
    error('Input must be a vector')
end
y = sum(x)/length(x);      % Actual computation
```

Try entering these commands in an M-file called `average.m`. The `average` function accepts a single input argument and returns a single output argument. To call the `average` function, enter

```
z = 1:99;

average(z)
ans =
    50
```

### Types of Functions

MATLAB provides the following types of functions. Each function type is described in more detail in a later section of this documentation:

- The “Primary M-File Functions” on page 5-14 is the first function in an M-file and typically contains the main program.
- “Subfunctions” on page 5-31 act as subroutines to the main function. You can also use them to define multiple functions within a single M-file.
- “Nested Functions” on page 5-15 are functions defined within another function. They can help to improve the readability of your program and also give you more flexible access to variables in the M-file.
- “Anonymous Functions” on page 5-3 provide a quick way of making a function from any MATLAB expression. You can compose anonymous functions either from within another function or at the MATLAB command prompt.
- “Overloaded Functions” on page 5-34 are useful when you need to create a function that responds to different types of inputs accordingly. They are similar to overloaded functions in any object-oriented language.
- “Private Functions” on page 5-33 give you a way to restrict access to a function. You can call them only from an M-file function in the parent directory.

You might also see the term *function functions* in the documentation. This is not really a separate function type. The term function functions refers to any functions that accept another function as an input argument. You can pass a function to another function using a function handle.

## Identifying Dependencies

Most any program you write will make calls to other functions and scripts. If you need to know what other functions and scripts your program is dependent upon, use one of the techniques described below.

### Simple Display of M-File Dependencies

For a simple display of all M-files referenced by a particular function, follow these steps:

- 1 Type `clear functions` to clear all functions from memory (see Note below).

---

**Note** `clear functions` does not clear functions locked by `mlock`. If you have locked functions (which you can check using `inmem`) unlock them with `munlock`, and then repeat step 1.

---

- 2 Execute the function you want to check. Note that the function arguments you choose to use in this step are important, because you can get different results when calling the same function with different arguments.
- 3 Type `inmem` to display all M-files that were used when the function ran. If you want to see what MEX-files were used as well, specify an additional output:

```
[mfiles, mexfiles] = inmem
```

### Detailed Display of M-File Dependencies

For a much more detailed display of dependent function information, use the `depfun` function. In addition to M-files, `depfun` shows which built-ins and classes a particular function depends on:

```
[list, builtins, classes] = depfun('strtok.m');  
  
list  
list =  
    'D:\matlabR14\toolbox\matlab\strfun\strtok.m'  
    'D:\matlabR14\toolbox\matlab\graphics\newplot.p'  
    'D:\matlabR14\toolbox\matlab\graphics\closereq.m'  
    'D:\matlabR14\toolbox\matlab\ops\@double\any.bi'  
    .  
    .  
    .
```

## Function Arguments

When calling a function, the caller provides the function with any data it needs by passing the data in an argument list. Data that needs to be returned to the caller is passed back in a list of return values.

Semantically speaking, MATLAB always passes argument data by value. (Internally, MATLAB optimizes away any unnecessary copy operations.)

If you pass data to a function that then modifies this data, you will need to update your own copy of the data. You can do this by having the function return the updated value as an output argument.

This section covers

- “Checking the Number of Input Arguments” on page 4-21
- “Passing Variable Numbers of Arguments” on page 4-23
- “Returning Output Arguments” on page 4-25

### Checking the Number of Input Arguments

The `nargin` and `nargout` functions enable you to determine how many input and output arguments a function is called with. You can then use conditional statements to perform different tasks depending on the number of arguments. For example,

```
function c = testarg1(a, b)
if (nargin == 1)
    c = a .^ 2;
elseif (nargin == 2)
    c = a + b;
end
```

Given a single input argument, this function squares the input value. Given two inputs, it adds them together.

Here is a more advanced example that finds the first token in a character string. A *token* is a set of characters delimited by white space or some other character. Given one input, the function assumes a default delimiter of white space; given two, it lets you specify another delimiter if desired. It also allows for two possible output argument lists:

```
function [token, remainder] = strtok(string, delimiters)
% Function requires at least one input argument
if nargin < 1
    error('Not enough input arguments.');
```

end

```
token = []; remainder = [];
len = length(string);
if len == 0
    return
end

% If one input, use white space delimiter
if (nargin == 1)
    delimiters = [9:13 32]; % White space characters
end
i = 1;

% Determine where non-delimiter characters begin
while (any(string(i) == delimiters))
    i = i + 1;
    if (i > len), return, end
end

% Find where token ends
start = i;
while (~any(string(i) == delimiters))
    i = i + 1;
    if (i > len), break, end
end
finish = i - 1;
token = string(start:finish);
```

```

% For two output arguments, count characters after
% first delimiter (remainder)
if (nargout == 2)
    remainder = string(finish+1:end);
end

```

The `strtok` function is a MATLAB M-file in the `strfun` directory.

---

**Note** The order in which output arguments appear in the function declaration line is important. The argument that the function returns in most cases appears first in the list. Additional, optional arguments are appended to the list.

---

## Passing Variable Numbers of Arguments

The `varargin` and `varargout` functions let you pass any number of inputs or return any number of outputs to a function. This section describes how to use these functions and also covers

- “Unpacking `varargin` Contents” on page 4-24
- “Packing `varargout` Contents” on page 4-24
- “`varargin` and `varargout` in Argument Lists” on page 4-25

MATLAB packs all specified input arguments into a *cell array*, a special kind of MATLAB array that consists of cells instead of array elements. Each cell can hold any size or kind of data — one might hold a vector of numeric data, another in the same array might hold an array of string data, and so on. For output arguments, your function code must pack them into a cell array so that MATLAB can return the arguments to the caller.

Here is an example function that accepts any number of two-element vectors and draws a line to connect them:

```

function testvar(varargin)
for k = 1:length(varargin)
    x(k) = varargin{k}(1); % Cell array indexing
    y(k) = varargin{k}(2);
end

```

```
xmin = min(0,min(x));  
ymin = min(0,min(y));  
axis([xmin fix(max(x))+3 ymin fix(max(y))+3])  
plot(x,y)
```

Coded this way, the `testvar` function works with various input lists; for example,

```
testvar([2 3],[1 5],[4 8],[6 5],[4 2],[2 3])  
testvar([-1 0],[3 -5],[4 2],[1 1])
```

### Unpacking varargin Contents

Because `varargin` contains all the input arguments in a cell array, it's necessary to use cell array indexing to extract the data. For example,

```
y(n) = varargin{n}(2);
```

Cell array indexing has two subscript components:

- The indices within curly braces `{}` specify which cell to get the contents of.
- The indices within parentheses `()` specify a particular element of that cell.

In the preceding code, the indexing expression `{i}` accesses the *n*th cell of `varargin`. The expression `(2)` represents the second element of the cell contents.

### Packing varargout Contents

When allowing any number of output arguments, you must pack all of the output into the `varargout` cell array. Use `nargout` to determine how many output arguments the function is called with. For example, this code accepts a two-column input array, where the first column represents a set of *x* coordinates and the second represents *y* coordinates. It breaks the array into separate `[xi yi]` vectors that you can pass into the `testvar` function shown in the earlier example:

```
function [varargout] = testvar2(arrayin)  
for k = 1:nargout  
    varargout{k} = arrayin(k,:);    % Cell array assignment  
end
```

The assignment statement inside the for loop uses cell array assignment syntax. The left side of the statement, the cell array, is indexed using curly braces to indicate that the data goes inside a cell. For complete information on cell array assignment, see “Cell Arrays” on page 2-66.

To call `testvar2`, type

```
a = [1 2; 3 4; 5 6; 7 8; 9 0];

[p1, p2, p3, p4, p5] = testvar2(a)
p1 =
     1     2
p2 =
     3     4
p3 =
     5     6
p4 =
     7     8
p5 =
     9     0
```

### **varargin and varargout in Argument Lists**

`varargin` or `varargout` must appear last in the argument list, following any required input or output variables. That is, the function call must specify the required arguments first. For example, these function declaration lines show the correct placement of `varargin` and `varargout`:

```
function [out1,out2] = example1(a,b,varargin)
function [i,j,varargout] = example2(x1,y1,x2,y2,flag)
```

### **Returning Output Arguments**

Any variables you place to the left of the equals sign in the function definition line are returned to the caller when the function terminates. If you pass any input variables that the function can modify, you will need to include the same variables as output arguments so that the caller receives the updated value.

For example, if the function `readText` shown below reads one line of a file each time it is called, then it must keep track of the offset into the file. But when `readText` terminates, its copy of the offset variable is cleared from memory.

To keep the offset value from being lost, `readText` must return this value to the caller:

```
[text, offset] = readText(filestart, offset)
```

---

**Note** If you include output arguments in your function definition line, make sure that your function code assigns a value to each one.

---

## Function Handles

A *function handle* is a MATLAB value that provides a means of calling a function indirectly. You can pass function handles in calls to other functions (often called *function functions*). You can also store function handles in data structures for later use (for example, as Handle Graphics callbacks).

### Constructing a Function Handle

Use the following syntax to construct a function handle, preceding the name of the function with an @ sign. Use only the function *name*, with no path information, after the @ sign:

```
fhandle = @functionname
```

MATLAB maps the handle to the function you specify and saves this mapping information in the handle. If there is more than one function with this name, MATLAB maps to the one function source it would dispatch to if you were actually calling the function.

A function handle retains that same mapping even if its corresponding function goes out of scope. For example, if, after creating the handle, you change the MATLAB path so that a different function of the same name now takes precedence, invoking the function handle still executes the code to which the handle was originally mapped.

### Handles to Anonymous Functions

Another way to construct a function handle is to create an anonymous function. For example,

```
sqr = @(x) x.^2;
```

creates an anonymous function that computes the square of its input argument *x*. The variable *sqr* contains a handle to the anonymous function. See “Anonymous Functions” on page 5-3 for more information.

### Arrays of Function Handles

To store function handles in an array, use a cell array:

```
trigFun = {@sin, @cos, @tan};  
  
plot(trigFun{2}(-pi:0.01:pi))
```

### Invalid or Obsolete Function Handles

If you create a handle to a function that is not on the MATLAB path, or if you load a handle to a function that is no longer on the path, MATLAB catches the error only when the handle is invoked. You can assign an invalid handle and use it in such operations as `func2str`. MATLAB catches and reports an error only when you attempt to use it in a runtime operation.

### Calling a Function Using Its Handle

To execute a function associated with a function handle, use the syntax shown here, treating the function handle `fhandle` as if it were a function name:

```
fhandle(arg1, arg2, ..., argn)
```

If the function being called takes no input arguments, then use empty parentheses after the function handle name:

```
fhandle()
```

### Simple Function Handle Example

The following example calls a function `plotFHandle`, passing it a handle for the MATLAB `sin` function. `plotFHandle` then calls the `plot` function, passing it some data and the function handle to `sin`. The `plot` function calls the function associated with the handle to compute its *y*-axis values:

```
function x = plotFHandle(fhandle, data)  
plot(data, fhandle(data))
```

Call `plotFHandle` with a handle to the `sin` function and the value shown below:

```
plotFHandle(@sin, -pi:0.01:pi)
```

## Functions That Operate on Function Handles

MATLAB provides the following functions for working with function handles. See the reference pages for these functions for more information.

Function	Description
<code>functions</code>	Return information describing a function handle.
<code>func2str</code>	Construct a function name string from a function handle.
<code>str2func</code>	Construct a function handle from a function name string.
<code>save</code>	Save a function handle from the current workspace to a MAT-file.
<code>load</code>	Load a function handle from a MAT-file into the current workspace.
<code>isa</code>	Determine if a variable contains a function handle.
<code>isequal</code>	Determine if two function handles are handles to the same function.

## Additional Information on Function Handles

This section covers the following topics:

- “Maximum Length of a Function Name” on page 4-29
- “How MATLAB Constructs a Function Handle” on page 4-30
- “Saving and Loading Function Handles” on page 4-30

### Maximum Length of a Function Name

Function names used in handles are unique up to `N` characters, where `N` is the number returned by the function `namelengthmax`. If the function name exceeds that length, MATLAB truncates the latter part of the name.

For function handles created for Java constructors, the length of any segment of the package name or class name must not exceed `namelengthmax` characters. (The term *segment* refers to any portion of the name that lies before, between, or after a dot. For example, `java.lang.String` has three segments). The overall length of the string specifying the package and class has no limit.

### How MATLAB Constructs a Function Handle

At the time you create a function handle, MATLAB maps the handle to one or more implementations of the function specified in the constructor statement:

```
fhandle = @functionname
```

In selecting which function(s) to map to, MATLAB considers

- **Scope** — The function named must be on the MATLAB path at the time the handle is constructed.
- **Precedence** — MATLAB selects which function(s) to map to according to the function precedence rules described under “How MATLAB Determines Which Method to Call” on page 8-67.
- **Overloading** — If additional M-files on the path overload the function for any of the standard MATLAB data types, such as `double` or `char`, then MATLAB maps the handle to these M-files as well.

M-files that overload a function for classes outside of the standard MATLAB data types are not mapped to the function handle at the time it is constructed. Function handles do operate on these types of overloaded functions, but MATLAB determines which implementation to call at the time of evaluation in this case.

### Saving and Loading Function Handles

You can save and load function handles in a MAT-file using the MATLAB `save` and `load` functions. If you load a function handle that you saved in an earlier MATLAB session, the following conditions could cause unexpected behavior:

- Any of the M-files that define the function have been moved, and thus no longer exist on the path stored in the handle.
- You load the function handle into an environment different from that in which it was saved. For example, the source for the function either doesn't exist or is located in a different directory than on the system on which the handle was saved.

In both of these cases, the function handle is now invalid because it no longer maps to any existing function code. Although the handle is invalid, MATLAB still performs the load successfully and without displaying a warning. Attempting to invoke the handle, however, results in an error.

## Calling Functions

This section describes the following topics on how to make function calls:

- “What Happens When You Call a Function” on page 4-31
- “Determining Which Function Is Called” on page 4-31
- “MATLAB Calling Syntax” on page 4-34
- “Passing Certain Argument Types” on page 4-37
- “Passing Arguments in Structures or Cell Arrays” on page 4-39
- “Calling External Functions” on page 4-40

### What Happens When You Call a Function

When you call a function M-file from either the command line or from within another M-file, MATLAB parses the function into pseudocode and stores it in memory. This prevents MATLAB from having to reparse a function each time you call it during a session. The pseudocode remains in memory until you clear it using the `clear` function, or until you quit MATLAB.

### Clearing Functions from Memory

You can use `clear` in any of the following ways to remove functions from the MATLAB workspace.

Syntax	Description
<code>clear functionname</code>	Remove specified function from workspace.
<code>clear functions</code>	Remove all compiled M-functions.
<code>clear all</code>	Remove all variables and functions.

### Determining Which Function Is Called

When more than one function has the same name, which one does MATLAB call? This section explains the process that MATLAB uses to make this decision. It covers the following topics:

- “Function Scope” on page 4-32
- “Precedence Order” on page 4-32

- “Multiple Implementation Types” on page 4-34
- “Querying Which Function MATLAB Will Call” on page 4-34

### **Function Scope**

Any functions you call must first be within the scope of (i.e., visible to) the calling function or your MATLAB session. MATLAB determines if a function is in scope by searching for the function’s executable file according to a certain order (see “Precedence Order” on page 4-32).

One key part of this search order is the MATLAB path. The path is an ordered list of directories that MATLAB defines on startup. You can add or remove any directories you want from the path. MATLAB searches the path for the given function name, starting at the first directory in the path string and continuing until either the function file is found or the list of directories is exhausted. If no function of that name is found, then the function is considered to be out of scope and MATLAB issues an error.

### **Precedence Order**

The function precedence order determines the precedence of one function over another based on the type of function and its location on the MATLAB path. MATLAB selects the correct function for a given context by applying the following function precedence rules in the order given here.

For items 3 through 7 in this list, the file MATLAB searches for can be any of five types: an M- or built-in (.bi) file, prepared M-file (P-Code), compiled C or Fortran file (MEX-file), or Simulink<sup>®</sup> model (MDL-file). See “Multiple Implementation Types” on page 4-34 for more on this.

#### **1 Variable**

Before assuming that a name should match a function, MATLAB checks the current workspace to see if it matches a variable name. If MATLAB finds a match, it stops the search.

#### **2 Subfunction**

Subfunctions take precedence over all other M-file functions and overloaded methods that are on the path and have the same name. Even if the function is called with an argument of type matching that of an overloaded method, MATLAB uses the subfunction and ignores the overloaded method.

### 3 Private function

Private functions are called if there is no subfunction of the same name within the current scope. As with subfunctions, even if the function is called with an argument of type matching that of an overloaded method, MATLAB uses the private function and ignores the overloaded method.

### 4 Class constructor

Constructor functions (functions having names that are the same as the @ directory, for example @polynom/polynom.m) take precedence over other MATLAB functions. Therefore, if you create an M-file called polynom.m and put it on your path before the constructor @polynom/polynom.m version, MATLAB will always call the constructor version.

### 5 Overloaded method

MATLAB calls an overloaded method if it is not superseded by a subfunction or private function. Which overloaded method gets called depends on the classes of the objects passed in the argument list.

### 6 Function in the current directory

A function in the current working directory is selected before one elsewhere on the path.

### 7 Function elsewhere on the path

Finally, a function elsewhere on the path is selected. A function in a directory that is toward the beginning of the path string is given higher precedence.

---

**Note** Because variables have the highest precedence, if you have created a variable of the same name as a function, MATLAB will not be able to run that function until you clear the variable from memory.

---

### Multiple Implementation Types

There are five file precedence types. MATLAB uses file precedence to select between identically named functions in the same directory. The order of precedence for file types is

- 1 Built-in (.bi) file
- 2 MEX-files
- 3 MDL (Simulink<sup>®</sup> model) file
- 4 P-code file
- 5 M-file

For example, if MATLAB finds a P-code and an M-file version of a method in a class directory, then the P-code version is used. It is, therefore, important to regenerate the P-code version whenever you edit the M-file.

### Querying Which Function MATLAB Will Call

You can determine which function MATLAB will call using the `which` command. For example,

```
which pie3
$matlabroot/toolbox/matlab/specgraph/pie3.m
```

However, if `p` is a portfolio object,

```
which pie3(p)
dir_on_your_path/@portfolio/pie3.m      % portfolio method
```

The `which` command determines which version of `pie3` MATLAB will call if you passed a portfolio object as the input argument. To see a list of all versions of a particular function that are on your MATLAB path, use the `-all` option. See the `which` reference page for more information on this command.

### MATLAB Calling Syntax

This section explains how to use the MATLAB command and function syntax:

- “MATLAB Command Syntax” on page 4-35
- “MATLAB Function Syntax” on page 4-35

- “Passing Arguments with Command and Function Syntax” on page 4-36

You can call function M-files from either the MATLAB command line or from within other M-files. Be sure to include all necessary arguments, enclosing input arguments in parentheses and output arguments in square brackets.

---

**Note** Function names are sensitive to case. When you call a function, use the correct combination of upper and lowercase letters so that the name is an exact match. Otherwise, you risk calling a different function that does match but is elsewhere on the path.

---

You often have the choice of using one of two syntaxes for a function call. You can use either a command or a function type of syntax. This is referred to in MATLAB as *command/function duality*.

### **MATLAB Command Syntax**

A function call made in command syntax consists of the function name followed by one or more arguments separated by spaces:

```
functionname in1 in2 ... inN
```

While the command syntax is simpler to write, it has the restriction that you may not assign any return values the function might generate. Attempting to do so generates an error.

Two examples of command syntax are

```
save mydata.mat x y z
clear length width depth
```

In the command syntax, MATLAB treats all arguments as string literals.

### **MATLAB Function Syntax**

Function calls written in the function syntax look essentially the same as those in many other programming languages. One difference is that, in MATLAB, functions can return more than one output value.

A function call with a single return value looks like this:

```
out = functionname(in1, in2, ..., inN)
```

If the function returns more than one value, separate the output variables with commas or spaces, and enclose them all in square brackets ([ ]):

```
[out1, out2, ..., outN] = functionname(in1, in2, ..., inN)
```

Here are two examples:

```
copyfile(srcfile, '..\mytests', 'writable')  
[x1, x2, x3, x4] = deal(A{:})
```

In the function syntax, MATLAB passes arguments to the function by value. See the examples under “Passing Arguments with Command and Function Syntax” on page 4-36.

### Passing Arguments with Command and Function Syntax

When you call a function using function syntax, MATLAB passes the *values* assigned to each variable in the argument list. For example, this expression passes the values assigned to A0, A1, and A2 to the `polyeig` function:

```
e = polyeig(A0, A1, A2)
```

Function calls written in command syntax pass all arguments as string literals. This expression passes the strings 'mydata.mat', 'x', 'y', and 'z' to the `save` function:

```
save mydata.mat x y z
```

The following examples show the difference between passing arguments in the two syntaxes.

**Passing Arguments – Example 1.** Calling `disp` with the function syntax, `disp(A)`, passes the value of variable `A` to the `disp` function:

```
A = pi;  
  
disp(A)                                % Function syntax  
3.1416
```

Calling it with the command syntax, `disp A`, passes the string 'A':

```
A = pi;  
  
disp A                                  % Command syntax  
A
```

**Passing Arguments — Example 2.** Passing two variables representing equal strings to the `strcmp` function using function and command syntaxes gives different results. The function syntax passes the values of the arguments. `strcmp` returns a 1, which means they are equal:

```
str1 = 'one';      str2 = 'one';

strcmp(str1, str2)           % Function syntax
ans =
    1           (equal)
```

The command syntax passes the names of the variables, 'str1' and 'str2', which are unequal:

```
str1 = 'one';      str2 = 'one';

strcmp str1 str2           % Command syntax
ans =
    0           (unequal)
```

## Passing Certain Argument Types

This section explains how to pass the following types of data in a function call:

- “Passing Strings” on page 4-37
- “Passing Filenames” on page 4-38
- “Passing Function Handles” on page 4-38

### Passing Strings

When using the function syntax to pass a string literal to a function, you must enclose the string in single quotes, ('string'). For example, to create a new directory called `myapptests`, use

```
mkdir('myapptests')
```

On the other hand, variables that contain strings do not need to be enclosed in quotes:

```
dirname = 'myapptests';
mkdir(dirname)
```

### Passing Filenames

You can specify a filename argument using the MATLAB command or function syntax. For example, either of the following are acceptable. (The `.mat` file extension is optional for `save` and `load`):

```
load mydata.mat           % Command syntax
load('mydata.mat')       % Function syntax
```

If you assign the output to a variable, you must use the function syntax:

```
savedData = load('mydata.mat')
```

Specify ASCII files as shown here. In this case, the file extension is required:

```
load mydata.dat -ascii    % Command syntax
load('mydata.dat', '-ascii') % Function syntax
```

**Determining Filenames at Run-Time.** There are several ways that your function code can work on specific files without your having to hard-code their filenames into the program. You can

- Pass the filename as an argument:

```
function myfun(datafile)
```

- Prompt for the filename using the input function:

```
filename = input('Enter name of file: ', 's');
```

- Browse for the file using the `uigetfile` function:

```
[filename, pathname] = uigetfile('*.mat', 'Select MAT-file');
```

### Passing Function Handles

The MATLAB function handle has several uses, the most common being a means of immediate access to the function it represents. You can pass function handles in argument lists to other functions, enabling the receiving function to make calls by means of the handle.

To pass a function handle, include its variable name in the argument list of the call:

```
fhandle = @humps;
x = fminbnd(fhandle, 0.3, 1);
```

The receiving function invokes the function being passed using the usual MATLAB calling syntax:

```
function [xf, fval, exitflag, output] = ...
    fminbnd(fhandle, ax, bx, options, varargin)
    .
    .
    .
113 fx = fhandle(x, varargin{:});
```

## Passing Arguments in Structures or Cell Arrays

Instead of requiring an additional argument for every value you want to pass in a function call, you can package them in a MATLAB structure or cell array.

### Passing Arguments in a Structure

Make each input you want to pass a separate field in the structure argument, using descriptive names for the fields. Structures allow you to change the number, contents, or order of the arguments without having to modify the function. They can also be useful when you have a number of functions that need similar information.

This example updates weather statistics from information in the following chart.

City	Temp.	Heat Index	Wind Speed	Wind Chill
Boston	43	32	8	37
Chicago	34	27	3	30
Lincoln	25	17	11	16
Denver	15	-5	9	0
Las Vegas	31	22	4	35
San Francisco	52	47	18	42

The information is stored in structure `W`. The structure has one field for each column of data:

```
W = struct('city', {'Bos', 'Chi', 'Lin', 'Dnv', 'Vgs', 'SFr'}, ...
          'temp', {43, 34, 25, 15, 31, 52}, ...
          'heatix', {32, 27, 17, -5, 22, 47}, ...
          'wspeed', {8, 3, 11, 9, 4, 18}, ...
          'wchill', {37, 30, 16, 0, 35, 42});
```

To update the data base, you can pass the entire structure, or just one field with its associated data. In the call shown here, `W.wchill` is a comma separated list:

```
updateStats(W.wchill);
```

### Passing Arguments in a Cell Array

You can also group arguments into cell arrays. The advantage over structures is that cell arrays are referenced by index, allowing you to loop through a cell array and access each argument passed in or out of the function. The disadvantage is that you don't have field names to describe each variable.

This example passes several attribute-value arguments to the plot function:

```
X = -pi:pi/10:pi;
Y = tan(sin(X)) - sin(tan(X));

C{1,1} = 'LineWidth';           C{2,1} = 2;
C{1,2} = 'MarkerEdgeColor';    C{2,2} = 'k';
C{1,3} = 'MarkerFaceColor';    C{2,3} = 'g';

plot(X, Y, '--rs', C{:})
```

### Calling External Functions

The MATLAB external interface offers a number of ways to run external functions from MATLAB. This includes programs written in C or Fortran, methods invoked on Java or COM (Component Object Model) objects, functions that interface with serial port hardware, and functions stored in shared libraries. The MATLAB External Interfaces documentation describes these various interfaces and how to call these external functions.

# Types of Functions

---

The basic function types that you can work with in MATLAB are

Overview of MATLAB Function Types (p. 5-2)	An introduction to the basic types of functions available with MATLAB
Anonymous Functions (p. 5-3)	Functions defined from a MATLAB expression and without requiring an M-file
Primary M-File Functions (p. 5-14)	The first, and often the main, function in an M-file
Nested Functions (p. 5-15)	Functions defined within the body of another function
Subfunctions (p. 5-31)	Any functions that follow the primary function in an M-file
Private Functions (p. 5-33)	Functions with restricted access, callable only from an M-file function in the parent directory
Overloaded Functions (p. 5-34)	Functions with multiple implementations that respond to different types of inputs accordingly

## Overview of MATLAB Function Types

There are essentially two ways to create a new function in MATLAB: in a command entered at run-time, or in a file saved to permanent storage.

The command-oriented function, called an *anonymous function*, is relatively brief in its content. It consists of a single MATLAB statement that can interact with multiple input and output arguments. The benefit of using anonymous functions is that you do not have to edit and maintain a file for functions that require only a brief definition.

There are several types of functions that are stored in files (called M-files). The most basic of these are *primary functions* and *subfunctions*. Primary functions are visible to other functions outside of their M-file, while subfunctions, generally speaking, are not. That is, you can call a primary function from an anonymous function or from a function defined in a separate M-file, whereas you can call a subfunction only from functions within the same M-file. (See the Description section of the `function_handle` reference page for information on making a subfunction externally visible.)

Two specific types of primary M-file functions are the *private* and *overloaded function*. Private functions are visible only to a limited group of other functions. This type of function can be useful if you want to limit access to a function, or when you choose not to expose the implementation of a function. Overloaded functions act the same way as overloaded functions in most computer languages. You can create multiple implementations of a function so that each responds accordingly to different types of inputs.

The last type of MATLAB function is the *nested function*. Nested functions are not an independent function type; they exist within the body of one of the other types of functions discussed here (with the exception of anonymous functions), and also within other nested functions.

One type of function that is not discussed in this chapter is the MATLAB built-in function. Built-ins are defined only as executables internal to MATLAB. See “Built-In Functions” on page 3-82 for more information.

## Anonymous Functions

Anonymous functions give you a quick means of creating simple functions without having to create M-files each time. You can construct an anonymous function either at the MATLAB command line or in any M-file function or script.

This section covers

- “Constructing an Anonymous Function” on page 5-3
- “Arrays of Anonymous Functions” on page 5-5
- “Variables Used in the Expression” on page 5-8
- “Examples of Anonymous Functions” on page 5-11

### Constructing an Anonymous Function

The syntax for creating an anonymous function from an expression is

```
fhandle = @(arglist) expr
```

Starting from the right of this syntax statement, the term `expr` represents the body of the function: the code that performs the main task your function is to accomplish. This consists of any single, valid MATLAB expression. Next is `arglist`, which is a comma-separated list of input arguments to be passed to the function. These two components are similar to the body and argument list components of any function.

Leading off the entire right side of this statement is an @ sign. The @ sign is the MATLAB operator that constructs a function handle. Creating a function handle for an anonymous function gives you a means of invoking the function. It is also useful when you want to pass your anonymous function in a call to some other function. The @ sign is a required part of an anonymous function definition.

---

**Note** Function handles not only provide access to anonymous functions. You can create a function handle to any MATLAB function. The constructor uses a different syntax: `fhandle = @functionname` (e.g., `fhandle = @sin`). To find out more about function handles, see “Function Handles” on page 4-27.

---

The syntax statement shown above constructs the anonymous function, returns a handle to this function, and stores the value of the handle in variable `fhandle`. You can use this function handle in the same way as any other MATLAB function handle.

### Simple Example

The statement below creates an anonymous function that finds the square of a number. When you call this function, MATLAB assigns the value you pass in to variable `x`, and then uses `x` in the equation `x.^2`:

```
sqr = @(x) x.^2;
```

The `@` operator constructs a function handle for this function, and assigns the handle to the output variable `sqr`. As with any function handle, you execute the function associated with it by specifying the variable that contains the handle, followed by a comma-separated argument list in parentheses. The syntax is

```
fhandle(arg1, arg2, ..., argN)
```

To execute the `sqr` function defined above, type

```
a = sqr(5)
a =
    25
```

Because `sqr` is a function handle, you can pass it in an argument list to other functions. The code shown here passes the `sqr` anonymous function to the MATLAB `quad` function to compute its integral from zero to one:

```
quad(sqr, 0, 1)
ans =
    0.3333
```

### A Two-Input Example

As another example, you could create the following anonymous function that uses two input arguments, `x` and `y`. (The example assumes that variables `A` and `B` are already defined):

```
sumAxB = @(x, y) (A*x + B*y);
```

```
whos sumAxBY
Name          Size          Bytes  Class

sumAxBY       1x1                16  function_handle
```

To call this function, assigning 5 to x and 7 to y, type

```
sumAxBY(5, 7)
```

### Evaluating With No Input Arguments

For anonymous functions that do not take any input arguments, construct the function using empty parentheses for the input argument list:

```
t = @() datestr(now);
```

Also use empty parentheses when invoking the function:

```
t()

ans =
04-Sep-2003 10:17:59
```

You must include the parentheses. If you type the function handle name with no parentheses, MATLAB just identifies the handle; it does not execute the related function:

```
t

t =
@() datestr(now)
```

### Arrays of Anonymous Functions

To store multiple anonymous functions in an array, use a cell array. The example shown here stores three simple anonymous functions in cell array A:

```
A = {@(x)x.^2, @(y)y+10, @(x,y)x.^2+y+10}
A =
    @(x)x.^2    @(y)y+10    @(x,y)x.^2+y+10
```

Execute the first two functions in the cell array by referring to them with the usual cell array syntax, `A{1}` and `A{2}`:

```
A{1}(4) + A{2}(7)
ans =
    33
```

Do the same with the third anonymous function that takes two input arguments:

```
A{3}(4, 7)
ans =
    33
```

### Space Characters in Anonymous Function Elements

Note that while using space characters in the definition of any function can make your code easier to read, spaces in the body of an anonymous function that is defined in a cell array can sometimes be ambiguous to MATLAB. To ensure accurate interpretation of anonymous functions in cell arrays, you can do any of the following:

- Remove all spaces from at least the body (not necessarily the argument list) of each anonymous function:

```
A = {@(x)x.^2, @(y)y+10, @(x, y)x.^2+y+10};
```

- Enclose in parentheses any anonymous functions that include spaces:

```
A = {(@(x)x .^ 2), (@(y) y +10), (@(x, y) x.^2 + y+10)};
```

- Assign each anonymous function to a variable, and use these variable names in creating the cell array:

```
A1 = @(x)x .^ 2; A2 = @(y) y +10; A3 = @(x, y)x.^2 + y+10;
A = {A1, A2, A3};
```

### Outputs from Anonymous Functions

As with other MATLAB functions, the number of outputs returned by an anonymous function depends mainly on how many variables you specify to the left of the equals (=) sign when you call the function.

For example, consider an anonymous function `getPersInfo` that returns a person's address, home phone, business phone, and date of birth, in that order. To get someone's address, you can call the function specifying just one output:

```
address = getPersInfo(name);
```

To get more information, specify more outputs:

```
[address, homePhone, busPhone] = getPersInfo(name);
```

Of course, you cannot specify more outputs than the maximum number generated by the function, which is four in this case.

### Example

The anonymous `getXLSData` function shown here calls the MATLAB `xlsread` function with a preset spreadsheet filename (`records.xls`) and a variable worksheet name (`worksheet`):

```
getXLSData = @(worksheet) xlsread('records.xls', worksheet);
```

The `records.xls` worksheet used in this example contains both numeric and text data. The numeric data is taken from instrument readings, and the text data describes the category that each numeric reading belongs to.

Because the MATLAB `xlsread` function is defined to return up to three values (numeric, text, and raw data), `getXLSData` can also return this same number of values, depending on how many output variables you specify to the left of the equals sign in the call. Call `getXLSData` a first time, specifying only a single (numeric) output, `dNum`:

```
dNum = getXLSData('Week 12');
```

Display the data that is returned using a `for` loop. You have to use generic names (`v1`, `v2`, `v3`) for the categories, due to the fact that the text of the real category names was not returned in the call:

```
for k = 1:length(dNum)
    disp(sprintf('%s    v1: %2.2f    v2: %d    v3: %d', ...
        datestr(clock, 'HH:MM'), dNum(k,1), dNum(k,2), ...
        dNum(k,3)));
end
```

Here is the output from the first call:

```
12:55   v1: 78.42   v2: 32   v3: 37
13:41   v1: 69.73   v2: 27   v3: 30
14:26   v1: 77.65   v2: 17   v3: 16
15:10   v1: 68.19   v2: 22   v3: 35
```

Now try this again, but this time specifying two outputs, numeric (dNum) and text (dTxt):

```
[dNum, dTxt] = getXLSData('Week 12');

for k = 1:length(dNum)
    disp(sprintf('%s   %s: %2.2f   %s: %d   %s: %d', ...
        datestr(clock, 'HH:MM'), dTxt{1}, dNum(k,1), ...
        dTxt{2}, dNum(k,2), dTxt{3}, dNum(k,3)));
end
```

This time, you can display the category names returned from the spreadsheet:

```
12:55   Temp: 78.42   HeatIndex: 32   WindChill: 37
13:41   Temp: 69.73   HeatIndex: 27   WindChill: 30
14:26   Temp: 77.65   HeatIndex: 17   WindChill: 16
15:10   Temp: 68.19   HeatIndex: 22   WindChill: 35
```

## Variables Used in the Expression

Anonymous functions commonly include two types of variables:

- Variables specified in the argument list. These often vary with each function call.
- Variables specified in the body of the expression. MATLAB captures these variables and holds them constant throughout the lifetime of the function handle.

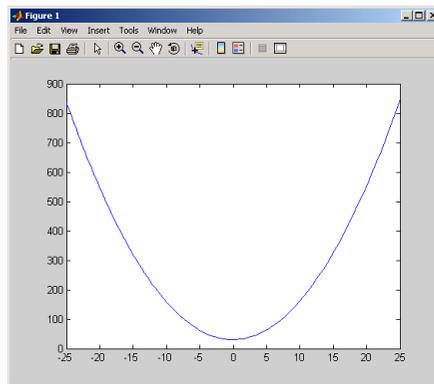
The latter variables must have a value assigned to them at the time you construct an anonymous function that uses them. Upon construction, MATLAB captures the current value for each variable specified in the body of that function. The function will continue to associate this value with the variable even if the value should change in the workspace or go out of scope.

The fact that MATLAB captures the values of these variables when the handle to the anonymous function is constructed enables you to execute an anonymous function from anywhere in the MATLAB environment, even outside the scope in which its variables were originally defined. But it also means that to supply new values for any variables specified within the expression, you must reconstruct the function handle.

### Changing Variables Used in an Anonymous Function

The second statement shown below constructs a function handle for an anonymous function called `parabola` that uses variables `a`, `b`, and `c` in the expression. Passing the function handle to the MATLAB `fplot` function plots it out using the initial values for these variables:

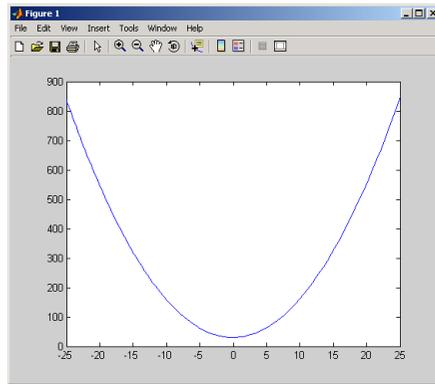
```
a = 1.3;    b = .2;    c = 30;  
parabola = @(x) a*x.^2 + b*x + c;  
fplot(parabola, [-25 25])
```



If you change the three variables in the workspace and replot the figure, the parabola remains unchanged because the `parabola` function is still using the initial values of `a`, `b`, and `c`:

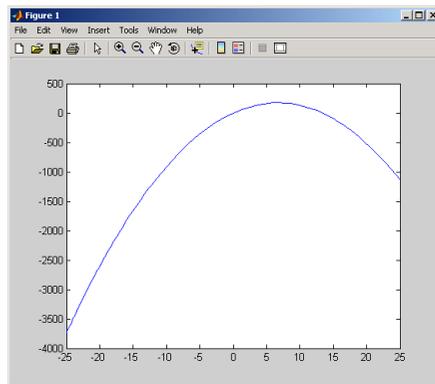
```
a = -3.9;    b = 52;    c = 0;
```

```
fplot(parabola, [-25 25])
```



To get the function to use the new values, you need to reconstruct the function handle, causing MATLAB to capture the updated variables. Replot using the new construct, and this time the parabola takes on the new values:

```
a = -3.9;    b = 52;    c = 0;  
parabola = @(x) a*x.^2 + b*x + c;  
fplot(parabola, [-25 25])
```



For the purposes of this example, there is no need to store the handle to the anonymous function in a variable (parabola, in this case). You can just construct and pass the handle right within the call to fplot. In this way, you update the values of a, b, and c on each call:

```
fplot(@(x) a*x.^2 + b*x + c, [-25 25])
```

## Examples of Anonymous Functions

This section shows a few examples of how you can use anonymous functions. These examples are intended to show you how to program with this type of function. For more mathematically oriented examples, see the MATLAB Mathematics documentation.

The examples in this section include

- “Example 1 — Passing a Function to quad” on page 5-11
- “Example 2 — Multiple Anonymous Functions” on page 5-12

### Example 1 — Passing a Function to quad

The equation shown here has one variable  $t$  that can vary each time you call the function, and two additional variables,  $g$  and  $\omega$ . Leaving these two variables flexible allows you to avoid having to hardcode values for them in the function definition:

$$x = g * \cos(\omega * t)$$

One way to program this equation is to write an M-file function, and then create a function handle for it so that you can pass the function to other functions, such as the MATLAB `quad` function as shown here. However, this requires creating and maintaining a new M-file for a purpose that is likely to be temporary, using a more complex calling syntax when calling `quad`, and passing the  $g$  and  $\omega$  parameters on every call. Here is the function M-file:

```
function f = vOut(t, g, omega)
f = g * cos(omega * t);
```

This code has to specify  $g$  and  $\omega$  on each call:

```
g = 2.5; omega = 10;

quad(@vOut, 0, 7, [], [], g, omega)
ans =
    0.1935

quad(@vOut, -5, 5, [], [], g, omega)
ans =
   -0.1312
```

You can simplify this procedure by setting the values for  $g$  and  $\omega$  just once at the start, constructing a function handle to an anonymous function that only lasts the duration of your MATLAB session, and using a simpler syntax when calling `quad`:

```
g = 2.5; omega = 10;

quad(@(t) (g * cos(omega * t)), 0, 7)
ans =
    0.1935

quad(@(t) (g * cos(omega * t)), -5, 5)
ans =
   -0.1312
```

To preserve an anonymous function from one MATLAB session to the next, save the function handle to a MAT-file

```
save anon.mat f
```

and then load it into the MATLAB workspace in a later session:

```
load anon.mat f
```

### Example 2 — Multiple Anonymous Functions

This example solves the following equation by combining two anonymous functions:

$$g(c) = \int_0^1 (x^2 + cx + 1) dx$$

The equivalent anonymous function for this expression is

```
g = @(c) (quad(@(x) (x.^2 + c*x + 1), 0, 1));
```

This was derived as follows. Take the parenthesized part of the equation (the integrand) and write it as an anonymous function. You don't need to assign the output to a variable as it will only be passed as input to the `quad` function:

```
@(x) (x.^2 + c*x + 1)
```

Next, evaluate this function from zero to one by passing the function handle, shown here as the entire anonymous function, to `quad`:

```
quad(@(x) (x.^2 + c*x + 1), 0, 1)
```

Supply the value for `c` by constructing an anonymous function for the entire equation and you are done:

```
g = @(c) (quad(@(x) (x.^2 + c*x + 1), 0, 1));
```

```
g(2)
```

```
ans =
```

```
2.3333
```

## Primary M-File Functions

The first function in any M-file is called the *primary function*. Following the primary function can be any number of subfunctions, which can serve as subroutines to the primary function.

Under most circumstances, the primary function is the only function in an M-file that you can call from the MATLAB command line or from another M-file function. You invoke this function using the name of the M-file in which it is defined.

For example, the average function shown here resides in the file `average.m`:

```
function y = average(x)
% AVERAGE Mean of vector elements.

y = sum(x)/length(x);      % Actual computation
```

You can invoke this function from the MATLAB command line with this command to find the average of three numbers:

```
average([12 60 42])
```

Note that it is customary to give the primary function the same name as the M-file in which it resides. If the function name differs from the filename, then you must use the filename to invoke the function.

---

## Nested Functions

You can define one or more functions within another function in MATLAB. These inner functions are said to be *nested* within the function that contains them. You can also nest functions within other nested functions.

This section covers the following topics on using nested functions in MATLAB:

- “Writing Nested Functions” on page 5-15
- “Calling Nested Functions” on page 5-17
- “Variable Scope in Nested Functions” on page 5-18
- “Using Function Handles with Nested Functions” on page 5-20
- “Examples of Nested Functions” on page 5-25

### Writing Nested Functions

To write a nested function, simply define one function within the body of another function in an M-file. Like any M-file function, a nested function contains any or all of the components described in “Basic Parts of an M-File” on page 4-8. In addition, you must always terminate a nested function with an end statement:

```
function x = A(p1, p2)
...
    function y = B(p3)
        ...
    end
...
end
```

---

**Note** M-file functions don’t normally require a terminating end statement. This rule does not hold, however, when you nest functions. If an M-file contains one or more nested functions, you must terminate *all* functions (including subfunctions) in the M-file with end, whether or not they contain nested functions.

---

### Example — More Than One Nested Function

This example shows function A and two additional functions nested inside A at the same level:

```
function x = A(p1, p2)
...
    function y = B(p3)
        ...
    end

    function z = C(p4)
        ...
    end
...
end
```

### Example — Multiply Nested Functions

This example shows multiply nested functions, C nested inside B, and B in A:

```
function x = A(p1, p2)
...
    function y = B(p3)
        ...
        function z = C(p4)
            ...
        end
    end
...
end
...
end
```

---

## Calling Nested Functions

You can call a nested function

- From the level immediately above it. (In the following code, function A can call B or D, but not C or E.)
- From a function nested at the same level within the same parent function. (Function B can call D, and D can call B.)
- From a function at any lower level. (Function C can call B or D, but not E.)

```
function A(x, y)           % Primary function
B(x, y);
D(y);

    function B(x, y)       % Nested in A
C(x);
D(y);

        function C(x)     % Nested in B
D(x);
end
end

    function D(x)         % Nested in A
E(x);

        function E       % Nested in E
...
end
end
end
```

You can also call a subfunction from any nested function in the same M-file.

---

**Note** If you construct a function handle for a nested function, you can call the nested function from any MATLAB function that has access to the handle. See “Using Function Handles with Nested Functions” on page 5-20.

---

## Variable Scope in Nested Functions

The scope of a variable is the range of functions that have direct access to the variable to set, modify, or acquire its value. When you define a local (i.e., nonglobal) variable within a function, its scope is normally restricted to that function alone. For example, subfunctions do not share variables with the primary function or with other subfunctions. This is because each function and subfunction stores its variables in its own separate workspace.

Like other functions, a nested function has its own workspace. But it also has access to the workspaces of all functions in which it is nested. So, for example, a variable that has a value assigned to it by the primary function can be read or overwritten by a function nested at any level within the primary. Similarly, a variable that is assigned in a nested function can be read or overwritten by any of the functions containing that function.

In the following two examples, variable `x` is stored in the workspace of the outer `varScope` function and can be read or written to by all functions nested within it.

<pre>function varScope1 x = 5; nestfun1      function nestfun1         nestfun2              function nestfun2                 x = x + 1             end         end     end end</pre>	<pre>function varScope2 nestfun1      function nestfun1         nestfun2              function nestfun2                 x = 5;             end         end     end     x = x + 1 end</pre>
--	--

As a rule, a variable used or defined within a nested function resides in the workspace of the outermost function that both contains the nested function and accesses that variable. The scope of this variable is then the function to which this workspace belongs, and all functions nested to any level within that function.

In the next example, the outer function, `varScope3`, does not access variable `x`. Following the rule just stated, `x` is unknown to the outer function and thus is not shared between the two nested functions. In fact, there are two separate `x` variables in this example: one in the function workspace of `nestfun1` and one in the function workspace of `nestfun2`. When `nestfun2` attempts to update `x`, it fails because `x` does not yet exist in this workspace:

```
function varScope3
  nestfun1
  nestfun2

  function nestfun1
    x = 5;
  end

  function nestfun2
    x = x + 1
  end
end
```

### The Scope of Output Variables

Variables containing values returned by a nested function are not in the scope of outer functions. In the two examples shown here, the one on the left fails in the second to last line because, although the *value* of `y` is returned by the nested function, the *variable* `y` is local to the nested function, and unknown to the outer function. The example on the right assigns the return value to a variable, `z`, and then displays the value of `z` correctly.

Incorrect	Correct
<pre>function varScope4   x = 5;   nestfun;    function y = nestfun     y = x + 1;   end    y end</pre>	<pre>function varScope5   x = 5;   z = nestfun;    function y = nestfun     y = x + 1;   end    z end</pre>

## Using Function Handles with Nested Functions

Every function has a certain *scope*, that is, a certain range of other functions to which it is visible. A function's scope determines which other functions can call it. You can call a function that is out of scope by providing an alternative means of access to it in the form of a function handle. (The function handle, however, must be within the scope of its related function when you construct the handle.) Any function that has access to a function handle can call the function with which the handle is associated.

---

**Note** Although you can call an out of scope function by means of a function handle, the handle itself must be within the scope of its related function at the time it is constructed.

---

The section on “Calling Nested Functions” on page 5-17 defines the scope of a nested function. As with other types of functions, you can make a nested function visible beyond its normal scope with a function handle. The following function `getCubeHandle` constructs a handle for nested function `findCube` and returns its handle, `h`, to the caller. The `@` sign placed before a function name (e.g., `@findCube`) is the MATLAB operator that constructs a handle for that function:

```
function h = getCubeHandle
h = @findCube;           % Function handle constructor

    function cube = findCube(X) % Nested function
        cube = X .^ 3;
    end
end
```

Call `getCubeHandle` to obtain the function handle to the nested function `findCube`. Assign the function handle value returned by `getCubeHandle` to an output variable, `cubeIt` in this case:

```
cubeIt = getCubeHandle;
```

You can now use this variable as a means of calling `findCube` from outside of its M-file:

```
cubeIt(8)
ans =
    512
```

---

**Note** When calling a function by means of its handle, use the same syntax as if you were calling a function directly. But instead of calling the function by its name (e.g., `strcmp(S1, S2)`), use the variable that holds the function handle (e.g., `fhandle(S1, S2)`).

---

## Function Handles and Nested Function Variables

One characteristic of nested functions that makes them different from other MATLAB functions is that they can share nonglobal variables with certain other functions within the same M-file. A nested function `nFun` can share variables with any outer function that contains `nFun`, and with any function nested within `nFun`. This characteristic has an impact on how certain variables are stored when you construct a handle for a nested function.

**Defining Variables When Calling Via Function Handle.** The example below shows a primary function `getHandle` that returns a function handle for the nested function `nestFun`. The `nestFun` function uses three different types of variables. The `VLoc` variable is local to the nested function, `VInp` is passed in when the nested function is called, and `VExt` is defined by the outer function:

```
function h = getHandle(X)
h = @nestFun;
VExt = someFun(X);

function nestFun(VInp)
VLoc = 173.5;
doSomeTask(VInp, VLoc, VExt);
end
end
```

As with any function, when you call `nestFun`, you must ensure that you supply the values for any variables it uses. This is a straightforward matter when calling the nested function directly (that is, calling it from `getHandle`). `VLoc` has a value assigned to it within `nestFun`, `VInp` has its value passed in, and `VExt` acquires its value from the workspace it shares with `getHandle`.

However, when you call `nestFun` using a function handle, only the nested function executes; the outer function, `getHandle`, does not. It might seem at first that the variable `VExt`, otherwise given a value by `getHandle`, has no value assigned to it in the case. What in fact happens though is that MATLAB stores variables such as `VExt` inside the function handle itself when it is being constructed. These variables are available for as long as the handle exists.

The `VExt` variable in this example is considered to be *externally scoped* with respect to the nested function. Externally scoped variables that are used in nested functions for which a function handle exists are stored within the function handle. So, function handles not only contain information about accessing a function. For nested functions, a function handle also stores the values of any externally scoped variables required to execute the function.

### Example Using Externally Scoped Variables

The `sCountFun` and `nCountFun` functions shown below return function handles for subfunction `subCount` and nested function `nestCount`, respectively. These two inner functions store a persistent value in memory (the value is retained in memory between function calls), and then increment this value on every subsequent call. `subCount` makes its count value persistent with an explicit persistent declaration. In `nestCount`, the count variable is externally scoped and thus is maintained in the function handle:

Using a Subfunction	Using a Nested Function
<pre>function h = sCountFun(X) h = @subCount; count = X  subCount(0, count);  function subCount(incr, ini) persistent count; initializing = nargin &gt; 1;  if initializing     count = ini; else     count = count + incr end</pre>	<pre>function h = nCountFun(X) h = @nestCount; count = X  function nestCount(incr)     count = count + incr end  end</pre>

When `sCountFun` executes, it passes the initial value for `count` to the `subCount` subfunction. Keep in mind that the `count` variable in `sCountFun` is not the same as the `count` variable in `subCount`; they are entirely independent of each other. Whenever `subCount` is called via its function handle, the value for `count` comes from its persistent place in memory.

In `nestCount`, the `count` variable again gets its value from the primary function when called from within the M-file. However, in this case the `count` variable in the primary and nested functions are one and the same. When `nestCount` is called by means of its function handle, the value for `count` is assigned from its storage within the function handle.

**Running the Example.** The `subCount` and `nestCount` functions increment a value in memory by another value that you pass as an input argument. Both of these functions give the same results.

Get the function handle to `nestCount`, and initialize the count value to a four-element vector:

```
h = nCountFun([100 200 300 400])
count =
    100    200    300    400
```

Increment the persistent vector by 25, and then by 42:

```
h(25)
count =
    125    225    325    425
```

```
h(42)
count =
    167    267    367    467
```

Now do the same using `sCountFun` and `subCount`, and verify that the results are the same.

---

**Note** If you construct a new function handle to `subCount` or `nestCount`, the former value for `count` is no longer retained in memory. It is replaced by the new value.

---

## Separate Instances of Externally Scoped Variables

The code shown below constructs two separate function handles to the same nested function, `nestCount`, that was used in the last example. It assigns the handles to fields `counter1` and `counter2` of structure `s`. These handles reference different instances of the `nestCount` function. Each handle also maintains its own separate value for the externally scoped count variable.

Call `nCountFun` twice to get two separate function handles to `nestCount`. Initialize the two instances of `count` to two different vectors:

```
s.counter1 = nCountFun([100 200 300 400]);
count =
    100    200    300    400
```

```
s.counter2 = nCountFun([-100 -200 -300 -400]);  
count =  
    -100  -200  -300  -400
```

Now call `nestCount` by means of each function handle to demonstrate that MATLAB increments the two count variables individually.

Increment the first counter:

```
s.counter1(25)  
count =  
    125    225    325    425  
s.counter1(25)  
count =  
    150    250    350    450
```

Now increment the second counter:

```
s.counter2(25)  
count =  
    -75   -175   -275   -375  
s.counter2(25)  
count =  
    -50   -150   -250   -350
```

Go back to the first counter and you can see that it keeps its own value for count:

```
s.counter1(25)  
count =  
    175    275    375    475
```

## Examples of Nested Functions

This section shows a few examples of how you can use nested functions. These examples are intended to show you how to program with this type of function. For more mathematically oriented examples, see the MATLAB Mathematics documentation.

The examples in this section include

- “Example 1 — Creating a Function Handle for a Nested Function” on page 5-26
- “Example 2 — Function-Generating Functions” on page 5-27

### **Example 1 — Creating a Function Handle for a Nested Function**

The following example constructs a function handle for a nested function and then passes the handle to the MATLAB `fplot` function to plot the parabola shape. The `makeParabola` function shown here constructs and returns a function handle `fhandle` for the nested parabola function. This handle gets passed to `fplot`:

```
function fhandle = makeParabola(a, b, c)
% MAKEPARABOLA returns a function handle with parabola
% coefficients.

fhandle = @parabola;          % @ is the function handle constructor

function y = parabola(x)
y = a*x.^2 + b*x + c;
end
end
```

Assign the function handle returned from the call to a variable (`h`) and evaluate the function at points 0 and 25:

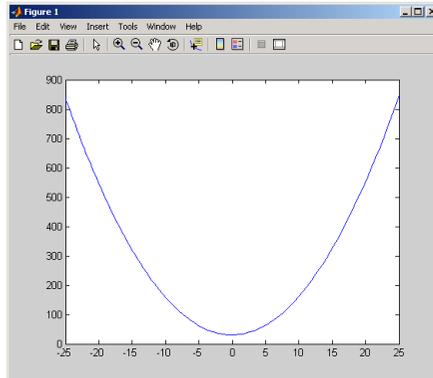
```
h = makeParabola(1.3, .2, 30)
h =
    @makeParabola/parabola

h(0)
ans =
    30

h(25)
ans =
    847.5000
```

Now pass the function handle `h` to the `fplot` function, evaluating the parabolic equation from  $x = -25$  to  $x = +25$ :

```
fplot(h, [-25 25])
```



## Example 2 — Function-Generating Functions

The fact that a function handle separately maintains a unique instance of the function from which it is constructed means that you can generate multiple handles for a function, each operating independently from the others. The function in this example makes IIR filtering functions by constructing function handles from nested functions. Each of these handles maintains its own internal state independent of the others.

The function `makeFilter` takes IIR filter coefficient vectors `a` and `b` and returns a filtering function in the form of a function handle. Each time a new input value  $x_n$  is available, you can call the filtering function to get the new output value  $y_n$ . Each filtering function created by `makeFilter` keeps its own private `a` and `b` vectors, in addition to its own private state vector, in the form of a transposed direct form II delay line:

```
function [filtfcn, statefcn] = makeFilter(b, a)
%   FILTFCN = MAKEFILTER(B, A) creates an IIR filtering
%   function and returns it in the form of a function handle,
%   FILTFCN. Each time you call FILTFCN with a new filter
%   input value, it computes the corresponding new filter
%   output value, updating its internal state vector at the
%   same time.
%
```

```
% [FILTFCN, STATEFCN] = MAKEFILTER(B, A) also returns a
% function (in the form of a function handle, STATEFCN)
% that can return the filter's internal state. The internal
% state vector is in the form of a transposed direct form
% II delay line.

% Initialize state vector. To keep this example a bit simpler,
% assume that a and b have the same length. Also assume that
% a(1) is 1.

v = zeros(size(a));

filtfcn = @iirFilter;
statefcn = @getState;

function yn = iirFilter(xn)
    % Update the state vector
    v(1) = v(2) + b(1) * xn;
    v(2:end-1) = v(3:end) + b(2:end-1) * xn - a(2:end-1) * v(1);
    v(end) = b(end) * xn - a(end) * v(1);

    % Output is the first element of the state vector.
    yn = v(1);
end

function vOut = getState
    vOut = v;
end
end
```

This sample session shows how `makeFilter` works. Make a filter that has a decaying exponential impulse response and then call it a few times in succession to see the output values change:

```
[filt1, state1] = makeFilter([1 0], [1 -.5]);

% First input to the filter is 1.
filt1(1)
ans =
    1
```

```
% Second input to the filter is 0.
filt1(0)
ans =
    0.5000

filt1(0)
ans =
    0.2500

% Show the filter's internal state.
state1()
ans =
    0.2500    0.1250

% Hit the filter with another impulse.
filt1(1)
ans =
    1.1250

% How did the state change?
state1()
ans =
    1.1250    0.5625

% Make an averaging filter.
filt2 = makeFilter([1 1 1]/3, [1 0 0]);

% Put a step input into filt2.
filt2(1)
ans =
    0.3333

filt2(1)
ans =
    0.6667

filt2(1)
ans =
    1
```

```
% The two filter functions can be used independently.
filt1(0)
ans =
    0.5625
```

As an extension of this example, suppose you were looking for a way to develop simulations of different filtering structures and compare them. This might be useful if you were interested in obtaining the range of values taken on by elements of the state vector, and how those values compare with a different filter structure. Here is one way you could capture the filter state at each step and save it for later analysis:

Call `makeFilter` with inputs `v1` and `v2` to construct function handles to the `iirFilter` and `getState` subfunctions:

```
[filtfcn, statefcn] = makeFilter(v1, v2);
```

Call the `iirFilter` and `getState` functions by means of their handles, passing in random values:

```
x = rand(1, 20);
for k = 1:20
    y(k) = filtfcn(x(k));
    states{k} = statefcn(); % Save the state at each step.
end
```

## Subfunctions

M-files can contain code for more than one function. Additional functions within the file are called *subfunctions*, and these are only visible to the primary function or to other subfunctions in the same file.

Each subfunction begins with its own function definition line. The functions immediately follow each other. The various subfunctions can occur in any order, as long as the primary function appears first:

```
function [avg, med] = newstats(u) % Primary function
% NEWSTATS Find mean and median with internal functions.
n = length(u);
avg = mean(u, n);
med = median(u, n);

function a = mean(v, n)           % Subfunction
% Calculate average.
a = sum(v)/n;

function m = median(v, n)        % Subfunction
% Calculate median.
w = sort(v);
if rem(n, 2) == 1
    m = w((n+1) / 2);
else
    m = (w(n/2) + w(n/2+1)) / 2;
end
```

The subfunctions `mean` and `median` calculate the average and median of the input list. The primary function `newstats` determines the length of the list and calls the subfunctions, passing to them the list length `n`.

Subfunctions cannot access variables used by other subfunctions, even within the same M-file, or variables used by the primary function of that M-file, unless you declare them as global within the pertinent functions, or pass them as arguments.

### Calling Subfunctions

When you call a function from within an M-file, MATLAB first checks the file to see if the function is a subfunction. It then checks for a private function (described in the following section) with that name, and then for a standard M-file or built-in function on your search path. Because it checks for a subfunction first, you can override existing M-files using subfunctions with the same name.

### Accessing Help for a Subfunction

You can write help for subfunctions using the same rules that apply to primary functions. To display the help for a subfunction, precede the subfunction name with the name of the M-file that contains the subfunction (minus file extension) and a forward slash.

For example, to get help on subfunction `mysubfun` in file `myfun.m`, type

```
help myfun/mysubfun
```

## Private Functions

Private functions are functions that reside in subdirectories with the special name `private`. These functions are called *private* because they are visible only to M-file functions and M-file scripts that meet these conditions:

- A function that calls a private function must be defined in an M-file that resides in the directory immediately above that `private` subdirectory.
- A script that calls a private function must itself be called from an M-file function that has access to the private function according to the above rule.

For example, assume the directory `newmath` is on the MATLAB search path. A subdirectory of `newmath` called `private` can contain functions that only the functions in `newmath` can call.

Because private functions are invisible outside the parent directory, they can use the same names as functions in other directories. This is useful if you want to create your own version of a particular function while retaining the original in another directory. Because MATLAB looks for private functions before standard M-file functions, it will find a private function named `test.m` before a nonprivate M-file named `test.m`.

Primary functions and subfunctions can also be implemented as private functions.

### Private Directories

You can create your own private directories simply by creating subdirectories called `private` using the standard procedures for creating directories or folders on your computer. Do not place these private directories on your path.

### Accessing Help for a Private Function

You can write help for private functions using the same rules that apply to primary functions. To display the help for a private function, precede the private function name with `private/`.

For example, to get help on private function `myprivfun`, type

```
help private/myprivfun
```

## Overloaded Functions

Overloaded functions are useful when you need to create a function that responds to different types of inputs accordingly. For instance, you might want one of your functions to accept both double-precision and integer input, but to handle each type somewhat differently. You can make this difference invisible to the user by creating two separate functions having the same name, and designating one to handle double types and one to handle integers.

See “Classes and Objects: An Overview” on page 8-2 for more information on overloading functions in MATLAB.

### Class Directories

MATLAB overloaded functions reside in subdirectories having a name starting with the symbol @ and followed by the name of a recognized MATLAB data type. For example, functions in the \@double directory execute when invoked with arguments of MATLAB type double. Those in an \@int32 directory execute when invoked with arguments of MATLAB type int32.

# Data Import and Export

---

Overview (p. 6-3)	Describes the import and export facilities for various data formats.
Using the Import Wizard (p. 6-7)	Shows how to import many types of binary data using this GUI-based interface.
Supported File Formats (p. 6-16)	Lists every file format and file extension supported for MATLAB import and export and the functions used with each type.
Saving and Loading MAT-Files (p. 6-18)	Introduces the MAT-file, a binary data file designed specifically to save and load MATLAB data.
Importing Text Data (p. 6-28)	Describes how to import ASCII text data into MATLAB using the Import Wizard as well as functions.
Exporting Text Data (p. 6-36)	Describes how to export ASCII text data.
Working with Graphics Files (p. 6-41)	Shows you how to import and export images stored in many different types of graphics files.
Working with Audio and Video Data (p. 6-44)	Describes how to import and export audio and video data.
Working with Spreadsheets (p. 6-48)	Explains how to interact with two popular types of spreadsheets.
Working with Scientific Data Formats (p. 6-54)	Shows you how to access data stored in scientific formats such as HDF4, HDF5, CDF, FITS, and multiband.
Importing HDF4 and HDF-EOS Data (p. 6-69)	Shows you how to access data stored in the HDF4 and HDF-EOS scientific formats using the MATLAB low-level command line interface.

Exporting MATLAB Data to an HDF4 File (p. 6-96)	Shows you how to export data from the MATLAB workspace into a file in the HDF4 scientific format using the MATLAB low-level command line interface.
Using Low-Level File I/O Functions (p. 6-105)	Describes how to use the MATLAB low-level file I/O functions, such as <code>fopen</code> , <code>fread</code> , and <code>fwrite</code> .
Exchanging Files over the Internet (p. 6-118)	Describes how to use the MATLAB URL, zip, and e-mail functions to exchange files over the Internet.

## Overview

MATLAB provides many ways to load data from disk files or the clipboard into the workspace, a process called *importing* data, and to save workspace variables to a disk file, a process called *exporting* data. Your choice of which import or export mechanism to use depends mainly on the format of the data being transferred: text, binary, or a standard format such as HDF.

---

**Note** If MATLAB does not support a high-level function that works with a data format, you can use the MATLAB low-level file I/O functions, if you know how the binary data is formatted in the file. See “Using Low-Level File I/O Functions” on page 6-105 for more information.

---

This overview covers the following topics:

- “Text Data” on page 6-3
- “Graphics Files” on page 6-4
- “Audio and Audio/Video Data” on page 6-4
- “Spreadsheets” on page 6-4
- “Scientific Formats” on page 6-4
- “The Internet” on page 6-5
- “Low-Level File I/O” on page 6-5
- “Large Data Sets” on page 6-5
- “Toolboxes for Importing Data” on page 6-6

For more information and examples on importing and exporting data, see Technical Note 1602:

<http://www.mathworks.com/support/tech-notes/1600/1602.html>

### Text Data

In text format, the data values are American Standard Code for Information Interchange (ASCII) codes that represent alphabetic and numeric characters. ASCII text data can be viewed in a text editor. For more information about working with text data in MATLAB, see

- “Importing Text Data” on page 6-28
- “Exporting Text Data” on page 6-36

These sections also describe how to import and export to XML documents.

### **Graphics Files**

MATLAB imports and exports images from many standard graphics file formats, including the Tagged Image File Format (TIFF), Graphics Interchange Format (GIF), Joint Photographic Experts Group (JPEG), and Portable Network Graphics (PNG) formats.

### **Audio and Audio/Video Data**

MATLAB provides functions to enable you to interact with the following types of audio and audio/video files:

- NeXT/SUN SPARCstation sound
- Microsoft WAVE sound
- Audio/Video Interleaved (AVI)
- Windows-compatible sound devices
- Audio player and recorder objects
- Linear audio signals

### **Spreadsheets**

You can use MATLAB to import and export data to the following types of spreadsheets:

- Microsoft Excel spreadsheets
- Lotus 123 spreadsheets

### **Scientific Formats**

Scientific data formats supported by MATLAB are

- Hierarchical Data Format (HDF)
- Common Data Format (CDF)
- Flexible Image Transport System (FITS)

- Band-Interleaved Data

## The Internet

From your MATLAB session, you can

- Send e-mail
- Download from the Internet
- Compress (zip) and uncompress (unzip) files
- Connect to an FTP server to perform remote file operations

## Low-Level File I/O

MATLAB also supports C-style, low-level I/O functions that you can use with any data format. For more information, see “Using Low-Level File I/O Functions” on page 6-105.

## Large Data Sets

An efficient way to read files with large data sets is to read the file in segments and process the data as you go. This method requires significantly less memory than if you were to try reading in the entire file at once. Using the `textscan` function, you can read a specified amount of data from a file, and maintain a pointer to the location in the file where your last read operation ended and your next read is to begin.

This example opens a large data file and reads the file a segment at a time in a for loop. The code calls `textscan` to read a particular pattern of data (as specified by format) 10,000 times for each segment. Following each read, the subfunction `process_data` processes the data collected in cell array `segarray`:

```
format = '%s %n %s %8.2f %8.2f %8.2f %8.2f %u8';
file_id = fopen('largefile.dat', 'r');

for k = 1:segcount
    segarray = textscan(file_id, format, 10000);
    process_data(segarray);
end

fclose(file_id);
```

### **Toolboxes for Importing Data**

In addition to MATLAB import functions, you can perform specialized import features using toolboxes. For example, use the Database Toolbox for importing data from relational databases. Refer to the documentation on the specific toolbox to see what import features are offered.

## Using the Import Wizard

The easiest way to import data into MATLAB is to use the Import Wizard. You do not need to know the format of the data to use this tool. You simply specify the file that contains the data and the Import Wizard processes the file contents automatically.

This section discusses how to import these types of data using the wizard:

- “Using the Import Wizard with Text Data” on page 6-7
- “Using the Import Wizard with Binary Data” on page 6-13

You can also use the Import Wizard to import HDF data. See “Using the HDF Import Tool” on page 6-69 for more information.

### Using the Import Wizard with Text Data

To import text data using the Import Wizard, perform these steps:

- 1** Start the Import Wizard by selecting the **Import Data** option on the MATLAB **File** menu. MATLAB displays a file selection dialog box. You can also use the `uiimport` function to start the Import Wizard.

To use the Import Wizard to import data from the clipboard, select the **Paste Special** option on the MATLAB **Edit** menu. Skip to step 3 to continue importing from the clipboard.

- 2** Specify the file you want to import in the file selection dialog box and click **Open**. The Import Wizard opens the file and attempts to process its contents.
- 3** Specify the character used to separate the individual data items. This character is called the delimiter or column separator. In many cases, the Import Wizard determines the delimiter automatically. However, you might need to specify the character used in your text file. For more information, see “Specifying the Delimiter” on page 6-8. Once the Import Wizard has correctly processed the data, click **Next**.

- 4 Select the variables that you want to import. By default, the Import Wizard puts all the numeric data in one variable and all the text data in other variables, but you can choose other options. For more information, see “Selecting the Variables to Import” on page 6-10.
- 5 Click **Finish** to import the data into the workspace.

### Specifying the Delimiter

When the Import Wizard opens a text file, or copies data from the clipboard, it displays a portion of the raw data in the preview pane of the dialog box. You can use this display to verify that the file contains the data you expected.

The Import Wizard also attempts to process the data, identifying the delimiter used in the data. The Import Wizard displays the variables it has created based on its interpretation of the delimiter, using tabbed panels to display multiple variables.

For example, in the following figure, the Import Wizard has opened this sample file, `grades.txt`:

John	85	90	95
Ann	90	92	98
Martin	100	95	97
Rob	77	86	93



**Specifying Other Delimiters.** If the Import Wizard cannot determine the delimiter used in the data, it displays a preview of the raw data, as before, but the variables it displays are not correct. If your data uses a character other than a comma, space, tab, or semicolon as a delimiter, you must specify it by clicking the **Other** button and entering the character in the text box. The Import Wizard immediately reprocesses the data, displaying the new variables it creates.

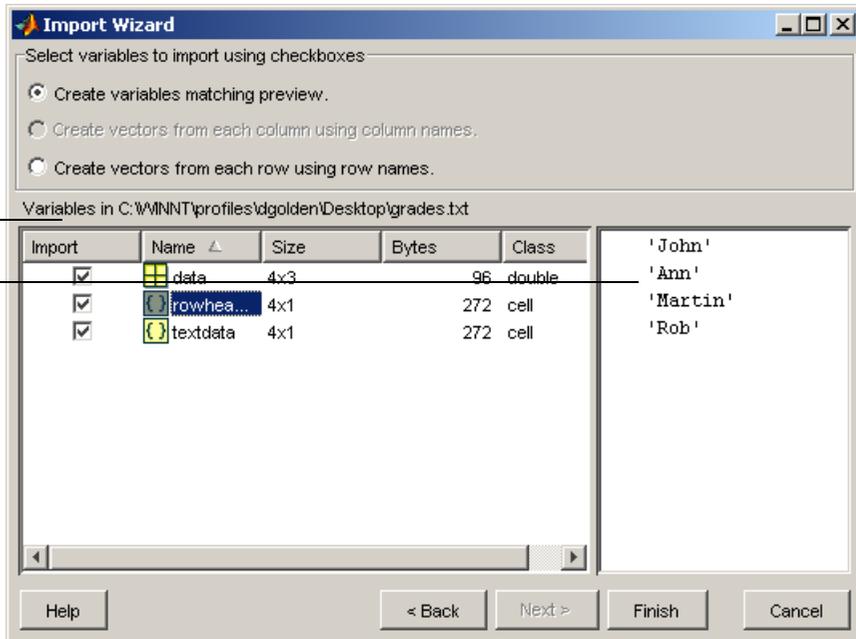
### Selecting the Variables to Import

The Import Wizard displays a list of the variables it has created from your data. To select a variable to import, click in the check box next to its name. By default, all variables are selected.

The Import Wizard displays the contents of the variable that is highlighted in the list in the right pane of the dialog box. To view the contents of one of the other variables, click it. Choose the variables you want to import and click **Next**.

List of variables to be imported.

Import Wizard displays content of the variable highlighted in the list.



**Changing the Variable Selection.** By default, the Import Wizard puts all the numeric data in the file into one variable. If the file contains text data, the Import Wizard puts it in a separate variable. If the file contains row or column headers, the Import Wizard puts them in separate variables, called rowheaders or colheaders, respectively.

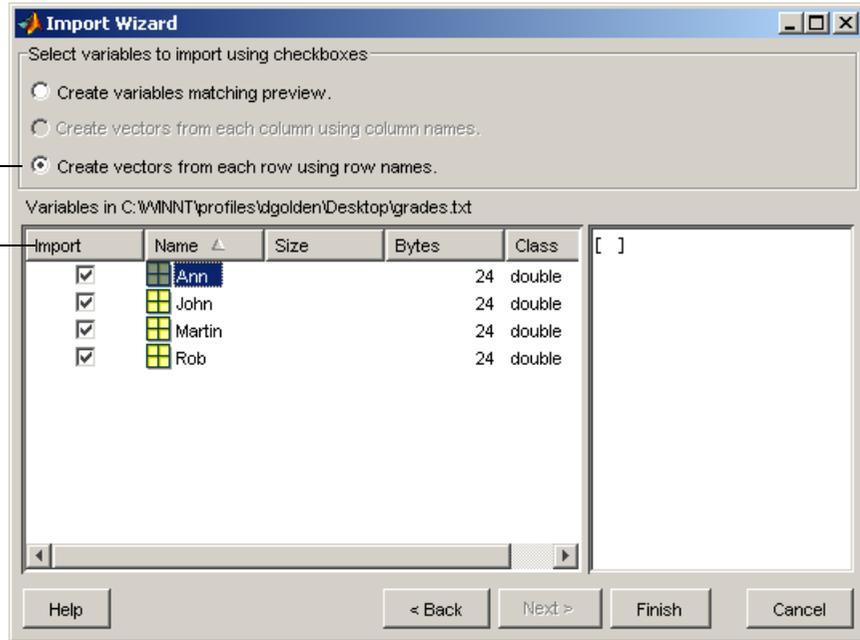
In some cases, it might be more convenient to create a vector from each row or column of data and use the row header or column header text as the name of each variable. To do this, click the appropriate button from the list of buttons at the top of the dialog box.

- Create variables matching preview.
- Create vectors from each column using column names.
- Create vectors from each row using row names.

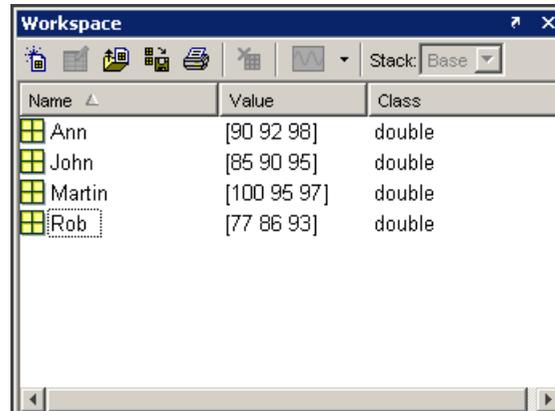
For example, it eases calculation of the student averages if you create a separate vector for each student that contains that student's grades. To create these variables, click the **Create vectors from each row using row names** button. When you click this option, the Import Wizard reprocesses the file, creating these new vectors.

Select this option to create vectors from row header.

List contains variables by row header.



When you are satisfied with the list of vectors to be imported, click **Next** to bring the data into the MATLAB workspace. This button also dismisses the Import Wizard. The MATLAB workspace now contains the vectors.



## Using the Import Wizard with Binary Data

To import binary data using the Import Wizard, perform these steps:

- 1 Start the Import Wizard by selecting the **Import Data** option on the MATLAB **File** menu. MATLAB displays a file selection dialog box. You can also use the `uiimport` function to start the Import Wizard.

To use the Import Wizard to import data from the clipboard, select the **Paste Special** option on the MATLAB **Edit** menu. You can also right-click in the MATLAB Command Window and choose **Paste Special** from the context menu. Skip to step 3 to continue importing from the clipboard.

- 2 Specify the file you want to import in the file selection dialog box and click **Open**. The Import Wizard opens the file and attempts to process its contents. See “Viewing the Variables” below for more information.
- 3 Select the variables that you want to import. (See “Selecting the Variables to Import” on page 6-10) By default, the Import Wizard creates variables depending on the type of data in the file.
- 4 Click **Finish** to import the data into the workspace.

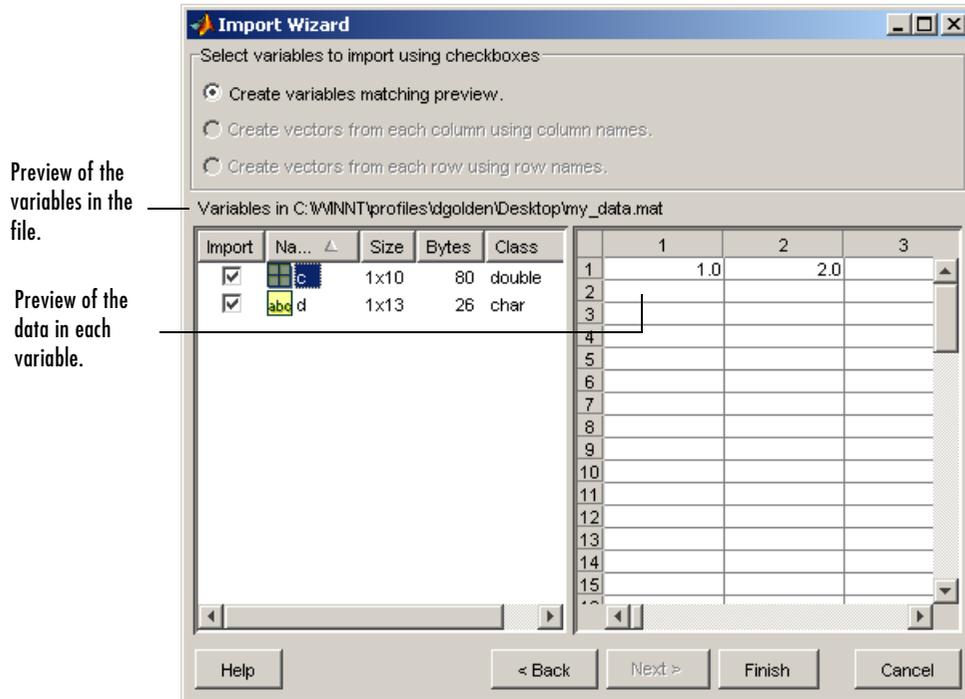
### Viewing the Variables

When the Import Wizard opens a binary data file, it attempts to process the data in the file, creating variables from the data it finds in the file.

For example, if you use the Import Wizard to import this sample MAT-file, `my_data.mat`,

```
C =  
    1  2  3  4  5  
    6  7  8  9 10  
  
D =  
    a test string
```

it creates two variables, listed in the preview pane. To select a variable to import, click in the check box next to its name. All variables are preselected by default.

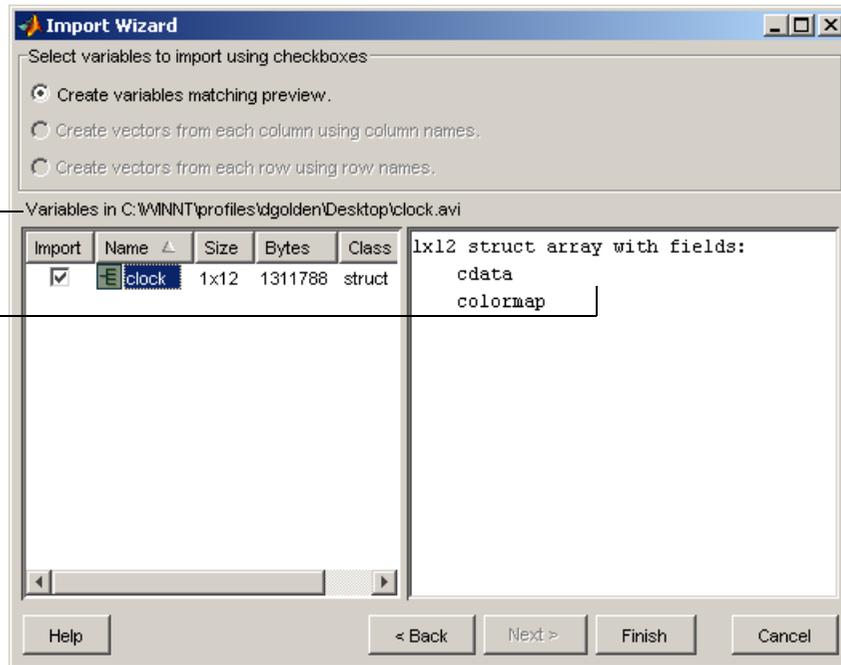


For other binary data types, such as images and sound files, the Import Wizard displays information about the data in the left pane and provides a preview button in the right pane of the dialog box. Click the preview button to view (or listen to) the data.

For example, when used to import a movie in Audio/Video Interleaved (AVI) format, the Import Wizard displays this dialog box.

Preview of the variables in the file.

Information about the data.



## Supported File Formats

The table below shows the file formats that you can read or write from MATLAB along with the functions that support each format.

<b>File Format</b>	<b>File Content</b>	<b>Extension</b>	<b>Functions</b>
MATLAB formatted	Saved MATLAB workspace	.mat	load, save
Text	Text	any	textscan
	Text	any	textread
	Delimited text	any	dlmread, dlmwrite
	Comma-separated numbers	.csv	csvread, csvwrite
Extended Markup Language	XML-formatted text	.xml	xmlread, xmlwrite
Audio	NeXT/SUN sound	.au	auread, auwrite
	Microsoft WAVE sound	.wav	wavread, wavwrite
Movie	Audio/video	.avi	aviread
Scientific data	Data in Common Data Format	.cdf	cdfread, cdfwrite
	Flexible Image Transport System data	.fits	fitsread
	Data in Hierarchical Data Format	.hdf	hdfread

<b>File Format</b>	<b>File Content</b>	<b>Extension</b>	<b>Functions</b>
Spreadsheet	Excel worksheet	.xls	xlsread, xlswrite
	Lotus 123 worksheet	.wk1	wk1read, wk1write
Graphics	TIFF image	.tiff	imread, imwrite
	PNG image	.png	same
	HDF image	.hdf	same
	BMP image	.bmp	same
	JPEG image	.jpeg	same
	GIF image	.gif	same
	PCX image	.pcx	same
	XWD image	.xwd	same
	Cursor image	.cur	same
	Icon image	.ico	same

## Saving and Loading MAT-Files

This section explains how to save the variables in your MATLAB session to a binary file called a MAT-file, and how to load them back into your MATLAB workspace. It covers the following:

- “Exporting Data to MAT-Files” on page 6-18
- “Importing Data from MAT-Files” on page 6-24

MAT-files are double-precision, binary, MATLAB format files. They can be created on one machine and later read by MATLAB on another machine with a different floating-point format, retaining as much accuracy and range as the different formats allow. They can also be manipulated by other programs external to MATLAB.

### Exporting Data to MAT-Files

This section covers

- “Using the save Function” on page 6-18
- “Saving Structures” on page 6-19
- “Appending to an Existing File” on page 6-20
- “Data Compression” on page 6-20
- “Unicode Character Encoding” on page 6-22
- “Optional Output Formats” on page 6-22
- “Storage Requirements” on page 6-23
- “Saving from External Programs” on page 6-24

### Using the save Function

To export workspace variables to a binary or ASCII file, use the save function. You can save all variables from the workspace in a single operation (if you omit the filename, MATLAB uses the name `matlab.mat`):

```
save filename
```

or save just those variables that you specify:

```
save filename var1 var2 ... varN
```

Use the wildcard character (\*) in the variable name to save those variables that match a specific pattern. For example, the following command saves all variables that start with `str`.

```
save strinfo str*
```

Use `whos -file` to examine what has been written to the MAT-file:

```
whos -file strinfo
Name              Size              Bytes  Class

str2              1x15              30    char array
strarray         2x5              678    cell array
strlen           1x1              8     double array
```

### Saving Structures

When saving a MATLAB structure, you have the option of saving the entire structure, saving each structure field as an individual variable in the MAT-file, or saving specific fields as individual variables.

For structure `S`,

```
S.a = 12.7; S.b = {'abc', [4 5; 6 7]}; S.c = 'Hello!';
```

Save the entire structure to `newstruct.mat` with the usual syntax:

```
save newstruct.mat S;
```

```
whos -file newstruct
Name      Size              Bytes  Class

S         1x1              550    struct array
```

Save the fields individually with the `-struct` option:

```
save newstruct.mat -struct S;
```

```
whos -file newstruct
Name      Size              Bytes  Class

a         1x1              8     double array
b         1x2             158    cell array
c         1x6              12    char array
```

Or save only selected fields using `-struct` and specifying each field name:

```
save newstruct.mat -struct S a c;
```

```
whos -file newstruct
Name          Size          Bytes  Class

a              1x1              8  double array
c              1x6             12  char array
```

### Appending to an Existing File

You can add new variables to those already stored in an existing MAT-file by using `save -append`. When you append to a MAT-file, MATLAB first looks in the designated file for each variable name specified in the argument list, or for all variables if no specific variable names are specified. Based on that information, MATLAB does both of the following:

- For each variable that already exists in the MAT-file, MATLAB overwrites its saved value with the new value taken from the workspace.
- For each variable not found in the MAT-file, MATLAB adds that variable to the file and stores its value from the workspace.

---

**Note** Saving with the `-append` option does not append additional elements to any arrays that are already saved in the MAT-file.

---

### Data Compression

MATLAB compresses the data that you save to a MAT-file. Data compression can save you a significant amount of storage space when you are working with large files or working over a network.

Data compression is optional, however, and you can disable it either for an individual save operation, or for all of your MATLAB sessions. Use the `-v6` option with the `save` function to turn off compression on a per-command basis:

```
save filename -v6
```

To disable data compression for all of your MATLAB sessions, open the **Preferences** dialog, select **General** and then **MAT-Files**, and click **Save uncompressed**. See “General Preferences for MATLAB” in the Desktop Tools and Development Environment documentation for more information.

---

**Note** You cannot read a compressed MAT-file with MATLAB Version 6 or 6.5. To write a MAT-file that you will be able to read with one of these versions, save to the file with data compression disabled.

---

Information returned by the command `whos -file` is independent of whether the variables in that file are compressed or not. The byte counts returned by this command represent the number of bytes data occupies in the MATLAB workspace, and not in the file the data was saved to.

**Evaluating When to Compress.** In general, data compression and decompression does slow down save and load operations to some extent. In most cases, however, the resulting reduction in file size is worth any additional time spent on performing the compression or decompression. Both of these factors, time and file size, are highly dependent on the type of data being saved and the size of the data set.

Concerning the type of data being saved, random data does not compress well and thus takes more time to store in a compressed format. Data that has more constant value takes less time to compress the same number of bytes.

Concerning data set size, the impact of data compression is very small for users who do not have large data sets. Users with large data sets need to be aware of the trade-off between compression and speed. Such users should override the default behavior, if necessary.

---

**Note** Compression and decompression during save and load is done transparently without the use of temporary files on disk. This is of significance to large dataset users in particular.

---

## Unicode Character Encoding

MATLAB saves character data to a MAT-file using Unicode character data encoding. As with data compression, Unicode encoding is optional. If you disable it, MATLAB writes the MAT-file using the default character set for your system. To disable Unicode character encoding on a per-command basis, use the `-v6` option with the `save` function:

```
save filename -v6
```

To disable Unicode encoding for all of your MATLAB sessions, open the **Preferences** dialog, select **General** and then **MAT-Files**, and click **Local Character Set**. See “General Preferences for MATLAB” in the Desktop Tools and Development Environment documentation for more information.

---

**Note** You cannot read a Unicode encoded MAT-file with MATLAB Version 6 or 6.5. To write a MAT-file that you will be able to read with one of these versions, save to the file with Unicode character encoding disabled.

---

For more information on how MATLAB saves specific ASCII data formats, and on preventing loss or corruption of character data, see “Writing Character Data” in the MATLAB External Interfaces documentation.

## Optional Output Formats

You can choose from any of the following formats for your output file. If you do not specify a format, MATLAB uses the binary MAT-file format.

Output File Format	Command
Binary MAT-file (default)	<code>save filename</code>
8-digit ASCII	<code>save filename -ascii</code>
8-digit ASCII, tab delimited	<code>save filename -ascii -tabs</code>
16-digit ASCII	<code>save filename -ascii -double</code>
16-digit ASCII, tab delimited	<code>save filename -ascii -double -tabs</code>
MATLAB Version 4 compatible	<code>save filename -v4</code>

**Saving in ASCII Format.** When saving in any of the ASCII formats, consider the following:

- Each variable to be saved must be either a two-dimensional double array or a two-dimensional character array. Saving a complex double array causes the imaginary part of the data to be lost, as MATLAB cannot load nonnumeric data ('i').
- To read the file with the MATLAB load function, make sure all the variables have the same number of columns. If you are using a program other than MATLAB to read the saved data, this restriction can be relaxed.
- Each MATLAB character in a character array is converted to a floating-point number equal to its internal ASCII code and written out as a floating-point number string. There is no information in the saved file that indicates whether the value was originally a number or a character.
- The values of all variables saved merge into a single variable that takes the name of the ASCII file (minus any extension). Therefore, it is advisable to save only one variable at a time.

**Saving in Version 4 Format.** With the -v4 option, you can save only those data constructs that are compatible with MATLAB Version 4. Therefore, you cannot save structures, cell arrays, multidimensional arrays, or objects. Variable names cannot exceed 19 characters in length. In addition, you must use filenames that are supported by MATLAB Version 4.

### Storage Requirements

The binary formats used to save depend on the size and type of each array. Arrays with any noninteger entries and arrays with 10,000 or fewer elements are saved in floating-point formats requiring 8 bytes per real element. Arrays with all integer entries and more than 10,000 elements are saved in the formats shown, requiring fewer bytes per element.

Element Range	Bytes per Element
0 to 255	1
0 to 65535	2
-32767 to 32767	2

$-2^{31}$ to $2^{31}-1$	4
Other	8

### **Saving from External Programs**

The MATLAB External Interfaces documentation provides details on reading and writing MAT-files from external C or Fortran programs. It is important to use recommended access methods, rather than rely upon the specific MAT-file format, which is likely to change in the future.

### **Importing Data from MAT-Files**

This section covers

- “Using the load Function” on page 6-24
- “Previewing MAT-File Contents” on page 6-25
- “Loading into a Structure” on page 6-25
- “Loading Binary Data” on page 6-26
- “Loading ASCII Data” on page 6-26

### **Using the load Function**

To import variables from a binary or ASCII file on your disk to your workspace, use the load function. You can load all variables from the workspace in a single operation (if you omit the filename, MATLAB loads from file `matlab.mat`):

```
load filename
```

or load just those variables that you specify:

```
load filename var1 var2 ... varN
```

Use the wildcard character (\*) in the variable name to load those variables that match a specific pattern. (This works for MAT-files only.) For example, the following command loads all variables that start with `str` from file `strinfo.mat`:

```
load strinfo str*
```

---

**Caution** When you import data into the MATLAB workspace, it overwrites any existing variable in the workspace with the same name.

---

### Previewing MAT-File Contents

To see what variables are stored in a MAT-file before actually loading the file into your workspace, use `whos -file filename`. This command returns the name, dimensions, size, and data type of all variables in the specified MAT-file.

You can use `whos -file` on binary MAT-files only:

```
whos -file mydata.mat
  Name          Size          Bytes  Class

  jsonArray     10x1                java.lang.Double[][]
  spArray        5x5                 84  double array (sparse)
  strArray       2x5                 678  cell array
  x              3x2x2              96  double array
  y              4x5                1230  cell array
```

### Loading into a Structure

To load MAT-file data into a MATLAB structure, specify an output variable in your load command. This example reads the data in `mydata.mat` into the fields of structure `S`:

```
S = load('mydata.mat')
S =
    x: [3x2x2 double]
    y: {4x5 cell}
  spArray: [5x5 double]
  strArray: {2x5 cell}
  jsonArray: [10x1 java.lang.Double[][]]
```

```
whos S
  Name          Size          Bytes  Class

  S             1x1                2840  struct array
```

### **Loading Binary Data**

MAT-files are double-precision binary MATLAB format files created by the `save` function and readable by the `load` function. They can be created on one machine and later read by MATLAB on another machine with a different floating-point format, retaining as much accuracy and range as the different formats allow. They can also be manipulated by other programs, external to MATLAB.

MAT-files can contain data in an uncompressed or a compressed form, or both. MATLAB knows which variables in the file have been compressed by looking at a tag that it attaches to each variable during the save operation. When loading data from a MAT-file into the workspace, MATLAB automatically handles the decompression of the appropriate data.

The External Interface libraries contain C- and Fortran-callable routines to read and write MAT-files from external programs.

### **Loading ASCII Data**

ASCII files must be organized as a rectangular table of numbers, with each number in a row separated by a blank or tab character, and with an equal number of elements in each row. MATLAB generates an error if the number of values differs between any two rows. ASCII files can contain MATLAB comments (lines that begin with `%`).

MATLAB returns all the data in the file as a single two-dimensional array of type `double`. The number of rows in the array is equal to the number of lines in the file, and the number of columns is equal to the number of values on a line.

In the workspace, MATLAB assigns the array to a variable named after the file being loaded (minus any file extension). For example, the command

```
load mydata.dat
```

reads all of the data from `mydata.dat` into the MATLAB workspace as a single array, and assigns it to a variable called `mydata`. In naming the variable, `load` precedes any leading underscores or digits in `filename` with an `X` and replaces any other nonalphanumeric characters with underscores.

For example, the command

```
load 10-May-data.dat
```

assigns the data in file 10-May-data.dat to a new workspace variable called X10\_May\_data.

## Importing Text Data

The section describes various ways to import text data into MATLAB. It covers these topics:

- “The MATLAB Import Wizard” on page 6-28
- “Using Import Functions with Text Data” on page 6-28
- “Importing Numeric Text Data” on page 6-31
- “Importing Delimited ASCII Data Files” on page 6-32
- “Importing Numeric Data with Text Headers” on page 6-32
- “Importing Mixed Alphabetic and Numeric Data” on page 6-33
- “Importing from XML Documents” on page 6-35

---

**Caution** When you import data into the MATLAB workspace, you overwrite any existing variable in the workspace with the same name.

---

### The MATLAB Import Wizard

The easiest way to import data into MATLAB is to use the Import Wizard. You do not need to know the format of the data to use this tool. You simply specify the file that contains the data and the Import Wizard processes the file contents automatically.

For more information, see “Using the Import Wizard with Text Data” on page 6-7.

### Using Import Functions with Text Data

To import text data from the command line or in an M-file, you must use one of the MATLAB import functions. Your choice of function depends on how the data in the text file is formatted.

The text data must be formatted in a uniform pattern of rows and columns, using a text character, called a *delimiter* or *column separator*, to separate each data item. The delimiter can be a space, comma, semicolon, tab, or any other character. The individual data items can be alphabetic or numeric characters or a mix of both.

The text file can also contain one or more lines of text, called *header lines*, or can use text headers to label each column or row. The following example illustrates a tab-delimited text file with header text and row and column headers.

Text header line	_____			
		Class Grades for Spring Term		
Column headers	_____	Grade1	Grade2	Grade3
	John	85	90	95
Row headers	_____	Ann	90	92
	Martin	100	95	97
	Rob	77	86	93
Tab-delimited data	_____			

To find out how your data is formatted, view it in a text editor. After you determine the format, find the sample in the table below that most closely resembles the format of your data. Then read the topic referred to in the table for information on how to import that format.

**Table 6-1: ASCII Data File Formats**

Data Format Sample	File Extension	Description
1 2 3 4 5 6 7 8 9 10	.txt .dat or other	See “Importing Numeric Text Data” on page 6-31 or “Using the Import Wizard with Text Data” on page 6-7 for information.
1; 2; 3; 4; 5 6; 7; 8; 9; 10 or 1, 2, 3, 4, 5 6, 7, 8, 9, 10	.txt .dat .csv or other	See “Importing Delimited ASCII Data Files” on page 6-32 or “Using the Import Wizard with Text Data” on page 6-7 for information.

**Table 6-1: ASCII Data File Formats (Continued)**

<b>Data Format Sample</b>	<b>File Extension</b>	<b>Description</b>
Ann Type1 12.34 45 Yes Joe Type2 45.67 67 No	.txt .dat or other	See “Importing Numeric Data with Text Headers” on page 6-32 for information.
Grade1 Grade2 Grade3 91.5 89.2 77.3 88.0 67.8 91.0 67.3 78.1 92.5	.txt .dat or other	See “Importing Numeric Data with Text Headers” on page 6-32 or “Using the Import Wizard with Text Data” on page 6-7 for information.

If you are familiar with MATLAB import functions but are not sure when to use them, see the following table, which compares the features of each function.

**Table 6-2: ASCII Data Import Function Features**

<b>Function</b>	<b>Data Type</b>	<b>Delimiters</b>	<b>Number of Return Values</b>	<b>Notes</b>
csvread	Numeric data	Commas only	One	Primarily used with spreadsheet data. See “Working with Spreadsheets” on page 6-48.
dlmread	Numeric data	Any character	One	Flexible and easy to use.
fscanf	Alphabetic and numeric; however, both types returned in a single return variable	Any character	One	Part of low-level file I/O routines. Requires use of fopen to obtain file identifier and fclose after read.
load	Numeric data	Spaces only	One	Easy to use. Use the functional form of load to specify the name of the output variable.

**Table 6-2: ASCII Data Import Function Features (Continued)**

Function	Data Type	Delimiters	Number of Return Values	Notes
textread	Alphabetic and numeric	Any character	Multiple values in cell arrays	Flexible, powerful, and easy to use. Use format string to specify conversions.
textscan	Alphabetic and numeric	Any character	Multiple values returned to one cell array	More flexible than textread. Also more format options.

## Importing Numeric Text Data

If your data file contains only numeric data, you can use many of the MATLAB import functions (listed in Table 6-2), depending on how the data is delimited. If the data is rectangular, that is, each row has the same number of elements, the simplest command to use is the `load` command. (The `load` function can also be used to import MAT-files, the MATLAB binary format for saving the workspace.)

For example, the file named `my_data.txt` contains two rows of numbers delimited by space characters:

```
1 2 3 4 5
6 7 8 9 10
```

When you use `load` as a command, it imports the data and creates a variable in the workspace with the same name as the filename, minus the file extension:

```
load my_data.txt;
whos
  Name           Size           Bytes    Class

  my_data        2x5             80      double array

my_data

my_data =
     1     2     3     4     5
     6     7     8     9    10
```

If you want to name the workspace variable something other than the file name, use the functional form of `load`. In the following example, the data from `my_data.txt` is loaded into the workspace variable `A`:

```
A = load('my_data.txt');
```

## Importing Delimited ASCII Data Files

If your data file uses a character other than a space as a delimiter, you have a choice of several import functions you can use. (See Table 6-2 for a complete list.) The simplest to use is the `d1mread` function.

For example, consider a file named `ph.dat` whose contents are separated by semicolons:

```
7.2;8.5;6.2;6.6  
5.4;9.2;8.1;7.2
```

To read the entire contents of this file into an array named `A`, enter

```
A = d1mread('ph.dat', ';');
```

You specify the delimiter used in the data file as the second argument to `d1mread`. Note that, even though the last items in each row are not followed by a delimiter, `d1mread` can still process the file correctly. `d1mread` ignores space characters between data elements. So, the preceding `d1mread` command works even if the contents of `ph.dat` are

```
7.2; 8.5;      6.2;6.6  
5.4; 9.2   ;8.1;7.2
```

## Importing Numeric Data with Text Headers

To import an ASCII data file that contains text headers, use the `textscan` function, specifying the `headerlines` parameter. `textscan` accepts a set of predefined parameters that control various aspects of the conversion. (For a complete list of these parameters, see the `textscan` reference page.) Using the `headerlines` parameter, you can specify the number of lines at the head of the file that `textscan` should ignore.

For example, the file `grades.dat` contains formatted numeric data with a one-line text header:

```
Grade1 Grade2 Grade3
  78.8   55.9   45.9
  99.5   66.8   78.0
  89.5   77.0   56.7
```

To import this data, first open the file and then use this `textscan` command to read the contents:

```
fid = fopen('grades.dat', 'r');
grades = textscan(fid, '%f %f %f', 3, 'headerlines', 1);

grades{:}
ans =
    78.8000
    99.5000
    89.5000

ans =
    55.9000
    66.8000
    77.0000

ans =
    45.9000
    78.0000
    56.7000

fclose(fid);
```

## Importing Mixed Alphabetic and Numeric Data

If your data file contains a mix of alphabetic and numeric ASCII data, use the `textscan` or `textread` function to import the data. `textscan` returns its output in a single cell array, while `textread` returns its output in separate variables and you can specify the data type of each variable. The `textscan` function offers better performance than `textread`, making it a better choice when reading large files.

This example uses `textread` to import the file `mydata.dat` that contains a mix of alphabetic and numeric data:

```
Sally    Type1 12.34 45 Yes
Larry    Type2 34.56 54 Yes
Tommy    Type1 67.89 23 No
```

---

**Note** To read an ASCII data file that contains numeric data with text column headers, see “Importing Numeric Data with Text Headers” on page 6-32.

---

To read the entire contents of the file `mydata.dat` into the workspace, specify the name of the data file and the format string as arguments to `textread`. In the format string, you include conversion specifiers that define how you want each data item to be interpreted. For example, specify `%s` for string data, `%f` for floating-point data, and so on. (For a complete list of format specifiers, see the `textread` reference page.)

For each conversion specifier in your format string, you must specify a separate output variable. `textread` processes each data item in the file as specified in the format string and puts the value in the output variable. The number of output variables must match the number of conversion specifiers in the format string.

In this example, `textread` reads the file `mydata.dat`, applying the format string to each line in the file until the end of the file:

```
[names, types, x, y, answer] = ...
    textread('mydata.dat', '%s %s %f %d %s', 3)
names =
    'Sally'
    'Larry'
    'Tommy'

types =
    'Type1'
    'Type2'
    'Type1'
```

```
x =  
    12.3400  
    34.5600  
    67.8900  
  
y =  
    45  
    54  
    23  
  
answer =  
    'Yes'  
    'Yes'  
    'No'
```

If your data uses a character other than a space as a delimiter, you must use the `textread` parameter `'delimiter'` to specify the delimiter. For example, if the file `mydata.dat` used a semicolon as a delimiter, you would use this command:

```
[names, types, x, y, answer]= ...  
    textread('mydata.dat', '%s %s %f %d %s', 'delimiter', ';')
```

See the `textread` reference page for more information about these optional parameters.

## Importing from XML Documents

With the `xmlread` function, you can read from a given URL or file, generating a Document Object Model (DOM) node to represent the parsed document.

MATLAB also provides these other XML functions:

- `xmlwrite` — Serializes a Document Object Model node to a file
- `xslt` — Transforms an XML document using an XSLT engine

See the reference pages for these functions for more information.

## Exporting Text Data

This section describes how to use MATLAB functions to export data in several common ASCII formats. For example, you can use these functions to export a MATLAB matrix as a text file where the rows and columns are represented as space-separated, numeric values. The function you use depends on the amount of data you want to export and its format. Topics covered include

- “Exporting Delimited ASCII Data Files” on page 6-37
- “Using the diary Function to Export Data” on page 6-39
- “Exporting to XML Documents” on page 6-40

If you are not sure which section describes your data, find the sample in the table below that most nearly matches the data format you want to create. Then read the section referred to in the table.

If you are familiar with MATLAB export functions but are not sure when to use them, see Table 6-4, which compares the features of each function.

---

**Note** If C or Fortran routines for writing data files in the form needed by other applications exist, create a MEX-file to write the data. See the MATLAB External Interfaces documentation for more information.

---

**Table 6-3: ASCII Data File Formats**

Data Format Sample	MATLAB Export Function
1 2 3 4 5 6 7 8 9 10	See “Exporting Delimited ASCII Data Files” on page 6-37 and “Using the diary Function to Export Data” on page 6-39 for information about these options.
1; 2; 3; 4; 5; 6; 7; 8; 9; 10;	See “Exporting Delimited ASCII Data Files” on page 6-37 for information. The example shows a semicolon-delimited file, but you can specify another character as the delimiter.

**Table 6-4: ASCII Data Export Function Features**

Function	Use With	Delimiters	Notes
csvwrite	Numeric data	Commas only	Primarily used with spreadsheet data. See “Working with Spreadsheets” on page 6-48.
diary	Numeric data or cell array	Spaces only	Can be used for small arrays. Requires editing of data file to remove extraneous text.
dlmwrite	Numeric data	Any character	Easy to use, flexible.
fprintf	Alphabetic and numeric data	Any character	Part of low-level file I/O routines. This function is the most flexible but also the most difficult to use. You must use fopen to obtain a file identifier before writing the data and fclose to close the file after writing the data.
save	Numeric data	Tabs or spaces	Easy to use; output values are high precision.

## Exporting Delimited ASCII Data Files

To export an array as a delimited ASCII data file, you can use either the save function, specifying the -ASCII qualifier, or the dlmwrite function. The save function is easy to use; however, the dlmwrite function provides more flexibility, allowing you to specify any character as a delimiter and to export subsets of an array by specifying a range of values.

### Using the save Function

To export the array A,

```
A = [ 1 2 3 4 ; 5 6 7 8 ];
```

use the save function, as follows:

```
save my_data.out A -ASCII
```

If you view the created file in a text editor, it looks like this:

```
1.0000000e+000 2.0000000e+000 3.0000000e+000 4.0000000e+000
5.0000000e+000 6.0000000e+000 7.0000000e+000 8.0000000e+000
```

By default, `save` uses spaces as delimiters but you can use tabs instead of spaces by specifying the `-tabs` option.

When you use `save` to write a character array to an ASCII file, it writes the ASCII equivalent of the characters to the file. If you write the character string 'hello' to a file, `save` writes the values

```
104 101 108 108 111
```

### Using the `dlmwrite` Function

To export an array in ASCII format and specify the delimiter used in the file, use the `dlmwrite` function.

For example, to export the array `A`,

```
A = [ 1 2 3 4 ; 5 6 7 8 ];
```

as an ASCII data file that uses semicolons as a delimiter, use this command:

```
dlmwrite('my_data.out',A, ';')
```

If you view the created file in a text editor, it looks like this:

```
1;2;3;4
5;6;7;8
```

Note that `dlmwrite` does not insert delimiters at the end of rows.

By default, if you do not specify a delimiter, `dlmwrite` uses a comma as a delimiter. You can specify a space (' ') as a delimiter or, if you specify empty quotes (''), no delimiter.

## Using the diary Function to Export Data

To export small numeric arrays or cell arrays, you can use the `diary` function. `diary` creates a verbatim copy of your MATLAB session in a disk file (excluding graphics).

For example, if you have the array `A` in your workspace,

```
A = [ 1 2 3 4; 5 6 7 8 ];
```

execute these commands at the MATLAB prompt to export this array using `diary`:

- 1 Turn on the diary function. You can optionally name the output file `diary` creates.

```
diary my_data.out
```

- 2 Display the contents of the array you want to export. This example displays the array `A`. You could also display a cell array or other MATLAB data type.

```
A =  
    1     2     3     4  
    5     6     7     8
```

- 3 Turn off the diary function.

```
diary off
```

`diary` creates the file `my_data.out` and records all the commands executed in the MATLAB session until it is turned off.

```
A =  
    1     2     3     4  
    5     6     7     8
```

```
diary off
```

- 4 Open the diary file `my_data.out` in a text editor and remove all the extraneous text.

### Exporting to XML Documents

With the `xmlwrite` function, you can serialize a Document Object Model (DOM) node to an XML file.

MATLAB also provides these other XML functions:

- `xmlread` — Imports from a given URL or file to a Document Object Model node
- `xslt` — Transforms an XML document using an XSLT engine

See the reference pages for these functions for more information.

## Working with Graphics Files

This section describes how to use MATLAB functions to import and export data in many standard graphics file formats. Topics covered include

- “Getting Information About Graphics Files” on page 6-41
- “Importing Graphics Data” on page 6-42
- “Exporting Graphics Data” on page 6-42

### Getting Information About Graphics Files

If you have a file in a standard graphics format, use the `imfinfo` function to get information about its contents. The `imfinfo` function returns a structure containing information about the file. The fields in the structure vary with the file format but `imfinfo` always returns some basic information including file name, last modification date, file size, and format.

This example returns information about a file in Joint Photographic Experts Group (JPEG) format:

```
info = imfinfo('ngc6543a.jpg')

info =

    Filename: [1x57 char]
    FileModDate: '01-Oct-1996 16:19:44'
    FileSize: 27387
    Format: 'jpg'
    FormatVersion: ''
    Width: 600
    Height: 650
    BitDepth: 24
    ColorType: 'truecolor'
    FormatSignature: ''
    NumberOfSamples: 3
    CodingMethod: 'Huffman'
    CodingProcess: 'Sequential'
    Comment: {[1x69 char]}
```

## Importing Graphics Data

To import data into the MATLAB workspace from a graphics file, use the `imread` function. Using this function, you can import data from files in many standard file formats, including the Tagged Image File Format (TIFF), Graphics Interchange Format (GIF), Joint Photographic Experts Group (JPEG), and Portable Network Graphics (PNG) formats. For a complete list of supported formats, see the `imread` reference page.

This example reads the image data stored in a file in JPEG format into the MATLAB workspace as the array `I`:

```
I = imread('ngc6543a.jpg');
```

`imread` represents the image in the workspace as a multidimensional array of class `uint8`. The dimensions of the array depend on the format of the data. For example, `imread` uses three dimensions to represent RGB color images:

```
whos I
  Name      Size                Bytes  Class
  I         650x600x3          1170000  uint8 array
```

```
Grand total is 1170000 elements using 1170000 bytes
```

## Exporting Graphics Data

To export data from the MATLAB workspace using one of the standard graphics file formats, use the `imwrite` function. Using this function, you can export data in formats such as the Tagged Image File Format (TIFF), Joint Photographic Experts Group (JPEG), and Portable Network Graphics (PNG). For a complete list of supported formats, see the `imwrite` reference page.

The following example writes a multidimensional array of `uint8` data `I` from the MATLAB workspace into a file in TIFF format. The class of the output image written to the file depends on the format specified. For most formats, if the input array is of class `uint8`, `imwrite` outputs the data as 8-bit values. See the `imwrite` reference page for details.

```
whos I
  Name      Size          Bytes  Class
  I         650x600x3      1170000  uint8 array

Grand total is 1170000 elements using 1170000 bytes

imwrite(I, 'my_graphics_file.tif','tif');
```

---

**Note** `imwrite` supports different syntaxes for several of the standard formats. For example, with TIFF file format, you can specify the type of compression used to store the image. See the `imwrite` reference page for details.

---

## Working with Audio and Video Data

This section describes the MATLAB functions that you can use to read, write, and get information about files that contain audio data, video data, or both. Topics covered include

- “Getting Information About Audio/Video Files” on page 6-44
- “Importing Audio/Video Data” on page 6-45
- “Exporting Audio/Video Data” on page 6-46

### Getting Information About Audio/Video Files

MATLAB includes several functions that you can use to get information about files that contain audio data, video data, or both. Some work only with specific file formats. One function, the `mmfileinfo` function, can retrieve information about many file formats.

#### Format-Specific Functions

MATLAB includes several functions that return information about files that contain audio and video data in specific formats.

- `aufinfo` — Returns a text description of the contents of a sound (AU) file
- `aviinfo` — Returns a structure containing information about the contents of an Audio/Video Interleaved (AVI) file
- `wavfinfo` — Returns a text description of the contents of a sound (WAV) file

#### Using the General Multimedia Information Function

MATLAB also includes a general-purpose, audio/video file information function named `mmfileinfo`. The `mmfileinfo` function returns information about both the audio data in a file as well as the video data in the file, if present.

---

**Note** `mmfileinfo` can be used only on Windows systems.

---

## Importing Audio/Video Data

MATLAB includes several functions that you can use to bring audio or video data into the MATLAB workspace. Some of these functions read audio or video data from files. Another way to import audio data into the MATLAB workspace is to record it using an audio input device, such as a microphone. The following sections describe

- “Reading Audio and Video Data from a File” on page 6-45
- “Recording Audio Data” on page 6-45

### Reading Audio and Video Data from a File

MATLAB includes several functions for reading audio or video data from a file. These files are format-specific.

- `auread` — Returns sound data from a sound (AU) file
- `aviread` — Returns AVI data as a MATLAB movie
- `wavread` — Returns sound data from a sound (WAV) file

### Recording Audio Data

To bring sound data into the MATLAB workspace by recording it from an audio input device, use the audio recorder object. This object represents the connection between MATLAB and an audio input device, such as a microphone, that is connected to your system. You use the `audiorecorder` function to create this object and then use methods and properties of the object to record the audio data.

On PCs running Windows, you can also use the `wavrecord` function to bring live audio data in WAV format into the MATLAB workspace.

Once you import audio data, MATLAB supports several ways to listen to the data. You can use an audio player object to play the audio data. Use the `audioplayer` function to create an audio player object.

You can also use the `sound` or `soundsc` function.

On PCs running Windows, you can use the `wavplay` function to listen to `.wav` files.

## Exporting Audio/Video Data

MATLAB includes several functions that you can use to export audio or video data from the MATLAB workspace. These functions write audio data to a file using specific file formats. The following sections describe

- “Exporting Audio Data” on page 6-46
- “Exporting Video Data in AVI Format” on page 6-46

This section also provides an example of writing video data to a file in “Example: Creating an AVI file” on page 6-47.

### Exporting Audio Data

In MATLAB, audio data is simply numeric data that you can export using standard MATLAB data export functions, such as `save`.

MATLAB also includes several functions that write audio data to files in specific file formats:

- `auwrite` — Exports sound data in AU file format
- `wavwrite` — Exports sound data in WAV file format

### Exporting Video Data in AVI Format

You can export MATLAB video data as an Audio/Video Interleaved (AVI) file. To do this, you use the `avifile` function to create an `avifile` object. Once you have the object, you can use AVI file object methods and properties to control various aspects of the data export process.

For example, in MATLAB, you can save a sequence of graphs as a movie that can then be played back using the `movie` function. You can export a MATLAB movie by saving it in MAT-file format, like any other MATLAB workspace variable. However, anyone who wants to view your movie must have MATLAB. (For more information about MATLAB movies, see the “Animation” section in the MATLAB Graphics documentation.)

To export a sequence of MATLAB graphs in a format that does not require MATLAB for viewing, save the figures in Audio/Video Interleaved (AVI) format. AVI is a file format that allows animation and video clips to be played on a PC running Windows or on UNIX systems.

---

**Note** To convert an existing MATLAB movie into an AVI file, use the `movie2avi` function.

---

### Example: Creating an AVI file

To export a sequence of MATLAB graphs as an AVI format movie, perform these steps:

- 1 Create an AVI file object, using the `avifile` function.

```
aviobj = avifile('mymovie.avi','fps',5);
```

AVI file objects support properties that let you control various characteristics of the AVI movie, such as colormap, compression, and quality. (See the `avifile` reference page for a complete list.) `avifile` uses default values for all properties, unless you specify a value. The example sets the value of the frames per second (`fps`) property.

- 2 Capture the sequence of graphs and put them into the AVI file, using the `addframe` function.

```
for k=1:25
    h = plot(fft(eye(k+16)));
    set(h,'EraseMode','xor');
    axis equal;
    frame = getframe(gca);
    aviobj = addframe(aviobj,frame);
end
```

The example uses a `for` loop to capture the series of graphs to be included in the movie. You typically use `addframe` to capture a sequence of graphs for AVI movies. However, because this particular MATLAB animation uses XOR graphics, you must call `getframe` to capture the graphs and then call `addframe` to add the captured frame to the movie.

- 3 Close the AVI file, using the `close` function.

```
aviobj = close(aviobj);
```

## Working with Spreadsheets

You can use MATLAB to import and export data to the following types of spreadsheets:

- “Microsoft Excel Spreadsheets” on page 6-48
- “Lotus 123 Spreadsheets” on page 6-51

### Microsoft Excel Spreadsheets

This section covers

- “Getting Information About the File” on page 6-48
- “Exporting to the File” on page 6-49
- “Importing from the File” on page 6-50

See the `xlsfinfo`, `xlswrite`, and `xlsread` reference pages for more detailed information and examples.

#### Getting Information About the File

Use the `xlsfinfo` function to determine if a file contains a readable Microsoft Excel spreadsheet.

Inputs to `xlsfinfo` are

- Name of the spreadsheet file

Outputs from `xlsfinfo` are

- String 'Microsoft Excel Spreadsheet' if the file contains an Excel worksheet readable with the `xlsread` function. Otherwise, it contains an empty string ('').
- Cell array of strings containing the names of each worksheet in the file.

**Example — Querying an XLS File.** This example returns information about spreadsheet file `tempdata.xls`:

```
[type, sheets] = xlsfinfo('tempdata.xls')  
  
type =  
Microsoft Excel Spreadsheet
```

```

sheets =
    'Locations'    'Rainfall'    'Temperatures'

```

## Exporting to the File

Use the `xlswrite` function to export a matrix to an Excel spreadsheet file. With `xlswrite`, you can export data from the workspace to any worksheet in the file, and to any location within that worksheet.

Inputs to `xlswrite` are

- Name of the spreadsheet file
- Matrix to be exported
- Name of the worksheet to receive the data
- Range of cells on the worksheet in which to write the data

Outputs from `xlswrite` are

- Pass or fail status
- Any warning or error message generated along with its message identifier

**Example — Writing To an XLS File.** This example writes a mix of text and numeric data to the file `tempdata.xls`. Call `xlswrite`, specifying a worksheet labeled `Temperatures`, and the region within the worksheet to write the data to. The 4-by-2 matrix is written to the rectangular region that starts at cell E1 in its upper-left corner:

```

d = {'Time', 'Temp'; 12 98; 13 99; 14 97}
d =
    'Time'    'Temp'
    [ 12]    [ 98]
    [ 13]    [ 99]
    [ 14]    [ 97]

xlswrite('tempdata.xls', d, 'Temperatures', 'E1');

```

**Adding a New Worksheet.** If the worksheet being written to does not already exist in the file, MATLAB displays the following warning:

```
Warning: Added specified worksheet.
```

You can disable these warnings with the command

```
warning off MATLAB:xlswrite:AddSheet
```

### Importing from the File

Use `xlsread` to import a matrix from an Excel spreadsheet file into the MATLAB workspace. You can import data from any worksheet in the file, and from any location within that worksheet. You can also optionally have `xlsread` open an Excel window showing the file and then interactively select the worksheet and range of data to be read by the function.

Inputs to `xlsread` are

- Name of the spreadsheet file
- Matrix to be imported
- Name of the worksheet from which to read the data
- Range of cells on the worksheet from which to read the data
- Keyword that opens an Excel window, enabling you to interactively select the worksheet and range of data to read
- Keyword that imports using basic import mode

Three separate outputs from `xlsread` are

- Numeric data
- String data
- Any unprocessed cell content

**Example — Reading from an XLS File.** Continuing with the previous example, to import only the numeric data, use `xlsread` with a single return argument. `xlsread` ignores any leading row or column of text in the numeric result:

```
ndata = xlsread('tempdata.xls', 'Temperatures')
ndata =
    12    98
    13    99
    14    97
```

To import both numeric data and text data, specify two return values for `xlsread`:

```
[ndata, headertext] = xlsread('tempdata.xls', 'Temperatures')

headertext =
    'Time'    'Temp'

ndata =
    12    98
    13    99
    14    97
```

## Lotus 123 Spreadsheets

This section covers

- “Getting Information About the File” on page 6-51
- “Exporting to the File” on page 6-52
- “Importing from the File” on page 6-53

See the `wk1finfo`, `wk1write`, and `wk1read` reference pages for more detailed information and examples.

### Getting Information About the File

Use the `wk1finfo` function to determine if a file contains a Lotus WK1 spreadsheet:

Inputs to `wk1finfo` are

- Name of the spreadsheet file

Outputs from `wk1finfo` are

- String 'WK1' if the file is a Lotus spreadsheet readable with the `wk1read` function. Otherwise, it contains an empty string ('').
- String 'Lotus 123 Spreadsheet'

**Example — Querying a WK1 File.** This example returns information about spreadsheet file `matA.wk1`:

```
[extens, type] = wk1finfo('matA.wk1')

extens =
    WK1
type =
    Lotus 123 Spreadsheet
```

### Exporting to the File

Use the `wk1write` function to export a matrix to a Lotus spreadsheet file. You have the choice of positioning the matrix starting at the first row and column of the spreadsheet, or at any other location in the file.

To export to a specific location in the file, use the second syntax, indicating a zero-based starting row and column.

Inputs to `wk1write` are

- Name of the spreadsheet file
- Matrix to be exported
- Location in the file in which to write the data

**Example — Writing to a WK1 File.** This example exports an 8-by-8 matrix to spreadsheet file `matA.wk1`:

```
A = [1:8; 11:18; 21:28; 31:38; 41:48; 51:58; 61:68; 71:78];
A =
     1     2     3     4     5     6     7     8
    11    12    13    14    15    16    17    18
    21    22    23    24    25    26    27    28
    31    32    33    34    35    36    37    38
    41    42    43    44    45    46    47    48
    51    52    53    54    55    56    57    58
    61    62    63    64    65    66    67    68
    71    72    73    74    75    76    77    78

wk1write('matA.wk1', A);
```

## Importing from the File

To import data from the spreadsheet into the MATLAB workspace, use `wk1read`. There are three ways to call `wk1read`. The first two shown here are similar to `wk1write`. The third enables you to select a range of values from the spreadsheet. You can specify the range argument with a one-based vector, spreadsheet notation (e.g., 'A1..B7'), or using a named range (e.g., 'Sales').

Inputs to `wk1read` are

- Name of the spreadsheet file
- Spreadsheet location from which to read the data
- Range of cells from which to read the data

Outputs from `wk1read` are

- Requested data from the spreadsheet

**Example — Reading from a WK1 File.** Read in a limited block of the spreadsheet data by specifying the upper-left row and column of the block using zero-based indexing:

```
M = wk1read('matA.wk1', 3, 2)
M =
    33    34    35    36    37    38
    43    44    45    46    47    48
    53    54    55    56    57    58
    63    64    65    66    67    68
    73    74    75    76    77    78
```

## Working with Scientific Data Formats

This section describes how to import and export data in several standard scientific data formats. Topics covered include

- “Working with Common Data Format (CDF) Files” on page 6-54
- “Working with Flexible Image Transport System (FITS) Files” on page 6-57
- “Working with Hierarchical Data Format (HDF5) Files” on page 6-59

For information about working with Hierarchical Data Format (HDF) files, see these sections:

- “Importing HDF4 and HDF-EOS Data” on page 6-69
- “Exporting MATLAB Data to an HDF4 File” on page 6-96

### Working with Common Data Format (CDF) Files

MATLAB includes functions that let you import and export data using the Common Data Format (CDF). CDF was created by the National Space Science Data Center (NSSDC) to provide a self-describing data storage and manipulation format that matches the structure of scientific data and applications (i.e., statistical and numerical methods, visualization, and management). The MATLAB CDF functions are described in the following sections:

- “Getting Information About CDF Files” on page 6-54
- “Importing Data from a CDF File” on page 6-55
- “Exporting Data to a CDF File” on page 6-57

#### Getting Information About CDF Files

To get information about the contents of a CDF file, use the `cdfinfo` function. The `cdfinfo` function returns a structure containing general information about the file and detailed information about the variables and attributes in the file.

This example returns information about the sample CDF file included with MATLAB. To determine the variables contained in the file, view the `Variables` field. This field contains a cell array that lists all the variables in the file with information that describes the variable, such as name, size, and data type. For an example, see “Importing Data from a CDF File” on page 6-55.

---

**Note** Before executing the `cdfinfo` function, make sure that your current working directory is writable because `cdfinfo` creates temporary files.

---

```
info = cdfinfo('example.cdf')

info =

    Filename: 'example.cdf'
  FileModDate: '09-Mar-2001 16:45:22'
    FileSize: 1240
      Format: 'CDF'
FormatVersion: '2.7.0'
  FileSettings: [1x1 struct]
    Subfiles: {}
    Variables: {5x6 cell}
GlobalAttributes: [1x1 struct]
VariableAttributes: [1x1 struct]
```

### Importing Data from a CDF File

To import data into the MATLAB workspace from a CDF file, use the `cdfread` function. Using this function, you can import all the data in the file, specific variables, or subsets of the data in a specific variable. This example illustrates how to use the `cdfread` function to read data associated with a particular variable:

- 1 Determine the names of variables in the CDF file. The information returned by `cdfinfo` indicates that the file contains five variables.

```
info = cdfinfo('example.cdf');

vars = info.Variables

vars =

Columns 1 through 5

    'Time'           [1x2 double]   [24]   'epoch'   'T/'
    'Longitude'     [1x2 double]   [ 1]   'int8'    'F/FT'
```

```
'Latitude'          [1x2 double] [ 1] 'int8'   'F/TF'  
'Data'             [1x3 double] [ 1] 'double' 'T/TTT'  
'multidimensional' [1x4 double] [ 1] 'uint8'  'T/TTTT'
```

Column 6

```
'Full'  
'Full'  
'Full'  
'Full'  
'Full'
```

- 2 Read the data associated with the Time variable. Variable names are case sensitive.

```
data = cdfread('example.cdf','variable','Time');
```

The return value data is a 24-by-1 cell array, where each cell contains a CDF epoch object.

### Representing CDF Time Values

CDF represents time differently than MATLAB. CDF represents date and time as the number of milliseconds since 1-Jan-0000. This is called an *epoch* in CDF terminology. MATLAB represents date and time as a serial date number, which is the number of days since 0-Jan-0000. To represent CDF dates, MATLAB uses an object called a CDF epoch object. To access the time information in a CDF object, use the object's `todatenum` method.

For example, this code extracts the date information from a CDF epoch object:

- 1 Extract the date information from the CDF epoch object returned in the cell array data (see “Importing Data from a CDF File” on page 6-55). Use the `todatenum` method of the CDF epoch object to get the date information, which is returned as a MATLAB serial date number.

```
m_date = todatenum(data{1});
```

- 2 View the MATLAB serial date number as a string.

```
datestr(m_date)
```

```
ans =
```

```
01-Jan-2001
```

### Exporting Data to a CDF File

To export data from the MATLAB workspace to a CDF file, use the `cdfwrite` function. Using this function, you can write variables and attributes to the file, specifying their names and associated values. See the `cdfwrite` reference page for more information.

This example shows how to write date information to a CDF file. Note how the example uses the CDF epoch object constructor, `cdfepoch`, to convert a MATLAB serial date number into a CDF epoch.

```
cdfwrite('myfile',{'Time_val',cdfepoch(now)});
```

## Working with Flexible Image Transport System (FITS) Files

MATLAB includes functions that let you import and export data using the Flexible Image Transport System (FITS) format. FITS is the standard data format used in astronomy, endorsed by both NASA and the International Astronomical Union (IAU). FITS is designed to store scientific data sets consisting of multidimensional arrays (1-D spectra, 2-D images, or 3-D data cubes) and two-dimensional tables containing rows and columns of data.

A data file in FITS format can contain multiple components, each marked by an ASCII text header followed by binary data. The first component in a FITS file is known as the *primary*, which can be followed by any number of other components, called *extensions*, in FITS terminology. For more information about the FITS standard, go to the official FITS Web site, [fits.gsfc.nasa.gov/](http://fits.gsfc.nasa.gov/).

The MATLAB FITS functions are described in the following sections:

- “Getting Information About CDF Files” on page 6-54
- “Importing Data from a CDF File” on page 6-55

## Getting Information About FITS Files

To get information about the contents of a FITS file, use the `fitsinfo` function. The `fitsinfo` function returns a structure containing the information about the file and detailed information about the data in the file.

This example returns information about a sample FITS file included with MATLAB. The structure returned contains fields for the primary component, `PrimaryData`, and all the extensions in the file, such as the `BinaryTable`, `Image`, and `AsciiTable` extensions.

```
info = fitsinfo('tst0012.fits')

info =

    Filename: 'tst0012.fits'
  FileModDate: '12-Mar-2001 18:37:46'
    FileSize: 109440
   Contents: {1x5 cell}
  PrimaryData: [1x1 struct]
  BinaryTable: [1x1 struct]
     Unknown: [1x1 struct]
       Image: [1x1 struct]
   AsciiTable: [1x1 struct]
```

## Importing Data from a FITS File

To import data into the MATLAB workspace from a FITS file, use the `fitsread` function. Using this function, you can import the data in the `PrimaryData` section of the file or you can import the data in any of the extensions in the file, such as the `Image` extension. This example illustrates how to use the `fitsread` function to read data from a FITS file:

- 1 Determine which extensions the FITS file contains, using the `fitsinfo` function.

```
info = fitsinfo('tst0012.fits')

info =

    Filename: 'tst0012.fits'
  FileModDate: '12-Mar-2001 18:37:46'
    FileSize: 109440
```

```
Contents: {1x5 cell}
PrimaryData: [1x1 struct]
BinaryTable: [1x1 struct]
    Unknown: [1x1 struct]
    Image: [1x1 struct]
AsciiTable: [1x1 struct]
```

The info structure shows that the file contains several extensions including the BinaryTable, AsciiTable, and Image extensions.

## 2 Read data from the file.

To read the PrimaryData in the file, specify the filename as the only argument:

```
pdata = fitsread('tst0012.fits');
```

To read any of the extensions in the file, you must specify the name of the extension as an optional parameter. This example reads the BinaryTable extension from the FITS file:

```
bindata = fitsread('tst0012.fits','bintable');
```

---

**Note** To read the BinaryTable extension using fitsread, you must specify the parameter 'bintable'. Similarly, to read the AsciiTable extension, you must specify the parameter 'table'. See the fitsread reference page for more information.

---

## Working with Hierarchical Data Format (HDF5) Files

Hierarchical Data Format, Version 5, (HDF5) is a general-purpose, machine-independent standard for storing scientific data in files, developed by the National Center for Supercomputing Applications (NCSA). For more information about the HDF5 file format, read the HDF5 documentation available at the NCSA Web site ([hdf.ncsa.uiuc.edu](http://hdf.ncsa.uiuc.edu)).

---

**Note** For information about importing HDF4 data, which is a completely separate, incompatible format, see “Importing HDF4 and HDF-EOS Data” on page 6-69.

---

This section describes how to import data or metadata from an HDF5 file. Topics covered include

- “Determining the Contents of an HDF5 File” on page 6-60
- “Importing Data from an HDF5 File” on page 6-64
- “Mapping HDF5 Data Types to MATLAB Data Types” on page 6-65

### Determining the Contents of an HDF5 File

HDF5 files can contain data and metadata. HDF5 files organize the data and metadata, called *attributes*, in a hierarchical structure, similar to the hierarchical structure of a file system.

In an HDF5 file, the directories in the hierarchy are called *groups*. A group can contain other groups, data sets, attributes, links, and data types. A data set is a collection of data, such as a multidimensional numeric array or string. An attribute is any data that is associated with another entity, such as a data set. A link is similar to a UNIX file system symbolic link. Links are a way to reference data without having to make a copy of the data.

Data types are a description of the data in the data set or attribute. Data types tell how to interpret the data in the data set. For example, a file might contain a data type called “Reading” that is comprised of three elements: a longitude value, a latitude value, and a temperature value.

To find the names of all the data sets and attributes contained in an HDF5 file, use the `hdf5info` function. For example, to find out what the sample HDF5 file, `example.h5`, contains, use this syntax:

```
fileinfo = hdf5info('example.h5');
```

`hdf5info` returns a structure that contains various information about the HDF5 file, including the name of the file and the version of the HDF5 library that MATLAB is using:

```
fileinfo =  
    Filename: 'example.h5'
```

```

LibVersion: '1.4.2'
Offset: 0
FileSize: 8172
GroupHierarchy: [1x1 struct]

```

## Exploring the Contents of an HDF5 File

To explore the hierarchical organization of the file, examine the `GroupHierarchy` field in the structure returned by `hdf5info`. The `GroupHierarchy` field is a structure that describes the top-level group in the file, called the root group. HDF5 uses the UNIX convention and names this top-level group `/` (forward slash).

The following example shows that the `GroupHierarchy` structure for the sample HDF5 file contains two groups and two attributes. The root group does not contain any data sets, data types, or links.

```

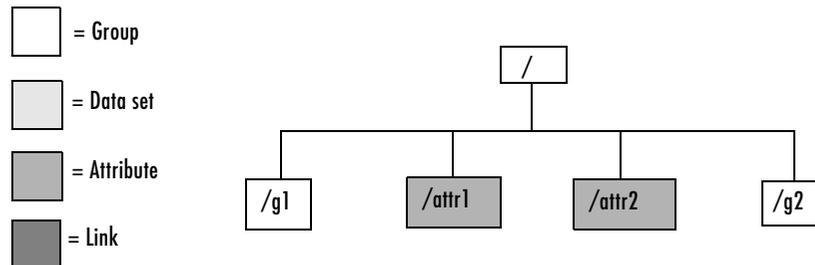
toplevel = fileinfo.GroupHierarchy

toplevel =

    Filename: 'C:\matlab\toolbox\matlab\demos\example.h5'
      Name: '/'
     Groups: [1x2 struct]
   Datasets: []
  Datatypes: []
     Links: []
  Attributes: [1x2 struct]

```

The following figure illustrates the organization of the root group.



### Organization of the Root Group of the Sample HDF5 File

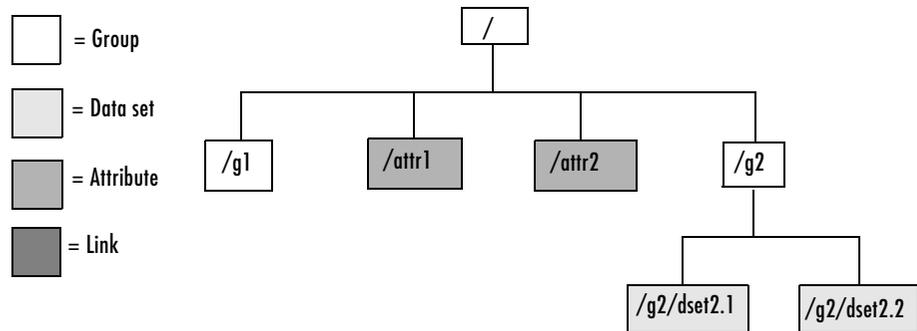
To explore the contents of the sample HDF5 file further, examine one of the two structures in the Groups field of the GroupHierarchy structure. Each structure in this field represents a group contained in the root group:

```
level2 = toplevel.Groups(2)

level2 =

    Filename: 'C:\matlab\toolbox\matlab\demos\example.h5'
    Name: '/g2'
    Groups: []
    Datasets: [1x2 struct]
    Datatypes: []
    Links: []
    Attributes: []
```

In the sample file, the group named /g2 contains two data sets. The following figure illustrates this part of the sample HDF5 file organization.



### Organization of the Data Set /g2 in the Sample HDF5 File

To get information about a data set, look at either of the structures returned in the Datasets field. These structures provide information about the data set, such as its name, dimensions, and data type.

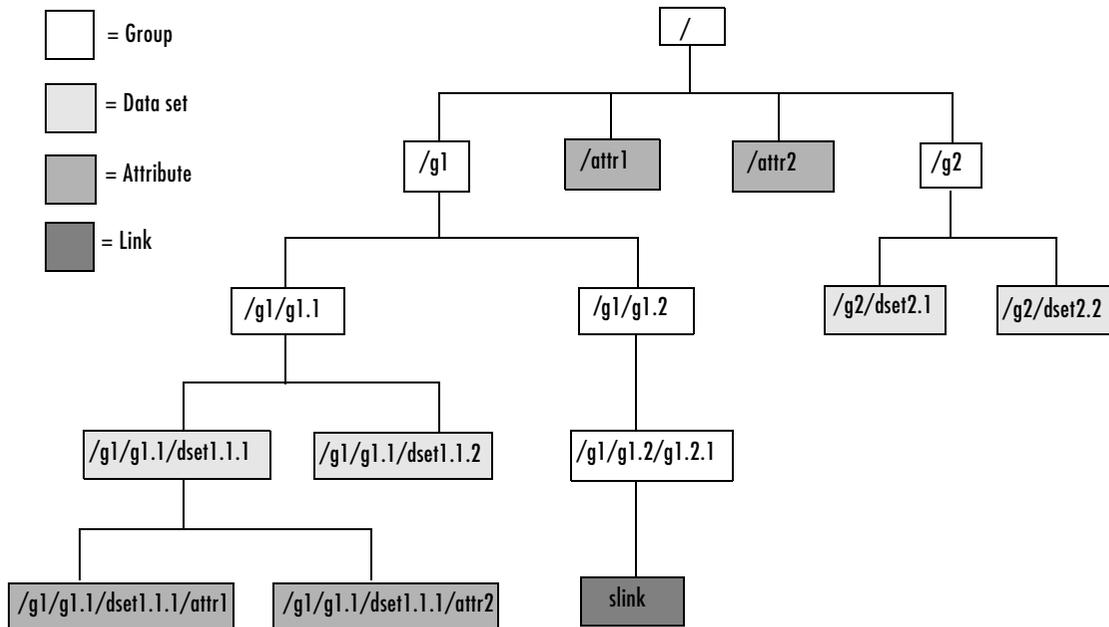
```
dataset1 = level2.Datasets(1)
```

```

dataset1 =
  Filename: 'L:\matlab\toolbox\matlab\demos\example.h5'
  Name: '/g2/dset2.1'
  Rank: 1
  Datatype: [1x1 struct]
  Dims: 10
  MaxDims: 10
  Layout: 'contiguous'
  Attributes: []
  Links: []
  Chunksize: []
  Fillvalue: []

```

By examining the structures at each level of the hierarchy, you can traverse the entire file. The following figure describes the hierarchical organization of the sample file `example.h5`.



**Hierarchical Structure of `example.h5` HDF5 File**

## Importing Data from an HDF5 File

To read data or metadata from an HDF5 file, use the `hdf5read` function. As arguments, you must specify the name of the HDF5 file and the name of the data set. For information about finding the name of a data set, see “Determining the Contents of an HDF5 File” on page 6-60.

For example, to read the data set, `/g2/dset2.1` from the HDF5 file `example.h5`, use this syntax:

```
data = hdf5read('example.h5', '/g2/dset2.1');
```

The return value contains the values in the data set, in this case a 1-by-10 vector of single-precision values:

```
data =  
  
Columns 1 through 8  
    1.0000    1.1000    1.2000    1.3000    1.4000    1.5000    1.6000    1.7000  
  
Columns 9 through 10  
    1.8000    1.9000
```

The `hdf5read` function maps HDF5 data types to appropriate MATLAB data types, whenever possible. If the HDF5 file contains data types that cannot be represented in MATLAB, `hdf5write` uses one of the predefined MATLAB HDF5 data type objects to represent the data.

For example, if an HDF5 data set contains four array elements, `hdf5read` can return the data as a 1-by-4 array of `hdf5.h5array` objects:

```
whos  
  
Name      Size      Bytes      Class  
  
data      1x4              hdf5.h5array  
  
Grand total is 4 elements using 0 bytes
```

For more information about the MATLAB HDF5 data type objects, see “Mapping HDF5 Data Types to MATLAB Data Types” on page 6-65.

## Exporting Data to HDF5 Files

To write data or metadata from the MATLAB workspace to an HDF5 file, use the `hdf5write` function. As arguments, specify:

- Name of an existing HDF5 file, or the name you want to assign to a new file.
- Name of an existing data set or attribute, or the name you want to assign to a new data set or attribute. To learn how to determine the name of data sets in an existing HDF5 file, see “Determining the Contents of an HDF5 File” on page 6-60.
- Data or metadata you want to write to the file. `hdf5write` converts MATLAB data types to the appropriate HDF5 data type automatically. For nonatomic data types, you can also create HDF5 objects to represent the data.

This example creates a 5-by-5 array of `uint8` values and then writes the array to an HDF5 file. By default, `hdf5write` overwrites the file, if it already exists. The example specifies an `hdf5write` mode option to append data to existing file.

- 1 Create a MATLAB variable in the workspace. This example creates a 5-by-5 array of `uint8` values.

```
testdata = uint8(magic(5))
```

- 2 Add the data to an existing HDF5 file. To add data to an existing file, you must specify 'append' mode. The file must already exist and it cannot already contain a data set with the same name.

```
hdf5write('myfile.h5', '/dataset1', testdata,...  
         'writemode','append')
```

## Mapping HDF5 Data Types to MATLAB Data Types

When the `hdf5read` function reads data from an HDF5 file into the MATLAB workspace, it maps HDF5 data types to MATLAB data types, depending on whether the data in the dataset is in an *atomic* data type or a *nonatomic composite* data type.

Atomic data types describe commonly used binary formats for numbers (integers and floating point) and characters (ASCII). Because different computing architectures and programming languages support different number and character representations, the HDF5 library provides the platform-independent data types, which it then maps to an appropriate data

type for each platform. For example, a computer may support 8-, 16-, 32-, and 64-bit signed integers, stored in memory in little endian byte order.

A composite data type is an aggregation of one or more atomic data types. Composite data types include structures, multidimensional arrays, and variable-length data types (one-dimensional arrays).

**Mapping Atomic Data Types.** If the data in the data set is stored in one of the HDF5 atomic data types, `hdf5read` uses the equivalent MATLAB data type to represent the data. Each data set contains a `Datatype` field that names the data type. For example, the data set `/g2/dset2.2` in the sample HDF5 file includes atomic data and data type information.

```
dtype = dataset1.Datatype
dtype =

      Name: []
      Class: 'H5T_IEEE_F32BE'
  Elements: []
```

The `H5T_IEEE_F32BE` class name indicates the data is a 4-byte, big endian, IEEE floating-point data type. (See the HDF5 specification for more information about atomic data types.)

**HDF5 Nonatomic Data Types.** If the data in the data set is stored in one of the HDF5 nonatomic data types, `hdf5read` represents the data set in MATLAB as an object. MATLAB supports the following objects to represent HDF5 nonatomic data types:

- `hdf5.h5array`
- `hdf5.h5enum`
- `hdf5.h5vlen`
- `hdf5.h5compound`
- `hdf5.h5string`

To access the data in the data set in the MATLAB workspace, you must access the `Data` field in the object. This example converts a simple MATLAB vector into an `h5array` object and then displays the fields in the object:

```
vec = [ 1 2 3];

hhh = hdf5.h5array(vec);

hhh:

    Name: ''
    Data: [1 2 3]

hhh.Data

ans =

     1     2     3
```

**Using HDF5 Data Type Objects.** If you are writing simple data sets, such as scalars, strings, or a simple compound data set, you can just pass the data directly to `hdf5write`. The `hdf5write` function can automatically map the MATLAB data types to appropriate HDF5 data types.

However, if your data is a complex data set, you must use one of the predefined MATLAB HDF5 objects to pass to the `hdf5write` function. The HDF5 objects are designed for situations where the mapping between MATLAB and HDF5 types is ambiguous.

For example, when passed a cell array of strings, the `hdf5write` function writes a data set made up of strings, not a data set of arrays containing strings. If that is not the mapping you intend, use HDF5 objects to specify the correct mapping.

In addition, note that HDF5 makes a distinction between the size of a data set and the size of a data type. In MATLAB, data types are always scalar. In HDF5, data types can have a size; that is, types can be either scalar (like MATLAB) or m-by-n.

In HDF5, a 5-by-5 data set containing a single `uint8` value in each element is distinct from a 1-by-1 data set containing a 5-by-5 array of `uint8` values. In the

first case, the data set contains 25 observations of a single value; in the second case, the data set contains a single observation with 25 values.

This example uses an HDF5 enumeration object for enumerated data:

- 1** Create an HDF5 enumerated object.

```
enum_obj = hdf5.h5enum;
```

- 2** Define the enumerated values and their corresponding names.

```
enum_obj.defineEnum({'RED' 'GREEN' 'BLUE'}, uint8([1 2 3]));
```

enum\_obj now contains the definition of the enumeration that associates the names RED, GREEN, and BLUE with the numbers 1, 2, and 3.

- 3** Add enumerated data to the object.

```
enum_obj.setData(uint8([2 1 3 3 2 3 2 1]));
```

In the HDF5 file, these numeric values map to the enumerated values GREEN, RED, BLUE, BLUE, GREEN, etc.

- 4** Write the enumerated data to a data set named objects in an HDF5 file.

```
hdf5write('myfile3.h5', '/g1/objects', enum_obj);
```

- 5** Read the enumerated data set from the file.

```
ddd = hdf5read('myfile3.h5', '/g1/objects')
```

```
hdf5.h5enum:
```

```
    Name: ''
    Data: [2 1 3 3 2 3 2 1]
 EnumNames: {'RED' 'GREEN' 'BLUE'}
 EnumValues: [1 2 3]
```

---

## Importing HDF4 and HDF-EOS Data

Hierarchical Data Format (HDF4) is a general-purpose, machine-independent standard for storing scientific data in files, developed by the National Center for Supercomputing Applications (NCSA). HDF-EOS is an extension of HDF4 that was developed by the National Aeronautics and Space Administration (NASA) for storage of data returned from the Earth Observing System (EOS). For more information about these file formats, read the HDF documentation at the NCSA Web site ([hdf.ncsa.uiuc.edu](http://hdf.ncsa.uiuc.edu)) and the HDF-EOS documentation at the NASA Web site ([hdfeos.gsfc.nasa.gov/hdfeos/index.cfm](http://hdfeos.gsfc.nasa.gov/hdfeos/index.cfm)).

---

**Note** For information about importing HDF5 data, which is a separate, incompatible format, see “Working with Hierarchical Data Format (HDF5) Files” on page 6-59.

---

HDF4 and HDF-EOS files can contain multidimensional numeric arrays or text data, called *data sets*, in HDF terminology. MATLAB provides three ways to import HDF4 or HDF-EOS data sets into the MATLAB workspace:

- “Using the HDF Import Tool” on page 6-69
- “Using the MATLAB `hdfread` Function” on page 6-84
- “Using the HDF4 Command-Line Interface” on page 6-88

For information about exporting data in HDF files, see “Exporting MATLAB Data to an HDF4 File” on page 6-96.

### Using the HDF Import Tool

The HDF Import Tool is a graphical user interface that you can use to navigate through HDF4 or HDF-EOS data sets and import data from them. Importing data using the HDF Import Tool involves these steps:

- “Step 1: Opening an HDF4 File in the HDF Import Tool” on page 6-70
- “Step 2: Selecting a Data Set in an HDF File” on page 6-71
- “Step 3: Specifying a Subset of the Data” on page 6-72
- “Step 4: Importing Data and Metadata” on page 6-73

The following sections provide more detail about each of these steps.

### Step 1: Opening an HDF4 File in the HDF Import Tool

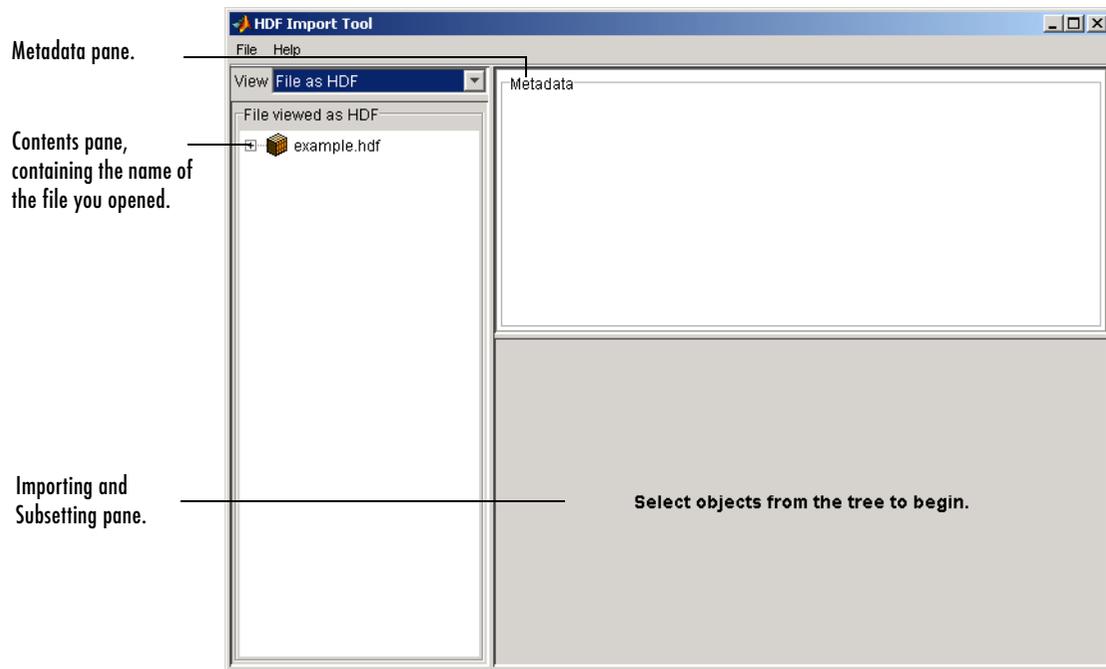
To open an HDF4 file, select the **Import Data** option from the MATLAB **File** menu. MATLAB displays a file selection dialog box. If you select an HDF4 file, the Import Wizard automatically starts the HDF Import Tool.

You can also open a file with the HDF Import Tool by entering the `hdftool` command at the MATLAB command line:

```
hdftool('example.hdf')
```

If you use the `hdftool` function without arguments, it starts the HDF Import Tool and automatically opens a file selection dialog box. You can open multiple files in the HDF Import Tool at the same time.

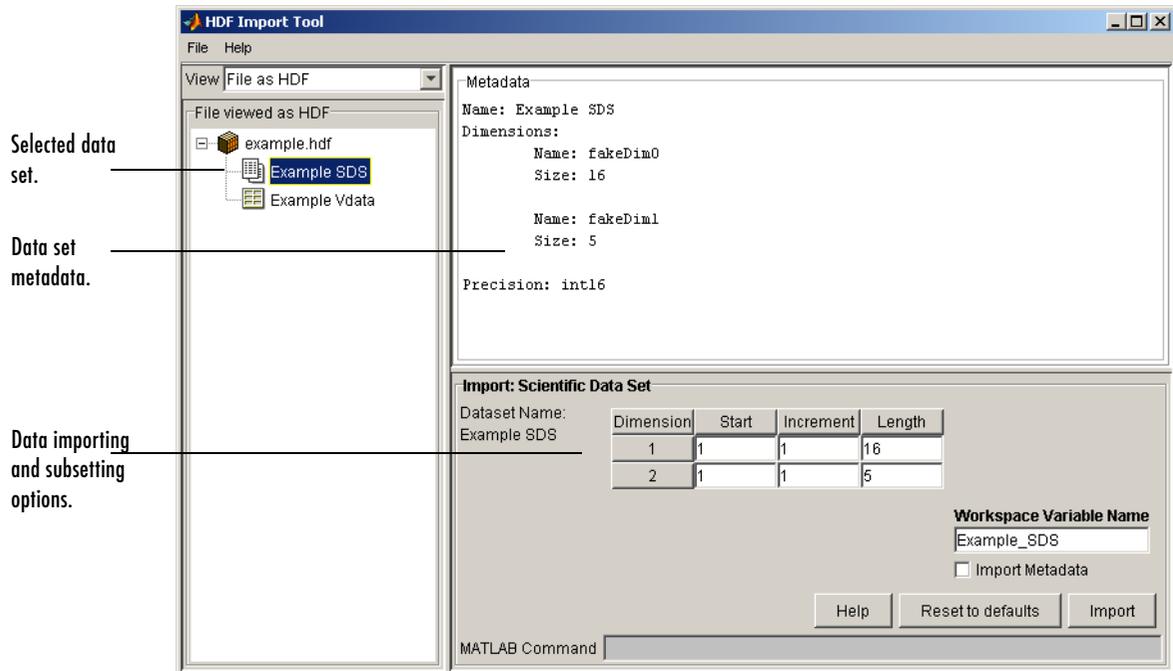
**Overview of the HDF Import Tool.** Initially, the HDF Import Tool window contains three panes: the **Contents** pane, **Metadata** pane, and the **Importing and Subsetting** pane. Initially, the **Contents** pane contains the name of the file you opened and the other panes are empty, as illustrated in the following figure.



## Step 2: Selecting a Data Set in an HDF File

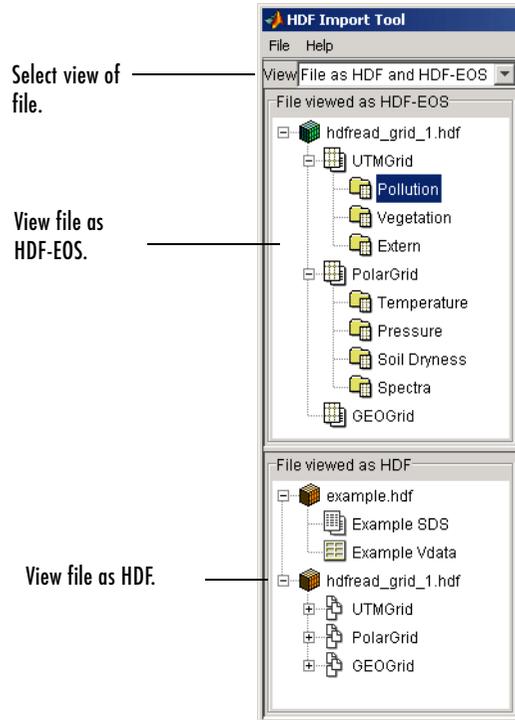
To select a data set to import, use the **Contents** pane to navigate through the file. Click the plus sign at the left of the filename to expand the hierarchical table of contents and see a list of the data sets in the file. (For HDF-EOS files, you can choose how the HDF Import Tool displays the contents. See “Specifying the View in the Contents Pane” on page 6-71 for more information.)

For example, this figure shows the expanded table of contents in the **Contents** pane, with the data set `Example SDS` selected. Note how the **Metadata** pane now displays information about the data set and the **Importing and Subsetting** pane displays subsetting options available for the data set.



**Specifying the View in the Contents Pane.** If you are opening an HDF-EOS file, you can specify whether you want to view the contents of the file as an HDF4 file or as an HDF-EOS file, or both, as illustrated in the following figure. Note, however, that the contents of HDF4 files cannot be viewed as HDF-EOS files.

**Note** Although HDF-EOS files can appear in both windows of the **Contents** pane, remember that you are getting two different views of one file.



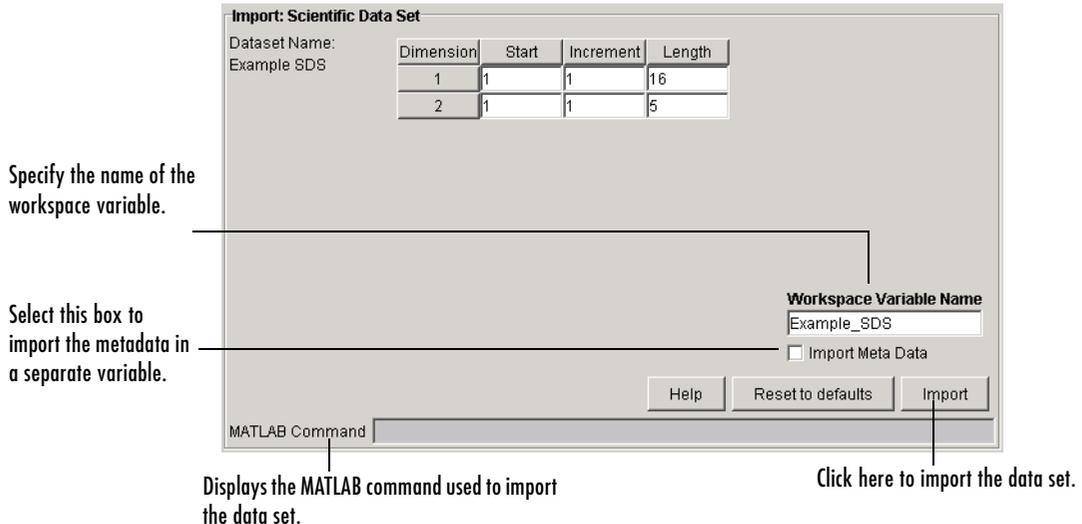
### Step 3: Specifying a Subset of the Data

When you select a data set, the **Importing and Subsetting** pane displays the subsetting options available for that type of data set. The subsetting options available depend on the type of data set. For more information, see “Using the HDF4 Import Tool Subsetting Options” on page 6-74.

### Step 4: Importing Data and Metadata

To import the data set you have selected, click the **Import** button in the bottom right corner of the **Importing and Subsetting** pane, as illustrated in the following figure. Using the **Importing and Subsetting** pane, you can

- Specify the name of the workspace variable — By default, the HDF Import Tool uses the name of the HDF4 data set as the name of the MATLAB workspace variable. In the figure, the variable name is `Example_SDS`. To specify another name, enter text in the **Workspace Variable Name** text box.
- Specify whether to import metadata associated with the data set — To import any metadata that might be associated with the data set, select the **Import Metadata** check box. To store the metadata, the HDF Import Tool creates a second variable in the workspace with the same name with “\_info” appended to it. For example, the name of the metadata variable in the figure is `Example_SDS_info`.
- Save the MATLAB import command syntax — The **MATLAB Command** text window displays the import command used to import the data set. This text is not editable, but you can select it to copy and paste it into the MATLAB Command Window or text editor to reuse it.



## Using the HDF4 Import Tool Subsetting Options

When you select a data set, the **Importing and Subsetting** pane displays the subsetting options available for that type of data set. For data sets that support multiple, mutually exclusive subsetting options, like HDF-EOS Grid data, the contents of the **Importing and Subsetting** pane change when you select one of the options. The following sections describe these subsetting options for all supported data set types. For general information about the tool, see “Using the HDF Import Tool” on page 6-69.

- “HDF Scientific Data (SD)” on page 6-74
- “HDF Vdata” on page 6-75
- “HDF-EOS Grid Data” on page 6-76
- “HDF-EOS Point Data” on page 6-80
- “HDF-EOS Swath Data” on page 6-81
- “HDF Raster Image Data” on page 6-84

---

**Note** To use these data subsetting options effectively, you must understand the HDF and HDF-EOS data formats. Use this documentation with the HDF documentation available at the NCSA Web site ([hdf.ncsa.uiuc.edu](http://hdf.ncsa.uiuc.edu)) and the HDF-EOS documentation at the NASA Web site ([hdfEOS.gsfc.nasa.gov/hdfEOS/index.cfm](http://hdfEOS.gsfc.nasa.gov/hdfEOS/index.cfm)).

---

### HDF Scientific Data (SD)

HDF Scientific Data (SD) data sets are multidimensional arrays. You can import a subset of an HDF SD data set by specifying the location, range, and values to be read from the data set.

Dimension	Start	Increment	Length
1	1	1	16
2	1	1	5

The HDF Import Tool displays the subsetting options available, where each row represents a dimension in the data set and each column represents these subsetting parameters:

- **Start** — Specifies the position on the dimension to begin reading. The default value is 1, which starts reading at the first element of each dimension. The values specified must not exceed the size of any dimension of the data set.
- **Increment** — Specifies the interval between the values to read. The default value is 1, which reads every element of the data set.
- **Length** — Specifies how much data to read along each dimension. The default value is the length of the dimension, which causes all the data to be read. The **Length** parameter automatically updates when you change the value of **Start** or **Increment**.

### HDF Vdata

HDF Vdata data sets are tables. You can import a subset of an HDF Vdata data set in two ways:

- By field name
- By record

**Fields.** Select a specific field you want to import.

The screenshot shows a window titled "Data Fields" with a list box containing three items: "Idx", "Temp", and "Dewpt". The "Dewpt" item is highlighted with a blue background, indicating it is the selected field.

**Records.** Specify the range of records you want to import.

The screenshot shows a window titled "First Record" with two input fields. The first field, labeled "First Record", contains the number "1". The second field, labeled "Number of Records", contains the number "10".

## HDF-EOS Grid Data

In HDF-EOS Grid data, a rectilinear grid overlays a map. The map uses a known map projection. The HDF Import Tool supports the following mutually exclusive subsetting options for Grid data:

- Direct Index
- Geographic Box
- Interpolation
- Pixels
- Tile
- Time
- User-Defined

**Direct Index.** You can import a subset of an HDF-EOS Grid data set by specifying the location, range, and values to be read along each dimension.

Dimension	Start	Increment	Length
1	1	1	16
2	1	1	5

Each row represents a dimension in the data set and each column represents these subsetting parameters:

- **Start** — Specifies the position on the dimension to begin reading. The default value is 1, which starts reading at the first element of each dimension. The values specified must not exceed the size of any dimension of the data set.
- **Increment** — Specifies the interval between the values to read. The default value is 1, which reads every element of the data set.
- **Length** — Specifies how much data to read along each dimension. The default value is the length of the dimension, which causes all the data to be read.

**Geographic Box.** You can import a subset of an HDF-EOS Grid data set by specifying the rectangular area of the grid that you are interested in.

Box Corners  
Units: *Decimal Degrees*

Corner 1  
Longitude      Latitude

Corner 2  
Longitude      Latitude

You define the rectangular area of interest by specifying two points that are two corners of the box:

- **Corner 1** — Specify longitude and latitude values in decimal degrees. Typically, **Corner 1** is the upper-left corner of the box.
- **Corner 2** — Specify longitude and latitude values in decimal degrees. Typically, **Corner 2** is the lower-right corner of the box.

When specifying geographic box subsetting, you can optionally further define the subset of data you are interested in by using Time parameters (see “Time” on page 6-79) and by specifying other User-Defined subsetting parameters (see “User-Defined” on page 6-84).

**Interpolation.** Interpolation is the process of estimating a pixel value at a location in between other pixels. In interpolation, the value of a particular pixel is determined by computing the weighted average of some set of pixels in the vicinity of the pixel.

Box Corners  
Units: *Decimal Degrees*

Corner 1  
Longitude      Latitude

Corner 2  
Longitude      Latitude

You define the region used for bilinear interpolation by specifying two points that are two corners of the interpolation area:

- **Corner 1** — Specify longitude and latitude values in decimal degrees. Typically, **Corner 1** is the upper-left corner of the box.
- **Corner 2** — Specify longitude and latitude values in decimal degrees. Typically, **Corner 2** is the lower-right corner of the box.

**Pixels.** You can import a subset of the pixels in a Grid data set by defining a rectangular area over the grid.

Box Corners  
Units: Decimal Degrees

Corner 1  
Longitude      Latitude

Corner 2  
Longitude      Latitude

You define the box by specifying two points that define two corners of the box:

- **Corner 1** — Specify longitude and latitude values in decimal degrees. Typically, **Corner 1** is the upper-left corner of the box.
- **Corner 2** — Specify longitude and latitude values in decimal degrees. Typically, **Corner 2** is the lower-right corner of the box.

**Tile.** In HDF-EOS Grid data, a rectilinear grid overlays a map. Each rectangle defined by the horizontal and vertical lines of the grid is referred to as a *tile*. HDF-EOS Grid data can be stored as tiles. If it is, you can import a subset of the Grid data set by specifying the coordinates of the tile you are interested in.

Tile Coordinates (eg 1,2,3...)

1,1

Tile coordinates are 1-based, with the upper-left corner of a two-dimensional data set identified as 1, 1. In a three-dimensional data set, this tile would be referenced as 1, 1, 1.

**Time.** You can import a subset of the Grid data set by specifying a time period.

Specify these values:

- **Start** — Specifies the start time.
- **Stop** — Specifies the endpoint in the time span.

---

**Note** The units used (hours, minutes, seconds) to specify the time are defined by the data set.

---

Along with these time parameters, you can optionally further define the subset of data to import by supplying user-defined parameters (see “User-Defined” on page 6-84).

**User-Defined.** You can import a subset of the Grid data set by specifying user-defined parameters.

Specify these values:

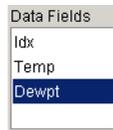
- **Dimension or Field Name** — Specifies the name of the dimension or field to be read from. Dimension names are prefixed with the characters DIM:.
- **Min** — Specifies the start of a range. For dimensions, **Min** represents the start of a range of *elements* to extract. For fields, **Min** represents the start of a range of *values* to extract.
- **Max** — Specifies the endpoint of a range. For dimensions, **Max** represents the end of a range of *elements* to extract. For fields, **Max** represents the end of a range of *values* to extract.

### HDF-EOS Point Data

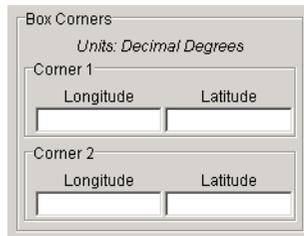
HDF-EOS Point data sets are tables. You can import a subset of an HDF-EOS Point data set by specifying any of these parameters:

- Field name
- Rectangular area of interest
- Record
- Time

**Fields.** Select a specific field you want to import.



**Rectangular Area.** You can import a subset of an HDF-EOS Point data set by specifying the rectangular area that you are interested in.



You define the rectangular area of interest by specifying two points that are two corners of the box:

- **Corner 1** — Specify longitude and latitude values in decimal degrees. Typically, **Corner 1** is the upper-left corner of the box.
- **Corner 2** — Specify longitude and latitude values in decimal degrees. Typically, **Corner 2** is the lower-right corner of the box.

**Records.** Specify the range of records you want to import.

A rectangular input field with a light gray border. The text "Record Numbers" is positioned to the left of the input area, which is currently empty.

**Time.** You can import a subset of the HDF-EOS Point data set by specifying a time period.

A rectangular form with a light gray border. The word "Time" is at the top left. Below it are two labels, "Start" and "Stop", each followed by an empty rectangular input field.

Specify these values:

- **Start** — Specifies the start time.
- **Stop** — Specifies the endpoint in the time span.

---

**Note** The units used (hours, minutes, seconds) to specify the time are defined by the data set.

---

### HDF-EOS Swath Data

HDF-EOS Swath data is data that is produced by a satellite as it traces a path over the earth. This path is called its ground track. The sensor aboard the satellite takes a series of scans perpendicular to the ground track. Swath data can also include a vertical measure as a third dimension. For example, this vertical dimension can represent the height above the Earth of the sensor.

The HDF Import Tool supports the following mutually exclusive subsetting options for Swath data:

- Direct indexing
- Geographic region
- Time
- User-Defined

**Direct Index.** You can import a subset of an HDF-EOS Swath data set by specifying the location, range, and values to be read along each dimension.

Dimension	Start	Increment	Length
1	1	1	16
2	1	1	5

Each row represents a dimension in the data set and each column represents these subsetting parameters:

- **Start** — Specifies the position on the dimension to begin reading. The default value is 1, which starts reading at the first element of each dimension. The values specified must not exceed the size of any dimension of the data set.
- **Increment** — Specifies the interval between the values to read. The default value is 1, which reads every element of the data set.
- **Length** — Specifies how much data to read along each dimension. The default value is the length of the dimension, which causes all the data to be read.

**Geographic Box.** You can import a subset of an HDF-EOS Swath data set by specifying the rectangular area of the grid that you are interested in. When you use this subsetting method, you can also specify the Cross Track Inclusion Mode and the Geolocation Mode.

Box Corners  
Units: Decimal Degrees

Corner 1  
Longitude      Latitude

Corner 2  
Longitude      Latitude

Define the area by specifying two points that specify two corners of the box:

- **Corner 1** — Specify longitude and latitude values in decimal degrees. Typically, **Corner 1** is the upper-left corner of the box.
- **Corner 2** — Specify longitude and latitude values in decimal degrees. Typically, **Corner 2** is the lower-right corner of the box.

For Swath data, you must also specify the **Cross Track Inclusion Mode**. This determines how much of the area of the geographic box that you define must fall within the boundaries of the swath.

 A screenshot of a software interface showing a dropdown menu for 'Cross Track Inclusion Mode'. The menu is currently set to 'AnyPoint' and is open, showing three options: 'AnyPoint', 'Midpoint', and 'Endpoint'. The 'Midpoint' option is highlighted with a blue background.

Select from these values:

- **AnyPoint** — Any part of the box overlaps with the swath.
- **Midpoint** — At least half of the box overlaps with the swath. This is the default.
- **Endpoint** — All of the area defined by the box overlaps with the swath.

For Swath data, you must also specify **Geolocation Mode**. This specifies whether geolocation fields and data must be in the same swath.

 A screenshot of a software interface showing a dropdown menu for 'Geolocation Mode'. The menu is currently set to 'Internal' and is open, showing two options: 'Internal' and 'External'. The 'Internal' option is highlighted with a blue background.

Select from these values:

- **Internal** — Geolocation fields and data fields must be in the same swath.
- **External** — Geolocation fields and data fields can be in different swaths.

**Time.** You can import a subset of the Swath data set by specifying a time period.

 A screenshot of a software interface showing a dialog box titled 'Time'. It contains two input fields labeled 'Start' and 'Stop' for specifying a time period. Both fields are currently empty.

Specify these values:

- **Start** — Specifies the start time.
- **Stop** — Specifies the endpoint in the time span.

---

**Note** The units used (hours, minutes, seconds) to specify the time are defined by the data set.

---

When you use this subsetting method, you must also specify the Cross Track Inclusion Mode and the Geolocation Mode. You can optionally also specify user-defined subsetting options.

**User-Defined.** You can import a subset of the Swath data set by specifying user-defined parameters.

Dimension or Field Name	Min	Max
DIM:Time		
DIM:Time		
DIM:Time		

Specify these values:

- **Dimension or Field Name** — Specifies the name of the dimension or field to be read from. Dimension names are prefixed with the characters DIM:.
- **Min** — Specifies the start of a range. For dimensions, **Min** represents the start of a range of *elements* to extract. For fields, **Min** represents the start of a range of *values* to extract.
- **Max** — Specifies the endpoint of a range. For dimensions, **Max** represents the end of a range of *elements* to extract. For fields, **Max** represents the end of a range of *values* to extract.

### HDF Raster Image Data

No subsetting options are available for HDF raster image data.

## Using the MATLAB `hdfread` Function

To import data from an HDF or HDF-EOS file, you can use the `hdfread` function. The `hdfread` function provides a programmatic way to import data from an HDF4 file, in contrast to the interactive HDF Import Tool, described in “Using the HDF Import Tool” on page 6-69. The `hdfread` function hides many of the details that you need to know if you use the low-level HDF

functions, described in “Using the HDF4 Command-Line Interface” on page 6-88.

To use the `hdfread` function to import data from an HDF4 file, you must know the name of the data set in the file that you want to read. To get this information, use the `hdfinfo` function. This section describes these high-level MATLAB HDF functions, including

- “Getting Information About an HDF4 File” on page 6-85
- “Importing Data from a CDF File” on page 6-55

To export data to an HDF4 file, you must use the low-level functions described in “Exporting MATLAB Data to an HDF4 File” on page 6-96.

### Getting Information About an HDF4 File

To get information about the contents of an HDF4 file, use the `hdfinfo` function. The `hdfinfo` function returns a structure that contains information about the file and the data in the file.

---

**Note** You can also use the HDF Import Tool to get information about the contents of an HDF4 file. See “Using the HDF Import Tool” on page 6-69 for more information.

---

This example returns information about a sample HDF4 file included with MATLAB:

```
info = hdfinfo('example.hdf')

info =

    Filename: 'example.hdf'
         SDS: [1x1 struct]
         Vdata: [1x1 struct]
```

To get information about the data sets stored in the file, look at the SDS field.

### Importing Data from a HDF4 File

To import data into the MATLAB workspace from an HDF4 file, use the `hdfread` function. To use this function, you must specify the data set that you want to read. Use the `hdfinfo` function to get this information.

The following example illustrates this process. This example also illustrates how to import a subset of the data in a data set. See “Reading a Subset of the Data in a Data Set” on page 6-87 for more information.

- 1 Determine the names of data sets in the HDF4 file, using the `hdfinfo` function.

```
info = hdfinfo('example.hdf')
```

```
info =
```

```
    Filename: 'example.hdf'  
         SDS: [1x1 struct]  
         Vdata: [1x1 struct]
```

To determine the names and other information about the data sets in the file, look at the contents of the SDS field. The `Name` field in the SDS structure gives the name of the data set.

```
dsets = info.SDS
```

```
dsets =
```

```
    Filename: 'example.hdf'  
         Type: 'Scientific Data Set'  
         Name: 'Example SDS'  
         Rank: 2  
    DataType: 'int16'  
Attributes: []  
         Dims: [2x1 struct]  
         Label: {}  
Description: {}  
         Index: 0
```

- 2** Read the data set from the HDF4 file, using the `hdfread` function. Specify the name of the data set as a parameter to the function. Note that the data set name is case sensitive. This example returns a 16-by-5 array:

```
dset = hdfread('example.hdf', 'Example SDS');
```

```
dset =
```

```

     3     4     5     6     7
     4     5     6     7     8
     5     6     7     8     9
     6     7     8     9    10
     7     8     9    10    11
     8     9    10    11    12
     9    10    11    12    13
    10    11    12    13    14
    11    12    13    14    15
    12    13    14    15    16
    13    14    15    16    17
    14    15    16    17    18
    15    16    17    18    19
    16    17    18    19    20
    17    18    19    20    21
    18    19    20    21    22

```

**Reading a Subset of the Data in a Data Set.** To read a subset of a data set, you can use the optional `'index'` parameter. The value of the index parameter is a cell array of three vectors that specify the location in the data set to start reading, the skip interval (e.g., read every other data item), and the amount of data to read (e.g., the length along each dimension). In HDF4 terminology, these parameters are called the *start*, *stride*, and *edge* values.

For example, this code

- Starts reading data at the third row, third column (`[3 3]`).
- Reads every element in the array(`[]`).
- Reads 10 rows and 2 columns (`[10 2]`).

```
subset = hdfread('Example.hdf','Example SDS',...
                'Index',{[3 3],[],[10 2]})
```

```
subset =  
  
      7      8  
      8      9  
      9     10  
     10     11  
     11     12  
     12     13  
     13     14  
     14     15  
     15     16  
     16     17
```

## Using the HDF4 Command-Line Interface

This section describes how to use MATLAB functions to access the HDF4 Application Programming Interfaces (APIs). These APIs are libraries of C routines that you can use to import data from an HDF4 file. For a complete list of the HDF APIs supported by MATLAB and the functions you use to access each one, see the hdf reference page.

Topics covered include

- “Understanding the HDF4 to MATLAB Syntax Mapping” on page 6-88
- “Example: Using the HDF4 SD API to Import Data” on page 6-89

---

**Note** You can also use the HDF Import Tool to get information about the contents of an HDF4 file. See “Using the HDF Import Tool” on page 6-69 for more information.

---

## Understanding the HDF4 to MATLAB Syntax Mapping

Each HDF4 API includes many individual routines that you use to read data from files, write data to files, and perform other related functions. For example, the HDF Scientific Data (SD) API includes separate C routines to open (SDopen), close (SDend), and read data (SDreaddata).

Instead of supporting each routine in the HDF APIs, MATLAB provides a single function that serves as a gateway to all the routines in a particular HDF

API. For example, the HDF Scientific Data (SD) API includes the C routine `SDend` to close an HDF file:

```
status = SDend(sd_id); /* C code */
```

To call this routine from MATLAB, use the MATLAB function associated with the SD API, `hdfsd`. You must specify the name of the routine, minus the API acronym, as the first argument and pass any other required arguments to the routine in the order they are expected. For example,

```
status = hdfsd('end',sd_id); % MATLAB code
```

**Handling HDF Routines with Output Arguments.** Some HDF API routines use output arguments to return data. Because MATLAB does not support output arguments, you must specify these arguments as return values.

For example, the `SDfileinfo` routine returns data about an HDF file in two output arguments, `ndatasets` and `nglobal_atts`:

```
status = SDfileinfo(sd_id, ndatasets, nglobal_atts); /* C code */
```

To call this routine from MATLAB, change the output arguments into return values:

```
[ndatasets, nglobal_atts, status] = hdfsd('fileinfo',sd_id);
```

Specify the return values in the same order as they appear as output arguments. The function `status` return value is always specified as the last return value.

### Example: Using the HDF4 SD API to Import Data

To illustrate using HDF4 API routines in MATLAB, this section describes how to import HDF4 Scientific Data (SD) into the MATLAB workspace. The following gives an overview of the steps required by the SD API to import data from an HDF file. The following sections walk you through a detailed example.

- “Step 1: Opening the HDF4 File” on page 6-90
- “Step 2: Retrieving Information About the HDF File” on page 6-91
- “Step 3: Retrieving Attributes from an HDF File (Optional)” on page 6-91
- “Step 4: Selecting the Data Sets to Import” on page 6-92
- “Step 5: Getting Information About a Data Set” on page 6-92

- “Step 6: Reading Data from the HDF File” on page 6-93
- “Step 7: Closing the HDF4 Data Set” on page 6-95
- “Step 8: Closing the HDF File” on page 6-95

---

**Note** The following sections, when referring to specific routines in the HDF SD API, use the C library name rather than the MATLAB function name. The MATLAB syntax is used in all examples.

---

**Step 1: Opening the HDF4 File.** To import an HDF SD data set, you must first open the file using the SD API routine `SDstart`. In MATLAB, you use the `hdfsd` function, specifying as arguments:

- Name of the SD API routine, `start` in this case.
- Name of the file you want to open.
- Mode in which you want to open it. The following table lists the file access modes supported by the `SDstart` routine. In MATLAB, you specify these modes as text strings. You can specify the full HDF mode name or one of the abbreviated forms listed in the table.

HDF File Creation Mode	HDF Mode Name	MATLAB String
Create a new file	'DFACC_CREATE'	'create'
Read access	'DFACC_RDONLY'	'read' or 'rdonly'
Read and write access	'DFACC_RDWR'	'rdwr' or 'write'

For example, this code opens the file `mydata.hdf` for read access:

```
sd_id = hdfsd('start', 'mydata.hdf', 'read');
```

If `SDstart` can find and open the file specified, it returns an HDF SD file identifier, named `sd_id` in the example. Otherwise, it returns `-1`.

**Step 2: Retrieving Information About the HDF File.** To get information about an HDF4 file, you must use the SD API routine `SDfileinfo`. This function returns the number of data sets in the file and the number of global attributes in the file, if any. (For more information about global attributes, see “Exporting MATLAB Data to an HDF4 File” on page 6-96.) In MATLAB, you use the `hdfsd` function, specifying the following arguments:

- Name of the SD API routine, `fileinfo` in this case
- SD file identifier, `sd_id`, returned by `SDstart`

In this example, the HDF4 file contains three data sets and one global attribute.

```
[ndatasets, nglobal_atts, stat] = hdfsd('fileinfo',sd_id)

ndatasets =
    3

nglobal_atts =
    1

status =
    0
```

**Step 3: Retrieving Attributes from an HDF File (Optional).** HDF files can optionally include information, called *attributes*, that describes the data the file contains. Attributes associated with an entire HDF file are called *global* attributes. Attributes associated with a data set are called *local* attributes. (You can also associate attributes with files or dimensions. For more information, see “Step 4: Writing Metadata to an HDF File” on page 6-100.)

To retrieve attributes from an HDF file, use the HDF4 API routine `SDreadattr`. In MATLAB, use the `hdfsd` function, specifying as arguments:

- Name of the SD API routine, `readattr` in this case.
- File identifier (`sd_id`) returned by `SDstart`, for global attributes, or the data set identifier for local attributes. (See “Step 4: Selecting the Data Sets to Import” on page 6-92 to learn how to get a data set identifier.)
- Index identifying the attribute you want to view. HDF uses zero-based indexing. If you know the name of an attribute but not its index, use the

SDfindattr routine to determine the index value associated with the attribute.

For example, this code returns the contents of the first global attribute, which is the character string `my global attribute`:

```
attr_idx = 0;
[attr, status] = hdfsd('readattr', sd_id, attr_idx);

attr =
    my global attribute
```

**Step 4: Selecting the Data Sets to Import.** To select a data set, use the SD API routine `SDselect`. In MATLAB, you use the `hdfsd` function, specifying as arguments:

- Name of the SD API routine, `select` in this case
- HDF SD file identifier (`sd_id`) returned by `SDstart`

If `SDselect` finds the specified data set in the file, it returns an HDF SD data set identifier, called `sds_id` in the example. If it cannot find the data set, it returns `-1`.

---

**Note** Do not confuse HDF SD *file* identifiers, named `sd_id` in the examples, with HDF SD *data set* identifiers, named `sds_id` in the examples.

---

```
sds_id = hdfsd('select',sd_id,1)
```

**Step 5: Getting Information About a Data Set.** To read a data set, you must get information about the data set, such as its name, size, and data type. In the HDF SD API, you use the `SDgetinfo` routine to gather this information. In MATLAB, use the `hdfsd` function, specifying as arguments:

- Name of the SD API routine, `getinfo` in this case
- HDF SD data set identifier (`sds_id`) returned by `SDselect`

This code retrieves information about the data set identified by `sds_id`:

```
[dsname, dsndims, dsdims, dstype, dsatts, stat] =
    hdfsd('getinfo',sds_id)
```

```
dsname =  
    A  
  
dsndims =  
    2  
  
dsdims =  
    5    3  
  
dstype =  
    double  
  
dsatts =  
    0  
  
stat =  
    0
```

**Step 6: Reading Data from the HDF File.** To read data from an HDF4 file, you must use the `SDreaddata` routine. In MATLAB, use the `hdf_sd` function, specifying as arguments:

- Name of the SD API function, `readdata` in this case.
- HDF SD data set identifier (`sds_id`) returned by `SDselect`.
- Location in the data set where you want to start reading data, specified as a vector of index values, called the *start* vector. To read from the beginning of a data set, specify zero for each element of the start vector. Use `SDgetinfo` to determine the dimensions of the data set.
- Number of elements along each dimension to skip between each read operation, specified as a vector of scalar values, called the *stride* vector. To read every element of a data set, specify 1 as the value for each element of the vector or specify an empty array (`[]`).
- Total number of elements to read along each dimension, specified as a vector of scalar values, called the *edges* vector. To read every element of a data set, set each element of the edges vector to the size of each dimension of the data set. Use `SDgetinfo` to determine these sizes.

---

**Note** `SDgetinfo` returns dimension values in row-major order, the ordering used by HDF. Because MATLAB stores data in column-major order, you must specify the dimensions in column-major order, that is, `[columns, rows]`. In addition, you must use zero-based indexing in these arguments.

---

For example, to read the entire contents of a data set, use this code:

```
[ds_name, ds_ndims, ds_dims, ds_type, ds_atts, stat] =
    hdfsd('getinfo',sds_id);

ds_start = zeros(1,ds_ndims); % Creates the vector [0 0]
ds_stride = [];
ds_edges = ds_dims;

[ds_data, status] =
    hdfsd('readdata',sds_id,ds_start,ds_stride,ds_edges);

disp(ds_data)
     1     2     3     4     5
     6     7     8     9    10
    11    12    13    14    15
```

To read less than the entire data set, use the `start`, `stride`, and `edges` vectors to specify where you want to start reading data and how much data you want to read. For example, this code reads the entire second row of the sample data set:

```
ds_start = [0 1] % Start reading at the first column, second row
ds_stride = []; % Read each element
ds_edges = [5 1]; % Read a 1-by-5 vector of data

[ds_data, status] =
    hdfsd('readdata',sds_id,ds_start,ds_stride,ds_edges);
```

**Step 7: Closing the HDF4 Data Set.** After writing data to a data set in an HDF file, you must close access to the data set. In the HDF SD API, you use the `SDendaccess` routine to close a data set. In MATLAB, use the `hdfsd` function, specifying as arguments:

- Name of the SD API routine, `endaccess` in this case
- HDF SD data set identifier (`sds_id`) returned by `SDselect`

For example, this code closes the data set:

```
stat = hdfsd('endaccess',sds_id);
```

You must close access to all the data sets in an HDF file before closing it.

**Step 8: Closing the HDF File.** After writing data to a data set and closing the data set, you must also close the HDF file. In the HDF SD API, you use the `SDend` routine. In MATLAB, use the `hdfsd` function, specifying as arguments:

- Name of the SD API routine, `end` in this case
- HDF SD file identifier (`sd_id`) returned by `SDstart`

For example, this code closes the data set:

```
stat = hdfsd('end',sd_id);
```

## Exporting MATLAB Data to an HDF4 File

To export data from MATLAB in an HDF file, you must use the functions in the HDF API associated with the HDF data type. Each API has a particular programming model, that is, a prescribed way to use the routines to write data sets to the file. (In HDF terminology, the numeric arrays stored in HDF files are called data sets.)

To illustrate this concept, this section describes the programming model of one particular HDF API: the HDF Scientific Data (SD) API. For a complete list of the HDF APIs supported by MATLAB, see the `hdf` reference page.

---

**Note** This section does not attempt to describe all HDF features and routines. To use the MATLAB HDF functions effectively, you must refer to the official NCSA documentation at the NCSA Web site (`hdf.ncsa.uiuc.edu`).

---

### Example: Exporting Data to an HDF4 File

The programming model for exporting HDF SD data involves these steps:

- “Step 1: Creating an HDF File” on page 6-96
- “Step 2: Creating an HDF Data Set” on page 6-97
- “Step 3: Writing MATLAB Data to an HDF File” on page 6-98
- “Step 4: Writing Metadata to an HDF File” on page 6-100
- “Step 5: Closing HDF Data Sets” on page 6-102
- “Step 6: Closing an HDF File” on page 6-102

#### Step 1: Creating an HDF File

To export MATLAB data in HDF format, you must first create an HDF file, or open an existing one. In the HDF SD API, you use the `SDstart` routine. In MATLAB, use the `hdfsd` function, specifying `start` as the first argument. As other arguments, specify

- A text string specifying the name you want to assign to the HDF file (or the name of an existing HDF file)
- A text string specifying the HDF SD interface file access mode

For example, this code creates an HDF file named `mydata.hdf`:

```
sd_id = hdfsd('start', 'mydata.hdf', 'DFACC_CREATE');
```

When you specify the `DFACC_CREATE` access mode, `SDstart` creates the file and initializes the HDF SD multifile interface, returning an HDF SD file identifier, named `sd_id` in the example.

If you specify `DFACC_CREATE` mode and the file already exists, `SDstart` fails, returning `-1`. To open an existing HDF file, you must use HDF read or write modes. For information about using `SDstart` in these modes, see “Step 1: Opening the HDF4 File” on page 6-90.

## Step 2: Creating an HDF Data Set

After creating the HDF file, or opening an existing one, you must create a data set in the file for each MATLAB array you want to export.

---

**Note** To write data to an existing data set, you can skip this step.

---

In the HDF SD API, you use the `SDcreate` routine to create data sets. In MATLAB, you use the `hdfsd` function, specifying as arguments:

- Name of the SD API routine, `create` in this case
- Valid HDF SD file identifier, `sd_id`, returned by `SDstart`
- Name you want assigned to the data set
- Data type of the data set. For information about specifying data types, see “Importing HDF4 and HDF-EOS Data” on page 6-69.
- Number of dimensions in the data set. This is called the *rank* of the data set in HDF terminology.
- Size of each dimension, specified as a vector

The values you assign to these arguments depend on the MATLAB array you want to export. For example, to export the following MATLAB 3-by-5 array of doubles,

```
A = [ 1 2 3 4 5 ; 6 7 8 9 10 ; 11 12 13 14 15 ];
```

you could set the values of these arguments as in this code fragment:

```
ds_name = 'A';
ds_type = 'double';
ds_rank = ndims(A);
ds_dims = fliplr(size(A));

sds_id = hdfsd('create',sd_id,ds_name,ds_type,ds_rank,ds_dims);
```

If `SDcreate` can successfully create the data set, it returns an HDF SD data set identifier, (`sds_id`). Otherwise, `SDcreate` returns -1.

---

**Note** In this example, note how the code fragment reverses the order of the values in the dimensions argument (`ds_dims`). This processing is necessary because the MATLAB `size` function returns the dimensions in column-major order and HDF expects to receive dimensions in row-major order.

---

Once you create a data set, you cannot change its characteristics. You can, however, modify the data it contains. To do this, initiate access to the data set, using `SDselect`, and write to the data set as described in “Step 3: Writing MATLAB Data to an HDF File” on page 6-98.

### Step 3: Writing MATLAB Data to an HDF File

After creating an HDF file and creating a data set in the file, you can write data to the entire data set or just a portion of the data set. In the HDF SD API, you use the `SDwritedata` routine. In MATLAB, use the `hdfsd` function, specifying specifying as arguments:

- Name of the SD API routine, `writedata` in this case
- Valid HDF SD data set identifier, `sds_id`, returned by `SDcreate`
- Location in the data set where you want to start writing data, called the *start* vector in HDF terminology
- Number of elements along each dimension to skip between each write operation, called the *stride* vector in HDF terminology
- Total number of elements to write along each dimension, called the *edges* vector in HDF terminology
- MATLAB array to be written

---

**Note** You must specify the values of the start, stride, and edges arguments in row-major order, rather than the column-major order used in MATLAB. Note how the example uses `fliplr` to reverse the order of the dimensions in the vector returned by the `size` function before assigning it as the value of the edges argument.

---

The values you assign to these arguments depend on the MATLAB array you want to export. For example, the following code fragment writes this MATLAB 3-by-5 array of doubles,

```
A = [ 1 2 3 4 5; 6 7 8 9 10; 11 12 13 14 15 ];
```

into an HDF file:

```
ds_start = zeros(1:ndims(A)); % Start at the beginning
ds_stride = []; % Write every element.
ds_edges = fliplr(size(A)); % Reverse the dimensions.

stat = hdfsd('writedata',sds_id,...
            ds_start, ds_stride, ds_edges, A)
```

If it can write the data to the data set, `SDwritedata` returns 0; otherwise, it returns -1.

---

**Note** `SDwritedata` queues write operations. To ensure that these queued write operations are executed, you must close the file, using the `SDend` routine. See “Step 6: Closing an HDF File” on page 6-102 for more information. As a convenience, MATLAB provides a function, `MLcloseall`, that you can use to close all open data sets and file identifiers with a single call. See “Using the MATLAB HDF Utility API” on page 6-103 for more information.

---

**Writing Data to Portions of Data Sets.** To write less than the entire data set, use the start, stride, and edges vectors to specify where you want to start writing data and how much data you want to write.

For example, the following code fragment uses `SDwritedata` to replace the values of the entire second row of the sample data set:

```
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15
```

with the vector `B`:

```
B = [ 9 9 9 9 9];
```

In the example, the start vector specifies that you want to start the write operation in the first column of the second row. Note how HDF uses zero-based indexing and specifies the column dimension first. In MATLAB, you would specify this location as  $(2, 1)$ . The `edges` argument specifies the dimensions of the data to be written. Note that the size of the array of data to be written must match the edge specification.

```
ds_start = [0 1]; % Start writing at the first column, second row.
ds_stride = []; % Write every element.
ds_edges = [5 1]; % Each row is a 1-by-5 vector.

stat = hdfsd('writedata',sds_id,ds_start,ds_stride,ds_edges,B);
```

#### Step 4: Writing Metadata to an HDF File

You can optionally include information in an HDF file, called attributes, that describes the file and its contents. Using the HDF SD API, you can associate attributes with three types of HDF objects:

- An entire HDF file — File attributes, also called *global* attributes, generally contain information pertinent to all the data sets in the file.
- A data set in an HDF file — Data set attributes, also called *local* attributes, describe individual data sets.
- A dimension of a data set — Dimension attributes provide information about one particular dimension of a data set.

**Associating Multiple Attributes with a Single Object.** You can associate multiple attributes with a single HDF object. HDF maintains an attribute index for each object. The attribute index is zero-based. The first attribute has index value 0, the second has index value 1, and so on. You access an attribute by its index value.

Each attribute has the format `name=value`, where `name` (called `label` in HDF terminology) is a text string up to 256 characters in length and `value` contains one or more entries of the same data type. A single attribute can have multiple values.

**Creating Attributes.** To create an attribute in the HDF SD API, use the `SDsetattr` routine. In MATLAB, use the `hdfsd` function, specifying `setattr` as the first argument. As other arguments, specify

- A valid HDF SD identifier associated with the object. This value can be a file identifier (`sd_id`), a data set identifier (`sds_id`), or a dimension identifier (`dim_id`).
- A text string that defines the name of the attribute. The SD interface supports predefined attributes that have reserved names and, in some cases, data types. For information about these attributes, see “Creating Predefined Attributes” on page 6-101.
- The attribute value.

For example, this code creates a global attribute, named `my_global_attr`, and associates it with the HDF file identified by `sd_id`:

```
status = hdfsd('setattr',sd_id,'my_global_attr','my_attr_val');
```

---

**Note** In the NCSA documentation, the `SDsetattr` routine has two additional arguments: data type and the number of values in the attribute. When calling this routine from MATLAB, you do not have to include these arguments. The MATLAB HDF function can determine the data type and size of the attribute from the value you specify.

---

**Creating Predefined Attributes.** Predefined attributes are identical to user-defined attributes except that the HDF SD API has already defined their names and data types. For example, the HDF SD API defines an attribute, named `coordsys`, in which you can specify the coordinate system used by the data set. Possible values of this attribute include the text strings `'cartesian'`, `'polar'`, and `'spherical'`.

Predefined attributes can be useful because they establish conventions that applications can depend on. The HDF SD API supports predefined attributes for data sets and dimensions only; there are no predefined attributes for files. For a complete list of the predefined attributes, see the NCSA documentation.

In the HDF SD API, you create predefined attributes the same way you create user-defined attributes, using the `SDsetattr` routine. In MATLAB, use the `hdfsd` function, specifying `setattr` as the first argument:

```
attr_name = 'cordsys';  
attr_value = 'polar';  
  
status = hdfsd('setattr',sds_id,attr_name,attr_value);
```

The HDF SD API also includes specialized functions for writing and reading the predefined attributes. These specialized functions, such as `SDsetdatastrs`, are sometimes easier to use, especially when you are reading or writing multiple related predefined attributes. You must use specialized functions to read or write the predefined dimension attributes.

### **Step 5: Closing HDF Data Sets**

After writing data to a data set in an HDF file, you must close access to the data set. In the HDF SD API, you use the `SDendaccess` routine to close a data set. In MATLAB, use the `hdfsd` function, specifying `endaccess` as the first argument. As the only other argument, specify a valid HDF SD data set identifier, `sds_id` in this example:

```
stat = hdfsd('endaccess',sds_id);
```

### **Step 6: Closing an HDF File**

After writing data to a data set and closing the data set, you must also close the HDF file. In the HDF SD API, you use the `SDend` routine. In MATLAB, use the `hdfsd` function, specifying `end` as the first argument. As the only other argument, specify a valid HDF SD file identifier, `sd_id` in this example:

```
stat = hdfsd('end',sd_id);
```

You must close access to all the data sets in an HDF file before closing it.

---

**Note** Closing an HDF file executes all the write operations that have been queued using `SDwritedata`. As a convenience, the MATLAB HDF Utility API provides a function, `MLcloseall`, that can close all open data set and file identifiers with a single call. See “Using the MATLAB HDF Utility API” on page 6-103 for more information.

---

## Using the MATLAB HDF Utility API

In addition to the standard HDF APIs, listed in “Using HDF5 Data Type Objects” on page 6-67, MATLAB supports utility functions that are designed to make using HDF in the MATLAB environment easier.

For example, the MATLAB HDF utility API includes a function, `MLlistinfo`, that you can use to view all types of open HDF identifiers, such as HDF SD file identifiers. MATLAB updates these lists whenever HDF identifiers are created or closed.

This code obtains a list of all open HDF file and data set identifiers, using the `MLlistinfo` function. In this example, only two identifiers are open:

```
hdfml('listinfo')
No open RI identifiers
No open GR identifiers
No open grid identifiers
No open grid file identifiers
No open annotation identifiers
No open AN identifiers
Open scientific dataset identifiers:
    262144
Open scientific data file identifiers:
    393216
No open Vdata identifiers
No open Vgroup identifiers
No open Vfile identifiers
No open point identifiers
No open point file identifiers
```

```
No open swath identifiers
No open swath file identifiers
No open access identifiers
No open file identifiers
```

### **Closing All Open HDF Identifiers**

To close all open HDF identifiers in a single call, use the `MLcloseall` function. This call closes all open HDF identifiers:

```
hdfml('closeall')
```

## Using Low-Level File I/O Functions

MATLAB includes a set of low-level file I/O functions that are based on the I/O functions of the ANSI Standard C Library. If you know C, you are probably familiar with these routines. This section covers

- “Opening Files” on page 6-106
- “Reading Binary Data” on page 6-108
- “Writing Binary Data” on page 6-110
- “Controlling Position in a File” on page 6-110
- “Reading Strings Line by Line from Text Files” on page 6-112
- “Reading Formatted ASCII Data” on page 6-114
- “Writing Formatted Text Files” on page 6-115
- “Closing a File” on page 6-117

The MATLAB file I/O functions use the same programming model as the C language routines. To read or write data, you perform these steps:

- 1** Open the file, using `fopen`. `fopen` returns a file identifier that you use with all the other low-level file I/O routines.
- 2** Operate on the file.
  - a** Read binary data, using `fread`.
  - b** Write binary data, using `fwrite`.
  - c** Read text strings from a file line-by-line, using `fgets/fgetl`.
  - d** Read formatted ASCII data, using `fscanf`.
  - e** Write formatted ASCII data, using `fprintf`.
- 3** Close the file, using `fclose`.

This section also describes how these functions affect the current position in the file where read or write operations happen and how you can change the position in the file.

---

**Note** While the MATLAB file I/O commands are modeled on the C language I/O routines, in some ways their behavior is different. For example, the `fread` function is *vectorized*; that is, it continues reading until it encounters a text string or the end of file. These sections, and the MATLAB reference pages for these functions, highlight any differences in behavior.

---

## Opening Files

Before reading or writing a text or binary file, you must open it with the `fopen` command.

```
fid = fopen('filename','permission')
```

### Specifying the Permission String

The permission string specifies the kind of access to the file you require. Possible permission strings include

- `r` for reading only
- `w` for writing only
- `a` for appending only
- `r+` for both reading and writing

---

**Note** Systems such as Microsoft Windows that distinguish between text and binary files might require additional characters in the permission string, such as `'rb'` to open a binary file for reading.

---

### Using the Returned File Identifier (`fid`)

If successful, `fopen` returns a nonnegative integer, called a *file identifier* (`fid`). You pass this value as an argument to the other I/O functions to access the open file. For example, this `fopen` statement opens the data file named `penny.dat` for reading:

```
fid = fopen('penny.dat','r')
```

If `fopen` fails, for example if you try to open a file that does not exist, `fopen`

- Assigns `-1` to the file identifier.
- Assigns an error message to an optional second output argument. Note that the error messages are system dependent and are not provided for all errors on all systems. The function `error` can also provide information about errors.

Test the file identifier each time you open a file in your code. For example, this code loops until a readable filename is entered:

```
fid=0;
while fid < 1
    filename=input('Open file: ', 's');
    [fid,message] = fopen(filename, 'r');
    if fid == -1
        disp(message)
    end
end
```

When you run this code, if you specify a file that doesn't exist, such as `nofile.mat`, at the `Open file:` prompt, the results are

```
Open file: nofile.mat
Sorry. No help in figuring out the problem . . .
```

If you specify a file that does exist, such as `goodfile.mat`, the code example returns the file identifier, `fid`, and exits the loop.

```
Open file: goodfile.mat
```

## Opening Temporary Files and Directories

The `tempdir` and `tempname` functions assist in locating temporary data on your system.

Function	Purpose
<code>tempdir</code>	Get temporary directory name.
<code>tempname</code>	Get temporary filename.

Use these functions to create temporary files. Some systems delete temporary files every time you reboot the system. On other systems, designating a file as temporary can mean only that the file is not backed up.

The `tempdir` function returns the name of the directory or folder that has been designated to hold temporary files on your system. For example, issuing `tempdir` on a UNIX system returns the `/tmp` directory.

MATLAB also provides a `tempname` function that returns a filename in the temporary directory. The returned filename is a suitable destination for temporary data. For example, if you need to store some data in a temporary file, then you might issue the following command first:

```
fid = fopen(tempname, 'w');
```

---

**Note** The filename that `tempname` generates is not guaranteed to be unique; however, it is likely to be so.

---

## Reading Binary Data

The `fread` function reads all or part of a binary file (as specified by a file identifier) and stores it in a matrix. In its simplest form, it reads an entire file and interprets each byte of input as the next element of the matrix. For example, the following code reads the data from a file named `nickel.dat` into matrix `A`:

```
fid = fopen('nickel.dat','r');  
A = fread(fid);
```

To echo the data to the screen after reading it, use `char` to display the contents of `A` as characters, transposing the data so it is displayed horizontally:

```
disp(char(A'))
```

The `char` function causes MATLAB to interpret the contents of `A` as characters instead of as numbers. Transposing `A` displays it in its more natural horizontal format.

## Controlling the Number of Values Read

`fread` accepts an optional second argument that controls the number of values read (if unspecified, the default is the entire file). For example, this statement reads the first 100 data values of the file specified by `fid` into the column vector `A`.

```
A = fread(fid,100);
```

Replacing the number 100 with the matrix dimensions `[10 10]` reads the same 100 elements into a 10-by-10 array.

## Controlling the Data Type of Each Value

An optional third argument to `fread` controls the data type of the input. The data type argument controls both the number of bits read for each value and the interpretation of those bits as character, integer, or floating-point values. MATLAB supports a wide range of precisions, which you can specify with MATLAB specific strings or their C or Fortran equivalents.

Some common precisions include

- 'char' and 'uchar' for signed and unsigned characters (usually 8 bits)
- 'short' and 'long' for short and long integers (usually 16 and 32 bits, respectively)
- 'float' and 'double' for single- and double-precision floating-point values (usually 32 and 64 bits, respectively)

---

**Note** The meaning of a given precision can vary across different hardware platforms. For example, a 'uchar' is not always 8 bits. `fread` also provides a number of more specific precisions, such as 'int8' and 'float32'. If in doubt, use precisions that are not platform dependent. See `fread` for a complete list of precisions.

---

For example, if `fid` refers to an open file containing single-precision floating-point values, then the following command reads the next 10 floating-point values into a column vector `A`:

```
A = fread(fid,10,'float');
```

## Writing Binary Data

The `fwrite` function writes the elements of a matrix to a file in a specified numeric precision, returning the number of values written. For instance, these lines create a 100-byte binary file containing the 25 elements of the 5-by-5 magic square, each stored as 4-byte integers:

```
fwriteid = fopen('magic5.bin','w');
count = fwrite(fwriteid,magic(5),'int32');
status = fclose(fwriteid);
```

In this case, `fwrite` sets the count variable to 25 unless an error occurs, in which case the value is less.

## Controlling Position in a File

Once you open a file with `fopen`, MATLAB maintains a file position indicator that specifies a particular location within a file. MATLAB uses the file position indicator to determine where in the file the next read or write operation will begin. The following sections describe how to

- Determine whether the file position indicator is at the end of the file
- Move to a specific location in the file
- Retrieve the current location of the file position indicator
- Reset the file position indicator to the beginning of the file

## Determining End-of-File

The `fseek` and `ftell` functions let you set and query the position in the file at which the next input or output operation takes place:

- The `fseek` function repositions the file position indicator, letting you skip over data or back up to an earlier part of the file.
- The `ftell` function gives the offset in bytes of the file position indicator for a specified file.

The syntax for `fseek` is

```
status = fseek(fid,offset,origin)
```

`fid` is the file identifier for the file. `offset` is a positive or negative offset value, specified in bytes. `origin` is one of the following strings that specify the location in the file from which to calculate the position.

```
'bof'    Beginning of file
'cof'    Current position in file
'eof'    End of file
```

### Understanding File Position

To see how `fseek` and `ftell` work, consider this short M-file:

```
A = 1:5;
fid = fopen('five.bin','w');
fwrite(fid, A,'short');
status = fclose(fid);
```

This code writes out the numbers 1 through 5 to a binary file named `five.bin`. The call to `fwrite` specifies that each numerical element be stored as a short. Consequently, each number uses two storage bytes.

Now reopen `five.bin` for reading:

```
fid = fopen('five.bin','r');
```

This call to `fseek` moves the file position indicator forward 6 bytes from the beginning of the file:

```
status = fseek(fid,6,'bof');
```

File Position	bof	1	2	3	4	5	6	7	8	9	10	eof
File Contents		0	1	0	2	0	3	0	4	0	5	
File Position Indicator								↑				

This call to `fread` reads whatever is at file positions 7 and 8 and stores it in variable `four`:

```
four = fread(fid,1,'short');
```

The act of reading advances the file position indicator. To determine the current file position indicator, call `ftell`:

```
position = ftell(fid)

position =

    8
```

File Position	bof	1	2	3	4	5	6	7	8	9	10	eof
File Contents		0	1	0	2	0	3	0	4	0	5	
File Position Indicator										↑		

This call to `fseek` moves the file position indicator back 4 bytes:

```
status = fseek(fid,-4,'cof');
```

File Position	bof	1	2	3	4	5	6	7	8	9	10	eof
File Contents		0	1	0	2	0	3	0	4	0	5	
File Position Indicator					↑							

Calling `fread` again reads in the next value (3):

```
three = fread(fid,1,'short');
```

## Reading Strings Line by Line from Text Files

MATLAB provides two functions, `fgetl` and `fgets`, that read lines from formatted text files and store them in string vectors. The two functions are almost identical; the only difference is that `fgets` copies the newline character to the string vector but `fgetl` does not.

The following M-file function demonstrates a possible use of `fgetl`. This function uses `fgetl` to read an entire file one line at a time. For each line, the function determines whether an input literal string (`literal`) appears in the line.

If it does, the function prints the entire line preceded by the number of times the literal string appears on the line.

```
function y = litcount(filename, literal)
% Search for number of string matches per line.

fid = fopen(filename, 'rt');
y = 0;
while feof(fid) == 0
    tline = fgetl(fid);
    matches = findstr(tline, literal);
    num = length(matches);
    if num > 0
        y = y + num;
        fprintf(1, '%d:%s\n', num, tline);
    end
end
fclose(fid);
```

For example, consider the following input data file called badpoem:

```
Oranges and lemons,
Pineapples and tea.
Orangutans and monkeys,
Dragonflys or fleas.
```

To find out how many times the string 'an' appears in this file, use `litcount`:

```
litcount('badpoem', 'an')
2: Oranges and lemons,
1: Pineapples and tea.
3: Orangutans and monkeys,
```

## Reading Formatted ASCII Data

The `fscanf` function is like the `fscanf` function in standard C. Both functions operate in a similar manner, reading data from a file and assigning it to one or more variables. Both functions use the same set of conversion specifiers to control the interpretation of the input data.

The conversion specifiers for `fscanf` begin with a `%` character; common conversion specifiers include

Conversion Specifier	Description
<code>%s</code>	Match a string.
<code>%d</code>	Match an integer in base 10 format.
<code>%g</code>	Match a double-precision floating-point value.

You can also specify that `fscanf` skip a value by specifying an asterisk in a conversion specifier. For example, `%*f` means skip the floating-point value in the input data; `%*d` means skip the integer value in the input data.

### Differences Between the MATLAB `fscanf` and the C `fscanf`

Despite all the similarities between the MATLAB and C versions of `fscanf`, there are some significant differences. For example, consider a file named `moon.dat` for which the contents are as follows:

```
3.654234533
2.71343142314
5.34134135678
```

The following code reads all three elements of this file into a matrix named `MyData`:

```
fid = fopen('moon.dat','r');
MyData = fscanf(fid,'%g');
status = fclose(fid);
```

Notice that this code does not use any loops. Instead, the `fscanf` function continues to read in text as long as the input format is compatible with the format specifier.

An optional size argument controls the number of matrix elements read. For example, if `fid` refers to an open file containing strings of integers, then this line reads 100 integer values into the column vector `A`:

```
A = fscanf(fid, '%5d', 100);
```

This line reads 100 integer values into the 10-by-10 matrix `A`:

```
A = fscanf(fid, '%5d', [10 10]);
```

A related function, `sscanf`, takes its input from a string instead of a file. For example, this line returns a column vector containing 2 and its square root:

```
root2 = num2str([2, sqrt(2)]);
rootvalues = sscanf(root2, '%f');
```

## Writing Formatted Text Files

The `fprintf` function converts data to character strings and outputs them to the screen or a file. A format control string containing conversion specifiers and any optional text specify the output format. The conversion specifiers control the output of array elements; `fprintf` copies text directly.

Common conversion specifiers include

Conversion Specifier	Description
<code>%e</code>	Exponential notation
<code>%f</code>	Fixed-point notation
<code>%g</code>	Automatically select the shorter of <code>%e</code> and <code>%f</code>

Optional fields in the format specifier control the minimum field width and precision. For example, this code creates a text file containing a short table of the exponential function:

```
x = 0:0.1:1;
y = [x; exp(x)];
```

The code below writes `x` and `y` into a newly created file named `exptable.txt`:

```
fid = fopen('exptable.txt','w');
fprintf(fid,'Exponential Function\n\n');
fprintf(fid,'%6.2f %12.8f\n',y);
status = fclose(fid);
```

The first call to `fprintf` outputs a title, followed by two carriage returns. The second call to `fprintf` outputs the table of numbers. The format control string specifies the format for each line of the table:

- A fixed-point value of six characters with two decimal places
- Two spaces
- A fixed-point value of twelve characters with eight decimal places

`fprintf` converts the elements of array `y` in column order. The function uses the format string repeatedly until it converts all the array elements.

Now use `fscanf` to read the exponential data file:

```
fid = fopen('exptable.txt','r');
title = fgetl(fid);
[table,count] = fscanf(fid,'%f %f',[2 11]);
table = table';
status = fclose(fid);
```

The second line reads the file title. The third line reads the table of values, two floating-point values on each line, until it reaches end of file. `count` returns the number of values matched.

A function related to `fprintf`, `sprintf`, outputs its results to a string instead of a file or the screen. For example,

```
root2 = sprintf('The square root of %f is %10.8e.\n',2,sqrt(2));
```

## Closing a File

When you finish reading or writing, use `fclose` to close the file. For example, this line closes the file associated with file identifier `fid`:

```
status = fclose(fid);
```

This line closes all open files:

```
status = fclose('all');
```

Both forms return 0 if the file or files were successfully closed or -1 if the attempt was unsuccessful.

MATLAB automatically closes all open files when you exit from MATLAB. It is still good practice, however, to close a file explicitly with `fclose` when you are finished using it. Not doing so can unnecessarily drain system resources.

---

**Note** Closing a file does not clear the file identifier variable `fid`. However, subsequent attempts to access a file through this file identifier variable will not work.

---

## Exchanging Files over the Internet

MATLAB provides functions for exchanging files over the Internet. You can exchange files using common protocols, such as File Transfer Protocol (FTP), Simple Mail Transport Protocol (SMTP), and HyperText Transfer Protocol (HTTP). In addition, you can create zip archives to minimize the transmitted file size, and you can save and work with Web pages.

For more information, see the following sections:

- “Downloading Web Content and Files” on page 6-118
- “Creating and Uncompressing Zip Archives” on page 6-120
- “Sending E-Mail” on page 6-121
- “Performing FTP File Operations” on page 6-123

### Downloading Web Content and Files

MATLAB provides two functions for downloading Web pages and files using HTTP: `urlread` and `urlwrite`. With the `urlread` function, you can read and save the contents of a Web page to a string variable in the MATLAB workspace. With the `urlwrite` function, you can save a Web page’s content to a file.

Because it creates a string variable in the workspace, the `urlread` function is useful for working with the contents of Web pages in MATLAB. The `urlwrite` function is useful for saving Web pages to a local directory.

---

**Note** When using `urlread`, remember that only the HTML in that specific Web page is retrieved. The hyperlink targets, images, and so on will not be retrieved.

---

If you need to pass parameters to a Web page, the `urlread` and `urlwrite` functions let you use HTTP post and get methods. For more information, see the `urlread` and `urlwrite` reference pages.

#### Example — Using the `urlread` Function

The following procedure demonstrates how to retrieve the contents of the Web page containing the Recent File list at the MATLAB Central File Exchange, <http://www.mathworks.com/matlabcentral/fileexchange/index.jsp>. It

assigns the results to a string variable, `recentFile`, and it uses the `strfind` function to search the retrieved content for a specific word:

- 1 Retrieve the Web page content with the `urlread` function:

```
recentFile =  
urlread('http://www.mathworks.com/matlabcentral/fileexchange/loadFileList.do?objectType=fileexchange&orderBy=date&srt3=0');
```

- 2 After retrieving the content, run the `strfind` function on the `recentFile` variable:

```
hits = findstr(recentFile, 'Simulink');
```

If the file contains the word `Simulink`, MATLAB will store the matches in the `hits` variable.

While you can manually pass arguments using the URL, the `urlread` function also lets you pass parameters to a Web page using standard HTTP methods, including `post` and `form`. Using the HTTP `get` method, which passes parameters in the URL, the following code queries Google for the word `Simulink`:

```
s =  
urlread('http://www.google.com/search','get',{'q','Simulink'})
```

For more information, see the `urlread` reference page.

### Example — Using the `urlwrite` Function

The following example builds on the procedure in the previous section. This example still uses `urlread` and checks for a specific word, but it also uses `urlwrite` to save the file if it contains any matches:

```
%The urlread function loads the contents of the Web page into the  
%MATLAB workspace.  
recentFile =  
urlread('http://www.mathworks.com/matlabcentral/fileexchange/loadFileList.do?objectType=fileexchange&orderBy=date&srt3=0');
```

```
%The findstr function searches for the word "Simulink".  
hits = findstr(recentFile, 'Simulink');
```

```
%The if statement checks for any hits.
if isempty(hits) == 0
    %If there are hits, the Web page will be saved locally using the
    %urlwrite function.

    urlwrite('http://www.mathworks.com/matlabcentral/fileexchange/lo
adFileList.do?objectType=fileexchange&orderBy=date&srt3=0','cont
ains_simulink.html');
end;
```

Note that the Web page is saved as contains\_simulink.html.

## Creating and Uncompressing Zip Archives

Using the zip and unzip functions, you can compress and uncompress files and directories. The zip function compresses files or directories into a zip archive. The unzip function uncompresses zip archives.

### Example – Using the zip Function

Again building on the example from previous sections, the following code creates a zip archive of the retrieved Web page:

```
%The urlread function loads the contents of the Web page into the
%MATLAB workspace.
recentFile =
urlread('http://www.mathworks.com/matlabcentral/fileexchange/loa
dFileList.do?objectType=fileexchange&orderBy=date&srt3=0');

%The findstr function searches for the word Simulink .
hits = findstr(recentFile,'Simulink');

%The if statement checks for any hits.
if isempty(hits) == 0
    %If there are hits, the Web page will be saved locally using the
    %urlwrite function.
    urlwrite('http://www.mathworks.com/matlabcentral/fileexchange/lo
adFileList.do?objectType=fileexchange&orderBy=date&srt3=0','cont
ains_simulink.html');
```

```
%The zip function creates a zip archive of the retrieved Web page.  
    zip('simulink_matches.zip','contains_simulink.html');  
end;
```

## Sending E-Mail

To send an e-mail from within MATLAB, use the `sendmail` function. You can also attach files to an e-mail, which lets you mail files directly from MATLAB. To use `sendmail`, you must first set up your e-mail address and your SMTP server information with the `setpref` function.

The `setpref` function defines two mail-related preferences:

- **E-mail address** — This preference sets your e-mail address that will appear on the message. Here is an example of the syntax:

```
setpref('Internet','E_mail','youraddress@yourserver.com');
```

- **SMTP server** — This preference sets your outgoing SMTP server address, which can be almost any e-mail server that supports the Post Office Protocol (POP) or the Internet Message Access Protocol (IMAP). Here is an example of the syntax:

```
setpref('Internet','SMTP_Server','mail.server.network');
```

You should be able to find your outgoing SMTP server address in your e-mail account settings in your e-mail client application. You can also contact your system administrator for the information.

---

**Note** The `sendmail` function does not support e-mail servers that require authentication.

---

Once you have properly configured MATLAB, you can use the `sendmail` function. The `sendmail` function requires at least two arguments: the recipient's e-mail address and the e-mail subject:

```
sendmail('recipient@someserver.com','Hello From MATLAB!');
```

You can supply multiple e-mail addresses using a cell array of strings, such as:

```
sendmail({'recipient@someserver.com',  
         'recipient2@someserver.com'}, 'Hello From MATLAB!');
```

You can also specify a message body with the `sendmail` function, such as:

```
sendmail('recepient@someserver.com', 'Hello From MATLAB!',  
        'Thanks for using sendmail.');
```

In addition, you can also attach files to an e-mail using the `sendmail` function, such as:

```
sendmail('recepient@somesever.com', 'Hello from MATLAB!', 'Thanks  
for using sendmail.', 'C:\yourFileSystem\message.txt');
```

You cannot attach a file without including a message. However, the message can be empty. You can also attach multiple files to an e-mail with the `sendmail` function, such as:

```
sendmail('recepient@somesever.com', 'Hello from MATLAB!', 'Thanks  
for using sendmail.', 'C:\yourFileSystem\message.txt',  
        'C:\yourFileSystem\message2.txt');
```

### Example – Using the `sendmail` Function

The following example sends e-mail with the retrieved Web page archive attached if it contains any matches for the specified word:

```
%The urlread function loads the contents of the Web page into the  
%MATLAB workspace.  
recentFile =  
urlread('http://www.mathworks.com/matlabcentral/fileexchange/load  
dFileList.do?objectType=fileexchange&orderBy=date&srt3=0');  
  
%The findstr function searches for the word "Simulink".  
hits = findstr(recentFile, 'Simulink');  
  
%The if statement checks for any hits.  
if isempty(hits) == 0  
  
%If there are hits, the Web page will be saved locally using the  
%urlwrite function.  
urlwrite('http://www.mathworks.com/matlabcentral/fileexchange/lo  
adFileList.do?objectType=fileexchange&orderBy=date&srt3=0', 'cont  
ains_simulink.html');
```

```
%The zip function creates a zip archive of the retrieved web page.
zip('simulink_matches.zip', 'contains_simulink.html');

%The setpref function supplies your e-mail address and SMTP server
%address to MATLAB.
setpref('Internet', 'SMTP_Server', 'mail.server.network');
setpref('Internet', 'E_mail', 'youraddress@yourserver.com');

%The sendmail function sends an e-mail with the zip archive of the
%retrieved Web page attached.
sendmail('youraddress@yourserver.com', 'New Simulink Files
Founds', 'New Simulink files have been uploaded to MATLAB Central.
See the attached zip archive.', 'simulink_matches.zip');
end;
```

## Performing FTP File Operations

From within MATLAB, you can connect to an FTP server to perform remote file operations. The following procedure uses a public MathWorks FTP server ([ftp.mathworks.com](http://ftp.mathworks.com)). To perform any file operation on an FTP server, follow these steps:

- 1 Connect to the server using the `ftp` function.

For example, you can create an FTP object for the public MathWorks FTP server with `tmw=ftp('ftp.mathworks.com')`.

- 2 Perform the file operations using appropriate MATLAB FTP functions as methods acting on the server object.

For example, you can display the file directories on the FTP server with `dir(tmw)`.

- 3 When you finish working on the server, close the connection object using the `close` function.

For example, you can disconnect from the FTP server with `close(tmw)`.

**Example — Retrieving a File from an FTP Server**

In this example, you retrieve the file `pub/pentium/Moler_1.txt`, which is on the MathWorks FTP server. You can run this example; the FTP server and content are valid.

- 1 Connect to the MathWorks FTP server using `ftp`. This creates the server object `tmw`:

```
tmw=ftp('ftp.mathworks.com');
```

- 2 List the contents of the server using the `dir` FTP function, which operates on the server object `tmw`:

```
dir(tmw)

.          incoming      pickup
README    matlab        pub
README.incoming  outgoing    pubs
```

- 3 Change to the `pub/pentium` directory by using the FTP `cd` function. As with all FTP functions, as part of the syntax you specify the server object you created using `ftp`, which in this case is `tmw`:

```
cd(tmw,'pub/pentium');
```

The server object `tmw` represents the current directory on the FTP server, which now is `pub/pentium`.

- 4 Now when you run

```
dir(tmw)
```

you see the contents of `pub/pentium`, rather than the top level contents as displayed previously when you ran `dir(tmw)`:

```
.          Intel_resp.txt      Intel_resp.txt
..         Intel_support.txt  Intel_support.txt
Andy_Grove.txt  Intel_white.ps    Intel_white.ps
Associated_Press.txt  MathWorks_press.txt  MathWorks_press.txt
CNN.html      Mathisen.txt      Mathisen.txt
Coe.txt       Moler_1.txt       Moler_1.txt
...
Intel_replace.txt  Myths.txt         Myths.txt
```

- 5** From the current directory on the FTP server, use `mget` to retrieve the file `Moler_1.txt` to the MATLAB current directory:

```
mget(tmw, 'Moler_1.txt');
```

- 6** Running a `dir` in MATLAB confirms that the file was downloaded.

```
dir
```

```
.          ..          Moler_1.txt
```

- 7** Close the FTP connection using `close`.

```
close(tmw);
```

### Summary of FTP Functions

The following table lists the available FTP functions. For more information, refer to the applicable reference pages.

Function	Description
<code>ascii</code>	Set FTP transfer type to ASCII (convert new lines).
<code>binary</code>	Set FTP transfer type to binary (transfer verbatim, default).
<code>cd</code>	Change current directory on FTP server.
<code>delete</code>	Delete file on FTP server.
<code>dir</code>	List contents of directory on FTP server.
<code>close</code>	Close connection with FTP server.
<code>ftp</code>	Connect to FTP server, creating an FTP object.
<code>mget</code>	Download file from FTP site.
<code>mkdir</code>	Create new directory on FTP server.
<code>mput</code>	Upload file or directory to FTP server.

<b>Function</b>	<b>Description</b>
rename	Rename file on FTP server.
rmdir	Remove directory on FTP server.

# Error Handling

---

In many cases, it is desirable to take specific actions when different kinds of errors occur. For example, you may want to prompt the user for more input, display extended error or warning information, or repeat a calculation using default values. The error handling capabilities in MATLAB let your application check for particular error conditions and execute appropriate code depending on the situation.

Checking for Errors with try-catch  
(p. 7-2)

Checking for errors that occur in your programs using the try and catch functions

Handling and Recovering from an Error (p. 7-4)

Reporting errors and identifying what caused them; regenerating errors

Message Identifiers (p. 7-8)

Attaching an identifier to an error or warning to enable you to better identify what caused it, and also for selective warning control

Warnings (p. 7-12)

Identifying warnings and identifying what caused them

Warning Control (p. 7-14)

Controlling the action taken when a warning is encountered

Debugging Errors and Warnings  
(p. 7-23)

Stopping code execution in the debugger on the occurrence of an error or warning

## Checking for Errors with try-catch

No matter how carefully you plan and test the programs you write, they may not always run as smoothly as expected when run under different conditions. It is always a good idea to include error checking in programs to ensure reliable operation under all conditions.

When you have statements in your code that could possibly generate unwanted results, put those statements into a try-catch block that will catch any errors and handle them appropriately. The example below shows a try-catch block within a sample function that multiplies two matrices:

```
function matrixMultiply(A, B)
try
    X = A * B
catch
    disp '** Error multiplying A * B'
end
```

A try-catch block is divided into two sections. The first begins with try and the second with catch. Terminate the block with end:

- All statements in the try segment are executed normally, just as if they were in the regular code flow. But if any of these operations result in an error, MATLAB skips the remaining statements in the try and jumps to the catch segment of the block.
- The catch segment handles the error. In this example, it displays a general error message. If there are different types of errors that can occur, you will want to identify which error has been caught and respond to that specific error. You can also try to recover from an error in the catch section.

When you execute the above example with inputs that are incompatible for matrix multiplication (e.g., the column dimension of A is not equal to the row dimension of B), MATLAB catches the error and displays the message generated in the catch section of the try-catch block.

```
A = [1 2 3; 6 7 2; 0 1 5];
B = [9 5 6; 0 4 9];

matrixMultiply(A, B)
** Error multiplying A * B
```

## Nested try-catch Blocks

You can also nest try-catch blocks, as shown here. You can use this to attempt to recover from an error caught in the first try section.

```
try
    statement1                                % Try to execute statement1
catch
    try
        statement2                            % Attempt to recover from error
    catch
        disp 'Operation failed'                % Handle the error
    end
end
end
```

## Handling and Recovering from an Error

The catch segment of a try-catch block needs to effectively handle any errors that may be caught by the preceding try. Frequently, you will want to simply report the error and stop execution. This prevents erroneous data from being propagated into the remainder of the program.

### Reporting an Error

To report an error and halt program execution, use the MATLAB `error` function. You determine what the error message will be by specifying it as an input to the error function in your code. For example,

```
if n < 1
    error('n must be 1 or greater.')
end
```

displays the message shown below when `n` is equal to zero.

```
??? n must be 1 or greater.
```

### Formatted Message Strings

The error message string that you specify can also contain formatting conversion characters, such as those used with the MATLAB `sprintf` function. Make the error string the first argument, and then add any variables used by the conversion as subsequent arguments.

```
error('formatted_errormsg', arg1, arg2, ...)
```

For example, if your program cannot find a specific file, you might report the error with

```
error('File %s not found', filename);
```

### Message Identifiers

Use a message identifier argument with `error` to attach a unique tag to that error message. MATLAB uses this tag to better identify the source of an error. The first argument in this example is the message identifier.

```
error('MATLAB:noSuchFile', 'File "%s" not found', ...
      filename);
```

See “Using Message Identifiers with `lasterr`” on page 7-9 for more information on how to use identifiers with errors.

### Formatted String Conversion

MATLAB converts special characters (like `\n` and `%d`) in the error message string only when you specify more than one input argument with error. In the single argument case shown below, `\n` is taken to mean backslash-n. It is not converted to a newline character.

```
error('In this case, the newline \n is not converted.')
??? In this case, the newline \n is not converted.
```

But, when more than one argument is specified, MATLAB does convert special characters. This is true regardless of whether the additional argument supplies conversion values or is a message identifier.

```
error('ErrorTests:convertTest', ...
      'In this case, the newline \n is converted.')
??? In this case, the newline
is converted.
```

### Identifying the Cause

Once an error has been caught, you will need to know the source of the error in order to handle it appropriately. The `lasterr` function returns information that enables you to identify the error that was most recently generated by MATLAB.

To return the most recent error message to the variable `errmsg`, type

```
errmsg = lasterr;
```

You can also change the text of the last error message with a new message or with an empty string as shown below. You might want to do this if a lower level routine detects an error that you don't want visible to the upper levels.

```
lasterr('newerrmsg');    % Replace last error with new string
lasterr('');            % Replace last error with empty string
```

### Example Using lasterr

The `matrixMultiply` function shown earlier in this section could fail for various reasons. If it is called with incompatible matrices, for example, `lasterr` returns the following string.

```
lasterr
ans =
    Error using ==> *
    Inner matrix dimensions must agree.
```

This example uses `lasterr` to determine the cause of an error in `matrixMultiply`.

```
function matrixMultiply(A, B)
try
    A * B
catch
    errmsg = lasterr;
    if(strfind(errmsg, 'Inner matrix dimensions'))
        disp('** Wrong dimensions for matrix multiply')
    else
        if(strfind(errmsg, 'not defined for variables of class'))
            disp('** Both arguments must be double matrices')
        end
    end
end
end
```

When calling the function with two matrices not compatible for matrix multiplication, you get the following error message.

```
A = [1 2 3; 6 7 2; 0 1 5];
B = [9 5 6; 0 4 9];
matrixMultiply(A, B)
** Wrong dimensions for matrix multiply
```

When calling the function with a cell array argument, you get a message that addresses that error.

```
C = {9 5 6; 0 4 9};
matrixMultiply(A, C)
** Both arguments must be double matrices
```

## Regenerating an Error

Use the `rethrow` function to regenerate an error that has previously been thrown. You may want to do this from the `catch` part of a `try-catch` block, for example, after performing some required cleanup tasks following an error.

This is how you would rethrow an error in a `try-catch` block:

```
try
    do_something
catch
    do_cleanup
    rethrow(lasterror)
end
```

`rethrow` generates an error based on the `err` input argument that you provide. This argument must be a MATLAB structure with at least the following two fields.

Field Name	Description
<code>message</code>	Text of the error message
<code>identifier</code>	Message identifier of the error message

If you simply want to regenerate the last error that occurred, the `lasterror` function returns a structure that can then be passed directly to `rethrow`.

**Note** `lasterror` is not the same as `lasterr`. The `lasterror` function returns a structure array that contains the message string and message identifier for the last error, and may contain more information in future versions of MATLAB. The `lasterr` function returns one or two character arrays that contain only the message string and identifier.

## Message Identifiers

A message identifier is a tag that you attach to an error or warning statement that makes that error or warning uniquely recognizable by MATLAB. You can use message identifiers with error reporting to better identify the source of an error, or with warnings to control any selected subset of the warnings in your programs. See the following topics for more information on how message identifiers are used:

- “Identifier Format” on page 7-8
- “Using Message Identifiers with lasterr” on page 7-9

### Identifier Format

The message identifier is a string that specifies a *component* and a *mnemonic* label for an error or warning. A simple identifier looks like this:

```
component:mnemonic
```

Some examples of message identifiers are

```
MATLAB:divideByZero  
Simulink:actionNotTaken  
TechCorp:notFoundInPath
```

Both the component and mnemonic fields must adhere to the following syntax rules:

- No white space (space or tab characters) is allowed in the entire identifier.
- The first character must be alphabetic, either uppercase or lowercase.
- The remaining characters can be alphanumeric or underscore.

There is no length limitation to either field.

### Component Field

The component field of a message identifier specifies a broad category under which various errors and warnings may be generated. Common components are a particular product or toolbox name, such as MATLAB or Control, or perhaps the name of your company, such as TechCorp in the example above.

You can also use this field to specify a multilevel component. The statement below has a three-level component followed by a mnemonic label:

```
TechCorp:TestEquipDiv:Waveform:obsoleteSyntax
```

One purpose of the component field is to enable you to guarantee the uniqueness of each identifier. Thus, while MATLAB uses the identifier `MATLAB:divideByZero` for its 'Divide by zero' warning, you could reuse the `divideByZero` mnemonic by using your own unique component. For example,

```
warning('TechCorp:divideByZero', 'A sprocket value was divided by
zero.')
```

### Mnemonic Field

The mnemonic field is a string normally used as a tag relating to the particular message. For example, when reporting an error resulting from the use of ambiguous syntax, a mnemonic such as `ambiguousSyntax` might be appropriate.

```
error('MATLAB:ambiguousSyntax', ...
      'Syntax %s could be ambiguous.\n', inputstr)
```

## Using Message Identifiers with `lasterr`

One use of message identifiers is to enable the `lasterr` and `lasterror` functions to better identify the source of an error. These functions return a message identifier, and you can use the combination of component and mnemonic parts of that identifier to identify both a broad category of errors and a specific instance within that category, respectively.

The first step in using this feature is to determine which of your error messages need this type of identification and then tag each with an identifier. You do this by specifying a message identifier argument (`msg_id`, below) along with the error message in your calls to `error`. Either form shown below is acceptable. The latter form uses formatting conversion characters as described in “Formatted Message Strings” on page 7-12.

```
error('msg_id', 'errmsg')
error('msg_id', 'formatted_errormsg', arg1, arg2, ...)
```

---

**Note** When you specify more than one input argument with `error`, MATLAB treats the `errmsg` string as if it were a `formatted_errormsg`. This is explained in “Formatted String Conversion” on page 7-5.

---

The message identifier must be the first argument and must be formatted according to the rules covered in “Message Identifiers” on page 7-8.

The message identifier is not a required argument for `error`. If you don’t need to return this type of information with `lasterr`, then you can omit the `msg_id` argument from the error function syntax shown previously:

```
error('errmsg')
```

### Returning a Message Identifier from `lasterr`

Use `lasterr` with one output to return just the string holding the error message from the most recently generated error.

```
errmsg = lasterr;
```

Use `lasterr` with two outputs to return both error message string and the message identifier for that error.

```
[errmsg, msg_id] = lasterr;
```

The following example performs an operation in the try segment of the try-catch block that results in an error. The first line of the catch segment retrieves both the error message string and message identifier for the error. The example then responds to the error in a manner that depends on the identifier returned.

```
try
    [d, x] = readimage(imageFile);
catch
    [errmsg, msg_id] = lasterr;
    switch (lower(msg_id))
    case 'matlab:nosuchfile'
        error('File "%s" does not exist.', filename);
    case 'myfileio:noaccess'
        error(['Can''t open file "%s" for reading\n', ...
            'You may not have read permission.'], filename);
```

```

        case 'myfileio:invformat'
            error('Unable to determine the file format.');
```

end

```

end
```

If the last error has no message identifier tag associated with it, then MATLAB returns an empty string in the second output argument.

```

error('This error has no message identifier.');
```

??? This error has no message identifier.

```

[errstr, msgid] = lasterr
errstr =
    This error has no message identifier.
msgid =
    ''
```

---

**Note** Both `lasterr` and `lasterror` return a message identifier. Although this section discusses only `lasterr`, you can use `lasterror` in the same way.

---

### Inputs to `lasterr`

In addition to returning information about the last error, `lasterr` also accepts inputs that modify the MATLAB copy of the last error. Use the command format shown below to change the error message string and message identifier returned by subsequent invocations of `lasterr`.

```

[lastErrMsg, lastMsgid] = lasterr('new_errmsg', 'new_msgid');
```

All `lasterr` input arguments are optional, but if you specify both an error message and message identifier input, they must appear in the order shown above.

## Warnings

Like error, the warning function alerts the user of unexpected conditions detected when running a program. However, warning does not halt the execution of the program. It displays the specified warning message and then continues.

### Reporting a Warning

Use warning in your code to generate a warning message during execution. Specify the message string as the input argument to warning. For example,

```
warning('Input must be a string')
```

Warnings also differ from errors in that you can disable any warnings that you don't want to see. You do this by invoking warning with certain control parameters. See “Warning Control” on page 7-14 for more information.

### Formatted Message Strings

The warning message string that you specify can also contain formatting conversion characters, such as those used with the MATLAB `sprintf` function. Make the warning string the first argument, and then add any variables used by the conversion as subsequent arguments.

```
warning('formatted_warningmsg', arg1, arg2, ...)
```

For example, if your program cannot process a given parameter, you might report a warning with

```
warning('Ambiguous parameter name, "%s".', param)
```

MATLAB converts special characters (like `\n` and `%d`) in the warning message string only when you specify more than one input argument with warning. See “Formatted String Conversion” on page 7-5 for information.

### Message Identifiers

Use a message identifier argument with warning to attach a unique tag to that warning message. MATLAB uses this tag to better identify the source of a warning. The first argument in this example is the message identifier.

```
warning('MATLAB:paramAmbiguous', ...  
       'Ambiguous parameter name, "%s".', param)
```

---

See “Warning Control Statements” on page 7-15 for more information on how to use identifiers with warnings.

## Identifying the Cause

The `lastwarn` function returns a string containing the last warning message issued by MATLAB. You can use this to enable your program to identify the cause of a warning that has just been issued. To return the most recent warning message to the variable `warnmsg`, type

```
warnmsg = lastwarn;
```

You can also change the text of the last warning message with a new message or with an empty string as shown here.

```
lastwarn('newwarnmsg'); % Replace last warning with new string  
lastwarn(' ');          % Replace last warning with empty string
```

## Warning Control

MATLAB gives you the ability to control what happens when a warning is encountered during M-file program execution. Options that are available include

- Display selected warnings
- Ignore selected warnings
- Stop in the debugger when a warning is invoked
- Display an M-stack trace after a warning is invoked

Depending on how you set up your warning controls, you can have these actions affect all warnings in your code, specific warnings that you select, or just the most recently invoked warning.

Setting up this system of warning control involves several steps.

- 1** Start by determining the scope of the control you will need for the warnings generated by your code. Do you want the above control operations to affect all the warnings in your code at once, or do you want to be able to control certain warnings separately?
- 2** If the latter is true, you will need to identify those warnings you want to selectively control. This requires going through your code and attaching unique *message identifiers* to those warnings. If, on the other hand, you don't require that fine a granularity of control, then the warning statements in your code need no message identifiers.
- 3** When you are ready to run your programs, use the MATLAB warning control statements to exercise the desired controls on all or selected warnings. Include message identifiers in these control statements when selecting specific warnings to act upon.

## Warning Statements

The warning statements that you put into your M-file code must contain the string that is to be displayed when the warning is incurred, and may also contain a message identifier. If you are not planning to use warning control or if you have no need to single out certain warnings for control, then you need to

specify the message string only. Use the syntax shown in the section on “Warnings” on page 7-12. Valid formats are

```
warning('warnmsg')  
warning('formatted_warnmsg', arg1, arg2, ...)
```

### Attaching an Identifier to the Warning Statement

If there are specific warnings that you want MATLAB to be able to apply control statements to, then you need to include a message identifier in the warning statement. The message identifier must be the first argument in the statement. Valid formats are

```
warning('msg_id', 'warnmsg')  
warning('msg_id', 'formatted_warnmsg', arg1, arg2, ...)
```

See “Message Identifiers” on page 7-8 for information on how to specify the `msg_id` argument.

---

**Note** When you specify more than one input argument with `warning`, MATLAB treats the `warnmsg` string as if it were a `formatted_warnmsg`. This is explained in “Formatted String Conversion” on page 7-5.

---

## Warning Control Statements

Once you have the warning statements in your M-file code and are ready to execute the code, you can indicate how you want MATLAB to act on these warnings by issuing control statements. These statements place the specified warning(s) into a desired state and have the format

```
warning state msg_id
```

Control statements can also return information on the state of selected warnings. This only happens if you assign the output to a variable, as shown below. See “Output from Control Statements” on page 7-17.

```
s = warning('state', 'msg_id');
```

### Warning States

There are three possible values for the state argument of a warning control statement.

State	Description
on	Enable the display of selected warning message.
off	Disable the display of selected warning message.
query	Display the current state of selected warning.

### Message Identifiers

In addition to the message identifiers already discussed, there are two other identifiers that you can use in control statements only.

Identifier	Description
<i>msg_id string</i>	Set selected warning to the specified state.
all	Set all warnings to the specified state.
last	Set only the last displayed warning to the specified state.

---

**Note** MATLAB starts up with all warnings enabled, except for those that are displayed in response to the command, `warning('query', 'all')`.

---

### Example 1 — Enabling a Selected Warning

Enable just the `actionNotTaken` warning from Simulink by first turning off all warnings and then setting just that warning to on.

```
warning off all
warning on Simulink:actionNotTaken
```

Use `query` to determine the current state of all warnings. It reports that you have set all warnings to off with the exception of `Simulink:actionNotTaken`.

```
warning query all
The default warning state is 'off'. Warnings not set to the
default are
```

```
State Warning Identifier
      on Simulink:actionNotTaken
```

### Example 2 — Disabling the Most Recent Warning

Evaluating `inv` on zero displays a warning message. Turn off the most recently invoked warning with `warning off last`.

```
inv(0)
Warning: Matrix is singular to working precision.
ans =
     Inf
```

```
warning off last
```

```
inv(0)                                % No warning is displayed this time
ans =
     Inf
```

## Output from Control Statements

The `warning` function, when used in a control statement, returns a MATLAB structure array containing the previous state of the selected warning(s). Use the following syntax to return this information in structure array, `s`:

```
s = warning('state', 'msg_id');
```

You must type the command using the MATLAB function format; that is, parentheses and quotation marks are required.

---

**Note** MATLAB does not display warning output if you do not assign the output to a variable.

---

This example turns off `divideByZero` warnings for the MATLAB component, and returns the identifier and previous state in a 1-by-1 structure array.

```
s = warning('off', 'MATLAB:divideByZero')
s =
    identifier: 'MATLAB:divideByZero'
    state: 'on'
```

You can use output variables with any type of warning control statement. If you just want to collect the information but don't want to change state, then simply perform a query on the warning(s). MATLAB returns the current state of those warnings selected by the message identifier.

```
s = warning('query', 'msg_id');
```

If you want to change state, but also save the former state so that you can restore it later, use the return structure array to save that state. The following example does an implicit query, returning state information in `s`, and then turns on all warnings.

```
s = warning('on', 'all');
```

See the section, “Saving and Restoring State” on page 7-20, for more information on restoring the former state of warnings.

### Output Structure Array

Each element of the structure array returned by `warning` contains two fields.

Field Name	Description
<code>identifier</code>	Message identifier string, 'all', or 'last'
<code>state</code>	State of warning(s) prior to invoking this control statement

If you query for the state of just one warning, using a message identifier or 'last' in the command, then MATLAB returns a one-element structure array. The `identifier` field contains the selected message identifier and the `state` field holds the current state of that warning:

```
s = warning('query','last')
s =
    identifier: 'MATLAB:divideByZero'
    state: 'on'
```

If you query for the state of all warnings, using 'all' in the command, MATLAB returns a structure array having one or more elements:

- The first element of the array always represents the default state. (This is the state set by the last warning on|off all command.)
- Each other element of the array represents a warning that is in a state that is different from the default.

```
warning off all
warning on MATLAB:divideByZero
warning on MATLAB:fileNotFound
```

```
s = warning('query', 'all')
s =
    3x1 struct array with fields:
    identifier
    state
```

```
s(1)
ans =
    identifier: 'all'
    state: 'off'
```

```
s(2)
ans =
    identifier: 'MATLAB:divideByZero'
    state: 'on'
```

```
s(3)
ans =
    identifier: 'MATLAB:fileNotFound'
    state: 'on'
```

## Saving and Restoring State

If you want to temporarily change the state of some warnings and then later return to your original settings, you can save the original state in a structure array and then restore it from that array. You can save and restore the state of all of your warnings or just one that you select with a message identifier.

To save the current warning state, assign the output of a warning control statement, as discussed in the last section, “Output from Control Statements” on page 7-17. The following statement saves the current state of all warnings in structure array `s`.

```
s = warning('query', 'all');
```

To restore state from `s`, use the syntax shown below. Note that the MATLAB function format (enclosing arguments in parentheses) is required.

```
warning(s)
```

### Example 1 — Performing an Explicit Query

Perform a query of all warnings to save the current state in structure array `s`.

```
s = warning('query', 'all');
```

Then, after doing some work that includes making changes to the state of some warnings, restore the original state of all warnings.

```
warning(s)
```

### Example 2 — Performing an Implicit Query

Turn on one particular warning, saving the previous state of this warning in `s`. Remember that this nonquery syntax (where state equals on or off) performs an implicit query prior to setting the new state.

```
s = warning('on', 'Control:parameterNotSymmetric');
```

Restore the state of that one warning when you are ready, with

```
warning(s)
```

## Debug, Backtrace, and Verbose Modes

In addition to warning messages, there are three *modes* that can be enabled or disabled with a warning control statement. These modes are shown here.

Mode	Description	Default
debug	Stop in the debugger when a warning is invoked.	off
backtrace	Display an M-stack trace after a warning is invoked.	off
verbose	Display a message on how to suppress the warning.	off

The syntax for using this type of control statement is as follows, where *state*, in this case, can be only on, off, or query.

```
warning state mode
```

Note that there is no need to include a message identifier with this type of control statement. All enabled warnings are affected by the this type of control statement.

### Example 1 — Entering Debug Mode on a Specific Warning

To enter debug mode whenever a Simulink `actionNotTaken` warning is invoked, first turn off all warnings and enable only this one type of warning using its message identifier. Then turn on debug mode for all enabled warnings. When you run your program, MATLAB will stop in debug mode just before this warning is executed. You will see the debug prompt (`K>>`) displayed.

```
warning off all
warning on Simulink:actionNotTaken
warning on debug
```

### Example 2 — Enabling Verbose Warnings

When you enable verbose warnings, MATLAB displays an extra line of information with each warning that tells you how to suppress it:

```
warning on verbose
```

```
A = 25/0;
```

```
Warning: Divide by zero.
```

```
(Type "warning off MATLAB:divideByZero" to suppress this  
warning.)
```

## Debugging Errors and Warnings

You can direct MATLAB to temporarily stop the execution of an M-file program in the event of a run-time error or warning, at the same time opening a debug window paused at the M-file line that generated the error or warning. This enables you to examine values internal to the program and determine the cause of the error.

Use the `dbstop` function to have MATLAB stop execution and enter debug mode when any M-file you subsequently run produces a run-time error or warning. There are three types of such breakpoints that you can set.

<b>Command</b>	<b>Description</b>
<code>dbstop if all error</code>	Stop on any error.
<code>dbstop if error</code>	Stop on any error not detected within a try-catch block.
<code>dbstop if warning</code>	Stop on any warning.

In all three cases, the M-file you are trying to debug must be in a directory that is on the search path or in the current directory.

You cannot resume execution after an error; use `dbquit` to exit from the Debugger. To resume execution after a warning, use `dbcont` or `dbstep`.



# Classes and Objects

---

Classes and Objects: An Overview (p. 8-2)	Using object-oriented programming in MATLAB
Designing User Classes in MATLAB (p. 8-9)	The basic set of methods that should be included in a class
Overloading Operators and Functions (p. 8-21)	Overloading the MATLAB operators and functions to change their behavior
Example — A Polynomial Class (p. 8-24)	Example that defines a new class to implement a MATLAB data type for polynomials
Building on Other Classes (p. 8-35)	Inheritance and aggregation
Example — Assets and Asset Subclasses (p. 8-38)	An example that uses simple inheritance
Example — The Portfolio Container (p. 8-54)	An example that uses aggregation
Saving and Loading Objects (p. 8-60)	Saving and retrieving user-defined objects to and from MAT-files
Example — Defining saveobj and loadobj for Portfolio (p. 8-61)	Defining methods that automatically execute on save and load
Object Precedence (p. 8-65)	Determining which operator or function to call in a given situation
How MATLAB Determines Which Method to Call (p. 8-67)	How function arguments and precedence determine which method to call

## Classes and Objects: An Overview

You can view classes as new data types having specific behaviors defined for the class. For example, a polynomial class might redefine the addition operator (+) so that it correctly performs the operation of addition on polynomials. Operations defined to work with objects of a particular class are known as *methods* of that class.

You can also view classes as new items that you can treat as single entities. An example is an arrow object that MATLAB can display on graphs (perhaps composed of MATLAB line and patch objects) and that has properties like a Handle Graphics object. You can create an arrow simply by instantiating the arrow class.

You can add classes to your MATLAB environment by specifying a MATLAB structure that provides data storage for the object and creating a class directory containing M-files that operate on the object. These M-files contain the methods for the class. The class directory can also include functions that define the way various MATLAB operators, including arithmetic operations, subscript referencing, and concatenation, apply to the objects. Redefining how a built-in operator works for your class is known as *overloading* the operator.

### Features of Object-Oriented Programming

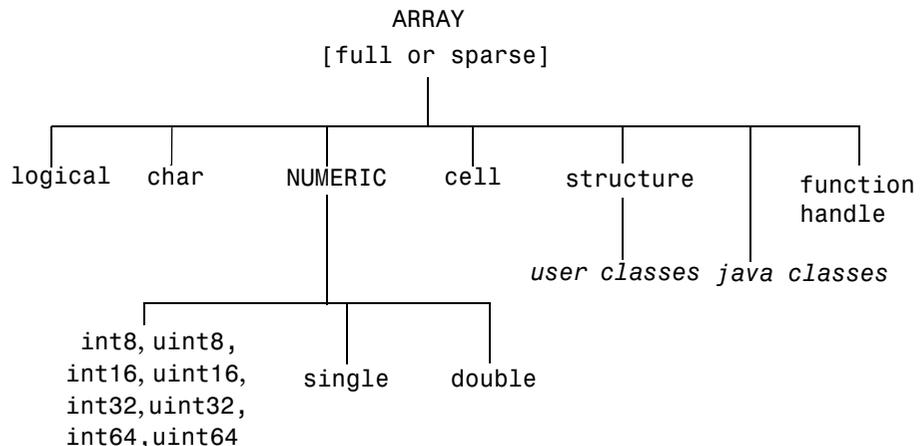
When using well-designed classes, object-oriented programming can significantly increase code reuse and make your programs easier to maintain and extend. Programming with classes and objects differs from ordinary structured programming in these important ways:

- **Function and operator overloading.** You can create methods that override existing MATLAB functions. When you call a function with a user-defined object as an argument, MATLAB first checks to see if there is a method defined for the object's class. If there is, MATLAB calls it, rather than the normal MATLAB function.
- **Encapsulation of data and methods.** Object properties are not visible from the command line; you can access them only with class methods. This protects the object properties from operations that are not intended for the object's class.

- **Inheritance.** You can create class hierarchies of parent and child classes in which the child class inherits data fields and methods from the parent. A child class can inherit from one parent (*single inheritance*) or many parents (*multiple inheritance*). Inheritance can span one or more generations. Inheritance enables sharing common parent functions and enforcing common behavior amongst all child classes.
- **Aggregation.** You can create classes using *aggregation*, in which an object contains other objects. This is appropriate when an object type is part of another object type. For example, a savings account object might be a part of a financial portfolio object.

## MATLAB Data Class Hierarchy

All MATLAB data types are designed to function as classes in object-oriented programming. The diagram below shows the fifteen fundamental data types (or classes) defined in MATLAB. You can add new data types to MATLAB by extending the class hierarchy.



The diagram shows a *user class* that inherits from the structure class. All classes that you create are structure based since this is the point in the class hierarchy where you can insert your own classes. (For more information about MATLAB data types, see Chapter 2, “Data Types.”)

## Creating Objects

You create an object by calling the class constructor and passing it the appropriate input arguments. In MATLAB, constructors have the same name as the class name. For example, the statement,

```
p = polynom([1 0 -2 -5]);
```

creates an object named `p` belonging to the class `polynom`. Once you have created a `polynom` object, you can operate on the object using methods that are defined for the `polynom` class. See “Example — A Polynomial Class” on page 8-24 for a description of the `polynom` class.

## Invoking Methods on Objects

Class methods are M-file functions that take an object as one of the input arguments. The methods for a specific class must be placed in the class directory for that class (the `@classname` directory). This is the first place that MATLAB looks to find a class method.

The syntax for invoking a method on an object is similar to a function call. Generally, it looks like

```
[out1,out2,...] = methodName(object,arg1,arg2,...);
```

For example, suppose a user-defined class called `polynom` has a `char` method defined for the class. This method converts a `polynom` object to a character string and returns the string. This statement calls the `char` method on the `polynom` object `p`.

```
s = char(p);
```

Using the `class` function, you can confirm that the returned value `s` is a character string.

```
class(s)
ans =
    char

s
s =
    x^3-2*x-5
```

You can use the `methods` command to produce a list of all of the methods that are defined for a class.

## Private Methods

Private methods can be called only by other methods of their class. You define private methods by placing the associated M-files in a private subdirectory of the `@classname` directory. In the example,

```
@classname/private/updateObj.m
```

the method `updateObj` has scope only within the `classname` class. This means that `updateObj` can be called by any method that is defined in the `@classname` directory, but it cannot be called from the MATLAB command line or by methods outside of the class directory, including parent methods.

Private methods and private functions differ in that private methods (in fact all methods) have an object as one of their input arguments and private functions do not. You can use private functions as helper functions, such as described in the next section.

## Helper Functions

In designing a class, you may discover the need for functions that perform support tasks for the class, but do not directly operate on an object. These functions are called *helper functions*. A helper function can be a subfunction in a class method file or a private function. When determining which version of a particular function to call, MATLAB looks for these functions in the order listed above. For more information about the order in which MATLAB calls functions and methods, see “How MATLAB Determines Which Method to Call” on page 8-67.

## Debugging Class Methods

You can use the MATLAB debugging commands with object methods in the same way that you use them with other M-files. The only difference is that you need to include the class directory name before the method name in the command call, as shown in this example using `dbstop`.

```
dbstop @polynom/char
```

While debugging a class method, you have access to all methods defined for the class, including inherited methods, private methods, and private functions.

## Changing Class Definition

If you change the class definition, such as the number or names of fields in a class, you must issue a

```
clear classes
```

command to propagate the changes to your MATLAB session. This command also clears all objects from the workspace. See the `clear` command help entry for more information.

## Setting Up Class Directories

The M-files defining the methods for a class are collected together in a directory referred to as the class directory. The directory name is formed with the class name preceded by the character `@`. For example, one of the examples used in this chapter is a class involving polynomials in a single variable. The name of the class, and the name of the class constructor, is `polynom`. The M-files defining a polynomial class would be located in directory with the name `@polynom`.

The class directories are subdirectories of directories on the MATLAB search path, but are not themselves on the path. For instance, the new `@polynom` directory could be a subdirectory of the MATLAB working directory or your own personal directory that has been added to the search path.

## Adding the Class Directory to the MATLAB Path

After creating the class directory, you need to update the MATLAB path so that MATLAB can locate the class source files. The class directory should not be directly on the MATLAB path. Instead, you should add the parent directory to the MATLAB path. For example, if the `@polynom` class directory is located at

```
c:\myClasses\@polynom
```

you add the class directory to the MATLAB path with the `addpath` command

```
addpath c:\myClasses;
```

## Using Multiple Class Directories

A MATLAB class can access methods in multiple `@classname` directories if all such directories are visible to MATLAB (i.e., the parent directories are on the MATLAB path or in the current directory). When you attempt to use a method of the class, MATLAB searches all the visible directories named `@classname` for the appropriate method.

For more information, see “How MATLAB Determines Which Method to Call” on page 8-67.

## Data Structure

One of the first steps in the design of a new class is the choice of the data structure to be used by the class. Objects are stored in MATLAB structures. The fields of the structure, and the details of operations on the fields, are visible only within the methods for the class. The design of the appropriate data structure can affect the performance of the code.

## Tips for C++ and Java Programmers

If you are accustomed to programming in other object-oriented languages, such as C++ or Java, you will find that the MATLAB programming language differs from these languages in some important ways:

- In MATLAB, method dispatching is not syntax based, as it is in C++ and Java. When the argument list contains objects of equal precedence, MATLAB uses the left-most object to select the method to call.
- In MATLAB, there is no equivalent to a destructor method. To remove an object from the workspace, use the `clear` function.
- Construction of MATLAB data types occurs at runtime rather than compile time. You register an object as belonging to a class by calling the `class` function.
- When using inheritance in MATLAB, the inheritance relationship is established in the child class by creating the parent object, and then calling the `class` function. For more information on writing constructors for inheritance relationships, see “Building on Other Classes” on page 8-35.
- When using inheritance in MATLAB, the child object contains a parent object in a property with the name of the parent class.

- In MATLAB, there is no passing of variables by reference. When writing methods that update an object, you must pass back the updated object and use an assignment statement. For instance, this call to the `set` method updates the `name` field of the object `A` and returns the updated object.

```
A = set(A, 'name', 'John Smith');
```

- In MATLAB, there is no equivalent to an abstract class.
- In MATLAB, there is no equivalent to the C++ scoping operator.
- In MATLAB, there is no virtual inheritance or virtual base classes.
- In MATLAB, there is no equivalent to C++ templates.

## Designing User Classes in MATLAB

This section discusses how to approach the design of a class and describes the basic set of methods that should be included in a class.

### The MATLAB Canonical Class

When you design a MATLAB class, you should include a standard set of methods that enable the class to behave in a consistent and logical way within the MATLAB environment. Depending on the nature of the class you are defining, you may not need to include all of these methods and you may include a number of other methods to realize the class's design goals.

This table lists the basic methods included in MATLAB classes.

Class Method	Description
class constructor	Creates an object of the class.
display	Called whenever MATLAB displays the contents of an object (e.g., when an expression is entered without terminating with a semicolon).
set and get	Accesses class properties.
subsref and subsasgn	Enables indexed reference and assignment for user objects.
end	Supports end syntax in indexing expressions using an object; e.g., <code>A(1:end)</code> .
subsindex	Supports using an object in indexing expressions.
converters like double and char	Methods that convert an object to a MATLAB data type.

The following sections discuss the implementation of each type of method, as well as providing references to examples used in this chapter.

## The Class Constructor Method

The @ directory for a particular class must contain an M-file known as the *constructor* for that class. The name of the constructor is the same as the name of the directory (excluding the @ prefix and .m extension) that defines the name of the class. The constructor creates the object by initializing the data structure and instantiating an object of the class.

### Guidelines for Writing a Constructor

Class constructors must perform certain functions so that objects behave correctly in the MATLAB environment. In general, a class constructor must handle three possible combinations of input arguments:

- No input arguments
- An object of the same class as an input argument
- The input arguments used to create an object of the class (typically data of some kind)

**No Input Arguments.** If there are no input arguments, the constructor should create a default object. Since there are no inputs, you have no data from which to create the object, so you simply initialize the object's data structures with empty or default values, call the `class` function to instantiate the object, and return the object as the output argument. Support for this syntax is required for two reasons:

- When loading objects into the workspace, the `load` function calls the class constructor with no arguments.
- When creating arrays of objects, MATLAB calls the class constructor to add objects to the array.

**Object Input Argument.** If the first input argument in the argument list is an object of the same class, the constructor should simply return the object. Use the `isa` function to determine if an argument is a member of a class. See "Overloading the + Operator" on page 8-30 for an example of a method that uses this constructor syntax.

**Data Input Arguments.** If the input arguments exist and are not objects of the same class, then the constructor creates the object using the input data. Of course, as in any function, you should perform proper argument checking in your constructor function. A typical approach is to use a `varargin` input

argument and a switch statement to control program flow. This provides an easy way to accommodate the three cases: no inputs, object input, or the data inputs used to create an object.

It is in this part of the constructor that you assign values to the object's data structure, call the `class` function to instantiate the object, and return the object as the output argument. If necessary, place the object in an object hierarchy using the `superiorto` and `inferiorto` functions.

### Using the `class` Function in Constructors

Within a constructor method, you use the `class` function to associate an object structure with a particular class. This is done using an internal class tag that is only accessible using the `class` and `isa` functions. For example, this call to the `class` function identifies the object `p` to be of type `polynom`.

```
p = class(p, 'polynom');
```

## Examples of Constructor Methods

See the following sections for examples of constructor methods:

- “The `Polynom` Constructor Method” on page 8-24
- “The `Asset` Constructor Method” on page 8-39
- “The `Stock` Constructor Method” on page 8-46
- “The `Portfolio` Constructor Method” on page 8-55

## Identifying Objects Outside the Class Directory

The `class` and `isa` functions used in constructor methods can also be used outside of the class directory. The expression

```
isa(a, 'classname');
```

checks whether `a` is an object of the specified class. For example, if `p` is a `polynom` object, each of the following expressions is true.

```
isa(pi, 'double');  
isa('hello', 'char');  
isa(p, 'polynom');
```

Outside of the class directory, the `class` function takes only one argument (it is only within the constructor that `class` can have more than one argument).

The expression

```
class(a)
```

returns a string containing the class name of `a`. For example,

```
class(pi),  
class('hello'),  
class(p)
```

return

```
'double',  
'char',  
'polynom'
```

Use the `whos` function to see what objects are in the MATLAB workspace.

```
whos
```

Name	Size	Bytes	Class
p	1x1	156	polynom object

## The display Method

MATLAB calls a method named `display` whenever an object is the result of a statement that is not terminated by a semicolon. For example, creating the variable `a`, which is a double, calls the MATLAB `display` method for doubles.

```
a = 5  
a =  
5
```

You should define a `display` method so MATLAB can display values on the command line when referencing objects from your class. In many classes, `display` can simply print the variable name, and then use the `char` converter method to print the contents or value of the variable, since MATLAB displays output as strings. You must define the `char` method to convert the object's data to a character string.

## Examples of display Methods

See the following sections for examples of display methods:

- “The Polynom display Method” on page 8-28
- “The Asset display Method” on page 8-44
- “The Stock display Method” on page 8-53
- “The Portfolio display Method” on page 8-56

## Accessing Object Data

You need to write methods for your class that provide access to an object’s data. Accessor methods can use a variety of approaches, but all methods that change object data always accept an object as an input argument and return a new object with the data changed. This is necessary because MATLAB does not support passing arguments by reference (i.e., pointers). Functions can change only their private, temporary copy of an object. Therefore, to change an existing object, you must create a new one, and then replace the old one.

The following sections provide more detail about implementation techniques for the `set`, `get`, `subsasgn`, and `subsref` methods.

## The set and get Methods

The `set` and `get` methods provide a convenient way to access object data in certain cases. For example, suppose you have created a class that defines an arrow object that MATLAB can display on graphs (perhaps composed of existing MATLAB `line` and `patch` objects).

To produce a consistent interface, you could define `set` and `get` methods that operate on arrow objects the way the MATLAB `set` and `get` functions operate on built-in graphics objects. The `set` and `get` verbs convey what operations they perform, but insulate the user from the internals of the object.

## Examples of set and get Methods

See the following sections for examples of set and get methods:

- “The Asset get Method” on page 8-41 and “The Asset set Method” on page 8-41
- “The Stock get Method” on page 8-48 and “The Stock set Method” on page 8-49

### Property Name Methods

As an alternative to a general set method, you can write a method to handle the assignment of an individual property. The method should have the same name as the property name.

For example, if you defined a class that creates objects representing employee data, you might have a field in an employee object called `salary`. You could then define a method called `salary.m` that takes an employee object and a value as input arguments and returns the object with the specified value set.

### Indexed Reference Using `subsref` and `subsasgn`

User classes implement new data types in MATLAB. It is useful to be able to access object data via an indexed reference, as is possible with the MATLAB built-in data types. For example, if `A` is an array of class `double`, `A(i)` returns the  $i^{\text{th}}$  element of `A`.

As the class designer, you can decide what an index reference to an object means. For example, suppose you define a class that creates polynomial objects and these objects contain the coefficients of the polynomial.

An indexed reference to a polynomial object,

`p(3)`

could return the value of the coefficient of  $x^3$ , the value of the polynomial at  $x = 3$ , or something different depending on the intended design.

You define the behavior of indexing for a particular class by creating two class methods – `subsref` and `subsasgn`. MATLAB calls these methods whenever a subscripted reference or assignment is made on an object from the class. If you do not define these methods for a class, indexing is undefined for objects of this class.

In general, the rules for indexing objects are the same as the rules for indexing structure arrays. For details, see “Structures” on page 2-49.

## Handling Subscripted Reference

The use of a subscript or field designator with an object on the right-hand side of an assignment statement is known as a *subscripted reference*. MATLAB calls a method named `subref` in these situations.

Object subscripted references can be of three forms—an array index, a cell array index, and a structure field name:

```
A(I)
A{I}
A.field
```

Each of these results in a call by MATLAB to the `subref` method in the class directory. MATLAB passes two arguments to `subref`.

```
B = subref(A,S)
```

The first argument is the object being referenced. The second argument, `S`, is a structure array with two fields:

- `S.type` is a string containing '()', '{}', or '.' specifying the subscript type. The parentheses represent a numeric array; the curly braces, a cell array; and the dot, a structure array.
- `S.subs` is a cell array or string containing the actual subscripts. A colon used as a subscript is passed as the string ':'.

For instance, the expression

```
A(1:2,:)
```

causes MATLAB to call `subref(A,S)`, where `S` is a 1-by-1 structure with

```
S.type = '()'
S.subs = {1:2, ':'}
```

Similarly, the expression

```
A{1:2}
```

uses

```
S.type = '{}'
S.subs = {1:2}
```

The expression

```
A.field
```

calls `subsref(A,S)` where

```
S.type = '.'  
S.subs = 'field'
```

These simple calls are combined for more complicated subscripting expressions. In such cases, `length(S)` is the number of subscripting levels. For example,

```
A(1,2).name(3:4)
```

calls `subsref(A,S)`, where `S` is a 3-by-1 structure array with the values:

```
S(1).type = '()'      S(2).type = '.'      S(3).type = '()  
S(1).subs = '{1,2}' S(2).subs = 'name'  S(3).subs = '{3:4}'
```

### How to Write `subsref`

The `subsref` method must interpret the subscripting expressions passed in by MATLAB. A typical approach is to use the `switch` statement to determine the type of indexing used and to obtain the actual indices. The following three code fragments illustrate how to interpret the input arguments. In each case, the function must return the value `B`.

For an array index:

```
switch S.type  
case '()'   
    B = A(S.subs{:});  
end
```

For a cell array:

```
switch S.type  
case '{}'   
    B = A(S.subs{:});           % A is a cell array  
end
```

For a structure array:

```
switch S.type
case '.'
    switch S.subs
    case 'field1'
        B = A.field1;
    case 'field2'
        B = A.field2;
    end
end
```

### Examples of the subsref Method

See the following sections for examples of the subsref method:

- “The Polynom subsref Method” on page 8-29
- “The Asset subsref Method” on page 8-42
- “The Stock subsref Method” on page 8-50
- “The Portfolio subsref Method” on page 8-64

## Handling Subscripted Assignment

The use of a subscript or field designator with an object on the left-hand side of an assignment statement is known as a *subscripted assignment*. MATLAB calls a method named `subsasgn` in these situations. Object subscripted assignment can be of three forms – an array index, a cell array index, and a structure field name.

```
A(I) = B
A{I} = B
A.field = B
```

Each of these results in a call to `subsasgn` of the form

```
A = subsasgn(A,S,B)
```

The first argument, `A`, is the object being referenced. The second argument, `S`, has the same fields as those used with `subsref`. The third argument, `B`, is the new value.

## Examples of the subsasgn Method

See the following sections for examples of the subsasgn method:

- “The Asset subsasgn Method” on page 8-43
- “The Stock subsasgn Method” on page 8-51

## Object Indexing Within Methods

If a subscripted reference is made within a class method, MATLAB uses its built-in subsref function to access data within the method’s own class. If the method accesses data from another class, MATLAB calls the overloaded subsref function in that class. The same holds true for subscripted assignment and subsasgn.

The following example shows a method, testref, that is defined in the class, employee. This method makes a reference to a field, address, in an object of its own class. For this, MATLAB uses the built-in subsref function. It also references the same field in another class, this time using the overloaded subsref of that class.

```
% ---- EMPLOYEE class method: testref.m ----  
function testref(myclass,otherclass)  
  
myclass.address           % use built-in subsref  
otherclass.address       % use overloaded subsref
```

The example creates an employee object and a company object.

```
empl = employee('Johnson','Chicago');  
comp = company('The MathWorks','Natick');
```

The employee class method, testref, is called. MATLAB uses an overloaded subsref only to access data outside of the method’s own class.

```
testref(empl,comp)  
ans =                               % built-in subsref was called  
    Chicago  
  
ans =                               % @company\subsref was called  
Executing @company\subsref ...  
    Natick
```

## Defining end Indexing for an Object

When you use `end` in an object indexing expression, MATLAB calls the object's `end` class method. If you want to be able to use `end` in indexing expressions involving objects of your class, you must define an `end` method for your class.

The `end` method has the calling sequence

```
end(a, k, n)
```

where `a` is the user object, `k` is the index in the expression where the `end` syntax is used, and `n` is the total number of indices in the expression.

For example, consider the expression

```
A(end-1, :)
```

MATLAB calls the `end` method defined for the object `A` using the arguments

```
end(A, 1, 2)
```

That is, the `end` statement occurs in the first index element and there are two index elements. The class method for `end` must then return the index value for the last element of the first dimension. When you implement the `end` method for your class, you must ensure it returns a value appropriate for the object.

## Indexing an Object with Another Object

When MATLAB encounters an object as an index, it calls the `subsindex` method defined for the object. For example, suppose you have an object `a` and you want to use this object to index into another object `b`.

```
c = b(a);
```

A `subsindex` method might do something as simple as convert the object to double format to be used as an index, as shown in this sample code.

```
function d = subsindex(a)
%SUBSINDEX
% convert the object a to double format to be used
% as an index in an indexing expression
d = double(a);
```

`subsindex` values are 0-based, not 1-based.

## Converter Methods

A converter method is a class method that has the same name as another class, such as `char` or `double`. Converter methods accept an object of one class as input and return an object of another class. Converters enable you to

- Use methods defined for another class
- Ensure that expressions involving objects of mixed class types execute properly

A converter function call is of the form

```
b = classname(a)
```

where `a` is an object of a class other than `classname`. In this case, MATLAB looks for a method called `classname` in the class directory for object `a`. If the input object is already of type `classname`, then MATLAB calls the constructor, which just returns the input argument.

## Examples of Converter Methods

See the following sections for examples of converter methods:

- “The Polynom to Double Converter” on page 8-26
- “The Polynom to Char Converter” on page 8-26

## Overloading Operators and Functions

In many cases, you may want to change the behavior of the MATLAB operators and functions for cases when the arguments are objects. You can accomplish this by *overloading* the relevant functions. Overloading enables a function to handle different types and numbers of input arguments and perform whatever operation is appropriate for the highest-precedence object. See “Object Precedence” on page 8-65 for more information on object precedence.

### Overloading Operators

Each built-in MATLAB operator has an associated function name (e.g., the + operator has an associated `plus.m` function). You can overload any operator by creating an M-file with the appropriate name in the class directory. For example, if either `p` or `q` is an object of type `classname`, the expression

```
p + q
```

generates a call to a function `@classname/plus.m`, if it exists. If `p` and `q` are both objects of different classes, then MATLAB applies the rules of precedence to determine which method to use.

### Examples of Overloaded Operators

See the following sections for examples of overloaded operators:

- “Overloading the + Operator” on page 8-30
- “Overloading the - Operator” on page 8-30
- “Overloading the \* Operator” on page 8-31

The following table lists the function names for most of the MATLAB operators.

<b>Operation</b>	<b>M-File</b>	<b>Description</b>
$a + b$	<code>plus(a,b)</code>	Binary addition
$a - b$	<code>minus(a,b)</code>	Binary subtraction
$-a$	<code>uminus(a)</code>	Unary minus
$+a$	<code>uplus(a)</code>	Unary plus
$a.*b$	<code>times(a,b)</code>	Element-wise multiplication
$a*b$	<code>mtimes(a,b)</code>	Matrix multiplication
$a./b$	<code>rdivide(a,b)</code>	Right element-wise division
$a.\backslash b$	<code>ldivide(a,b)</code>	Left element-wise division
$a/b$	<code>mrdivide(a,b)</code>	Matrix right division
$a\backslash b$	<code>mldivide(a,b)</code>	Matrix left division
$a.^b$	<code>power(a,b)</code>	Element-wise power
$a^b$	<code>mpower(a,b)</code>	Matrix power
$a < b$	<code>lt(a,b)</code>	Less than
$a > b$	<code>gt(a,b)</code>	Greater than
$a \leq b$	<code>le(a,b)</code>	Less than or equal to
$a \geq b$	<code>ge(a,b)</code>	Greater than or equal to
$a \neq b$	<code>ne(a,b)</code>	Not equal to
$a == b$	<code>eq(a,b)</code>	Equality
$a \& b$	<code>and(a,b)</code>	Logical AND
$a   b$	<code>or(a,b)</code>	Logical OR
$\sim a$	<code>not(a)</code>	Logical NOT

<b>Operation</b>	<b>M-File</b>	<b>Description</b>
a:d:b a:b	colon(a,d,b) colon(a,b)	Colon operator
a'	ctranspose(a)	Complex conjugate transpose
a.'	transpose(a)	Matrix transpose
command window output	display(a)	Display method
[a b]	horz- cat(a,b,...)	Horizontal concatenation
[a; b]	vert- cat(a,b,...)	Vertical concatenation
a(s1,s2,...sn)	subsref(a,s)	Subscripted reference
a(s1,...,sn) = b	subsasgn(a,s,b)	Subscripted assignment
b(a)	subsindex(a)	Subscript index

## Overloading Functions

You can overload any function by creating a function of the same name in the class directory. When a function is invoked on an object, MATLAB always looks in the class directory before any other location on the search path. To overload the `plot` function for a class of objects, for example, simply place your version of `plot.m` in the appropriate class directory.

### Examples of Overloaded Functions

See the following sections for examples of overloaded functions:

- “Overloading Functions for the Polynom Class” on page 8-31
- “The Portfolio `pie3` Method” on page 8-57

## Example – A Polynomial Class

This example implements a MATLAB data type for polynomials by defining a new class called `polynom`. The class definition specifies a structure for data storage and defines a directory (`@polynom`) of methods that operate on `polynom` objects.

### Polynom Data Structure

The `polynom` class represents a polynomial with a row vector containing the coefficients of powers of the variable, in decreasing order. Therefore, a `polynom` object `p` is a structure with a single field, `p.c`, containing the coefficients. This field is accessible only within the methods in the `@polynom` directory.

### Polynom Methods

To create a class that is well behaved within the MATLAB environment and provides useful functionality for a polynomial data type, the `polynom` class implements the following methods:

- A constructor method `polynom.m`
- A `polynom` to double converter
- A `polynom` to char converter
- A display method
- A `subsref` method
- Overloaded `+`, `-`, and `*` operators
- Overloaded `roots`, `polyval`, `plot`, and `diff` functions

### The Polynom Constructor Method

Here is the `polynom` class constructor, `@polynom/polynom.m`.

```
function p = polynom(a)
%POLYNOM Polynomial class constructor.
%   p = POLYNOM(v) creates a polynomial object from the vector v,
%   containing the coefficients of descending powers of x.
```

```
if nargin == 0
    p.c = [];
    p = class(p, 'polynom');
elseif isa(a, 'polynom')
    p = a;
else
    p.c = a(:).';
    p = class(p, 'polynom');
end
```

### Constructor Calling Syntax

You can call the `polynom` constructor method with one of three different arguments:

- No input argument — If you call the constructor function with no arguments, it returns a polynomial object with empty fields.
- Input argument is an object — If you call the constructor function with an input argument that is already a polynomial object, MATLAB returns the input argument. The `isa` function (pronounced “is a”) checks for this situation.
- Input argument is a coefficient vector — If the input argument is a variable that is not a polynomial object, reshape it to be a row vector and assign it to the `.c` field of the object’s structure. The `class` function creates the polynomial object, which is then returned by the constructor.

An example use of the `polynom` constructor is the statement

```
p = polynom([1 0 -2 -5])
```

This creates a polynomial with the specified coefficients.

### Converter Methods for the Polynomial Class

A converter method converts an object of one class to an object of another class. Two of the most important converter methods contained in MATLAB classes are `double` and `char`. Conversion to `double` produces the MATLAB traditional matrix, although this may not be appropriate for some classes. Conversion to `char` is useful for producing printed output.

### The Polynom to Double Converter

The double converter method for the polynom class is a very simple M-file, @polynom/double.m, which merely retrieves the coefficient vector.

```
function c = double(p)
% POLYNOM/DOUBLE Convert polynom object to coefficient vector.
% c = DOUBLE(p) converts a polynomial object to the vector c
% containing the coefficients of descending powers of x.
c = p.c;
```

On the object p,

```
p = polynom([1 0 -2 -5])
```

the statement

```
double(p)
```

returns

```
ans =
     1     0    -2    -5
```

### The Polynom to Char Converter

The converter to char is a key method because it produces a character string involving the powers of an independent variable, x. Therefore, once you have specified x, the string returned is a syntactically correct MATLAB expression, which you can then evaluate.

Here is @polynom/char.m.

```
function s = char(p)
% POLYNOM/CHAR
% CHAR(p) is the string representation of p.c
if all(p.c == 0)
    s = '0';
else
    d = length(p.c) - 1;
    s = [];
```

```

for a = p.c;
    if a ~= 0;
        if ~isempty(s)
            if a > 0
                s = [s ' + '];
            else
                s = [s ' - '];
                a = -a;
            end
        end
        if a ~= 1 | d == 0
            s = [s num2str(a)];
            if d > 0
                s = [s '*''];
            end
        end
        if d >= 2
            s = [s 'x^' int2str(d)];
        elseif d == 1
            s = [s 'x'];
        end
        end
        d = d - 1;
    end
end

```

### Evaluating the Output

If you create the polynom object p

```
p = polynom([1 0 -2 -5]);
```

and then call the char method on p

```
char(p)
```

MATLAB produces the result

```
ans =
    x^3 - 2*x - 5
```

The value returned by `char` is a string that you can pass to `eval` once you have defined a scalar value for `x`. For example,

```
x = 3;

eval(char(p))
ans =
    16
```

See “The Polynom `subsref` Method” on page 8-29 for a better method to evaluate the polynomial.

## The Polynom `display` Method

Here is `@polynom/display.m`. This method relies on the `char` method to produce a string representation of the polynomial, which is then displayed on the screen. This method produces output that is the same as standard MATLAB output. That is, the variable name is displayed followed by an equal sign, then a blank line, then a new line with the value.

```
function display(p)
% POLYNOM/DISPLAY Command window display of a polynom
disp(' ');
disp([inputname(1), ' = '])
disp(' ');
disp(['   ' char(p)])
disp(' ');
```

The statement

```
p = polynom([1 0 -2 -5])
```

creates a polynom object. Since the statement is not terminated with a semicolon, the resulting output is

```
p =
    x^3 - 2*x - 5
```

## The Polynom subsref Method

Suppose the design of the polynom class specifies that a subscripted reference to a polynom object causes the polynomial to be evaluated with the value of the independent variable equal to the subscript. That is, for a polynom object  $p$ ,

```
p = polynom([1 0 -2 -5]);
```

the following subscripted expression returns the value of the polynomial at  $x = 3$  and  $x = 4$ .

```
p([3 4])
ans =
    16    51
```

### subsref Implementation Details

This implementation takes advantage of the char method already defined in the polynom class to produce an expression that can then be evaluated.

```
function b = subsref(a,s)
% SUBSREF
switch s.type
case '('
    ind = s.subs{:};
    for k = 1:length(ind)
        b(k) = eval(strrep(char(a),'x',num2str(ind(k))));
    end
otherwise
    error('Specify value for x as p(x)')
end
```

Once the polynomial expression has been generated by the char method, the strrep function is used to swap the passed in value for the character  $x$ . The eval function then evaluates the expression and returns the value in the output argument.

## Overloading Arithmetic Operators for polynom

Several arithmetic operations are meaningful on polynomials and should be implemented for the polynom class. When overloading arithmetic operators, keep in mind what data types you want to operate on. In this section, the plus, minus, and mtimes methods are defined for the polynom class to handle

addition, subtraction, and multiplication on `polynom/polynom` and `polynom/double` combinations of operands.

### Overloading the + Operator

If either `p` or `q` is a `polynom`, the expression

$$p + q$$

generates a call to a function `@polynom/plus.m`, if it exists (unless `p` or `q` is an object of a higher precedence, as described in “Object Precedence” on page 8-65).

The following M-file redefines the + operator for the `polynom` class.

```
function r = plus(p,q)
% POLYNOM/PLUS Implement p + q for polynoms.
p = polynom(p);
q = polynom(q);
k = length(q.c) - length(p.c);
r = polynom([zeros(1,k) p.c] + [zeros(1,-k) q.c]);
```

The function first makes sure that both input arguments are polynomials. This ensures that expressions such as

$$p + 1$$

that involve both a `polynom` and a `double`, work correctly. The function then accesses the two coefficient vectors and, if necessary, pads one of them with zeros to make them the same length. The actual addition is simply the vector sum of the two coefficient vectors. Finally, the function calls the `polynom` constructor a third time to create the properly typed result.

### Overloading the – Operator

You can implement the overloaded minus operator (-) using the same approach as the plus (+) operator. MATLAB calls `@polynom/minus.m` to compute `p-q`.

```
function r = minus(p,q)
% POLYNOM/MINUS Implement p - q for polynoms.
p = polynom(p);
q = polynom(q);
k = length(q.c) - length(p.c);
r = polynom([zeros(1,k) p.c] - [zeros(1,-k) q.c]);
```

## Overloading the \* Operator

MATLAB calls the method `@polynom/mtimes.m` to compute the product  $p*q$ . The letter `m` at the beginning of the function name comes from the fact that it is overloading the MATLAB *matrix* multiplication. Multiplication of two polynomials is simply the convolution of their coefficient vectors.

```
function r = mtimes(p,q)
% POLYNOM/MTIMES    Implement p * q for polynoms.
p = polynom(p);
q = polynom(q);
r = polynom(conv(p.c,q.c));
```

## Using the Overloaded Operators

Given the `polynom` object

```
p = polynom([1 0 -2 -5])
```

MATLAB calls these two functions `@polynom/plus.m` and `@polynom/mtimes.m` when you issue the statements

```
q = p+1
r = p*q
```

to produce

```
q =
    x^3 - 2*x - 4
```

```
r =
    x^6 - 4*x^4 - 9*x^3 + 4*x^2 + 18*x + 20
```

## Overloading Functions for the Polynom Class

MATLAB already has several functions for working with polynomials represented by coefficient vectors. They should be overloaded to also work with the new `polynom` object. In many cases, the overloading methods can simply apply the original function to the coefficient field.

### Overloading roots for the Polynom Class

The method @polynom/roots.m finds the roots of polynom objects.

```
function r = roots(p)
% POLYNOM/ROOTS. ROOTS(p) is a vector containing the roots of p.
r = roots(p.c);
```

The statement

```
roots(p)
```

results in

```
ans =
    2.0946
   -1.0473 + 1.1359i
   -1.0473 - 1.1359i
```

### Overloading polyval for the Polynom Class

The function polyval evaluates a polynomial at a given set of points.

@polynom/polyval.m uses nested multiplication, or Horner's method to reduce the number of multiplication operations used to compute the various powers of x.

```
function y = polyval(p,x)
% POLYNOM/POLYVAL POLYVAL(p,x) evaluates p at the points x.
y = 0;
for a = p.c
    y = y.*x + a;
end
```

### Overloading plot for the Polynom Class

The overloaded plot function uses both root and polyval. The function selects the domain of the independent variable to be slightly larger than an interval containing all real roots. Then polyval is used to evaluate the polynomial at a few hundred points in the domain.

```
function plot(p)
% POLYNOM/PLOT PLOT(p) plots the polynom p.
r = max(abs(roots(p)));
x = (-1.1:0.01:1.1)*r;
y = polyval(p,x);
```

```

plot(x,y);
title(char(p))
grid on

```

### Overloading diff for the Polynom Class

The method @polynom/diff.m differentiates a polynomial by reducing the degree by 1 and multiplying each coefficient by its original degree.

```

function q = diff(p)
% POLYNOM/DIFF DIFF(p) is the derivative of the polynom p.
c = p.c;
d = length(c) - 1; % degree
q = polynom(p.c(1:d).*(d:-1:1));

```

### Listing Class Methods

The function call

```
methods('classname')
```

or its command form

```
methods classname
```

shows all the methods available for a particular class. For the polynom example, the output is

```

methods polynom

Methods for class polynom:

char    display    minus    plot    polynom    roots
diff    double    mtimes    plus    polyval    subsref

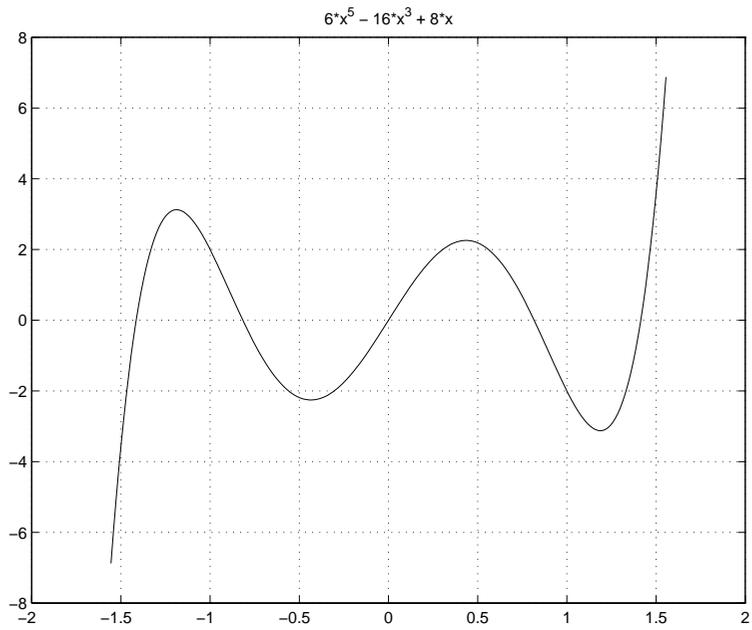
```

Plotting the two polynom objects x and p calls most of these methods.

```

x = polynom([1 0]);
p = polynom([1 0 -2 -5]);
plot(diff(p*p + 10*p + 20*x) - 20)

```



## Building on Other Classes

A MATLAB object can *inherit* properties and behavior from another MATLAB object. When one object (the child) inherits from another (the parent), the child object includes all the fields of the parent object and can call the parent's methods. The parent methods can access those fields that a child object inherited from the parent class, but not fields new to the child class.

Inheritance is a key feature of object-oriented programming. It makes it easy to reuse code by allowing child objects to take advantage of code that exists for parent objects. Inheritance enables a child object to behave exactly like a parent object, which facilitates the development of related classes that behave similarly, but are implemented differently.

There are two kinds of inheritance:

- Simple inheritance, in which a child object inherits characteristics from one parent class.
- Multiple inheritance, in which a child object inherits characteristics from more than one parent class.

This section also discusses a related topic, *aggregation*. Aggregation allows one object to contain another object as one of its fields.

### Simple Inheritance

A class that inherits attributes from a single parent class, and adds new attributes of its own, uses simple inheritance. Inheritance implies that objects belonging to the child class have the same fields as the parent class, as well as additional fields. Therefore, methods associated with the parent class can operate on objects belonging to the child class. The methods associated with the child class, however, cannot operate on objects belonging to the parent class. You cannot access the parent's fields directly from the child class; you must use access methods defined for the parent.

The constructor function for a class that inherits the behavior of another has two special characteristics:

- It calls the constructor function for the parent class to create the inherited fields.
- The calling syntax for the `class` function is slightly different, reflecting both the child class and the parent class.

The general syntax for establishing a simple inheritance relationship using the `class` function is

```
childObj = class(childObj, 'childClass', parentObj);
```

Simple inheritance can span more than one generation. If a parent class is itself an inherited class, the child object will automatically inherit from the grandparent class.

### Visibility of Class Properties and Methods

The parent class does not have knowledge of the child properties or methods. The child class cannot access the parent properties directly, but must use parent access methods (e.g., `get` or `subsref` method) to access the parent properties. From the child class methods, this access is accomplished via the parent field in the child structure. For example, when a constructor creates a child object `c`,

```
c = class(c, 'childClassname', parentObject);
```

MATLAB automatically creates a field, `c.parentClassname`, in the object's structure that contains the parent object. You could then have a statement in the child's `display` method that calls the parent's `display` method.

```
display(c.parentClassname)
```

See “Designing the Stock Class” on page 8-45 for examples that use simple inheritance.

## Multiple Inheritance

In the multiple inheritance case, a class of objects inherits attributes from more than one parent class. The child object gets fields from all the parent classes, as well as fields of its own.

Multiple inheritance can encompass more than one generation. For example, each of the parent objects could have inherited fields from multiple grandparent objects, and so on. Multiple inheritance is implemented in the constructors by calling `class` with more than three arguments.

```
obj = class(structure, 'classname', parent1, parent2, ...)
```

You can append as many parent arguments as desired to the `class` input list.

Multiple parent classes can have associated methods of the same name. In this case, MATLAB calls the method associated with the parent that appears first in the `class` function call in the constructor function. There is no way to access subsequent parent function of this name.

## Aggregation

In addition to standard inheritance, MATLAB objects support *containment* or *aggregation*. That is, one object can contain (embed) another object as one of its fields. For example, a rational object might use two polynom objects, one for the numerator and one for the denominator.

You can call a method for the contained object only from within a method for the outer object. When determining which version of a function to call, MATLAB considers only the outermost containing class of the objects passed as arguments; the classes of any contained objects are ignored.

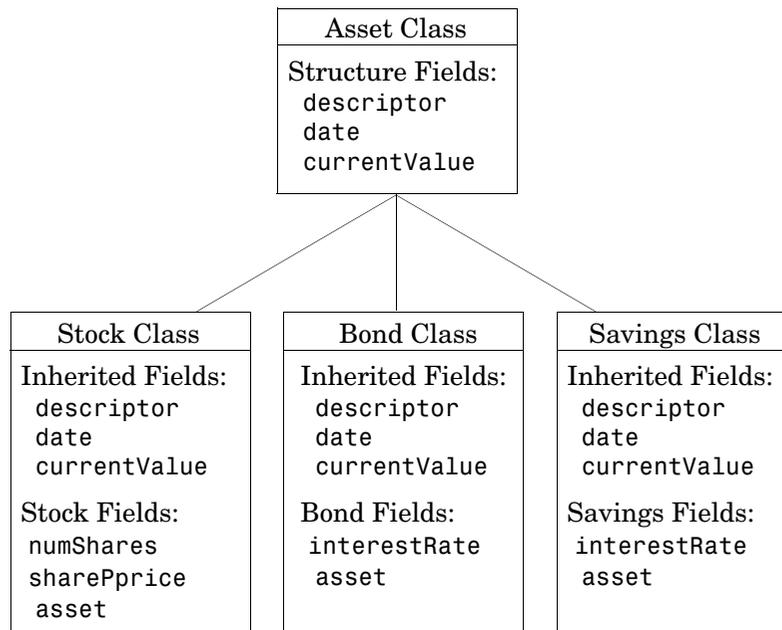
See “Example — The Portfolio Container” on page 8-54 for an example of aggregation.

## Example – Assets and Asset Subclasses

As an example of simple inheritance, consider a general asset class that can be used to represent any item that has monetary value. Some examples of an asset are: stocks, bonds, savings accounts, and any other piece of property. In designing this collection of classes, the asset class holds the data that is common to all of the specialized asset subclasses. The individual asset subclasses, such as the stock class, inherit the asset properties and contribute additional properties. The subclasses are “kinds of” assets.

### Inheritance Model for the Asset Class

An example of a simple inheritance relationship using an asset parent class is shown in this diagram.



As shown in the diagram, the stock, bond, and savings classes inherit structure fields from the asset class. In this example, the asset class is used to provide storage for data common to all subclasses and to share asset methods with these subclasses. This example shows how to implement the

asset and stock classes. The bond and savings classes can be implemented in a way that is very similar to the stock class, as would other types of asset subclasses.

## Asset Class Design

The asset class provides storage and access for information common to all asset children. It is not intended to be instantiated directly, so it does not require an extensive set of methods. To serve its purpose, the class needs to contain the following methods:

- Constructor
- get and set
- subsref and subsasgn
- display

## Other Asset Methods

The asset class provides inherited data storage for its child classes, but is not instanced directly. The set, get, and display methods provide access to the stored data. It is not necessary to implement the full complement of methods for asset objects (such as converters, end, and subsindex) since only the child classes access the data.

## The Asset Constructor Method

The asset class is based on a structure array with four fields:

- descriptor — Identifier of the particular asset (e.g., stock name, savings account number, etc.)
- date — The date the object was created (calculated by the date command)
- type — The type of asset (e.g., savings, bond, stock)
- currentValue — The current value of the asset (calculated from subclass data)

This information is common to asset child objects (stock, bond, and savings), so it is handled from the parent object to avoid having to define the same fields in each child class. This is particularly helpful as the number of child classes increases.

```
function a = asset(varargin)
% ASSET Constructor function for asset object
% a = asset(descriptor, currentValue)
switch nargin
case 0
% if no input arguments, create a default object
    a.descriptor = 'none';
    a.date = date;
    a.type = 'none';
    a.currentValue = 0;
    a = class(a,'asset');
case 1
% if single argument of class asset, return it
    if (isa(varargin{1},'asset'))
        a = varargin{1};
    else
        error('Wrong argument type')
    end
case 3
% create object using specified values
    a.descriptor = varargin{1};
    a.date = date;
    a.type = varargin{2};
    a.currentValue = varargin{3};
    a = class(a,'asset');
otherwise
    error('Wrong number of input arguments')
end
```

The function uses a switch statement to accommodate three possible scenarios:

- Called with no arguments, the constructor returns a default asset object.
- Called with one argument that is an asset object, the object is simply returned.
- Called with two arguments (subclass descriptor, and current value), the constructor returns a new asset object.

The asset constructor method is not intended to be called directly; it is called from the child constructors since its purpose is to provide storage for common data.

## The Asset get Method

The asset class needs methods to access the data contained in asset objects. The following function implements a get method for the class. It uses capitalized property names rather than literal field names to provide an interface similar to other MATLAB objects.

```
function val = get(a, propName)
% GET Get asset properties from the specified object
% and return the value
switch propName
case 'Descriptor'
    val = a.descriptor;
case 'Date'
    val = a.date;
case 'CurrentValue'
    val = a.currentValue;
otherwise
    error([propName, ' Is not a valid asset property'])
end
```

This function accepts an object and a property name and uses a switch statement to determine which field to access. This method is called by the subclass get methods when accessing the data in the inherited properties. See “The Stock get Method” on page 8-48 for an example.

## The Asset set Method

The asset class set method is called by subclass set methods. This method accepts an asset object and variable length argument list of property name/property value pairs and returns the modified object.

```
function a = set(a,varargin)
% SET Set asset properties and return the updated object
propertyArgIn = varargin;
while length(propertyArgIn) >= 2,
    prop = propertyArgIn{1};
    val = propertyArgIn{2};
```

```
propertyArgIn = propertyArgIn(3:end);
switch prop
case 'Descriptor'
    a.descriptor = val;
case 'Date'
    a.date = val;
case 'CurrentValue'
    a.currentValue = val;
otherwise
    error('Asset properties: Descriptor, Date, CurrentValue')
end
end
```

Subclass set methods call the asset set method and require the capability to return the modified object since MATLAB does not support passing arguments by reference. See “The Stock set Method” on page 8-49 for an example.

## The Asset subsref Method

The subsref method provides access to the data contained in an asset object using one-based numeric indexing and structure field name indexing. The outer switch statement determines if the index is a numeric or field name syntax. The inner switch statements map the index to the appropriate value.

MATLAB calls subsref whenever you make a subscripted reference to an object (e.g., `A(i)`, `A{i}`, or `A.fieldname`).

```
function b = subsref(a,index)
%SUBSREF Define field name indexing for asset objects
switch index.type
case '('
    switch index.subs{:}
    case 1
        b = a.descriptor;
    case 2
        b = a.date;
    case 3
        b = a.currentValue;
    otherwise
        error('Index out of range')
    end
end
```

```

case '.'
    switch index.subs
    case 'descriptor'
        b = a.descriptor;
    case 'date'
        b = a.date;
    case 'currentValue'
        b = a.currentValue;
    otherwise
        error('Invalid field name')
    end
case '{} '
    error('Cell array indexing not supported by asset objects')
end

```

See the “The Stock subsref Method” on page 8-50 for an example of how the child subsref method calls the parent subsref method.

## The Asset subsasgn Method

The subsasgn method is the assignment equivalent of the subsref method. This version enables you to change the data contained in an object using one-based numeric indexing and structure field name indexing. The outer switch statement determines if the index is a numeric or field name syntax. The inner switch statements map the index value to the appropriate value in the stock structure.

MATLAB calls subsasgn whenever you execute an assignment statement (e.g., `A(i) = val`, `A{i} = val`, or `A.fieldname = val`).

```

function a = subsasgn(a,index,val)
% SUBSASGN Define index assignment for asset objects
switch index.type
case '('
    switch index.subs{:}
    case 1
        a.descriptor = val;
    case 2
        a.date = val;

```

```
        case 3
            a.currentValue = val;
        otherwise
            error('Index out of range')
        end
    case '.'
        switch index.subs
        case 'descriptor'
            a.descriptor = val;
        case 'date'
            a.date = val;
        case 'currentValue'
            a.currentValue = val;
        otherwise
            error('Invalid field name')
        end
    end
end
```

The `subsasgn` method enables you to assign values to the asset object data structure using two techniques. For example, suppose you have a child stock object `s`. (If you want to run this statement, you first need to create a stock constructor method.)

```
s = stock('XYZ',100,25);
```

Within stock class methods, you could change the descriptor field with either of the following statements

```
s.asset(1) = 'ABC';
```

or

```
s.asset.descriptor = 'ABC';
```

See the “The Stock `subsasgn` Method” on page 8-51 for an example of how the child `subsasgn` method calls the parent `subsasgn` method.

## The Asset display Method

The asset display method is designed to be called from child-class `display` methods. Its purpose is to display the data it stores for the child object. The method simply formats the data for display in a way that is consistent with the formatting of the child’s `display` method.

```

function display(a)
% DISPLAY(a) Display an asset object
stg = sprintf(...
    'Descriptor: %s\nDate: %s\nType: %s\nCurrent Value:%9.2f',...
    a.descriptor,a.date,a.type,a.currentValue);
disp(stg)

```

The stock class display method can now call this method to display the data stored in the parent class. This approach isolates the stock display method from changes to the asset class. See “The Stock display Method” on page 8-53 for an example of how this method is called.

## The Asset fieldcount Method

The asset fieldcount method returns the number of fields in the asset object data structure. fieldcount enables asset child methods to determine the number of fields in the asset object during execution, rather than requiring the child methods to have knowledge of the asset class. This allows you to make changes to the number of fields in the asset class data structure without having to change child-class methods.

```

function numFields = fieldcount(assetObj)
% Determines the number of fields in an asset object
% Used by asset child class methods
numFields = length(fieldnames(assetObj));

```

The struct function converts an object to its equivalent data structure, enabling access to the structure’s contents.

## Designing the Stock Class

A stock object is designed to represent one particular asset in a person’s investment portfolio. This object contains two properties of its own and inherits three properties from its parent asset object.

Stock properties:

- NumberShares — The number of shares for the particular stock object.
- SharePrice — The value of each share.

Asset properties:

- **Descriptor** — The identifier of the particular asset (e.g., stock name, savings account number, etc.).
- **Date** — The date the object was created (calculated by the `date` command).
- **CurrentValue** — The current value of the asset.

Note that the property names are not actually the same as the field names of the structure array used internally by stock and asset objects. The property name interface is controlled by the stock and asset `set` and `get` methods and is designed to resemble the interface of other MATLAB object properties.

The asset field in the stock object structure contains the parent asset object and is used to access the inherited fields in the parent structure.

### **Stock Class Methods**

The stock class implements the following methods:

- **Constructor**
- `get` and `set`
- `subsref` and `subsasgn`
- `display`

### **The Stock Constructor Method**

The stock constructor creates a stock object from three input arguments:

- The stock name
- The number of shares
- The share price

The constructor must create an asset object from within the stock constructor to be able to specify it as a parent to the stock object. The stock constructor must, therefore, call the asset constructor. The `class` function, which is called to create the stock object, defines the asset object as the parent.

Keep in mind that the asset object is created in the temporary workspace of the stock constructor function and is stored as a field (`.asset`) in the stock structure. The stock object inherits the asset fields, but the asset object is not returned to the base workspace.

```

function s = stock(varargin)
% STOCK Stock class constructor.
% s = stock(descriptor, numShares, sharePrice)
switch nargin
case 0
% if no input arguments, create a default object
    s.numShares = 0;
    s.sharePrice = 0;
    a = asset('none',0);
    s = class(s, 'stock',a);
case 1
% if single argument of class stock, return it
    if (isa(varargin{1},'stock'))
        s = varargin{1};
    else
        error('Input argument is not a stock object')
    end
case 3
% create object using specified values
    s.numShares = varargin{2};
    s.sharePrice = varargin{3};
    a = asset(varargin{1},'stock',varargin{2} * varargin{3});
    s = class(s,'stock',a);
otherwise
    error('Wrong number of input arguments')
end

```

### Constructor Calling Syntax

The stock constructor method can be called in one of three ways:

- No input argument — If called with no arguments, the constructor returns a default object with empty fields.
- Input argument is a stock object — If called with a single input argument that is a stock object, the constructor returns the input argument. A single argument that is not a stock object generates an error.
- Three input arguments — If there are three input arguments, the constructor uses them to define the stock object.

Otherwise, if none of the above three conditions are met, return an error.

For example, this statement creates a stock object to record the ownership of 100 shares of XYZ corporation stocks with a price per share of 25 dollars.

```
XYZStock = stock('XYZ',100,25);
```

## The Stock get Method

The get method provides a way to access the data in the stock object using a “property name” style interface, similar to Handle Graphics. While in this example the property names are similar to the structure field name, they can be quite different. You could also choose to exclude certain fields from access via the get method or return the data from the same field for a variety of property names, if such behavior suits your design.

```
function val = get(s,propName)
% GET Get stock property from the specified object
% and return the value. Property names are: NumberShares
% SharePrice, Descriptor, Date, CurrentValue
switch propName
case 'NumberShares'
    val = s.numShares;
case 'SharePrice'
    val = s.sharePrice;
case 'Descriptor'
    val = get(s.asset,'Descriptor'); % call asset get method
case 'Date'
    val = get(s.asset,'Date');
case 'CurrentValue'
    val = get(s.asset,'CurrentValue');
otherwise
    error([propName , 'Is not a valid stock property'])
end
```

Note that the asset object is accessed via the stock object’s asset field (s.asset). MATLAB automatically creates this field when the class function is called with the parent argument.

## The Stock set Method

The set method provides a “property name” interface like the get method. It is designed to update the number of shares, the share value, and the descriptor. The current value and the date are automatically updated.

```
function s = set(s,varargin)
% SET Set stock properties to the specified values
% and return the updated object
propertyArgIn = varargin;
while length(propertyArgIn) >= 2,
    prop = propertyArgIn{1};
    val = propertyArgIn{2};
    propertyArgIn = propertyArgIn(3:end);
    switch prop
    case 'NumberShares'
        s.numShares = val;
    case 'SharePrice'
        s.sharePrice = val;
    case 'Descriptor'
        s.asset = set(s.asset,'Descriptor',val);
    otherwise
        error('Invalid property')
    end
end
s.asset = set(s.asset,'CurrentValue',...
             s.numShares * s.sharePrice,'Date',date);
```

Note that this function creates and returns a new stock object with the new values, which you then copy over the old value. For example, given the stock object,

```
s = stock('XYZ',100,25);
```

the following set command updates the share price.

```
s = set(s,'SharePrice',36);
```

It is necessary to copy over the original stock object (i.e., assign the output to s) because MATLAB does not support passing arguments by reference. Hence the set method actually operates on a copy of the object.

## The Stock subsref Method

The subsref method defines subscripted indexing for the stock class. In this example, subsref is implemented to enable numeric and structure field name indexing of stock objects.

```
function b = subsref(s,index)
% SUBSREF Define field name indexing for stock objects
fc = fieldcount(s.asset);
switch index.type
case '()'
    if (index.subs{:} <= fc)
        b = subsref(s.asset,index);
    else
        switch index.subs{:} - fc
        case 1
            b = s.numShares;
        case 2
            b = s.sharePrice;
        otherwise
            error(['Index must be in the range 1 to ',num2str(fc + 2)])
        end
    end
case '.'
    switch index.subs
    case 'numShares'
        b = s.numShares;
    case 'sharePrice'
        b = s.sharePrice;
    otherwise
        b = subsref(s.asset,index);
    end
end
end
```

The outer switch statement determines if the index is a numeric or field name syntax.

The fieldcount asset method determines how many fields there are in the asset structure, and the if statement calls the asset subsref method for indices 1 to fieldcount. See “The Asset fieldcount Method” on page 8-45 and “The Asset subsref Method” on page 8-42 for a description of these methods.

Numeric indices greater than the number returned by `fieldcount` are handled by the inner switch statement, which maps the index value to the appropriate field in the stock structure.

Field-name indexing assumes field names other than `numShares` and `sharePrice` are asset fields, which eliminates the need for knowledge of asset fields by child methods. The `asset subsref` method performs field-name error checking.

See the `subsref` help entry for general information on implementing this method.

## The Stock `subsasgn` Method

The `subsasgn` method enables you to change the data contained in a stock object using numeric indexing and structure field name indexing. MATLAB calls `subsasgn` whenever you execute an assignment statement (e.g., `A(i) = val`, `A{i} = val`, or `A.fieldname = val`).

```
function s = subsasgn(s,index,val)
% SUBSASGN Define index assignment for stock objects
fc = fieldcount(s.asset);
switch index.type
case '()'
    if (index.subs{:} <= fc)
        s.asset = subsasgn(s.asset,index,val);
    else
        switch index.subs{:}-fc
        case 1
            s.numShares = val;
        case 2
            s.sharePrice = val;
        otherwise
            error(['Index must be in the range 1 to ',num2str(fc + 2)])
        end
    end
case '.'
    switch index.subs
    case 'numShares'
        s.numShares = val;
```

```
        case 'sharePrice'  
            s.sharePrice = val;  
        otherwise  
            s.asset = subsasgn(s.asset,index,val);  
        end  
    end
```

The outer switch statement determines if the index is a numeric or field name syntax.

The `fieldcount` asset method determines how many fields there are in the asset structure and the `if` statement calls the asset `subsasgn` method for indices 1 to `fieldcount`. See “The Asset `fieldcount` Method” on page 8-45 and “The Asset `subsasgn` Method” on page 8-43 for a description of these methods.

Numeric indices greater than the number returned by `fieldcount` are handled by the inner switch statement, which maps the index value to the appropriate field in the stock structure.

Field-name indexing assumes field names other than `numShares` and `sharePrice` are asset fields, which eliminates the need for knowledge of asset fields by child methods. The asset `subsasgn` method performs field-name error checking.

The `subsasgn` method enables you to assign values to stock object data structure using two techniques. For example, suppose you have a stock object

```
s = stock('XYZ',100,25)
```

You could change the descriptor field with either of the following statements

```
s(1) = 'ABC';
```

or

```
s.descriptor = 'ABC';
```

See the `subsasgn` help entry for general information on assignment statements in MATLAB.

## The Stock display Method

When you issue the statement (without terminating with a semicolon)

```
XYZStock = stock('XYZ',100,25)
```

MATLAB looks for a method in the @stock directory called `display`. The `display` method for the stock class produces this output.

```
Descriptor: XYZ  
Date: 17-Nov-1998  
Type: stock  
Current Value: 2500.00  
Number of shares: 100  
Share price: 25.00
```

Here is the stock display method.

```
function display(s)  
% DISPLAY(s) Display a stock object  
display(s.asset)  
stg = sprintf('Number of shares: %g\nShare price: %3.2f\n',...  
    s.numShares,s.sharePrice);  
disp(stg)
```

First, the parent asset object is passed to the asset display method to display its fields (MATLAB calls the asset display method because the input argument is an asset object). The stock object's fields are displayed in a similar way using a formatted text string.

Note that if you did not implement a stock class display method, MATLAB would call the asset display method. This would work, but would display only the descriptor, date, type, and current value.

## Example – The Portfolio Container

Aggregation is the containment of one class by another class. The basic relationship is: each contained class “is a part of” the container class.

For example, consider a financial portfolio class as a container for a set of assets (stocks, bonds, savings, etc.). Once the individual assets are grouped, they can be analyzed, and useful information can be returned. The contained objects are not accessible directly, but only via the portfolio class methods.

See “Example — Assets and Asset Subclasses” on page 8-38 for information about the assets collected by this portfolio class.

### Designing the Portfolio Class

The portfolio class is designed to contain the various assets owned by a given individual and provide information about the status of his or her investment portfolio. This example implements a somewhat over-simplified portfolio class that

- Contains an individual’s assets
- Displays information about the portfolio contents
- Displays a 3-D pie chart showing the relative mix of asset types in the portfolio

### Required Portfolio Methods

The portfolio class implements only three methods:

- `portfolio` — The portfolio constructor.
- `display` — Displays information about the portfolio contents.
- `pie3` — Overloaded version of `pie3` function designed to take a single portfolio object as an argument.

Since a portfolio object contains other objects, the portfolio class methods can use the methods of the contained objects. For example, the portfolio `display` method calls the stock class `display` method, and so on.

## The Portfolio Constructor Method

The portfolio constructor method takes as input arguments a client's name and a variable length list of asset subclass objects (stock, bond, and savings objects in this example). The portfolio object uses a structure array with the following fields:

- name — The client's name.
- indAssets — The array of asset subclass objects (stock, bond, savings).
- totalValue — The total value of all assets. The constructor calculates this value from the objects passed in as arguments.
- accountNumber — The account number. This field is assigned a value only when you save a portfolio object (see “Saving and Loading Objects” on page 8-60).

```
function p = portfolio(name,varargin)
% PORTFOLIO Create a portfolio object containing the
% client's name and a list of assets
switch nargin
case 0
    % if no input arguments, create a default object
    p.name = 'none';
    p.totalValue = 0;
    p.indAssets = {};
    p.accountNumber = '';
    p = class(p,'portfolio');
case 1
    % if single argument of class portfolio, return it
    if isa(name,'portfolio')
        p = name;
    else
        disp([inputname(1) ' is not a portfolio object'])
        return
    end
otherwise
    % create object using specified arguments
    p.name = name;
    p.totalValue = 0;
```

```
    for k = 1:length(varargin)
        p.indAssets(k) = {varargin{k}};
        assetValue = get(p.indAssets{k}, 'CurrentValue');
        p.totalValue = p.totalValue + assetValue;
    end
    p.accountNumber = '';
    p = class(p, 'portfolio');
end
```

### Constructor Calling Syntax

The portfolio constructor method can be called in one of three different ways:

- No input arguments — If called with no arguments, it returns an object with empty fields.
- Input argument is an object — If the input argument is already a portfolio object, MATLAB returns the input argument. The `isa` function checks for this case.
- More than two input arguments — If there are more than two input arguments, the constructor assumes the first is the client's name and the rest are asset subclass objects. A more thorough implementation would perform more careful input argument checking, for example, using the `isa` function to determine if the arguments are the correct class of objects.

### The Portfolio display Method

The portfolio `display` method lists the contents of each contained object by calling the object's `display` method. It then lists the client name and total asset value.

```
function display(p)
% DISPLAY Display a portfolio object
for k = 1:length(p.indAssets)
    display(p.indAssets{k})
end
stg = sprintf('\nAssets for Client: %s\nTotal Value: %9.2f\n',...
p.name,p.totalValue);
disp(stg)
```

## The Portfolio pie3 Method

The portfolio class overloads the MATLAB pie3 function to accept a portfolio object and display a 3-D pie chart illustrating the relative asset mix of the client's portfolio. MATLAB calls the @portfolio/pie3.m version of pie3 whenever the input argument is a single portfolio object.

```
function pie3(p)
% PIE3 Create a 3-D pie chart of a portfolio
stockAmt = 0; bondAmt = 0; savingsAmt = 0;
for k = 1:length(p.indAssets)
    if isa(p.indAssets{k}, 'stock')
        stockAmt = stockAmt + ...
            get(p.indAssets{k}, 'CurrentValue');
    elseif isa(p.indAssets{k}, 'bond')
        bondAmt = bondAmt + ...
            get(p.indAssets{k}, 'CurrentValue');
    elseif isa(p.indAssets{k}, 'savings')
        savingsAmt = savingsAmt + ...
            get(p.indAssets{k}, 'CurrentValue');
    end
end
i = 1;
if stockAmt ~= 0
    label(i) = {'Stocks'};
    pieVector(i) = stockAmt;
    i = i + 1;
end
if bondAmt ~= 0
    label(i) = {'Bonds'};
    pieVector(i) = bondAmt;
    i = i + 1;
end
if savingsAmt ~= 0
    label(i) = {'Savings'};
    pieVector(i) = savingsAmt;
end
pie3(pieVector, label)
set(gcf, 'Renderer', 'zbuffer')
set(findobj(gca, 'Type', 'Text'), 'FontSize', 14)
cm = gray(64);
```

```
colormap(cm(48:end, :))
stg(1) = {'Portfolio Composition for ', p.name}};
stg(2) = {'Total Value of Assets: $', num2str(p.totalValue)};
title(stg, 'FontSize', 12)
```

There are three parts in the overloaded `pie3` method.

- The first uses the asset subclass get methods to access the `CurrentValue` property of each contained object. The total value of each class is summed.
- The second part creates the pie chart labels and builds a vector of graph data, depending on which objects are present.
- The third part calls the MATLAB `pie3` function, makes some font and colormap adjustments, and adds a title.

## Creating a Portfolio

Suppose you have implemented a collection of asset subclasses in a manner similar to the stock class. You can then use a portfolio object to present the individual's financial portfolio. For example, given the following assets

```
XYZStock = stock('XYZ', 200, 12);
SaveAccount = savings('Acc # 1234', 2000, 3.2);
Bonds = bond('U.S. Treasury', 1600, 12);
```

create a portfolio object:

```
p = portfolio('Gilbert Bates', XYZStock, SaveAccount, Bonds)
```

The portfolio `display` method summarizes the portfolio contents (because this statement is not terminated by a semicolon).

```
Descriptor: XYZ
Date: 24-Nov-1998
Current Value: 2400.00
Type: stock
Number of shares: 200
Share price: 12.00

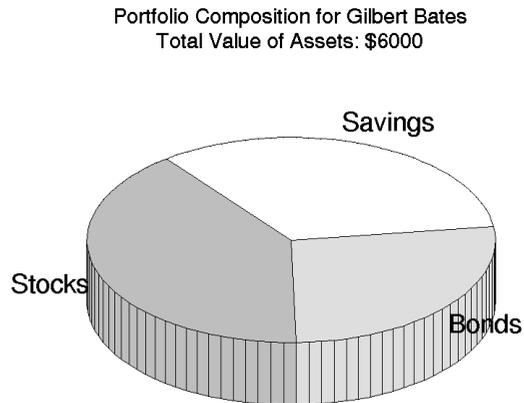
Descriptor: Acc # 1234
Date: 24-Nov-1998
Current Value: 2000.00
Type: savings
Interest Rate: 3.2%
```

Descriptor: U.S. Treasury  
Date: 24-Nov-1998  
Current Value: 1600.00  
Type: bond  
Interest Rate: 12%

Assets for Client: Gilbert Bates  
Total Value: 6000.00

The portfolio pie3 method displays the relative mix of assets using a pie chart.

pie3(p)



## **Saving and Loading Objects**

You can use the MATLAB `save` and `load` commands to save and retrieve user-defined objects to and from `.mat` files, just like any other variables.

When you load objects, MATLAB calls the object's class constructor to register the object in the workspace. The constructor function for the object class you are loading must be able to be called with no input arguments and return a default object. See “Guidelines for Writing a Constructor” on page 8-10 for more information.

### **Modifying Objects During Save or Load**

When you issue a `save` or `load` command on objects, MATLAB looks for class methods called `saveobj` and `loadobj` in the class directory. You can overload these methods to modify the object before the save or load operation. For example, you could define a `saveobj` method that saves related data along with the object or you could write a `loadobj` method that updates objects to a newer version when this type of object is loaded into the MATLAB workspace.

## Example — Defining saveobj and loadobj for Portfolio

In the section “Example — The Portfolio Container” on page 8-54, portfolio objects are used to collect information about a client’s investment portfolio. Now suppose you decide to add an account number to each portfolio object that is saved. You can define a portfolio saveobj method to carry out this task automatically during the save operation.

Suppose further that you have already saved a number of portfolio objects without the account number. You want to update these objects during the load operation so that they are still valid portfolio objects. You can do this by defining a loadobj method for the portfolio class.

### Summary of Code Changes

To implement the account number scenario, you need to add or change the following functions:

- `portfolio` — The portfolio constructor method needs to be modified to create a new field, `accountNumber`, which is initialized to the empty string when an object is created.
- `saveobj` — A new portfolio method designed to add an account number to a portfolio object during the save operation, only if the object does not already have one.
- `loadobj` — A new portfolio method designed to update older versions of portfolio objects that were saved before the account number structure field was added.
- `suboref` — A new portfolio method that enables subscripted reference to portfolio objects outside of a portfolio method.
- `getAccountNumber` — a MATLAB function that returns an account number that consists of the first three letters of the client’s name.

### New Portfolio Class Behavior

With the additions and changes made in this example, the portfolio class now

- Includes a field for an account number
- Adds the account number when a portfolio object is saved for the first time

- Automatically updates the older version of portfolio objects when you load them into the MATLAB workspace

## The saveobj Method

MATLAB looks for the portfolio `saveobj` method whenever the `save` command is passed a portfolio object. If `@portfolio/saveobj` exists, MATLAB passes the portfolio object to `saveobj`, which must then return the modified object as an output argument. The following implementation of `saveobj` determines if the object has already been assigned an account number from a previous save operation. If not, `saveobj` calls `getAccountNumber` to obtain the number and assigns it to the `accountNumber` field.

```
function b = saveobj(a)
if isempty(a.accountNumber)
    a.accountNumber = getAccountNumber(a);
end
b = a;
```

## The loadobj Method

MATLAB looks for the portfolio `loadobj` method whenever the `load` command detects portfolio objects in the `.mat` file being loaded. If `loadobj` exists, MATLAB passes the portfolio object to `loadobj`, which must then return the modified object as an output argument. The output argument is then loaded into the workspace.

If the input object does not match the current definition as specified by the constructor function, then MATLAB converts it to a structure containing the same fields and the object's structure with all the values intact (that is, you now have a structure, not an object).

The following implementation of `loadobj` first uses `isa` to determine whether the input argument is a portfolio object or a structure. If the input is an object, it is simply returned since no modifications are necessary. If the input argument has been converted to a structure by MATLAB, then the new `accountNumber` field is added to the structure and is used to create an updated portfolio object.

```

function b = loadobj(a)
% loadobj for portfolio class
if isa(a,'portfolio')
    b = a;
else % a is an old version
    a.accountNumber = getAccountNumber(a);
    b = class(a,'portfolio');
end

```

## Changing the Portfolio Constructor

The portfolio structure array needs an additional field to accommodate the account number. To create this field, add the line

```
p.accountNumber = '';
```

to @portfolio/portfolio.m in both the zero argument and variable argument sections.

## The getAccountNumber Function

In this example, getAccountNumber is a MATLAB function that returns an account number composed of the first three letters of the client name prepended to a series of digits. To illustrate implementation techniques, getAccountNumber is not a portfolio method so it cannot access the portfolio object data directly. Therefore, it is necessary to define a portfolio subref method that enables access to the name field in a portfolio object's structure.

For this example, getAccountNumber simply generates a random number, which is formatted and concatenated with elements 1 to 3 from the portfolio name field.

```

function n = getAccountNumber(p)
% provides a account number for object p
n = [upper(p.name(1:3)) strcat(num2str(round(rand(1,7)*10))'')'];

```

Note that the portfolio object is indexed by field name, and then by numerical subscript to extract the first three letters. The subref method must be written to support this form of subscripted reference.

## The Portfolio subsref Method

When MATLAB encounters a subscripted reference, such as that made in the `getAccountNumber` function

```
p.name(1:3)
```

MATLAB calls the `portfolio subsref` method to interpret the reference. If you do not define a `subsref` method, the above statement is undefined for portfolio objects (recall that here `p` is an object, not just a structure).

The `portfolio subsref` method must support field-name and numeric indexing for the `getAccountNumber` function to access the portfolio name field.

```
function b = subsref(p,index)
% SUBSREF Define field name indexing for portfolio objects
switch index(1).type
case '.'
    switch index(1).subs
    case 'name'
        if length(index)== 1
            b = p.name;
        else
            switch index(2).type
            case '()'
                b = p.name(index(2).subs{:});
            end
        end
    end
end
end
```

Note that the `portfolio` implementation of `subsref` is designed to provide access to specific elements of the `name` field; it is not a general implementation that provides access to all structure data, such as the `stock` class implementation of `subsref`.

See the `subsref` help entry for more information about indexing and objects.

## Object Precedence

Object precedence is a means to resolve the question of which of possibly many versions of an operator or function to call in a given situation. Object precedence enables you to control the behavior of expressions containing different classes of objects. For example, consider the expression

$$\text{objectA} + \text{objectB}$$

Ordinarily, MATLAB assumes that the objects have equal precedence and calls the method associated with the leftmost object. However, there are two exceptions:

- User-defined classes have precedence over MATLAB built-in classes.
- User-defined classes can specify their relative precedence with respect to other user-defined classes using the `inferiorto` and `superiorto` functions.

For example, in the section “Example — A Polynomial Class” on page 8-24 the polynomial class defines a `plus` method that enables addition of polynomial objects. Given the polynomial object `p`

```
p = polynomial([1 0 -2 -5])
p =
    x^3-2*x-5
```

The expression,

```
1 + p
ans =
    x^3-2*x-4
```

calls the polynomial `plus` method (which converts the double, `1`, to a polynomial object, and then adds it to `p`). The user-defined polynomial class has precedence over the MATLAB double class.

## Specifying Precedence of User-Defined Classes

You can specify the relative precedence of user-defined classes by calling the `inferiorto` or `superiorto` function in the class constructor.

The `inferiorto` function places a class below other classes in the precedence hierarchy. The calling syntax for the `inferiorto` function is

```
inferiorto('class1', 'class2', ...)
```

You can specify multiple classes in the argument list, placing the class below many other classes in the hierarchy.

Similarly, the `superiorto` function places a class above other classes in the precedence hierarchy. The calling syntax for the `superiorto` function is

```
superiorto('class1', 'class2', ...)
```

### Location in the Hierarchy

If *objectA* is above *objectB* in the precedence hierarchy, then the expression

```
objectA + objectB
```

calls `@classA/plus.m`. Conversely, if *objectB* is above *objectA* in the precedence hierarchy, then MATLAB calls `@classB/plus.m`.

See “How MATLAB Determines Which Method to Call” on page 8-67 for related information.

## How MATLAB Determines Which Method to Call

In MATLAB, functions exist in directories in the computer's file system. A directory may contain many functions (M-files). Function names are unique only within a single directory (e.g., more than one directory may contain a function called `pie3`). When you type a function name on the command line, MATLAB must search all the directories it is aware of to determine which function to call. This list of directories is called the *MATLAB path*.

When looking for a function, MATLAB searches the directories in the order they are listed in the path, and calls the first function whose name matches the name of the specified function.

If you write an M-file called `pie3.m` and put it in a directory that is searched before the `specgraph` directory that contains the MATLAB `pie3` function, then MATLAB uses your `pie3` function instead (note that this is not true for built-in functions like `plot`, which are always found first).

Object-oriented programming allows you to have many methods (MATLAB functions located in class directories) with the same name and enables MATLAB to determine which method to use based on the type or class of the variables passed to the function. For example, if `p` is a portfolio object, then

```
pie3(p)
```

calls `@portfolio/pie3.m` because the argument is a portfolio object.

### Selecting a Method

When you call a method for which there are multiple versions with the same name, MATLAB determines the method to call by:

- Looking at the classes of the objects in the argument list to determine which argument has the highest object precedence; the class of this object controls the method selection and is called the *dispatch type*.
- Applying the *function precedence order* to determine which of possibly several implementations of a method to call. This order is determined by the location and type of function.

### Determining the Dispatch Type

MATLAB first determines which argument controls the method selection. The class type of this argument then determines the class in which MATLAB searches for the method. The controlling argument is either

- The argument with the highest precedence, or
- The leftmost of arguments having equal precedence

User-defined objects take precedence over the MATLAB built-in classes such as `double` or `char`. You can set the relative precedence of user-defined objects with the `inferioorto` and `superiorto` functions, as described in “Object Precedence” on page 8-65.

MATLAB searches for functions by name. When you call a function, MATLAB knows the name, number of arguments, and the type of each argument. MATLAB uses the dispatch type to choose among multiple functions of the same name, but does not consider the number of arguments.

### Function Precedence Order

The function precedence order determines the precedence of one function over another based on the type of function and its location on the MATLAB path. MATLAB selects the correct function for a given context by applying the following function precedence rules, in the order given:

#### 1 Subfunctions

Subfunctions take precedence over all other M-file functions and overloaded methods that are on the path and have the same name. Even if the function is called with an argument of type matching that of an overloaded method, MATLAB uses the subfunction and ignores the overloaded method.

#### 2 Private functions

Private functions are called if there is no subfunction of the same name within the current scope. As with subfunctions, even if the function is called with an argument of type matching that of an overloaded method, MATLAB uses the private function and ignores the overloaded method.

### 3 Class constructor functions

Constructor functions (functions having names that are the same as the @ directory, for example @polynom/polynom.m) take precedence over other MATLAB functions. Therefore, if you create an M-file called polynom.m and put it on your path before the constructor @polynom/polynom.m version, MATLAB will always call the constructor version.

### 4 Overloaded methods

MATLAB calls an overloaded method if it is not masked by a subfunction or private function.

### 5 Current directory

A function in the current working directory is selected before one elsewhere on the path.

### 6 Elsewhere on path

Finally, a function anywhere else on the path is selected.

## Selecting Methods from Multiple Directories

There may be a number of directories on the path that contain methods with the same name. MATLAB stops searching when it finds the first implementation of the method on the path, regardless of the implementation type (MEX-file, P-code, M-file).

## Selecting Methods from Multiple Implementation Types

There are five file precedence types. MATLAB uses file precedence to select between identically named functions in the same directory. The order of precedence for file types is

- 1 Built-in (.bi) file
- 2 MEX-files
- 3 MDL (Simulink® model) file
- 4 P-code file

## 5 M-file

For example, if MATLAB finds a P-code and an M-file version of a method in a class directory, then the P-code version is used. It is, therefore, important to regenerate the P-code version whenever you edit the M-file.

### Querying Which Method MATLAB Will Call

You can determine which method MATLAB will call using the `which` command. For example,

```
which pie3
your_matlab_path/toolbox/matlab/specgraph/pie3.m
```

However, if `p` is a portfolio object,

```
which pie3(p)
dir_on_your_path/@portfolio/pie3.m      % portfolio method
```

The `which` command determines which version of `pie3` MATLAB will call if you passed a portfolio object as the input argument. To see a list of all versions of a particular function that are on your MATLAB path, use the `-all` option. See the `which` reference page for more information on this command.

# Scheduling Program Execution with Timers

---

MATLAB includes a timer object that you can use to schedule the execution of MATLAB commands. This chapter describes how you can create timer objects, start a timer running, and specify the processing that you want performed when a timer fires. A timer is said to *fire* when the amount of time specified by the timer object elapses and the timer object executes the commands you specify.

Using a MATLAB Timer Object (p. 9-2)	Step-by-step procedure for using a timer object with a simple example
Creating Timer Objects (p. 9-4)	Using the <code>timer</code> function to create a timer object
Working with Timer Object Properties (p. 9-6)	Setting timer object properties and retrieving the values of timer object properties
Starting and Stopping Timers (p. 9-9)	Using the <code>start</code> or <code>startat</code> function to start timer objects; using the <code>stop</code> function to stop them, and blocking the command line
Creating and Executing Callback Functions (p. 9-13)	Creating a callback function and specifying it as the value of a timer object callback property
Timer Object Execution Modes (p. 9-18)	Using the <code>ExecutionMode</code> property to control when a timer object executes
Deleting Timer Objects from Memory (p. 9-22)	Using the <code>delete</code> function to delete a timer object
Finding All Timer Objects in Memory (p. 9-23)	Using the <code>timerfind</code> and <code>timerfindall</code> functions to determine if timer objects exist in memory

## Using a MATLAB Timer Object

To use a timer, perform these steps:

**1** Create a timer object.

You use the `timer` function to create a timer object. See “Creating Timer Objects” on page 9-4 for more information.

**2** Specify which MATLAB commands you want executed when the timer fires and control other aspects of timer object behavior.

You use timer object properties to specify this information. To learn about all the properties supported by the timer object, see “Working with Timer Object Properties” on page 9-6. (You can also set timer object properties when you create them, in step 1.)

**3** Start the timer object.

After you create the timer object, you must start it, using either the `start` or `startat` function. See “Starting and Stopping Timers” on page 9-9 for more information.

**4** Delete the timer object when you are done with it.

After you are finished using a timer object, you should delete it from memory. See “Deleting Timer Objects from Memory” on page 9-22 for more information.

---

**Note** The specified execution time and the actual execution of a timer can vary because timer objects work in the MATLAB single-threaded execution environment. The length of this time lag is dependent on what other processing MATLAB is performing. To force the execution of the callback functions in the event queue, include a call to the `drawnow` function in your code. The `drawnow` function flushes the event queue. For more information about callback processing, see “The Event Queue” in the MATLAB Programming GUIs documentation.

---

## Example: Displaying a Message

The following example sets up a timer object that executes a MATLAB command string after 10 seconds elapse. The example creates a timer object, specifying the values of two timer object properties, `TimerFcn` and `StartDelay`. `TimerFcn` specifies the timer callback function. This is the MATLAB command string or M-file that you want to execute when the timer fires. In the example, the timer callback function sets the value of the MATLAB workspace variable `stat` and executes the MATLAB `disp` command. The `StartDelay` property specifies how much time elapses before the timer fires.

After creating the timer object, the example uses the `start` function to start the timer object. (The additional commands in this example are included to illustrate the timer but are not required for timer operation.)

```
t = timer('TimerFcn', 'stat=false; disp(''Timer!'')',...
          'StartDelay',10);
start(t)

stat=true;
while(stat==true)
    disp('.')
    pause(1)
end
```

When you execute this code, it produces this output:

```
.
.
.
.
.
.
.
.
.
Timer!
```

```
delete(t) % Always delete timer objects after using them.
```

## Creating Timer Objects

To use a timer in MATLAB, you must create a timer object. The timer object represents the timer in MATLAB, supporting various properties and functions that control its behavior.

To create a timer object, use the `timer` function. This creates a valid timer object with default values for most properties. The following shows an example of the default timer object and its summary display:

```
t = timer

Timer Object: timer-1

Timer Settings
  ExecutionMode: singleShot
        Period: 1
  BusyMode: drop
  Running: off

Callbacks
  TimerFcn: ''
  ErrorFcn: ''
  StartFcn: ''
  StopFcn: ''
```

MATLAB names the timer object `timer-1`. (See “Timer Object Naming” on page 9-5 for more information.)

To specify the value of timer object properties after you create it, you can use the `set` function. This example sets the value of the `TimerFcn` property and the `StartDelay` property. For more information about timer object properties, see “Working with Timer Object Properties” on page 9-6.

```
set(t, 'TimerFcn', 'disp(''Hello World!'')', 'StartDelay', 5)
```

You can also set timer object properties when you create the timer object by specifying property name and value pairs as arguments to the `timer` function. The following example sets the same properties at object creation time:

```
t = timer('TimerFcn', 'disp(''Hello World!'')', 'StartDelay', 5);
```

Always delete timer objects when you are done using them. See “Deleting Timer Objects from Memory” on page 9-22 for more information.

## Timer Object Naming

MATLAB assigns a name to each timer object you create. This name has the form 'timer-*i*', where *i* is a number representing the total number of timer objects created this session.

For example, the first time you call the timer function to create a timer object, MATLAB names the object timer-1. If you call the timer function again to create another timer object, MATLAB names the object timer-2.

MATLAB keeps incrementing the number associated with each timer object it creates, even if you delete the timer objects you already created. For example, if you delete the first two timer objects and create a new object, MATLAB names it timer-3, even though the other two timer objects no longer exist in memory. To reset the numeric part of timer object names to 1, execute the `clear classes` command.

## Working with Timer Object Properties

The timer object supports many properties that provide information about the current state of the timer object and control aspects of its functioning. The following sections describe

- “Retrieving the Value of Timer Object Properties” on page 9-6
- “Setting the Value of Timer Object Properties” on page 9-7

To get information about timer object properties, see the `timer` function reference page.

### Retrieving the Value of Timer Object Properties

To retrieve the value of a timer object property, you can use the `get` function or use subscripts (dot notation) to access the field.

The following example uses the `set` function to retrieve the value of the `ExecutionMode` property:

```
t = timer;  
  
tmode = get(t, 'ExecutionMode')  
  
tmode =  
  
singleShot
```

The following example uses dot notation to retrieve the value of the `ExecutionMode` property:

```
tmode = t.ExecutionMode  
  
tmode =  
  
singleShot
```

To view a list of all the properties of a timer object, use the `get` function, specifying the timer object as the only argument:

```
get(t)
  AveragePeriod: NaN
    BusyMode: 'drop'
    ErrorFcn: ''
  ExecutionMode: 'singleShot'
  InstantPeriod: NaN
    Name: 'timer-4'
ObjectVisibility: 'on'
  Period: 1
  Running: 'off'
  StartDelay: 0
  StartFcn: ''
  StopFcn: ''
  Tag: ''
  TasksExecuted: 0
  TasksToExecute: Inf
  TimerFcn: ''
  Type: 'timer'
  UserData: []
```

## Setting the Value of Timer Object Properties

To set the value of a timer object property, use the `set` function or subscripted assignment (dot notation). You can also set timer object properties when you create the timer object. For more information, see “Creating Timer Objects” on page 9-4.

The following example uses both methods to assign values to timer object properties. The example creates a timer that, once started, displays a message every second until you stop it with the `stop` command.

- 1 Create a timer object.

```
t = timer;
```

- 2 Assign values to timer object properties using the `set` function.

```
set(t, 'ExecutionMode', 'fixedRate', 'BusyMode', 'drop', 'Period', 1);
```

- 3 Assign a value to the timer object TimerFcn property using dot notation.

```
t.TimerFcn = 'disp(''Processing...'').'
```

- 4 Start the timer object. It displays a message at 1-second intervals.

```
start(t)
```

- 5 Stop the timer object.

```
stop(t)
```

- 6 Delete timer objects after you are done using them.

```
delete(t)
```

### Viewing a List of All Settable Properties

To view a list of all timer object properties that can have values assigned to them (in contrast to the read-only properties), use the `set` function, specifying the timer object as the only argument.

The display includes the values you can use to set the property if, like the `BusyMode` property, the property accepts an enumerated list of values.

```
t = timer;  
  
set(t)  
BusyMode: [ {drop} | queue | error ]  
ErrorFcn: string -or- function handle -or- cell array  
ExecutionMode: [{singleShot} | fixedSpacing | fixedDelay | fixedRate]  
Name  
ObjectVisibility: [ {on} | off ]  
Period  
StartDelay  
StartFcn: string -or- function handle -or- cell array  
StopFcn: string -or- function handle -or- cell array  
Tag  
TasksToExecute  
TimerFcn: string -or- function handle -or- cell array  
UserData
```

## Starting and Stopping Timers

After you create a timer object, you can start the timer by calling either the `start` or `startat` function. This section describes

- “Starting a Timer” on page 9-9
- “Starting a Timer at a Specified Time” on page 9-10
- “Stopping Timer Objects” on page 9-10
- “Blocking the MATLAB Command Line” on page 9-11

---

**Note** Because the timer works within the MATLAB single-threaded environment, it cannot guarantee execution times or execution rates.

---

### Starting a Timer

To start a timer object, call the `start` function, specifying the timer object as the only argument. The `start` function starts a timer object running; the amount of time the timer runs is specified in seconds in the `StartDelay` property.

This example creates a timer object that displays a greeting after 5 seconds elapse.

- 1 Create a timer object, specifying values for timer object properties.

```
t = timer('TimerFcn','disp(''Hello World!'')','StartDelay', 5);
```

- 2 Start the timer object.

```
start(t)
```

- 3 Delete the timer object after you are finished using it.

```
delete(t);
```

## Starting a Timer at a Specified Time

To start a timer object and specify when you want the timer to fire as a date and time, rather than as the number of seconds to elapse, use the `startat` function.

The `startat` function starts a timer object running and lets you specify the date, hour, minute, and second when you want the timer to execute. You specify the time as a MATLAB serial date number or as a specially formatted date text string. See the `startat` function reference page for details.

This example creates a timer object that displays a message after an hour has elapsed. The `startat` function starts the timer object running and calculates the value of the `StartDelay` property based on the time you specify.

```
t2=timer('TimerFcn','disp(''It has been an hour now.'')');  
  
startat(t2,now+1/24);
```

## Stopping Timer Objects

Once started, the timer object stops running if one of the following conditions apply:

- The timer function callback (`TimerFcn`) has been executed the number of times specified in the `TasksToExecute` property.
- An error occurred while executing a timer function callback (`TimerFcn`).

You can also stop a timer object by using the `stop` function, specifying the timer object as the only argument. The following example illustrates stopping a timer object:

**1** Create a timer object.

```
t = timer('TimerFcn','disp(''Hello World!'')','StartDelay', 100);
```

**2** Start it running.

```
start(t)
```

- 3 Check the state of the timer object after starting it.

```
get(t, 'Running')
```

```
ans =
```

```
on
```

- 4 Stop the timer using the stop command and check the state again. When a timer stops, the value of the Running property of the timer object is set to 'off'.

```
stop(t)
```

```
get(t, 'Running')
```

```
ans =
```

```
off
```

- 5 Delete the timer object when you are finished using it.

```
delete(t)
```

---

**Note** The timer object can execute a callback function that you specify when it starts or stops. See “Creating and Executing Callback Functions” on page 9-13.

---

## Blocking the MATLAB Command Line

By default, when you use the start or startat function to start a timer object, the function returns control to the command line immediately. For some applications, you might prefer to block the command line until the timer fires. To do this, call the wait function right after calling the start or startat function.

- 1 Create a timer object.

```
t = timer('StartDelay', 5, 'TimerFcn', 'disp(''Hello World!'')');
```

- 2** Start the timer object running.

```
start(t)
```

- 3** After the start function returns, call the wait function immediately. The wait function blocks the command line until the timer object fires.

```
wait(t)
```

- 4** Delete the timer object after you are finished using it.

```
delete(t)
```

## Creating and Executing Callback Functions

This section describes

- “Associating Commands with Timer Object Events” on page 9-13
- “Creating Callback Functions” on page 9-14
- “Specifying the Value of Callback Function Properties” on page 9-16

---

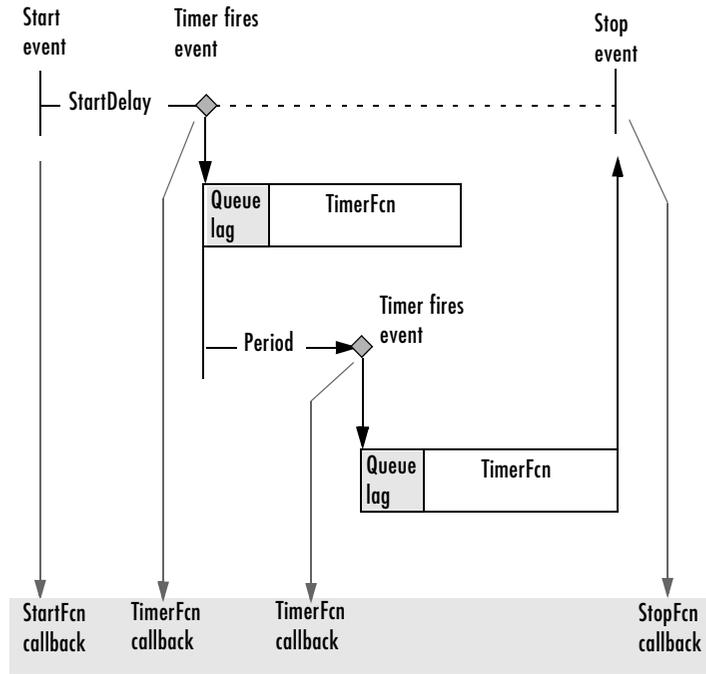
**Note** Callback function execution might be delayed if the callback involves a CPU-intensive task such as updating a figure.

---

### Associating Commands with Timer Object Events

The timer object supports properties that let you specify the MATLAB commands that execute when a timer fires, and for other timer object events, such as starting, stopping, or when an error occurs. These are called *callbacks*. To associate MATLAB commands with a timer object event, set the value of the associated timer object callback property.

The following diagram shows when the events occur during execution of a timer object and give the names of the timer object properties associated with each event. For example, to associate MATLAB commands with a start event, assign a value to the `StartFcn` callback property. Error callbacks can occur at any time.



### Timer Object Events and Related Callback Function

## Creating Callback Functions

When the time period specified by a timer object elapses, the timer object executes one or more MATLAB functions of your choosing. You can specify the functions directly as the value of the `TimerFcn` callback property. You can also put the commands in an M-file function and specify the M-file function as the value of the `callback` property.

### Specifying Callback Functions Directly

This example creates a timer object that displays a greeting after 5 seconds. The example specifies the value of the `TimerFcn` callback property directly, putting the commands in a text string.

```
t = timer('TimerFcn','disp(''Hello World!''),'StartDelay',5);
```

---

**Note** When you specify the callback commands directly as the value of the callback function property, the commands are evaluated in the MATLAB workspace.

---

### Putting Commands in a Callback Function

Instead of specifying MATLAB commands directly as the value of a callback property, you can put the commands in an M-file and specify the M-file as the value of the callback property.

When you create a callback function, the first two arguments must be a handle to the timer object and an event structure. An event structure contains two fields: Type and Data. The Type field contains a text string that identifies the type of event that caused the callback. The value of this field can be any of the following strings: 'StartFcn', 'StopFcn', 'TimerFcn', or 'ErrorFcn'. The Data field contains the time the event occurred.

In addition to these two required input arguments, your callback function can accept application-specific arguments. To receive these input arguments, you must use a cell array when specifying the name of the function as the value of a callback property. For more information, see “Specifying the Value of Callback Function Properties” on page 9-16.

### Example: Writing a Callback Function

This example implements a simple callback function that displays the type of event that triggered the callback and the time the callback occurred. To illustrate passing application-specific arguments, the example callback function accepts as an additional argument a text string and includes this text string in the display output. To see this function used with a callback property, see “Specifying the Value of Callback Function Properties” on page 9-16.

```
function my_callback_fcn(obj, event, string_arg)

txt1 = ' event occurred at ';
txt2 = string_arg;

event_type = event.Type;
event_time = datestr(event.Data.time);
```

```
msg = [event_type txt1 event_time];
disp(msg)
disp(txt2)
```

## Specifying the Value of Callback Function Properties

You associate a callback function with a specific event by setting the value of the appropriate callback property. You can specify the callback function as a text string, cell array, or function handle. To access the object and event arguments, you must specify the function as a cell array or as a function handle. If your callback function accepts additional arguments, you must use a cell array.

The following table shows the syntax for several sample callback functions and describes how you call them.

Callback Function Syntax	How to Specify as a Property Value
function myfile	set(h, 'StartFcn', 'myfile')
function myfile(obj, event)	set(h, 'StartFcn', @myfile)
function myfile(obj, event, arg1, arg2)	set(h, 'StartFcn', {'myfile', 5, 6})
function myfile(obj, event, arg1, arg2)	set(h, 'StartFcn', {@myfile, 5, 6})

This example illustrates several ways you can specify the value of timer object callback function properties, some with arguments and some without. To see the code of the callback function, `my_callback_fcn`, see “Example: Writing a Callback Function” on page 9-15.

**1** Create a timer object.

```
t = timer('StartDelay', 4, 'Period', 4, 'TasksToExecute', 2, ...
         'ExecutionMode', 'fixedRate');
```

**2** Specify the value of the `StartFcn` callback. Note that the example specifies the value in a cell array because the callback function needs to access arguments passed to it.

```
t.StartFcn = {'my_callback_fcn', 'My start message'};
```

- 3 Specify the value of the `StopFcn` callback. The example specifies the callback function by its handle, rather than as a text string. Again, the value is specified in a cell array because the callback function needs to access the arguments passed to it.

```
t.StopFcn = { @my_callback_fcn, 'My stop message'};
```

- 4 Specify the value of the `TimerFcn` callback. The example specifies the MATLAB commands in a text string.

```
t.TimerFcn = 'disp(''Hello World!'')';
```

- 5 Start the timer object.

```
start(t)
```

The example outputs the following.

```
StartFcn event occurred at 10-Mar-2004 17:16:59
Start message
Hello World!
Hello World!
StopFcn event occurred at 10-Mar-2004 17:16:59
Stop message
```

- 6 Delete the timer object after you are finished with it.

```
delete(t)
```

## Timer Object Execution Modes

The timer object supports several execution modes that determine how it schedules the timer callback function (`TimerFcn`) for execution. You specify the execution mode by setting the value of the `ExecutionMode` property. This section describes

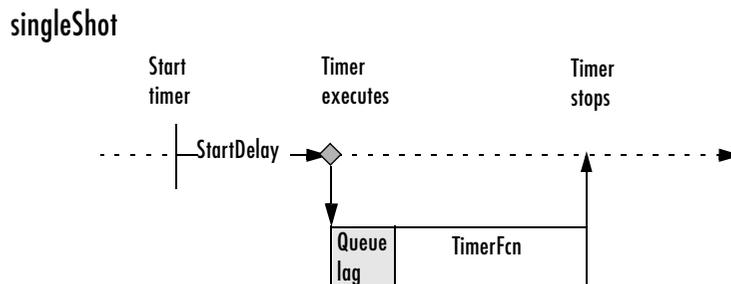
- “Executing a Timer Callback Function Once” on page 9-18
- “Executing a Timer Callback Function Multiple Times” on page 9-19
- “Handling Callback Function Queuing Conflicts” on page 9-20

### Executing a Timer Callback Function Once

To execute a timer callback function once, set the `ExecutionMode` property to `'singleShot'`. This is the default execution mode.

In this mode, the timer object starts the timer and, after the time period specified in the `StartDelay` property elapses, adds the timer callback function (`TimerFcn`) to the MATLAB execution queue. When the timer callback function finishes, the timer stops.

The following figure graphically illustrates the parts of timer callback execution for a `singleShot` execution mode. The shaded area in the figure, labelled `queue lag`, represents the indeterminate amount of time between when the timer adds a timer callback function to the MATLAB execution queue and when the function starts executing. The duration of this lag is dependent on what other processing MATLAB happens to be doing at the time.



**Timer Callback Execution (singleShot Execution Mode)**

## Executing a Timer Callback Function Multiple Times

The timer object supports three multiple-execution modes:

- 'fixedRate'
- 'fixedDelay'
- 'fixedSpacing'

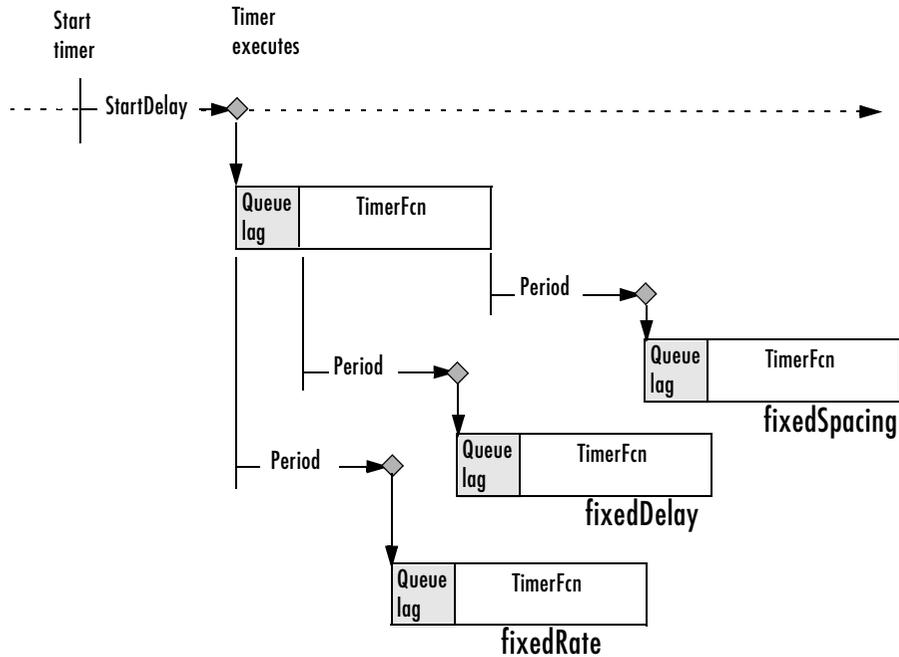
In many ways, these execution modes operate the same:

- The `TasksToExecute` property specifies the number of times you want the timer to execute the timer callback function (`TimerFcn`).
- The `Period` property specifies the amount of time between executions of the timer callback function.
- The `BusyMode` property specifies how the timer object handles queuing of the timer callback function when the previous execution of the callback function has not completed. See “Handling Callback Function Queuing Conflicts” on page 9-20 for more information.

The execution modes differ only in where they start measuring the time period between executions. The following table describes these differences.

Execution Mode	Description
'fixedRate'	Time period between executions begins immediately after the timer callback function is added to the MATLAB execution queue.
'fixedDelay'	Time period between executions begins when the timer function callback actually starts executing, after any time lag due to delays in the MATLAB execution queue.
'fixedSpacing'	Time period between executions begins when the timer callback function finishes executing.

The following figure illustrates the difference between these modes. Note that the amount of time between executions (specified by the `Period` property) remains the same. Only the point at which execution begins is different.



### Differences Between Execution Modes

## Handling Callback Function Queuing Conflicts

At busy times, in multiple-execution scenarios, the timer may need to add the timer callback function (`TimerFcn`) to the MATLAB execution queue before the previously queued execution of the callback function has completed. You can determine how the timer object handles this scenario by using the `BusyMode` property.

If you specify `'drop'` as the value of the `BusyMode` property, the timer object skips the execution of the timer function callback if the previously scheduled callback function has not already completed.

If you specify 'queue', the timer object waits until the currently executing callback function finishes before queuing the next execution of the timer callback function.

---

**Note** In 'queue' mode, the timer object tries to make the average time between executions equal the amount of time specified in the `Period` property. If the timer object has to wait longer than the time specified in the `Period` property between executions of the timer function callback, it shortens the time period for subsequent executions to make up the time.

---

If the `BusyMode` property is set to 'error', the timer object stops and executes the timer object error callback function (`ErrorFcn`), if one is specified.

## Deleting Timer Objects from Memory

When you are finished with a timer object, delete it from memory using the `delete` function:

```
delete(t)
```

When you delete a timer object, workspace variables that referenced the object remain. Deleted timer objects are invalid and cannot be reused. Use the `clear` command to remove workspace variables that reference deleted timer objects.

## Testing the Validity of a Timer Object

To test if a timer object has been deleted, use the `isvalid` function. The `isvalid` function returns logical 0 (false) for deleted timer objects:

```
isvalid(t)
ans =

    0
```

## Deleting All Existing Timer Objects

To remove all timer objects from memory, enter

```
delete(timerfind)
```

For information about the `timerfind` function, see “Finding All Timer Objects in Memory” on page 9-23.

## Finding All Timer Objects in Memory

To find all the timer objects that exist in memory, use the `timerfind` function. This function returns an array of timer objects. If you leave off the semicolon, and there are multiple timer objects in the array, `timerfind` displays summary information in a table:

```
t1 = timer;
t2 = timer;
t3 = timer;
t_array = timerfind
```

Timer Object Array

Index:	ExecutionMode:	Period:	TimerFcn:	Name:
1	singleShot	1	''	timer-3
2	singleShot	1	''	timer-4
3	singleShot	1	''	timer-5

Using `timerfind` to determine all the timer objects that exist in memory can be helpful when deleting timer objects. See “Deleting All Existing Timer Objects” on page 9-22 for an example.

### Finding Invisible Timer Objects

If you set the value of a timer object’s `ObjectVisibility` property to `'off'`, the timer object does not appear in listings of existing timer objects returned by `timerfind`. The `ObjectVisibility` property provides a way for application developers to prevent end-user access to the timer objects created by their application.

Objects that are not visible are still valid. If you have access to the object (for example, from within the M-file that created it), you can set its properties. To retrieve a list of all the timer objects in memory, including invisible ones, use the `timerfindall` function.



# Improving Performance and Memory Usage

---

Analyzing Your Program's Performance (p. 10-2)	What tools are provided by MATLAB to measure the performance of your programs and identify where the bottlenecks are
Techniques for Improving Performance (p. 10-4)	How to improve M-file performance by vectorizing loops, preallocating arrays, etc.
Making Efficient Use of Memory (p. 10-12)	Conserving memory; platform-specific memory handling; "Out of Memory" errors
Resolving "Out of Memory" Errors (p. 10-16)	What to do when if you get a "Out of Memory" error

## Analyzing Your Program's Performance

The M-file Profiler graphical user interface and the stopwatch timer functions enable you to get back information on how your program is performing and help you identify areas that need improvement. The Profiler can be more useful in measuring relative execution time and in identifying specific performance bottlenecks in your code, while the stopwatch functions tend to be more useful for providing absolute time measurements.

### The M-File Profiler Utility

A good first step to speeding up your programs is to find out where the bottlenecks are. This is where you need to concentrate your attention to optimize your code.

MATLAB provides the M-file Profiler, a graphical user interface that shows you where your program is spending its time during execution. Use the Profiler to help you determine where you can modify your code to make performance improvements.

To start the Profiler, type `profile viewer` or select **Desktop -> Profiler** in the MATLAB Command Window. See “Profiling for Improving Performance” in the MATLAB Desktop Tools and Development Environment documentation, and the `profile` function reference page.

### Stopwatch Timer Functions

If you just need to get an idea of how long your program (or a portion of it) takes to run, or to compare the speed of different implementations of a program, you can use the stopwatch timer functions, `tic` and `toc`. Invoking `tic` starts the timer, and the first subsequent `toc` stops it and reports the time elapsed between the two.

Use `tic` and `toc` as shown here:

```
tic
    run the program section to be timed
toc
```

Keep in mind that `tic` and `toc` measure overall elapsed time. Make sure that no other applications are running in the background on your system that could affect the timing of your MATLAB programs.

### Measuring Smaller Programs

Shorter programs sometimes run too fast to get useful data from `tic` and `toc`. When this is the case, try measuring the program running repeatedly in a loop, and then average to find the time for a single run:

```
tic
  for k = 1:100
    run the program
  end
toc
```

## Techniques for Improving Performance

This section covers the following suggestions on how you can improve the performance of your MATLAB programs:

- “Vectorizing Loops” on page 10-4
- “Preallocating Arrays” on page 10-7
- “Coding Loops in a MEX-File” on page 10-9
- “Assigning to Variables” on page 10-9
- “Operating on Real Data” on page 10-10
- “Using Appropriate Logical Operators” on page 10-10
- “Overloading Built-In Functions” on page 10-11
- “Functions Are Generally Faster Than Scripts” on page 10-11
- “Load and Save Are Faster Than File I/O Functions” on page 10-11
- “Avoid Large Background Processes” on page 10-11

### Vectorizing Loops

MATLAB is a matrix language, which means it is designed for vector and matrix operations. You can often speed up your M-file code by using vectorizing algorithms that take advantage of this design. *Vectorization* means converting `for` and `while` loops to equivalent vector or matrix operations.

#### Simple Example of Vectorizing

Here is one way to compute the sine of 1001 values ranging from 0 to 10:

```
i = 0;
for t = 0:.01:10
    i = i + 1;
    y(i) = sin(t);
end
```

A vectorized version of the same code is

```
t = 0:.01:10;
y = sin(t);
```

The second example executes much faster than the first and is the way MATLAB is meant to be used. Test this on your system by creating M-file scripts that contain the code shown, and then using the `tic` and `toc` functions to time the M-files.

### Advanced Example of Vectorizing

`repmat` is an example of a function that takes advantage of vectorization. It accepts three input arguments: an array `A`, a row dimension `M`, and a column dimension `N`.

`repmat` creates an output array that contains the elements of array `A`, replicated and “tiled” in an `M`-by-`N` arrangement:

```
A = [1 2 3; 4 5 6];

B = repmat(A,2,3);
B =
     1     2     3     1     2     3     1     2     3
     4     5     6     4     5     6     4     5     6
     1     2     3     1     2     3     1     2     3
     4     5     6     4     5     6     4     5     6
```

`repmat` uses vectorization to create the indices that place elements in the output array:

```
function B = repmat(A, M, N)

% Step 1 Get row and column sizes
[m,n] = size(A);

% Step 2 Generate vectors of indices from 1 to row/column size
mind = (1:m)';
nind = (1:n)';

% Step 3 Create index matrices from vectors above
mind = mind(:,ones(1, M));
nind = nind(:,ones(1, N));
```

```
% Step 4 Create output array  
B = A(mind,nind);
```

**Step 1**, above, obtains the row and column sizes of the input array.

**Step 2** creates two column vectors. `mind` contains the integers from 1 through the row size of A. The `nind` variable contains the integers from 1 through the column size of A.

**Step 3** uses a MATLAB vectorization trick to replicate a single column of data through any number of columns. The code is

```
B = A(:,ones(1,nCols))
```

where `nCols` is the desired number of columns in the resulting matrix.

**Step 4** uses array indexing to create the output array. Each element of the row index array, `mind`, is paired with each element of the column index array, `nind`, using the following procedure:

- 1 The first element of `mind`, the row index, is paired with each element of `nind`. MATLAB moves through the `nind` matrix in a columnwise fashion, so `mind(1,1)` goes with `nind(1,1)`, and then `nind(2,1)`, and so on. The result fills the first row of the output array.
- 2 Moving columnwise through `mind`, each element is paired with the elements of `nind` as above. Each complete pass through the `nind` matrix fills one row of the output array.

### Functions Used in Vectorizing

Some of the most commonly used functions for vectorizing are as follows

Function	Description
<code>all</code>	Test to determine if all elements are nonzero
<code>any</code>	Test for any nonzeros
<code>cumsum</code>	Find cumulative sum
<code>diff</code>	Find differences and approximate derivatives

<b>Function</b>	<b>Description</b>
<code>find</code>	Find indices and values of nonzero elements
<code>ind2sub</code>	Convert from linear index to subscripts
<code>ipermute</code>	Inverse permute dimensions of a multidimensional array
<code>logical</code>	Convert numeric values to logical
<code>ndgrid</code>	Generate arrays for multidimensional functions and interpolation
<code>permute</code>	Rearrange dimensions of a multidimensional array
<code>prod</code>	Find product of array elements
<code>repmat</code>	Replicate and tile an array
<code>reshape</code>	Change the shape of an array
<code>shiftdim</code>	Shift array dimensions
<code>sort</code>	Sort array elements in ascending or descending order
<code>squeeze</code>	Remove singleton dimensions from an array
<code>sub2ind</code>	Convert from subscripts to linear index
<code>sum</code>	Find the sum of array elements

## Preallocating Arrays

`for` and `while` loops that incrementally increase, or *grow*, the size of a data structure each time through the loop can adversely affect performance and memory use. Repeatedly resizing arrays often requires that MATLAB spend extra time looking for larger contiguous blocks of memory and then moving the array into those blocks. You can often improve on code execution time by preallocating the maximum amount of space that would be required for the array ahead of time.

The following code creates a scalar variable `x`, and then gradually increases the size of `x` in a `for` loop instead of preallocating the required amount of memory at the start:

```
x = 0;
for k = 2:1000
    x(k) = x(k-1) + 5;
end
```

Change the first line to preallocate a 1-by-1000 block of memory for `x` initialized to zero. This time there is no need to repeatedly reallocate memory and move data as more values are assigned to `x` in the loop:

```
x = zeros(1, 1000);
for k = 2:1000
    x(k) = x(k-1) + 5;
end
```

### Preallocation Functions

Preallocation makes it unnecessary for MATLAB to resize an array each time you enlarge it. Use the appropriate preallocation function for the kind of array you are working with.

Array Type	Function	Examples
Numeric	<code>zeros</code>	<code>y = zeros(1, 100);</code>
Cell	<code>cell</code>	<code>B = cell(2, 3);</code> <code>B{1,3} = 1:3;</code> <code>B{2,2} = 'string';</code>

### Preallocating a Nondouble Matrix

When you preallocate a block of memory to hold a matrix of some type other than double, avoid using the method

```
A = int8(zeros(100));
```

This statement preallocates a 100-by-100 matrix of `int8` first by creating a full matrix of doubles, and then converting each element to `int8`. This costs time and uses memory unnecessarily.

The next statement shows how to do this more efficiently:

```
A = zeros(100, 'int8');
```

## Coding Loops in a MEX-File

If there are instances where you cannot vectorize and must use a for or while loop, consider coding the loop in a MEX-file. In this way, the loop executes much more quickly since the instructions in the loop do not have to be interpreted each time they execute.

See “Introducing MEX-Files” in the External Interfaces documentation.

## Assigning to Variables

For best performance, keep the following suggestions in mind when assigning values to variables.

### Changing a Variable’s Data Type or Dimension

Changing the data type or array shape of an existing variable slows MATLAB down as it must take extra time to process this. When you need to store data of a different type, it is advisable to create a new variable.

This code changes the type for X from double to char, which has a negative impact on performance:

```
X = 23;
.
other code
.
X = 'A';           % X changed from type double to char
.
other code
```

### Assigning Real and Complex Numbers

Assigning a complex number to a variable that already holds a real number impacts the performance of your program. Similarly, you should not assign a real value to a variable that already holds a complex value.

## Operating on Real Data

When operating on real (i.e., noncomplex) numbers, it is more efficient to use MATLAB functions that have been designed specifically for real numbers. The following functions return numeric values that are real.

Function	Description
reallog	Find natural logarithm for nonnegative real arrays
realpow	Find array power for real-only output
realsqrt	Find square root for nonnegative real arrays

## Using Appropriate Logical Operators

When performing a logical AND or OR operation, you have a choice of two operators of each type.

Operator	Description
&,	Perform logical AND and OR on arrays element by element
&&,	Perform logical AND and OR on scalar values with short-circuiting

In `if` and `while` statements, it is more efficient to use the short-circuiting operators, `&&` for logical AND and `||` for logical OR. This is because these operators often don't have to evaluate the entire logical expression. For example, MATLAB evaluates only the first part of this expression whenever the number of input arguments is less than three:

```
if (nargin >= 3) && (ischar(varargin{3}))
```

See “Short-Circuit Operators” in the MATLAB documentation for a discussion on short-circuiting with `&&` and `||`.

## **Overloading Built-In Functions**

Overloading MATLAB built-in functions on any of the standard MATLAB data types can negatively affect performance. For example, if you overload the `plus` function to handle any of the integer data types differently, you may hinder certain optimizations in the MATLAB built-in function code for `plus`, and thus may slow down any programs that make use of this overload.

## **Functions Are Generally Faster Than Scripts**

Your code executes more quickly if it is implemented in a function rather than a script.

## **Load and Save Are Faster Than File I/O Functions**

If you have a choice of whether to use `load` and `save` instead of the low-level MATLAB file I/O routines such as `fread` and `fwrite`, choose the former. `load` and `save` have been optimized to run faster and reduce memory fragmentation.

## **Avoid Large Background Processes**

Avoid running large processes in the background at the same time you are executing your program in MATLAB. This frees more CPU time for your MATLAB session.

## Making Efficient Use of Memory

This section discusses how to conserve memory and improve memory use. Topics include

- “Memory Management Functions” on page 10-12
- “Preallocating Arrays to Reduce Fragmentation” on page 10-13
- “Enlarging Arrays with `repmat`” on page 10-13
- “Working with Variables” on page 10-13
- “Converting Full Matrices into Sparse” on page 10-15
- “Structure of Arrays vs. Array of Structures” on page 10-15
- “Working with Large Amounts of Data” on page 10-15

For more information on memory management, see Technical Note 1106: “Memory Management Guide” at the following URL:

<http://www.mathworks.com/support/tech-notes/1100/1106.shtml>

### Memory Management Functions

The following functions can help you to manage memory use in MATLAB:

- `whos` shows how much memory has been allocated for variables in the workspace.
- `pack` saves existing variables to disk, and then reloads them contiguously. This reduces the chances of running into problems due to memory fragmentation.  
See “Compressing Data in Memory” on page 10-16.
- `clear` removes variables from memory. One way to increase the amount of available memory is to periodically clear variables from memory that you no longer need.
- `save` selectively stores variables to the disk. This is a useful technique when you are working with large amounts of data. Save data to the disk periodically, and then use the `clear` function to remove the saved data from memory.
- `load` reloads a data file saved with the `save` function.

- `quit` exits MATLAB and returns all allocated memory to the system. This can be useful on UNIX systems as UNIX does not free up memory allocated to an application (e.g., MATLAB) until the application exits.

## Preallocating Arrays to Reduce Fragmentation

In the course of a MATLAB session, memory can become fragmented due to dynamic memory allocation and deallocation. `for` and `while` loops that incrementally increase, or *grow*, the size of a data structure each time through the loop can add to this fragmentation as they have to repeatedly find and allocate larger blocks of memory to store the data.

To make more efficient use of your memory, preallocate a block of memory large enough to hold the matrix at its final size before entering the loop. When you preallocate memory for a potentially large array, MATLAB “grabs” sufficient contiguous space for the data at the beginning of the computation. Once you have this space, you can add elements to the array without having to continually allocate new space for it in memory.

For more information on preallocation, see “Preallocating Arrays” on page 10-7.

## Enlarging Arrays with `repmat`

In cases where you cannot preallocate, see if you can increase the size of your array using the `repmat` function. `repmat` tries to allocate a contiguous block of memory for your expanding array.

## Working with Variables

To conserve memory when creating variables,

- Allocate your larger matrices first, as explained in “Allocating Large Matrices Earlier” on page 10-14.
- Avoid creating large temporary variables, and clear temporary variables when they are no longer needed.
- When working with arrays of fixed size, preallocate them rather than having MATLAB resize the array each time you enlarge it.
- Set variables equal to the empty matrix `[]` to free memory, or clear the variables using the `clear` function.

- Reuse variables as much as possible, but keeping in mind the guidelines stated in “Assigning to Variables” on page 10-9.

## Allocating Large Matrices Earlier

MATLAB uses a heap method of memory management. It requests memory from the operating system when there is not enough memory available in the MATLAB heap to store the current variables. It reuses memory as long as the size of the memory segment required is available in the MATLAB heap.

For example, on one machine these statements use approximately 15.4 MB of RAM:

```
a = rand(1e6,1);  
b = rand(1e6,1);
```

This statement uses approximately 16.4 MB of RAM:

```
c = rand(2.1e6,1);
```

These statements use approximately 32.4 MB of RAM. This is because MATLAB is not able to fit a 2.1 MB array in the space previously occupied by two 1-MB arrays:

```
a = rand(1e6,1);  
b = rand(1e6,1);  
clear  
c = rand(2.1e6,1);
```

The simplest way to prevent overallocation of memory is to allocate the largest vectors first. These statements use only about 16.4 MB of RAM:

```
c = rand(2.1e6,1);  
clear  
a = rand(1e6,1);  
b = rand(1e6,1);
```

## Clearing Unused Variables from Memory

If you use `pack` and there is still not enough free memory to proceed, you probably need to remove some of the variables you are no longer using from memory. Use `clear` to do this.

## Converting Full Matrices into Sparse

Matrices with values that are mostly zero are best stored in sparse format. Sparse matrices can use less memory and may also be faster to manipulate than full matrices. You can convert a full matrix to sparse using the `sparse` function.

Compare two 1000-by-1000 matrices: `X`, a matrix of doubles with 2/3 of its elements equal to zero; and `Y`, a sparse copy of `X`. As shown below, approximately half as much memory is required for the sparse matrix:

```
whos
  Name      Size      Bytes  Class
-----
  X         1000x1000  8000000  double array
  Y         1000x1000  4004000  double array (sparse)
```

## Structure of Arrays vs. Array of Structures

If your MATLAB application needs to store a large amount of data, and the definition of that data lends itself to being stored in either a structure of arrays or an array of structures, the former is preferable. A structure of arrays requires significantly less memory than an array of structures, and also has a corresponding speed benefit.

## Working with Large Amounts of Data

If your program generates very large amounts of data, consider writing the data to disk periodically. After saving that portion of the data, use the `clear` function to remove the variable from memory and continue with the data generation.

## Resolving “Out of Memory” Errors

MATLAB generates an Out of Memory message whenever it requests a segment of memory from the operating system that is larger than what is currently available. This section covers the following topics providing suggestions on how to resolve such errors:

- “General Suggestions For Reclaiming Memory” on page 10-16
- “Compressing Data in Memory” on page 10-16
- “Increasing System Swap Space” on page 10-17
- “Freeing Up System Resources on Windows Systems” on page 10-18
- “Reloading Variables on UNIX Systems” on page 10-18

### General Suggestions For Reclaiming Memory

When you see the Out of Memory message, use any of the techniques discussed under “Making Efficient Use of Memory” on page 10-12 to help optimize the available memory. If the Out of Memory message still appears, you can try any of the following:

- Compress data to reduce memory fragmentation
- If possible, break large matrices into several smaller matrices so that less memory is used at any one time.
- If possible, reduce the size of your data.
- Make sure that there are no external constraints on the memory accessible to MATLAB. (On UNIX systems, use the `limit` command to check).
- Increase the size of the swap file. We recommend that your machine be configured with twice as much swap space as you have RAM. See “Increasing System Swap Space” on page 10-17, below.
- Add more memory to the system.

### Compressing Data in Memory

Since MATLAB uses a heap method of memory management, extended MATLAB sessions may cause memory to become fragmented. When memory is fragmented, there may be plenty of free space, but not enough contiguous memory to store a new large variable. If you get the Out of Memory message

from MATLAB, the `pack` function may be able to compress some of your data in memory, thus freeing up larger contiguous blocks.

---

**Note** Because of time considerations, you should not use `pack` within loops or M-file functions.

---

## Increasing System Swap Space

How you set the swap space for your computer depends on what operating system you are running on.

### UNIX

Information about swap space can be procured by typing `psstat -s` at the UNIX command prompt. For detailed information on changing swap space, ask your system administrator.

### Linux

Swap space can be changed by using the `mkswap` and `swapon` commands. For more information on the above commands, type `man` followed by the command name at the Linux prompt.

### Windows NT

Follow the steps shown here:

- 1 Right-click the **My Computer** icon, and select **Properties**.
- 2 Select the **Performance** tab and click the **Change** button to change the amount of virtual memory.

### Windows 2000

Follow the steps shown here:

- 1 Right-click the **My Computer** icon, and select **Properties**.
- 2 Select the **Advanced** tab and click the **Performance Options** button.
- 3 Click the **Change** button to change the amount of virtual memory.

### Windows XP

Follow the steps shown here:

- 1 Right-click on the **My Computer** icon, and select **Properties**.
- 2 Select the **Advanced** tab and click the **Performance Settings** button.
- 3 Click the **Change** button to change the amount of virtual memory.

### Freeing Up System Resources on Windows Systems

There are no functions implemented to manipulate the way MATLAB handles Microsoft Windows system resources. Windows uses system resources to track fonts, windows, and screen objects. Resources can be depleted by using multiple figure windows, multiple fonts, or several UI controls. One way to free up system resources is to close all inactive windows. Windows icons still use resources.

### Reloading Variables on UNIX Systems

On UNIX systems, MATLAB does not return memory to the operating system even after variables have been cleared. This is due to the manner in which UNIX manages memory. UNIX does not accept memory back from a program until the program has terminated. So, the amount of memory used in a MATLAB session is not returned to the operating system until you exit MATLAB.

To free up the memory used in your MATLAB session, save your workspace variables, exit MATLAB, and then load your variables back in.

# Programming Tips

---

This chapter is a categorized compilation of tips for the MATLAB<sup>®</sup> programmer. Each item is relatively brief to help you to browse through them and find information that is useful. Many of the tips include a reference to specific MATLAB documentation that gives you more complete coverage of the topic. You can find information on the following topics:

Command and Function Syntax (p. 11-3)	Syntax, command shortcuts, command recall, etc.
Help (p. 11-6)	Getting help on MATLAB functions and your own
Development Environment (p. 11-10)	Useful features in the development environment
M-File Functions (p. 11-12)	M-file structure, getting information about a function
Function Arguments (p. 11-14)	Various ways to pass arguments, useful functions
Program Development (p. 11-16)	Suggestions for creating and modifying program code
Debugging (p. 11-19)	Using the debugging environment and commands
Variables (p. 11-23)	Variable names, global and persistent variables
Strings (p. 11-27)	String concatenation, string conversion, etc.
Evaluating Expressions (p. 11-30)	Use of <code>eval</code> , short-circuiting logical expressions, etc.
MATLAB Path (p. 11-32)	Precedence rules, making file changes visible to MATLAB, etc.
Program Control (p. 11-36)	Using program control statements like <code>if</code> , <code>switch</code> , <code>try</code>
Save and Load (p. 11-40)	Saving MATLAB data to a file, loading it back in
Files and Filenames (p. 11-43)	Naming M-files, passing filenames, etc.
Input/Output (p. 11-46)	Reading and writing various types of files
Starting MATLAB (p. 11-49)	Getting MATLAB to start up faster

Operating System Compatibility  
(p. 11-50)

Interacting with the operating system

Demos (p. 11-52)

Learning about the demos supplied with MATLAB

For More Information (p. 11-53)

Other valuable resources for information

For suggestions on how to improve the performance of your MATLAB programs, and how to write programs that use memory more efficiently, see Chapter 10, “Improving Performance and Memory Usage.”

## Command and Function Syntax

This section covers the following topics:

- “Syntax Help” on page 11-3
- “Command and Function Syntaxes” on page 11-3
- “Command Line Continuation” on page 11-3
- “Completing Commands Using the Tab Key” on page 11-4
- “Recalling Commands” on page 11-4
- “Clearing Commands” on page 11-5
- “Suppressing Output to the Screen” on page 11-5

### Syntax Help

For help about the general syntax of MATLAB functions and commands, type

```
help syntax
```

### Command and Function Syntaxes

You can enter MATLAB commands using either a *command* or *function* syntax. It is important to learn the restrictions and interpretation rules for both.

```
functionname arg1 arg2 arg3           % Command syntax
functionname('arg1','arg2','arg3')    % Function syntax
```

**For more information:** See “Calling Functions” in the MATLAB Programming documentation.

### Command Line Continuation

You can continue most statements to one or more additional lines by terminating each incomplete line with an ellipsis (...). Breaking down a statement into a number of lines can sometimes result in a clearer programming style.

```
sprintf ('Example %d shows a command coded on %d lines.\n', ...
        exampleNumber, ...
        numberOfLines)
```

Note that you cannot continue an incomplete string to another line.

```
disp 'This statement attempts to continue a string ...  
    to another line, resulting in an error.'
```

**For more information:** See “Entering Long Lines” in the MATLAB Desktop Tools and Development Environment documentation.

## Completing Commands Using the Tab Key

You can save some typing when entering commands by entering only the first few letters of the command, variable, property, etc. followed by the **Tab** key. Typing the second line below (with **T** representing **Tab**) yields the expanded, full command shown in the third line:

```
f = figure;  
set(f, 'papTuT', 'cT')           % Type this line.  
set(f, 'paperunits', 'centimeters') % This is what you get.
```

If there are too many matches for the string you are trying to complete, you will get no response from the first **Tab**. Press **Tab** again to see all possible choices:

```
set(f, 'paTT  
PaperOrientation  PaperPositionMode  PaperType          Parent  
PaperPosition    PaperSize          PaperUnits
```

**For more information:** See “Tab Completion” in the MATLAB Desktop Tools and Development Environment documentation

## Recalling Commands

Use any of the following methods to simplify recalling previous commands to the screen:

- To recall an earlier command to the screen, press the up arrow key one or more times, until you see the command you want. If you want to modify the recalled command, you can edit its text before pressing **Enter** or **Return** to execute it.
- To recall a specific command by name without having to scroll through your earlier commands one by one, type the starting letters of the command, followed by the up arrow key.

- Open the Command History window (**View -> Command History**) to see all previous commands. Double-click on the one you want to execute.

**For more information:** See “Recalling Previous Lines” and “Command History” in the MATLAB Desktop Tools and Development Environment documentation.

## Clearing Commands

If you have typed a command that you then decide not to execute, you can clear it from the Command Window by pressing the Escape (**Esc**) key.

## Suppressing Output to the Screen

To suppress output to the screen, end statements with a semicolon. This can be particularly useful when generating large matrices.

```
A = magic(100);    % Create matrix A, but do not display it.
```

## Help

This section covers the following topics:

- “Using the Help Browser” on page 11-6
- “Help on Functions from the Help Browser” on page 11-7
- “Help on Functions from the Command Window” on page 11-7
- “Topical Help” on page 11-7
- “Paged Output” on page 11-8
- “Writing Your Own Help” on page 11-8
- “Help for Subfunctions and Private Functions” on page 11-9
- “Help for Methods and Overloaded Functions” on page 11-9

### Using the Help Browser

Open the Help browser from the MATLAB Command Window using one of the following:

- Click on the question mark symbol in the toolbar.
- Select **Help -> MATLAB Help** from the menu.
- Type the word `doc` at the command prompt.

Some of the features of the Help browser are listed below.

<b>Feature</b>	<b>Description</b>
<b>Product Filter</b>	Establish which products to find help on.
<b>Contents</b>	Look up topics in the Table of Contents.
<b>Index</b>	Look up help using the documentation Index.
<b>Search</b>	Search the documentation for one or more words.
<b>Demos</b>	See what demos are available; run selected demos.
<b>Favorites</b>	Save bookmarks for frequently used Help pages.

**For more information:** See “Finding Information with the Help Browser” in the MATLAB Desktop Tools and Development Environment documentation.

## Help on Functions from the Help Browser

To find help on any function from the Help browser, do either of the following:

- Select the **Contents** tab of the Help browser, open the **Contents** entry labeled MATLAB, and find the two subentries shown below. Use one of these to look up the function you want help on.
  - Functions — Categorical List
  - Functions — Alphabetical List
- Type `doc functionname` at the command line.

## Help on Functions from the Command Window

Several types of help on functions are available from the Command Window:

- To list all categories that you can request help on from the Command Window, just type
 

```
help
```
- To see a list of functions for one of these categories, along with a brief description of each function, type `help category`. For example,
 

```
help datafun
```
- To get help on a particular function, type `help functionname`. For example,
 

```
help sortrows
```

## Topical Help

In addition to the help on individual functions, you can get help on any of the following topics by typing `help topicname` at the command line.

Topic Name	Description
<code>arith</code>	Arithmetic operators
<code>relop</code>	Relational and logical operators

<b>Topic Name</b>	<b>Description</b>
punct	Special character operators
slash	Arithmetic division operators
paren	Parentheses, braces, and bracket operators
precedence	Operator precedence
datatypes	MATLAB data types, their associated functions, and operators that you can overload
lists	Comma separated lists
strings	Character strings
function_handle	Function handles and the @ operator
debug	Debugging functions
java	Using Java from within MATLAB
fileformats	A list of readable file formats
changeNotification	Windows directory change notification

## **Paged Output**

Before displaying a lengthy section of help text or code, put MATLAB into its paged output mode by typing `more on`. This breaks up any ensuing display into pages for easier viewing. Turn off paged output with `more off`.

Page through the displayed text using the space bar key. Or step through line by line using **Enter** or **Return**. Discontinue the display by pressing the **Q** key or **Ctrl+C**.

## **Writing Your Own Help**

Start each program you write with a section of text providing help on how and when to use the function. If formatted properly, the MATLAB `help` function displays this text when you enter

```
help functionname
```

MATLAB considers the first group of consecutive lines immediately following the function definition line that begin with % to be the help section for the function. The first line without % as the left-most character ends the help.

**For more information:** See “Help Text” in the MATLAB Desktop Tools and Development Environment documentation.

## Help for Subfunctions and Private Functions

You can write help for subfunctions using the same rules that apply to main functions. To display the help for the subfunction `mysubfun` in file `myfun.m`, type

```
help myfun/mysubfun
```

To display the help for a private function, precede the function name with `private/`. To get help on private function `myprivfun`, type

```
help private/myprivfun
```

## Help for Methods and Overloaded Functions

You can write help text for object-oriented class methods implemented with M-files. Display help for the method by typing

```
help classname/methodname
```

where the file `methodname.m` resides in subdirectory `@classname`.

For example, if you write a `plot` method for a class named `polynom`, (where the `plot` method is defined in the file `@polynom/plot.m`), you can display this help by typing

```
help polynom/plot
```

You can get help on overloaded MATLAB functions in the same way. To display the help text for the `eq` function as implemented in `matlab/iofun/@serial`, type

```
help serial/eq
```

## Development Environment

This section covers the following topics:

- “Workspace Browser” on page 11-10
- “Using the Find and Replace Utility” on page 11-10
- “Commenting Out a Block of Code” on page 11-11
- “Creating M-Files from Command History” on page 11-11
- “Editing M-Files in EMACS” on page 11-11

### Workspace Browser

The Workspace browser is a graphical interface to the variables stored in the MATLAB base and function workspaces. You can view, modify, save, load, and create graphics from workspace data using the browser. Select **View -> Workspace** to open the browser.

To view function workspaces, you need to be in debug mode.

**For more information:** See “MATLAB Workspace” in the MATLAB Desktop Tools and Development Environment documentation.

### Using the Find and Replace Utility

Find any word or phrase in a group of files using the Find and Replace utility. Click on **View -> Current Directory**, and then click on the binoculars icon at the top of the **Current Directory** window.

When entering search text, you don’t need to put quotes around a phrase. In fact, parts of words, like `win` for `windows`, will not be found if enclosed in quotes.

**For more information:** See “Finding and Replacing a String” in the MATLAB Desktop Tools and Development Environment documentation.

## Commenting Out a Block of Code

To comment out a block of text or code within the MATLAB editor,

- 1 Highlight the block of text you would like to comment out.
- 2 Holding the mouse over the highlighted text, select **Text -> Comment** (or **Uncomment**, to do the reverse) from the toolbar. (You can also get these options by right-clicking the mouse.)

**For more information:** See “Commenting” in the MATLAB Desktop Tools and Development Environment documentation.

## Creating M-Files from Command History

If there is part of your current MATLAB session that you would like to put into an M-file, this is easily done using the Command History window:

- 1 Open this window by selecting **View -> Command History**.
- 2 Use **Shift+Click** or **Ctrl+Click** to select the lines you want to use. MATLAB highlights the selected lines.
- 3 Right-click once, and select **Create M-File** from the menu that appears. MATLAB creates a new Editor window displaying the selected code.

## Editing M-Files in EMACS

If you use Emacs, you can download editing modes for editing M-files with GNU-Emacs or with early versions of Emacs from the MATLAB Central Web site:

<http://www.mathworks.com/matlabcentral/>

At this Web site, select **File Exchange**, and then **Utilities -> Emacs**.

**For more information:** See “General Preferences for the Editor/Debugger” in the MATLAB Desktop Tools and Development Environment documentation.

## M-File Functions

This section covers the following topics:

- “M-File Structure” on page 11-12
- “Using Lowercase for Function Names” on page 11-12
- “Getting a Function’s Name and Path” on page 11-13
- “What M-Files Does a Function Use?” on page 11-13
- “Dependent Functions, Built-Ins, Classes” on page 11-13

### M-File Structure

An M-File consists of the components shown here:

```
function [x, y] = myfun(a, b, c) % Function definition line
% H1 line   A one-line summary of the function's purpose.
% Help text One or more lines of help text that explain
%   how to use the function. This text is displayed when
%   the user types "help functionname".

% The Function body normally starts after the first blank line.
% Comments   Description (for internal use) of what the function
%   does, what inputs are expected, what outputs are generated.
%   Typing "help functionname" does not display this text.

x = prod(a, b); % Start of Function code
```

**For more information:** See “Basic Parts of an M-File” on page 4-8 of the MATLAB Programming documentation.

### Using Lowercase for Function Names

Function names appear in uppercase in MATLAB help text only to make the help easier to read. In practice, however, it is usually best to use lowercase when calling functions.

For M-file functions, case requirements depend on the case sensitivity of the operating system you are using. As a rule, naming and calling functions using lowercase generally makes your M-files more portable from one operating system to another.

## Getting a Function's Name and Path

To obtain the name of an M-file that is currently being executed, use the following function in your M-file code.

```
mfilename
```

To include the path along with the M-file name, use

```
mfilename('fullpath')
```

**For more information:** See the `mfilename` function reference page.

## What M-Files Does a Function Use?

For a simple display of all M-files referenced by a particular function, follow the steps below:

- 1 Type `clear functions` to clear all functions from memory (see Note below).
- 2 Execute the function you want to check. Note that the function arguments you choose to use in this step are important, since you can get different results when calling the same function with different arguments.
- 3 Type `inmem` to display all M-Files that were used when the function ran. If you want to see what MEX-files were used as well, specify an additional output, as shown here:

```
[mfiles, mexfiles] = inmem
```

---

**Note** `clear functions` does not clear functions locked by `mlock`. If you have locked functions, (which you can check using `inmem`), unlock them with `munlock`, and then repeat step 1.

---

## Dependent Functions, Built-Ins, Classes

For a much more detailed display of dependent function information, use the `depfun` function. In addition to M-files, `depfun` shows which built-ins and classes a particular function depends on.

## Function Arguments

This section covers the following topics:

- “Getting the Input and Output Arguments” on page 11-14
- “Variable Numbers of Arguments” on page 11-14
- “String or Numeric Arguments” on page 11-15
- “Passing Arguments in a Structure” on page 11-15
- “Passing Arguments in a Cell Array” on page 11-15

### Getting the Input and Output Arguments

Use `nargin` and `nargout` to determine the number of input and output arguments in a particular function call. Use `nargchk` and `nargoutchk` to verify that your function is called with the required number of input and output arguments.

```
function [x, y] = myplot(a, b, c, d)
    disp(nargchk(2, 4, nargin))           % Allow 2 to 4 inputs
    disp(nargoutchk(0, 2, nargout))     % Allow 0 to 2 outputs

    x = plot(a, b);
    if nargin == 4
        y = myfun(c, d);
    end
```

### Variable Numbers of Arguments

You can call functions with fewer input and output arguments than you have specified in the function definition, but not more. If you want to call a function with a variable number of arguments, use the `varargin` and `varargout` function parameters in the function definition.

This function returns the size vector and, optionally, individual dimensions:

```
function [s, varargout] = mysize(x)
    nout = max(nargout, 1) - 1;
    s = size(x);
    for k = 1:nout
        varargout(k) = {s(k)};
    end
```

Try calling it with

```
[s, rows, cols] = mysize(rand(4, 5))
```

## String or Numeric Arguments

If you are passing only string arguments into a function, you can use MATLAB command syntax. All arguments entered in command syntax are interpreted as strings.

```
strcmp string1 string1  
ans =  
    1
```

When passing numeric arguments, it is best to use function syntax unless you want the number passed as a string. The right-hand example below passes the number 75 as the string, '75'.

```
isnumeric(75)                isnumeric 75  
ans =                          ans =  
    1                            0
```

**For more information:** See “Passing Arguments” in the MATLAB Programming documentation.

## Passing Arguments in a Structure

Instead of requiring an additional argument for every value you want to pass in a function call, you can package them in a MATLAB structure and pass the structure. Make each input you want to pass a separate field in the structure argument, using descriptive names for the fields.

Structures allow you to change the number, contents, or order of the arguments without having to modify the function. They can also be useful when you have a number of functions that need similar information.

## Passing Arguments in a Cell Array

You can also group arguments into cell arrays. The disadvantage over structures is that you don't have fieldnames to describe each variable. The advantage is that cell arrays are referenced by index, allowing you to loop through a cell array and access each argument passed in or out of the function.

## Program Development

This section covers the following topics:

- “Planning the Program” on page 11-16
- “Using Pseudo-Code” on page 11-16
- “Selecting the Right Data Structures” on page 11-16
- “General Coding Practices” on page 11-17
- “Naming a Function Uniquely” on page 11-17
- “The Importance of Comments” on page 11-17
- “Coding in Steps” on page 11-18
- “Making Modifications in Steps” on page 11-18
- “Functions with One Calling Function” on page 11-18
- “Testing the Final Program” on page 11-18

### Planning the Program

When planning how to write a program, take the problem you are trying to solve and break it down into a series of smaller, independent tasks. Implement each task as a separate function. Try to keep functions fairly short, each having a single purpose.

### Using Pseudo-Code

You may find it helpful to write the initial draft of your program in a structured format using your own natural language. This *pseudo-code* is often easier to think through, review, and modify than using a formal programming language, yet it is easily translated into a programming language in the next stage of development.

### Selecting the Right Data Structures

Look at what data types and data structures are available to you in MATLAB and determine which of those best fit your needs in storing and passing your data.

**For more information:** See “Data Types” in the MATLAB Programming documentation.

## General Coding Practices

A few suggested programming practices:

- Use descriptive function and variable names to make your code easier to understand.
- Order subfunctions alphabetically in an M-file to make them easier to find.
- Precede each subfunction with a block of help text describing what that subfunction does. This not only explains the subfunctions, but also helps to visually separate them.
- Don't extend lines of code beyond the 80th column. Otherwise, it will be hard to read when you print it out.
- Use full Handle Graphics<sup>®</sup> property and value names. Abbreviated names are often allowed, but can make your code unreadable. They also could be incompatible in future releases of MATLAB.

## Naming a Function Uniquely

To avoid choosing a name for a new function that might conflict with a name already in use, check for any occurrences of the name using this command:

```
which -all functionname
```

**For more information:** See the which function reference page.

## The Importance of Comments

Be sure to document your programs well to make it easier for you or someone else to maintain them. Add comments generously, explaining each major section and any smaller segments of code that are not obvious. You can add a block of comments as shown here.

```
%-----  
% This function computes the ... <and so on>  
%-----
```

**For more information:** See “Comments” in the MATLAB Programming documentation.

## **Coding in Steps**

Don't try to write the entire program all at once. Write a portion of it, and then test that piece out. When you have that part working the way you want, then write the next piece, and so on. It's much easier to find programming errors in a small piece of code than in a large program.

## **Making Modifications in Steps**

When making modifications to a working program, don't make widespread changes all at one time. It's better to make a few small changes, test and debug, make a few more changes, and so on. Tracking down a difficult bug in the small section that you've changed is much easier than trying to find it in a huge block of new code.

## **Functions with One Calling Function**

If you have a function that is called by only one other function, put it in the same M-file as the calling function, making it a subfunction.

**For more information:** See "Subfunctions" in the MATLAB Programming documentation.

## **Testing the Final Program**

One suggested practice for testing a new program is to step through the program in the MATLAB debugger while keeping a record of each line that gets executed on a printed copy of the program. Use different combinations of inputs until you have observed that every line of code is executed at least once.

## Debugging

This section covers the following topics:

- “The MATLAB Debug Functions” on page 11-19
- “More Debug Functions” on page 11-19
- “The MATLAB Graphical Debugger” on page 11-20
- “A Quick Way to Examine Variables” on page 11-20
- “Setting Breakpoints from the Command Line” on page 11-21
- “Finding Line Numbers to Set Breakpoints” on page 11-21
- “Stopping Execution on an Error or Warning” on page 11-21
- “Locating an Error from the Error Message” on page 11-21
- “Using Warnings to Help Debug” on page 11-22
- “Making Code Execution Visible” on page 11-22
- “Debugging Scripts” on page 11-22

### The MATLAB Debug Functions

For a brief description of the main debug functions in MATLAB, type

```
help debug
```

**For more information:** See “Debugging M-Files” in the MATLAB Desktop Tools and Development Environment documentation.

### More Debug Functions

Other functions you may find useful in debugging are listed below.

Function	Description
echo	Display function or script code as it executes.
disp	Display specified values or messages.
sprintf, fprintf	Display formatted data of different types.
whos	List variables in the workspace.

<b>Function</b>	<b>Description</b>
size	Show array dimensions.
keyboard	Interrupt program execution and allow input from keyboard.
return	Resume execution following a keyboard interruption.
warning	Display specified warning message.
error	Display specified error message.
lasterr	Return error message that was last issued.
lasterror	Return last error message and related information.
lastwarn	Return warning message that was last issued.

## **The MATLAB Graphical Debugger**

Learn to use the MATLAB graphical debugger. You can view the function and its calling functions as you debug, set and clear breakpoints, single-step through the program, step into or over called functions, control visibility into all workspaces, and find and replace strings in your files.

Start out by opening the file you want to debug using **File -> Open** or the open function. Use the debugging functions available on the toolbar and pull-down menus to set breakpoints, run or step through the program, and examine variables.

**For more information:** See “Debugging M-Files” and “Using Debugging Features” in the MATLAB Desktop Tools and Development Environment documentation.

## **A Quick Way to Examine Variables**

To see the value of a variable from the Editor/Debugger window, hold the mouse cursor over the variable name for a second or two. You will see the value of the selected variable displayed.

## Setting Breakpoints from the Command Line

You can set breakpoints with `dbstop` in any of the following ways:

- Break at a specific M-file line number.
- Break at the beginning of a specific subfunction.
- Break at the first executable line in an M-file.
- Break when a warning, or error, is generated.
- Break if any infinite or NaN values are encountered.

**For more information:** See “Setting Breakpoints” in the MATLAB Desktop Tools and Development Environment documentation.

## Finding Line Numbers to Set Breakpoints

When debugging from the command line, a quick way to find line numbers for setting breakpoints is to use `dbtype`. The `dbtype` function displays all or part of an M-file, also numbering each line. To display `copyfile.m`, use

```
dbtype copyfile
```

To display only lines 70 through 90, use

```
dbtype copyfile 70:90
```

## Stopping Execution on an Error or Warning

Use `dbstop if error` to stop program execution on any error and enter debug mode. Use `warning debug` to stop execution on any warning and enter debug mode.

**For more information:** See “Debug, Backtrace, and Verbose Modes” in the MATLAB Programming documentation.

## Locating an Error from the Error Message

Click on the underlined text in an error message, and MATLAB opens the M-file being executed in its editor and places the cursor at the point of error.

**For more information:** See “Types of Errors” in the MATLAB Desktop Tools and Development Environment documentation.

## Using Warnings to Help Debug

You can detect erroneous or unexpected behavior in your programs by inserting warning messages that MATLAB will display under the conditions you specify. See the section on “Warning Control” in the MATLAB “Programming and Data Types” documentation to find out how to selectively enable warnings.

**For more information:** See the warning function reference page.

## Making Code Execution Visible

An easy way to see the end result of a particular line of code is to edit the program and temporarily remove the terminating semicolon from that line. Then, run your program and the evaluation of that statement is displayed on the screen.

**For more information:** See “Finding Errors” in the MATLAB Desktop Tools and Development Environment documentation.

## Debugging Scripts

Scripts store their variables in a workspace that is shared with the caller of the script. So, when you debug a script from the command line, the script uses variables from the base workspace. To avoid errors caused by workspace sharing, type `clear all` before starting to debug your script to clear the base workspace.

# Variables

This section covers the following topics:

- “Rules for Variable Names” on page 11-23
- “Making Sure Variable Names Are Valid” on page 11-23
- “Don’t Use Function Names for Variables” on page 11-24
- “Checking for Reserved Keywords” on page 11-24
- “Avoid Using i and j for Variables” on page 11-24
- “Avoid Overwriting Variables in Scripts” on page 11-25
- “Persistent Variables” on page 11-25
- “Protecting Persistent Variables” on page 11-25
- “Global Variables” on page 11-26

## Rules for Variable Names

Although variable names can be of any length, MATLAB uses only the first N characters of the name, (where N is the number returned by the function `namelengthmax`), and ignores the rest. Hence, it is important to make each variable name unique in the first N characters to enable MATLAB to distinguish variables. Also note that variable names are case sensitive.

```
N = namelengthmax
N =
    63
```

**For more information:** See “Naming Variables” in the MATLAB Programming documentation.

## Making Sure Variable Names Are Valid

Before using a new variable name, you can check to see if it is valid with the `isvarname` function. Note that `isvarname` does not consider names longer than `namelengthmax` characters to be valid.

For example, the following name cannot be used for a variable since it begins with a number.

```
isvarname 8thColumn
ans =
    0
```

**For more information:** See “Naming Variables” in the MATLAB Programming documentation.

## Don't Use Function Names for Variables

When naming a variable, make sure you are not using a name that is already used as a function name. If you do define a variable with a function name, you won't be able to call that function until you `clear` the variable from memory. (If it's a MATLAB built-in function, then you will still be able to call that function but you must do so using `builtin`.)

To test whether a proposed variable name is already used as a function name, use

```
which -all name
```

## Checking for Reserved Keywords

MATLAB reserves certain keywords for its own use and does not allow you to override them. Attempts to use these words may result in any one of a number of error messages, some of which are shown here:

```
Error: Expected a variable, function, or constant, found "=".
Error: "End of Input" expected, "case" found.
Error: Missing operator, comma, or semicolon.
Error: "identifier" expected, "=" found.
```

Use the `iskeyword` function with no input arguments to list all reserved words.

## Avoid Using `i` and `j` for Variables

MATLAB uses the characters `i` and `j` to represent imaginary units. Avoid using `i` and `j` for variable names if you intend to use them in complex arithmetic.

If you want to create a complex number without using `i` and `j`, you can use the `complex` function.

## Avoid Overwriting Variables in Scripts

MATLAB scripts store their variables in a workspace that is shared with the caller of the script. When called from the command line, they share the base workspace. When called from a function, they share that function's workspace. If you run a script that alters a variable that already exists in the caller's workspace, that variable is overwritten by the script.

**For more information:** See “Scripts” in the MATLAB Programming documentation.

## Persistent Variables

To get the equivalent of a static variable in MATLAB, use `persistent`. When you declare a variable to be persistent within a function, its value is retained in memory between calls to that function. Unlike global variables, persistent variables are known only to the function in which they are declared.

**For more information:** See “Persistent Variables” in the MATLAB Programming documentation.

## Protecting Persistent Variables

You can inadvertently clear persistent variables from memory by either modifying the function in which the variables are defined, or by clearing the function with one of the following commands:

```
clear all
clear functions
```

Locking the M-file in memory with `mlock` prevents any persistent variables defined in the file from being reinitialized.

## Global Variables

Use global variables sparingly. The global workspace is shared by all of your functions and also by your interactive MATLAB session. The more global variables you use, the greater the chances of unintentionally reusing a variable name, thus leaving yourself open to having those variables change in value unexpectedly. This can be a difficult bug to track down.

**For more information:** See “Global Variables” in the MATLAB Programming documentation.

# Strings

This section covers the following topics:

- “Creating Strings with Concatenation” on page 11-27
- “Comparing Methods of Concatenation” on page 11-27
- “Store Arrays of Strings in a Cell Array” on page 11-28
- “Converting Between Strings and Cell Arrays” on page 11-28
- “Search and Replace Using Regular Expressions” on page 11-29

## Creating Strings with Concatenation

Strings are often created by concatenating smaller elements together (e.g., strings, values, etc.). Two common methods of concatenating are to use the MATLAB concatenation operator (`[]`) or the `sprintf` function. The second and third line below illustrate both of these methods. Both lines give the same result:

```
numChars = 28;  
s = ['There are ' int2str(numChars) ' characters here']  
s = sprintf('There are %d characters here\n', numChars)
```

**For more information:** See “Creating Character Arrays” and “Numeric/String Conversion” in the MATLAB Programming documentation.

## Comparing Methods of Concatenation

When building strings with concatenation, `sprintf` is often preferable to `[]` because

- It is easier to read, especially when forming complicated expressions
- It gives you more control over the output format
- It often executes more quickly

You can also concatenate using the `strcat` function. However, for simple concatenations, `sprintf` and `[]` are faster.

## Store Arrays of Strings in a Cell Array

It is usually best to store an array of strings in a cell array instead of a character array, especially if the strings are of different lengths. Strings in a character array must be of equal length, which often requires padding the strings with blanks. This is not necessary when using a cell array of strings that has no such requirement.

The `cellRecord` below does not require padding the strings with spaces:

```
charRecord = ['Allison Jones'; 'Development  '; 'Phoenix      '];  
cellRecord = {'Allison Jones'; 'Development'; 'Phoenix'};
```

**For more information:** See “Cell Arrays of Strings” in the MATLAB Programming documentation.

## Converting Between Strings and Cell Arrays

You can convert between standard character arrays and cell arrays of strings using the `cellstr` and `char` functions:

```
charRecord = ['Allison Jones'; 'Development  '; 'Phoenix      '];  
cellRecord = cellstr(charRecord);
```

Also, a number of the MATLAB string operations can be used with either character arrays, or cell arrays, or both:

```
cellRecord2 = {'Brian Lewis'; 'Development'; 'Albuquerque'};  
strcmp(charRecord, cellRecord2)  
ans =  
     0  
     1  
     0
```

**For more information:** See “Converting to a Cell Array of Strings” and “String Comparisons” in the MATLAB Programming documentation.

## Search and Replace Using Regular Expressions

Using regular expressions in MATLAB offers a very versatile way of searching for and replacing characters or phrases within a string. See the help on these functions for more information.

Function	Description
regexp	Match regular expression.
regexp_i	Match regular expression, ignoring case.
regexprep	Replace string using regular expression.

**For more information:** See “Regular Expressions” in the MATLAB Programming documentation.

## Evaluating Expressions

This section covers the following topics:

- “Find Alternatives to Using `eval`” on page 11-30
- “Assigning to a Series of Variables” on page 11-30
- “Short-Circuit Logical Operators” on page 11-31
- “Changing the Counter Variable within a for Loop” on page 11-31

### Find Alternatives to Using `eval`

While the `eval` function can provide a convenient solution to certain programming challenges, it is best to limit its use. The main reason is that code that uses `eval` is often difficult to read and hard to debug. A second reason is that `eval` statements cannot always be translated into C or C++ code by the MATLAB Compiler.

If you are evaluating a function, it is more efficient to use `feval` than `eval`. The `feval` function is made specifically for this purpose and is optimized to provide better performance.

**For more information:** See MATLAB Technical Note 1103, “What Is the EVAL Function, When Should I Use It, and How Can I Avoid It?” at URL <http://www.mathworks.com/support/tech-notes/1100/1103.html>.

### Assigning to a Series of Variables

One common pattern for creating variables is to use a variable name suffixed with a number (e.g., `phase1`, `phase2`, `phase3`, etc.). We recommend using a cell array to build this type of variable name series, as it makes code more readable and executes more quickly than some other methods. For example:

```
for k = 1:800
    phase{k} = expression;
end
```

## Short-Circuit Logical Operators

MATLAB has logical AND and OR operators (&& and ||) that enable you to partially evaluate, or *short-circuit*, logical expressions. Short-circuit operators are useful when you want to evaluate a statement only when certain conditions are satisfied.

In this example, MATLAB does not execute the function `myfun` unless its M-file exists on the current path.

```
comp = (exist('myfun.m') == 2) && (myfun(x) >= y)
```

**For more information:** See “Short-Circuit Operators” in the MATLAB Programming documentation.

## Changing the Counter Variable within a for Loop

You cannot change the value of the loop counter variable (e.g., the variable `k` in the example below) in the body of a for loop. For example, this loop executes just 10 times, even though `k` is set back to 1 on each iteration.

```
for k = 1:10
    disp(sprintf('Pass %d', k))
    k = 1;
end
```

Although MATLAB does allow you to use a variable of the same name as the loop counter within a loop, this is not a recommended practice.

## MATLAB Path

This section covers the following topics:

- “Precedence Rules” on page 11-32
- “File Precedence” on page 11-33
- “Adding a Directory to the Search Path” on page 11-33
- “Handles to Functions Not on the Path” on page 11-33
- “Making Toolbox File Changes Visible to MATLAB” on page 11-34
- “Making Nontoolbox File Changes Visible to MATLAB” on page 11-35
- “Change Notification on Windows” on page 11-35

### Precedence Rules

When MATLAB is given a name to interpret, it determines its usage by checking the name against each of the entities listed below, and in the order shown:

- 1 Variable
- 2 Subfunction
- 3 Private function
- 4 Class constructor
- 5 Overloaded method
- 6 M-file in the current directory
- 7 M-file on the path, or MATLAB built-in function

Regarding the last item on this list, all built-in functions have a `.bi` file on the path that is used internally by MATLAB. If you have an M-file on the path that has the same name as a MATLAB built-in, MATLAB selects the function that has its M- or `.bi` file in the directory closest to the beginning of the path string.

If you have two or more M-files on the path that have the same name, MATLAB selects the function that has its M-file in the directory closest to the beginning of the path string.

**For more information:** See “Function Precedence Order” in the MATLAB Programming documentation.

## File Precedence

If you refer to a file by its filename only (leaving out the file extension), and there is more than one file of this name in the directory, MATLAB selects the file to use according to the following precedence:

- 1 MEX-file
- 2 MDL-file (Simulink<sup>®</sup> model)
- 3 P-Code file
- 4 M-file

**For more information:** See “Selecting Methods from Multiple Implementation Types” in the MATLAB Programming documentation.

## Adding a Directory to the Search Path

To add a directory to the search path, use either of the following:

- At the toolbar, select **File -> Set Path**.
- At the command line, use the `addpath` function.

You can also add a directory and all of its subdirectories in one operation by either of these means. To do this from the command line, use `genpath` together with `addpath`. The online help for the `genpath` function shows how to do this.

This example adds `/control` and all of its subdirectories to the MATLAB path:

```
addpath(genpath('K:/toolbox/control'))
```

**For more information:** See “Search Path” in the MATLAB Desktop Tools and Development Environment documentation.

## Handles to Functions Not on the Path

You cannot create function handles to functions that are not on the MATLAB path. But you can achieve essentially the same thing by creating the handles

through a script file placed in the same off-path directory as the functions. If you then run the script, using `run path/script`, you will have created the handles that you need.

For example,

- 1 Create a script in this off-path directory that constructs function handles and assigns them to variables. That script might look something like this:

```
File E:/testdir/createFhandles.m
    fhset = @setItem
    fhsort = @sortItems
    fhdel = @deleteItem
```

- 2 Run the script from your current directory to create the function handles:

```
run E:/testdir/createFhandles
```

- 3 You can now execute one of the functions by means of its handle.

```
fhset(item, value)
```

## **Making Toolbox File Changes Visible to MATLAB**

Unlike functions in user-supplied directories, M-files (and MEX-files) in the `$MATLAB/toolbox` directories are not time-stamp checked, so MATLAB does not automatically see changes to them. If you modify one of these files, and then rerun it, you may find that the behavior does not reflect the changes that you made. This is most likely because MATLAB is still using the previously loaded version of the file.

To force MATLAB to reload a function from disk, you need to explicitly clear the function from memory using `clear functionname`. Note that there are rare cases where `clear` will not have the desired effect, (for example, if the file is locked, or if it is a class constructor and objects of the given class exist in memory).

Similarly, MATLAB does not automatically detect the presence of new files in `$MATLAB/toolbox` directories. If you add (or remove) files from these directories, use `rehash toolbox` to force MATLAB to see your changes. Note that if you use the MATLAB Editor to create files, these steps are unnecessary, as the Editor automatically informs MATLAB of such changes.

## Making Nontoolbox File Changes Visible to MATLAB

For M-files outside of the toolbox directories, MATLAB sees the changes made to these files by comparing timestamps and reloads any file that has changed the next time you execute the corresponding function.

If MATLAB does not see the changes you make to one of these files, try clearing the old copy of the function from memory using `clear functionname`. You can verify that MATLAB has cleared the function using `inmem` to list all functions currently loaded into memory.

## Change Notification on Windows

If MATLAB, running on Windows, is unable to see new files or changes you have made to an existing file, the problem may be related to operating system change notification handles.

Type the following for more information:

```
help changeNotification
help changeNotificationAdvanced
```

## Program Control

This section covers the following topics:

- “Using break, continue, and return” on page 11-36
- “Using switch Versus if” on page 11-37
- “MATLAB case Evaluates Strings” on page 11-37
- “Multiple Conditions in a case Statement” on page 11-37
- “Implicit Break in switch-case” on page 11-38
- “Variable Scope in a switch” on page 11-38
- “Catching Errors with try-catch” on page 11-38
- “Nested try-catch Blocks” on page 11-39
- “Forcing an Early Return from a Function” on page 11-39

### Using break, continue, and return

It’s easy to confuse the break, continue, and return functions as they are similar in some ways. Make sure you use these functions appropriately.

Function	Where to Use It	Description
break	for or while loops	Exits the loop in which it appears. In nested loops, control passes to the next outer loop.
continue	for or while loops	Skips any remaining statements in the current loop. Control passes to next iteration of the same loop.
return	Anywhere	Immediately exits the function in which it appears. Control passes to the caller of the function.

## Using switch Versus if

It is possible, but usually not advantageous, to implement switch-case statements using if-elseif instead. See pros and cons in the table.

<b>switch-case Statements</b>	<b>if-elseif Statements</b>
Easier to read.	Can be difficult to read.
Can compare strings of different lengths.	You need strcmp to compare strings of different lengths.
Test for equality only.	Test for equality or inequality.

## MATLAB case Evaluates Strings

A useful difference between switch-case statements in MATLAB and C is that you can specify string values in MATLAB case statements, which you cannot do in C.

```
switch(method)
    case 'linear'
        disp('Method is linear')
    case 'cubic'
        disp('Method is cubic')
end
```

## Multiple Conditions in a case Statement

You can test against more than one condition with switch. The first case below tests for either a linear or bilinear method by using a cell array in the case statement.

```
switch(method)
    case {'linear', 'bilinear'}
        disp('Method is linear or bilinear')
    case (<and so on>)
end
```

## Implicit Break in switch-case

In C, if you don't end each case with a break statement, code execution falls through to the following case. In MATLAB, case statements do not fall through; only one case may execute. Using break within a case statement is not only unnecessary, it is also invalid and generates a warning.

In this example, if result is 52, only the first disp statement executes, even though the second is also a valid match:

```
switch(result)
    case 52
        disp('result is 52')
    case {52, 78}
        disp('result is 52 or 78')
end
```

## Variable Scope in a switch

Since MATLAB executes only one case of any switch statement, variables defined within one case are not known in the other cases of that switch statement. The same holds true for if-elseif statements.

In these examples, you get an error when choice equals 2, because x is undefined.

SWITCH-CASE	IF-ELSEIF
<pre>switch choice     case 1         x = -pi:0.01:pi;     case 2         plot(x, sin(x)); end</pre>	<pre>if choice == 1     x = -pi:0.01:pi; elseif choice == 2     plot(x, sin(x)); end</pre>

## Catching Errors with try-catch

When you have statements in your code that could possibly generate unwanted results, put those statements into a try-catch block that will catch any errors and handle them appropriately.

The example below shows a try-catch block within a function that multiplies two matrices. If a statement in the try segment of the block fails, control passes to the catch segment. In this case, the catch statements check the error message that was issued (returned by `lasterr`) and respond appropriately.

```
try
    X = A * B
catch
    errmsg = lasterr;
    if(strfind(errmsg, 'Inner matrix dimensions'))
        disp('** Wrong dimensions for matrix multiply')
    end
end
```

**For more information:** See “Checking for Errors with try-catch” in the MATLAB Programming documentation.

## Nested try-catch Blocks

You can also nest try-catch blocks, as shown here. You can use this to attempt to recover from an error caught in the first try section:

```
try
    statement1                % Try to execute statement1
catch
    try
        statement2           % Attempt to recover from error
    catch
        disp 'Operation failed' % Handle the error
    end
end
```

## Forcing an Early Return from a Function

To force an early return from a function, place a return statement in the function at the point where you want to exit. For example,

```
if <done>
    return
end
```

## Save and Load

This section covers the following topics:

- “Saving Data from the Workspace” on page 11-40
- “Loading Data into the Workspace” on page 11-40
- “Viewing Variables in a MAT-File” on page 11-41
- “Appending to a MAT-File” on page 11-41
- “Save and Load on Startup or Quit” on page 11-42
- “Saving to an ASCII File” on page 11-42

### Saving Data from the Workspace

To save data from your workspace, you can do any of the following:

- Copy from the MATLAB Command Window and paste into a text file.
- Record part of your session in a diary file, and then edit the file in a text editor.
- Save to a binary or ASCII file using the save function.
- Save spreadsheet, scientific, image, or audio data with appropriate function.
- Save to a file using low-level file I/O functions (`fwrite`, `fprintf`, ...).

**For more information:** See “Saving the Current Workspace” in the MATLAB Desktop Tools and Development Environment documentation, “Using the diary Function to Export Data” on page 6-39, and “Using Low-Level File I/O Functions” on page 6-105.

### Loading Data into the Workspace

Similarly, to load new or saved data into the workspace, you can do any of the following:

- Enter or paste data at the command line.
- Create a script file to initialize large matrices or data structures.
- Read a binary or ASCII file using load.
- Load spreadsheet, scientific, image, or audio data with appropriate function.
- Load from a file using low-level file I/O functions (`fread`, `fscanf`, ...).

**For more information:** See “Loading a Saved Workspace and Importing Data” in the MATLAB Development Environment documentation, and “Using Low-Level File I/O Functions” on page 6-105.

## Viewing Variables in a MAT-File

To see what variables are saved in a MAT-file, use `who` or `whos` as shown here (the `.mat` extension is not required). `who` returns a cell array and `whos` returns a structure array.

```
mydataVariables = who('-file', 'mydata.mat');
```

## Appending to a MAT-File

To save additional variables to an existing MAT-file, use

```
save matfilename -append
```

Any variables you save that do not yet exist in the MAT-file are added to the file. Any variables you save that already exist in the MAT-file overwrite the old values.

---

**Note** Saving with the `-append` switch does not append additional elements to an array that is already saved in a MAT-file. See the example below.

---

In this example, the second save operation does not concatenate new elements to vector `A`, (making `A` equal to `[1 2 3 4 5 6 7 8]`) in the MAT-file. Instead, it replaces the 5 element vector, `A`, with a 3 element vector, also retaining all other variables that were stored on the first save operation.

```
A = [1 2 3 4 5];    B = 12.5;    C = rand(4);
save savefile;
A = [6 7 8];
save savefile A -append;
```

## Save and Load on Startup or Quit

You can automatically save your variables at the end of each MATLAB session by creating a `finish.m` file to save the contents of your base workspace every time you quit MATLAB. Load these variables back into your workspace at the beginning of each session by creating a `startup.m` file that uses the `load` function to load variables from your MAT-file.

**For more information:** See the `startup` and `finish` function reference pages.

## Saving to an ASCII File

When you save matrix data to an ASCII file using `save -ascii`, MATLAB combines the individual matrices into one collection of numbers. Variable names are not saved. If this is not acceptable for your application, use `fprintf` to store your data instead.

**For more information:** See “Exporting Delimited ASCII Data Files” on page 6-37.

## Files and Filenames

This section covers the following topics:

- “Naming M-files” on page 11-43
- “Naming Other Files” on page 11-43
- “Passing Filenames as Arguments” on page 11-44
- “Passing Filenames to ASCII Files” on page 11-44
- “Determining Filenames at Run-Time” on page 11-44
- “Returning the Size of a File” on page 11-45

### Naming M-files

M-file names must start with an alphabetic character, may contain any alphanumeric characters or underscores, and must be no longer than the maximum allowed M-file name length (returned by the function `namelengthmax`).

```
N = namelengthmax
N =
    63
```

Since variables must obey similar rules, you can use the `isvarname` function to check whether a filename (minus its `.m` file extension) is valid for an M-file.

```
isvarname mfilename
```

### Naming Other Files

The names of other files that MATLAB interacts with (e.g., MAT, MEX, and MDL-files) follow the same rules as M-files, but may be of any length.

Depending on your operating system, you may be able to include certain non-alphanumeric characters in your filenames. Check your operating system manual for information on valid filename restrictions.

## Passing Filenames as Arguments

In MATLAB commands, you can specify a filename argument using the MATLAB command or function syntax. For example, either of the following are acceptable. (The `.mat` file extension is optional for `save` and `load`).

```
load mydata.mat           % Command syntax
load('mydata.mat')       % Function syntax
```

If you assign the output to a variable, you must use the function syntax.

```
savedData = load('mydata.mat')
```

## Passing Filenames to ASCII Files

ASCII files are specified as follows. Here, the file extension is required.

```
load mydata.dat -ascii    % Command syntax
load('mydata.dat', '-ascii') % Function syntax
```

## Determining Filenames at Run-Time

There are several ways that your function code can work on specific files without you having to hard-code their filenames into the program. You can

- Pass the filename in as an argument

```
function myfun(datafile)
```

- Prompt for the filename using the `input` function

```
filename = input('Enter name of file: ', 's');
```

- Browse for the file using the `uigetfile` function

```
[filename, pathname] = uigetfile('*.mat', 'Select MAT-file');
```

**For more information:** See the `input` and `uigetfile` function reference pages.

## Returning the Size of a File

Two ways to have your program determine the size of a file are shown here.

### METHOD #1

```
s = dir('myfile.dat');  
filesize = s.bytes
```

### METHOD #2

```
fid = fopen('myfile.dat');  
fseek(fid, 0, 'eof');  
filesize = ftell(fid)  
fclose(fid);
```

The `dir` function also returns the filename (`s.name`), last modification date (`s.date`), and whether or not it's a directory (`s.isdir`).

(The second method requires read access to the file.)

**For more information:** See the `fopen`, `fseek`, `ftell`, and `fclose` function reference pages.

## Input/Output

This section covers the following topics:

- “File I/O Function Overview” on page 11-46
- “Common I/O Functions” on page 11-46
- “Readable File Formats” on page 11-46
- “Using the Import Wizard” on page 11-47
- “Loading Mixed Format Data” on page 11-47
- “Reading Files with Different Formats” on page 11-47
- “Reading ASCII Data into a Cell Array” on page 11-48
- “Interactive Input into Your Program” on page 11-48

For more information and examples on importing and exporting data, see Technical Note 1602:

<http://www.mathworks.com/support/tech-notes/1600/1602.html>

### File I/O Function Overview

For a good overview of MATLAB file I/O functions, use the online “Functions — Categorical List” reference. In the Help browser **Contents**, select **MATLAB -> Functions — Categorical List**, and then click **File I/O**.

### Common I/O Functions

The most commonly used, high-level, file I/O functions in MATLAB are `save` and `load`. For help on these, type `doc save` or `doc load`.

Functions for I/O to text files with delimited values are `textread`, `dlmread`, `dlmwrite`. Functions for I/O to text files with comma-separated values are `csvread`, `csvwrite`.

**For more information:** See “Text Files” in the MATLAB “Functions — Categorical List” reference documentation.

### Readable File Formats

Type `doc fileformats` to see a list of file formats that MATLAB can read, along with the associated MATLAB functions.

## Using the Import Wizard

A quick method of importing text or binary data from a file (e.g., Excel files) is to use the MATLAB Import Wizard. Open the Import Wizard with the command, `uiimport filename` or by selecting **File -> Import Data** at the Command Window.

Specify or browse for the file containing the data you want to import and you will see a preview of what the file contains. Select the data you want and click **Finish**.

**For more information:** See “Using the Import Wizard with Text Data” on page 6-7 and “Using the Import Wizard with Binary Data” on page 6-13.

## Loading Mixed Format Data

To load data that is in mixed formats, use `textread` instead of `load`. The `textread` function lets you specify the format of each piece of data.

If the first line of file `mydata.dat` is

```
Sally    12.34 45
```

Read the first line of the file as a free format file using the `%` format:

```
[names, x, y] = textread('mydata.dat', '%s %f %d', 1)
```

returns

```
names =  
    'Sally'  
x =  
    12.340000000000000  
y =  
    45
```

## Reading Files with Different Formats

Attempting to read data from a file that was generated on a different platform may result in an error because the binary formats of the platforms may differ. Using the `fopen` function, you can specify a machine format when you open the file to avoid these errors.

## Reading ASCII Data into a Cell Array

A common technique used to read an ASCII data file into a cell array is

```
[a,b,c,d] = textread('data.txt', '%s %s %s %s');  
mydata = cellstr([a b c d]);
```

**For more information:** See the `textread` and `cellstr` function reference pages.

## Interactive Input into Your Program

Your program can accept interactive input from users during execution. Use the `input` function to prompt the user for input, and then read in a response. When executed, `input` causes the program to display your prompt, pause while a response is entered, and then resume when the **Enter** key is pressed.

## Starting MATLAB

### Getting MATLAB to Start Up Faster

Here are some things that you can do to make MATLAB start up faster.

- Make sure toolbox path caching is enabled.
- Make sure that the system on which MATLAB is running has enough RAM.
- Choose only the windows you need in the MATLAB desktop.
- Close the Help browser before exiting MATLAB. When you start your next session, MATLAB will not open the Help browser, and thus will start faster.
- If disconnected from the network, check the `LM_LICENSE_FILE` variable. See <http://www.mathworks.com/support/solutions/data/25731.html> for a more detailed explanation.

**For more information:** See “Reduced Startup Time with Toolbox Path Caching” in the MATLAB Desktop Tools and Development Environment documentation.

## Operating System Compatibility

This section covers the following topics:

- “Executing O/S Commands from MATLAB” on page 11-50
- “Searching Text with grep” on page 11-50
- “Constructing Paths and Filenames” on page 11-50
- “Finding the MATLAB Root Directory” on page 11-51
- “Temporary Directories and Filenames” on page 11-51

### Executing O/S Commands from MATLAB

To execute a command from your operating system prompt without having to exit MATLAB, precede the command with the MATLAB ! operator.

On Windows, you can add an ampersand (&) to the end of the line to make the output appear in a separate window.

**For more information:** See “Running External Programs” in the MATLAB Desktop Tools and Development Environment documentation, and the `system` and `dos` function reference pages.

### Searching Text with grep

`grep` is a powerful tool for performing text searches in files on UNIX systems. To `grep` from within MATLAB, precede the command with an exclamation point (!`grep`).

For example, to search for the word `warning`, ignoring case, in all M-files of the current directory, you would use

```
!grep -i 'warning' *.m
```

### Constructing Paths and Filenames

Use the `fullfile` function to construct path names and filenames rather than entering them as strings into your programs. In this way, you always get the correct path specification, regardless of which operating system you are using at the time.

## Finding the MATLAB Root Directory

The `matlabroot` function returns the location of the MATLAB installation on your system. Use `matlabroot` to create a path to MATLAB and toolbox directories that does not depend on a specific platform or MATLAB version.

The following example uses `matlabroot` with `fullfile` to return a platform-independent path to the general toolbox directory:

```
fullfile(matlabroot, 'toolbox', 'matlab', 'general')
```

## Temporary Directories and Filenames

If you need to locate the directory on your system that has been designated to hold temporary files, use the `tempdir` function. `tempdir` returns a string that specifies the path to this directory.

To create a new file in this directory, use the `tempname` function. `tempname` returns a string that specifies the path to the temporary file directory, plus a unique filename.

For example, to store some data in a temporary file, you might issue the following command first.

```
fid = fopen(tempname, 'w');
```

## Demos

### Demos Available with MATLAB

MATLAB comes with a wide array of visual demonstrations to help you see the extent of what you can do with the product. To start running any of the demos, simply type `demo` at the MATLAB command prompt. Demos cover the following major areas:

- MATLAB
- Toolboxes
- Simulink
- Blocksets
- Real-Time Workshop®
- Stateflow®

**For more information:** See “Running Demonstrations” in the MATLAB Desktop Tools and Development Environment documentation, and the demo function reference page.

## For More Information

### Current CSSM

`news:comp.soft-sys.matlab`

### Archived CSSM

`http://mathforum.org/epigone/comp.soft-sys.matlab/`

### MATLAB Technical Support

`http://www.mathworks.com/support/`

### Tech Notes

`http://www.mathworks.com/support/tech-notes/list\_all.html`

### MATLAB Central

`http://www.mathworks.com/matlabcentral/`

### MATLAB Newsletters (Digest, News & Notes)

`http://www.mathworks.com/company/newsletters/index.html`

### MATLAB Documentation

`http://www.mathworks.com/access/helpdesk/help/helpdesk.html`

### MATLAB Index of Examples

`http://www.mathworks.com/access/helpdesk/help/techdoc/  
demo\_example.shtml`



# External Interfaces

---

Finding the Documentation in Online Help (p. A-2)	A summary of what information is available on the MATLAB external interfaces
Reference Documentation (p. A-5)	A summary of the MATLAB functions that support external interfaces

## Finding the Documentation in Online Help

MATLAB provides interface capabilities that enable you to communicate between MATLAB and the following programs and devices:

- MAT-files that store data from your MATLAB sessions
- Generic dynamic link library (DLL) files
- External C and Fortran programs
- Object-oriented technologies like Java and COM
- Web services, such as SOAP and WSDL
- Hardware devices on your computer's serial port

These interfaces, also referred to as the MATLAB Application Program Interface (API), are documented in full in the MathWorks book, "External Interfaces," and in the online help for MATLAB.

Use the following path in the left pane of the MATLAB Help browser to locate the help sections listed below:

```
MATLAB -> External Interfaces
```

### Importing and Exporting Data

This section describes how to use MAT-files to import data to and export data from the MATLAB environment. MAT-files provide a convenient mechanism for moving your MATLAB data between different platforms in a highly portable manner. In addition, they provide a means to import and export your data to other stand-alone MATLAB applications.

### MATLAB Interface to Generic DLLs

A shared library is a collection of functions that are available for use by one or more applications running on a system. On Windows systems, the library is precompiled into a dynamic link library (.dll) file. At run-time, the library is loaded into memory and made accessible to all applications. The MATLAB Interface to Generic DLLs enables you to interact with functions in dynamic link libraries directly from MATLAB.

---

## **Calling C and Fortran Programs from MATLAB**

MATLAB provides an interface to external programs written in the C and Fortran languages that enables you to interact with data and programs external to the MATLAB environment. This section explains how to call your own C or Fortran subroutines from MATLAB as if they were built-in functions.

### **Creating C Language MEX-Files**

MATLAB callable C and Fortran programs are referred to as MEX-files. This section explains how to create and work with C MEX-files.

### **Creating Fortran MEX-Files**

This section explains how to create and work with Fortran MEX-files.

## **Calling MATLAB from C and Fortran Programs**

You can employ MATLAB as a computational engine that responds to calls from your C and Fortran programs. This section describes the MATLAB functions that enable you to

- Start and end a MATLAB process
- Send commands to and exchange data with MATLAB
- Compile and link MATLAB engine programs

## **Calling Java from MATLAB**

This section describes how to use the MATLAB interface to Java classes and objects. This MATLAB capability enables you to

- Bring Java classes into the MATLAB environment
- Construct objects from those classes
- Work with Java arrays in MATLAB
- Call methods on Java objects, passing MATLAB or Java data types

## COM and DDE Support

MATLAB has interfaces that enable you to interact with Component Object Model (COM) and Dynamic Data Exchange (DDE). This section explains how to

- Integrate COM control components into a COM control container such as a MATLAB figure window
- Use COM Automation components to control or be controlled by MATLAB
- Enable access between MATLAB and other Windows applications using DDE

## Using Web Services in MATLAB

Web services are designed to let applications running on disparate operating systems and development platforms communicate with each other. MATLAB acts as a Web service client by sending requests to a server and handling the responses. MATLAB implements the following Web service technologies:

- Simple Object Access Protocol (SOAP)
- Web Services Description Language (WSDL)

## Serial Port I/O

This section describes the MATLAB serial port interface which provides direct access to peripheral devices such as modems, printers, and scientific instruments that you connect to your computer's serial port. This interface is established through a serial port object that enables you to

- Configure serial port communications
- Use serial port control pins
- Write data to and read data from the device
- Execute an action when a particular event occurs
- Create a record of your serial port session

## Reference Documentation

The online help provides a detailed description of each of the MATLAB functions available in the MATLAB external interfaces. Use the following path in the left pane of the MATLAB Help browser to locate the help sections listed below:

MATLAB -> External Interfaces Reference

### **Generic DLL Interface Functions**

Functions you can use to interact with external shared library (.dll) files.

### **C MAT-File Functions**

Functions that enable you to incorporate and use MATLAB data in your own C programs.

### **C MX-Functions**

Array access and creation functions that you use in your C files to manipulate MATLAB arrays.

### **C MEX-Functions**

Functions that you use in your C files to perform operations in the MATLAB environment.

### **C Engine Functions**

Functions that enable you to call MATLAB from your own C programs.

### **Fortran MAT-File Functions**

Functions that enable you to incorporate and use MATLAB data in your own Fortran programs.

### **Fortran MX-Functions**

Array access and creation functions that you use in your Fortran files to manipulate MATLAB arrays.

### **Fortran MEX-Functions**

Functions that you use in your Fortran files to perform operations in the MATLAB environment.

**Fortran Engine Functions**

Functions that enable you to call MATLAB from your own Fortran programs.

**Java Interface Functions**

Functions that enable you to create and interact with Java classes and objects from MATLAB.

**COM Functions**

Functions that create Component Object Model objects and manipulate their interfaces.

**DDE Functions**

Dynamic Data Exchange functions that enable MATLAB to access other Windows applications and vice versa.

**Serial Port I/O Functions**

Functions that enable you to interact with devices connected to your computer's serial port.

## Symbols

- ! symbol
    - for entering a shell escape function 3-75
  - % symbol
    - for specifying character conversions 3-76
    - for writing single-line comments 3-76
    - for writing the H1 help line 4-10
  - %{ and %} symbols
    - for writing multiple-line comments 3-77
  - ( ) symbol
    - for indexing into an array 3-76
    - for specifying function input arguments 3-76
  - (space) character
    - for separating array row elements 3-78
    - for separating function return values 3-79
  - \* symbol
    - for filename wildcards 3-70
  - , symbol
    - for separating array indices 3-72
    - for separating array row elements 3-72
    - for separating input or output arguments 3-73
    - for separating MATLAB commands 3-73
  - . symbol
    - for defining a structure field 3-73
    - for specifying object methods 3-74
  - .( ) symbol
    - for creating a dynamic structure field 3-75
  - .. symbol
    - for referring to a parent directory 3-74
  - ... symbol
    - for continuing a command line 3-74
  - : symbol
    - for converting to a column vector 3-72
    - for generating a numeric sequence 3-71
    - for preserving array shape on assignment 3-72
    - for specifying an indexing range 3-72
    - generating a numeric sequence 1-10
  - ; symbol
    - for separating rows of an array 3-77
    - for suppressing command output 3-77
  - @ symbol
    - for class directories 3-71
    - for constructing function handles 3-70
  - [ ] symbol
    - for concatenating arrays 3-79
    - for constructing an array 3-79
    - for specifying function return values 3-80
  - { } symbol
    - for constructing a cell array 3-73
    - for indexing into a cell array 3-73
  - ' symbol
    - for constructing a character array 3-78
- ## A
- access modes
    - HDF files 6-90
  - accuracy of calculations 3-10
  - addition operator 3-12
  - aggregation 8-37
  - and (M-file function equivalent for &) 3-16
  - anonymous functions 5-3
    - changing variables 5-9
    - constructing 5-3
    - evaluating variables 5-8
    - example — multiple anonymous functions 5-12
    - example — passing a function to quad 5-11
    - in cell arrays 5-5
    - using space characters in 5-6
    - with no input arguments 5-5
  - answer, assigned to ans 3-10

- arguments
    - checking number of 4-21
    - function 4-10
    - order in argument list 4-25
    - order of outputs 4-23
    - passing 4-36
    - passing variable number 4-23
  - arithmetic operators
    - overloading 8-29
  - arrays
    - cell array of strings 2-27
    - concatenating diagonally 1-40
    - deleting rows and columns 1-30
    - diagonal 1-39
    - dimensions
      - inverse permutation 1-61
    - empty 1-41
    - expanding 1-26
    - flipping 1-32
    - functions
      - changing indexing style 1-72
      - creating a matrix 1-69
      - determining data type 1-70
      - finding matrix structure or shape 1-70
      - modifying matrix shape 1-69
      - multidimensional arrays 1-72
      - sorting and shifting 1-71
    - functions for diagonals 1-71
    - getting dimensions of 1-23
    - linear indexing 1-18
    - logical indexing 1-21
    - multidimensional 1-48
    - numeric
      - converting to cell array 2-78
    - of strings 2-26
    - reshaping 1-31
    - rotating 1-32
    - shifting 1-35
    - sorting column data 1-37
    - sorting row data 1-37
    - sorting row vectors 1-38
    - storage 1-18
    - transposing 1-32
  - ASCII data
    - definition of 6-3
    - exporting 6-36
    - exporting delimited data 6-37
    - exporting with diary function 6-39
    - formats 6-28
    - importing 6-28
    - importing delimited files 6-32
    - importing mixed alphabetic and numeric data 6-33
    - importing space-delimited data 6-31
    - reading formatted text 6-114
    - saving 6-37
    - specifying delimiter used in file 6-32
    - using the Import Wizard with 6-7
    - with text headers 6-32
    - writing 6-115
  - assignment statements
    - building structure arrays with 2-50
    - local and global variables 3-7
  - attributes
    - retrieving from HDF files 6-91
    - writing to an HDF file 6-100
- B**
- base (numeric), converting 2-35
  - base date 2-42
  - binary data
    - controlling data type of values read 6-109
    - using the Import Wizard 6-7

- using the Import Wizard with 6-13
  - writing to 6-110
- binary files
- importing 6-13
- binary from decimal conversion 2-35
- blanks
- finding in string arrays 2-32
  - removing from strings 2-27
- break 3-66
- built-in functions 3-82
- forcing a built-in call 3-83
  - identifying 3-82
- ## C
- C++ and MATLAB OOP 8-7
- caching
- MATLAB directory 4-14
- callback functions
- creating 9-14
  - specifying 9-16
- calling context 4-17
- calling MATLAB functions
- storing as pseudocode 4-15
- canonical class 8-9
- case 3-62
- case conversion 2-38, 2-39
- cat 1-52
- catch 3-67
- cell
- building nested arrays with 2-76
  - preallocating empty arrays with 2-70
- cell arrays 2-66
- accessing a subset of cells 2-71
  - accessing data 2-70
  - adding cells to 2-68
  - applying functions to 2-74
- cell indexing 2-67
- accessing a subset of cells 2-71
- celldisp 2-68
- cellplot 2-68
- char data type 6-109
- cell indexing 2-67
- concatenating 2-69
- content indexing 2-68
- converting to numeric array 2-78
- creating 2-67
- using assignments 2-67
  - with cells function 2-70
  - with curly braces 2-69
- deleting cells 2-72
- deleting dimensions 2-72
- displaying 2-68
- expanding 2-68
- flat 2-76
- functions 2-79
- growing 1-27, 1-29
- indexing 2-68
- multidimensional 1-66
- nested 2-76
- building with the cells function 2-76
  - indexing 2-77
- of strings 2-27
- comparing strings 2-31
  - functions 2-29
- of structures 2-78
- organizing data 2-75
- preallocating 2-70, 10-8
- replacing comma-separated list with 2-72
- reshaping 2-72
- visualizing 2-68
- with anonymous function elements 5-5

- character arrays
  - categorizing characters of 2-32
  - comparing 2-30
  - comparing values on cell arrays 2-31
  - concatenating 3-22
  - conversion 2-34
  - converting to cell arrays 2-27
  - creating 2-25
  - delimiting character 2-33
  - evaluating 3-22
  - finding a substring 2-33
  - functions 2-39
  - functions that create 2-38
  - functions that modify 2-38
  - in cell arrays 2-27
  - padding for equal row length 2-26
  - removing trailing blanks 2-27
  - representation 2-25
  - scalar 2-31
  - searching and replacing 2-33
  - searching or comparing 2-39
  - token 2-33
  - two-dimensional 2-26
  - using relational operators on 2-31
- characters
  - corresponding ASCII values 2-36
  - finding in string 2-32
  - used as delimiters 6-28
- characters and strings 2-25
- class 8-11
- class directories 8-6
- classes
  - clearing definition 8-6
  - constructor method 8-10
  - debugging 8-5
  - designing 8-9
  - matlab 2-82
  - methods required by MATLAB 8-9
  - object-oriented methods 8-2
  - overview 8-2
- clear 4-31, 10-12
- clipboard
  - importing binary data 6-7
- closing
  - files 6-117
- colon operator 1-10
  - for multidimensional array subscripting 1-55
  - scalar expansion with 1-51
- column separators
  - defined 6-28
- command/function duality 4-35
- comma-separated lists 3-54
  - FFT example 3-58
  - generating from cell array 3-54
  - generating from structure 3-55
  - replacing with cell array 2-72
  - usage 3-56
    - concatenation 3-57
    - constructing arrays 3-56
    - displaying arrays 3-56
    - function call arguments 3-57
    - function return values 3-58
- comments
  - in code 4-12
  - in scripts and functions 4-8
- comparing
  - strings 2-30
- complex conjugate transpose operator 3-12
- complex number functions 2-18
- complex numbers 2-11
  - creating 2-11
- computational functions
  - applying to cell arrays 2-74
  - applying to multidimensional arrays 1-62

- applying to structure fields 2-56
  - in M-file 4-8
  - computer 3-10
  - computer type 3-10
  - concatenating
    - cell arrays 2-69
    - diagonal matrices 1-40
    - matrices 1-7
    - strings 3-22
  - concatenation 1-3
    - of unlike data types 1-11
  - concatenation functions 1-8
  - conditional statements 4-21
  - constructor methods 8-10
    - guidelines 8-10
    - using `class` in 8-11
  - containment 8-37
  - content indexing 2-68
    - to access cell contents 2-70
  - Contents.m file 4-15
  - continue 3-66
  - control statements
    - break 3-66
    - case 3-62
    - catch 3-67
    - continue 3-66
    - else 3-60
    - elseif 3-60
    - for 3-64
    - if 3-60
    - otherwise 3-62
    - return 3-68
    - switch 3-62
    - try 3-67
    - while 3-65
  - conv 2-73
  - converter methods 8-20
  - converting
    - cases of strings 2-38, 2-39
    - dates 2-41
    - numbers 2-34
    - strings 2-34
  - converting numeric and string data types 2-40
  - converting numeric to string 2-34
  - converting string to numeric 2-36
  - cos 4-16
  - creating
    - cell array 2-67
    - multidimensional array 1-50
    - string array 2-27
    - strings 2-25
    - structure array 2-50
    - timer objects 9-4
  - cross 1-62
  - curly braces
    - for cell array indexing 2-67
    - in cell array syntax 2-69
    - to build cell arrays 2-69
    - to nest cell arrays 2-76
- D**
- data
    - binary, dependence upon array size and type 6-23
  - data class hierarchy 8-3
  - data organization
    - cell arrays 2-75
    - multidimensional arrays 1-64
    - structure arrays 2-59
  - data sets
    - See* HDF data sets

- data types 2-2
  - cell arrays 2-66
  - cell arrays of strings 2-27
  - combining unlike data types 1-11
  - complex numbers 2-11
  - dates and times 2-41
  - determining 2-39
  - double precision 6-109
  - floating point 2-6
    - double-precision 2-7
    - single-precision 2-9
  - infinity 2-12
  - integers 2-4
  - java classes 2-83
  - logical 2-20
  - logicals 2-20
  - NaN 2-12
  - numeric 2-4
  - precision 6-109
  - reading files 6-109
  - specifying for input 6-109
  - structure arrays 2-49
  - user-defined classes 8-3
- date 2-46
- date and time functions 2-47
- datenum 2-43
- dates
  - base 2-42
  - conversions 2-43
  - handling and converting 2-41
  - number 2-42
  - string, vector of input 2-44
- dates and times 2-41
- datestr 2-43
- datevec 2-43
- deblank 2-27
- debugging
  - errors and warnings 7-23
- debugging class methods 8-5
- decimal representation
  - to binary 2-35
  - to hexadecimal 2-35
- delaying program execution
  - using timers 9-2
- deleting
  - cells from cell array 2-72
  - fields from structure arrays 2-56
  - matrix rows and columns 1-30
- deleting array elements 1-30
- deletion operator 1-30
- delimiter in string 2-33
- delimiters
  - defined 6-28
- diagonal matrices 1-39
- diary 6-39
- dim argument for cat 1-52
- dimensions
  - deleting 2-72
  - permuting 1-60
  - removing singleton 1-59
- directories
  - adding to path 8-6
  - class 8-6
  - Contents.m file 4-15
  - help for 4-15
  - MATLAB
    - caching 4-14
    - private functions for 5-33
    - private methods for 8-5
    - temporary 6-108
- disp 2-63
- dispatch type 8-68

- display method 8-12
    - examples 8-13
  - displaying
    - cell arrays 2-68
    - field names for structure array 2-51
  - division operators
    - left division 3-12
    - matrix left division 3-13
    - matrix right division 3-13
    - right division 3-12
  - double precision 6-109
  - double-precision matrix 2-4
  - downloading files 6-118
  - duality, command/function 4-35
  - dynamic field names in structure arrays 2-54
- E**
- Earth Observing System (EOS) 6-69
  - editor
    - accessing 4-13
    - for creating M-files 4-13
  - eig 1-63
  - element-by-element organization for structures 2-61
  - else 3-60
  - else, elseif 3-61
  - elseif 3-60
  - empty arrays
    - and if statement 3-62
    - and relational operators 3-14
    - and while loops 3-66
  - empty matrices 1-41
  - end 1-20
  - end method 8-19
  - end of file 6-111
  - EOS (Earth Observing System)
    - sources of information 6-69
  - eps 3-10
  - epsilon 3-10
  - equal to operator 3-14
  - error 4-18
  - error handling
    - debugging 7-23
    - error recovery 7-4
    - formatted message strings 7-4
    - identifying errors 7-5
    - message identifiers 7-4
    - regenerating errors 7-7
    - reporting errors 7-4
    - with try-catch 7-2
  - error messages
    - formatted message strings 7-4
  - evaluating
    - string containing function name 3-23
    - string containing MATLAB expression 3-22
  - examples
    - checking number of function arguments 4-22
    - container class 8-54
    - for 3-64
    - function 4-17
    - if 3-61
    - inheritance 8-38
    - M-file for structure array 2-57
    - polynomial class 8-24
    - script 4-16
    - switch 3-63
    - vectorization 10-4
    - while 3-65
  - expanding
    - cell arrays 2-68
    - structure arrays 2-50
  - expanding cell arrays 1-27, 1-29

expanding structure arrays 1-27, 1-29

exporting

ASCII data 6-36

in HDF4 format 6-96

in HDF5 format 6-65

overview 6-3

expressions

involving empty arrays 3-14

most recent answer 3-10

overloading 8-21

scalar expansion with 3-13

external program, running from MATLAB 3-23

## F

`fclose` 6-117

`feof` 6-110

`fid`

*See* file identifiers

field names

dynamic 2-54

`fieldnames` 2-51

fields 2-49, 2-50

accessing data within 2-53

adding to structure array 2-56

applying functions to 2-56

all like-named fields 2-57

assigning data to 2-50

deleting from structures 2-56

indexing within 2-54

names 2-51

size 2-55

writing M-files for 2-57

`fields` 2-51

file exchange

over Internet 6-118

file formats

importing and exporting 6-3

file identifiers

clearing 6-117

defined 6-106

file operations

FTP 6-123

filenames

wildcards 3-70

files

ASCII

reading 6-112

reading formatted text 6-114

writing 6-115

beginning of 6-111

binary

controlling data type values read 6-109

data types 6-109

reading 6-108

writing to 6-110

closing 6-117

current position 6-111

end of 6-111

failing to open 6-107

file identifiers (FID) 6-106

MAT 6-26

opening 6-106

permissions 6-106

position 6-110

specifying delimiter used in ASCII files 6-32

temporary 6-108

`find` function

and subscripting 3-17

finding

substring within a string 2-33

flipping matrices 1-32

float 6-109

- floating point 2-6
- floating point, double-precision 2-7
  - converting to 2-8
  - creating 2-8
  - maximum and minimum values 2-7
- floating point, single-precision 2-9
  - converting to 2-10
  - creating 2-10
  - maximum and minimum values 2-9
- floating-point functions 2-17
- floating-point numbers
  - largest 3-10
  - smallest 3-10
- floating-point precision 6-109
- floating-point relative accuracy 3-10
- flow control
  - break 3-66
  - case 3-62
  - catch 3-67
  - continue 3-66
  - else 3-60
  - elseif 3-60
  - for 3-64
  - if 3-60
  - otherwise 3-62
  - return 3-68
  - switch 3-62
  - try 3-67
  - while 3-65
- fopen 6-106
  - failing 6-107
- for 2-78, 3-64
  - example 3-64
  - indexing 3-65
  - nested 3-64
  - syntax 3-64
- format for numeric values 2-14
- fread 6-108
- frewind 6-110
- fseek 6-110
- ftell 6-110
- FTP file operations 6-123
- function definition line
  - for subfunction 5-31
  - in an M-file 4-8
  - syntax 4-9
- function handles
  - example 4-28
  - for nested functions 5-20
  - maximum name length 4-29
  - naming 4-29
  - operations on 4-29
  - overview of 4-27
- function workspace 4-17
- functions
  - applying
    - to multidimensional structure arrays 1-68
    - to structure contents 2-56
  - applying to cell arrays 2-74
  - arguments
    - passing variable number of 4-23
  - body 4-8, 4-11
  - built-in 3-82
    - forcing a built-in call 3-83
    - identifying 3-82
  - calling
    - command syntax 4-35
    - function syntax 4-35
    - passing arguments 4-36
  - calling context 4-17
  - cell arrays 2-79
  - cell arrays of strings 2-29
  - changing indexing style 1-72
  - character arrays 2-39

- clearing from memory 4-31
- comments 4-8
- comparing character arrays 2-39
- complex number 2-18
- computational, applying to structure fields 2-56
- creating a matrix 1-69
- creating arrays with 1-52
- creating matrices 1-4
- date and time 2-47
- determining data type 1-70
- example 4-17
- executing function name string 3-23
- finding matrix structure or shape 1-70
- floating-point 2-17
- for diagonal matrices 1-71
- infinity 2-18
- integer 2-16
- logical array 2-21
- matrix concatenation 1-8
- M-file 3-81
- modifying character arrays 2-38
- modifying matrix shape 1-69
- multidimensional arrays 1-72
- multiple output arguments 4-10
- NaN 2-18
- numeric and string conversion 2-40
- numeric to string conversion 2-34
- output formatting 2-19
- overloaded 3-83
- overloading 8-23
- primary 5-31
- searching character arrays 2-39
- sorting and shifting 1-71
- sparse matrix 1-46
- storing as pseudocode 4-15
- string to numeric conversion 2-36

- structures 2-65
- that determine data type 2-39
- type identification 2-18
- types of 4-18
  - anonymous 5-3
  - nested 5-15
  - overloaded 5-34
  - primary 5-14
  - private 5-33
  - subfunctions 5-31
- fwrite 6-110

## G

- get method 8-13
- global attributes
  - HDF files 6-91
- global variables 3-3
  - alternatives 3-4
  - creating 3-3
  - displaying 3-4
  - suggestions for use 3-4
- graphics files
  - getting information about 6-41
  - importing and exporting 6-41
- greater than operator 3-14
- greater than or equal to operator 3-14
- growing an array 1-26
- growing cell array 1-27, 1-29
- growing structure arrays 1-27, 1-29

## H

- H1 line 4-8, 4-10
  - and help command 4-8
  - and lookfor command 4-8

- HDF (Hierarchical Data Format)
    - exporting in HDF4 format 6-96
    - exporting in HDF5 format 6-65
    - importing into MATLAB 6-69
    - importing subsets of data 6-72
    - MATLAB utility API 6-103
    - output arguments 6-89
    - programming model 6-88
    - selecting data sets to import 6-71
  - HDF data sets
    - accessing 6-92
    - associating attributes with 6-100
    - closing access 6-95
    - creating 6-97
    - getting information about 6-92
    - reading 6-93
    - using predefined attributes 6-101
  - HDF files
    - access modes 6-90
    - associating attributes with 6-100
    - closing 6-95
    - closing all open identifiers 6-104
    - creating 6-96
    - getting information about 6-91
    - listing open identifiers 6-103
    - opening 6-90
    - reading global attributes 6-91
    - writing data 6-98
  - HDF5
    - reading files 6-59
  - HDF-EOS
    - Earth Observing System 6-69
  - help
    - M-file 4-11
  - help
    - and H1 line 4-8
  - help text 4-8
  - hexadecimal, converting from decimal 2-35
  - Hierarchical Data Format
    - reading HDF5 files 6-59
  - hierarchy of data classes 8-3
- I**
- if 3-60
    - and empty arrays 3-62
    - example 3-61
    - nested 3-61
  - imaginary unit 3-10
  - Import Data option 6-7
  - import functions
    - comparison of features 6-30
  - Import Wizard
    - importing binary data 6-7
    - with ASCII data 6-7
    - with binary data 6-13
  - importing
    - ASCII data 6-28
    - HDF data 6-69
      - from the command line 6-88
    - Import Wizard 6-7
    - overview 6-3
    - selecting HDF data sets 6-71
    - subsets of HDF data 6-72
  - indexed reference 8-14
  - indexing
    - cell array 2-67
    - content 2-68
    - for loops 3-65
    - multidimensional arrays 1-54
    - nested cell arrays 2-77
    - nested structure arrays 2-64
    - structures within cell arrays 2-79
    - within structure fields 2-54

- indices, how MATLAB calculates 1-57
- Inf 3-10
- inferiorto 8-66
- inferiorto function 8-66
- infinity 2-12
  - functions 2-18
  - represented in MATLAB 3-10
- inheritance
  - example class 8-38
  - multiple 8-37
  - simple 8-35
- integer data type 6-115
- integer functions 2-16
- integers 2-4
  - creating 2-5
  - largest system can represent 3-10
  - smallest system can represent 3-10
- Internet functions 6-118
- intmax 3-10
- intmin 3-10
- inverse permutation of array dimensions 1-61
- ipermute 1-61
- isa 8-11

## J

- Java and MATLAB OOP 8-7

## K

- keywords 3-9
  - checking for 11-24

## L

- lasterr
  - using with message identifiers 7-9

- less than operator 3-13
- less than or equal to operator 3-13
- load 10-12
- loading data
  - overview 6-3
- loading objects 8-60
- loadobj example 8-62
- local variables 3-2
- logical array functions 2-21
- logical data type 2-20
- logical expressions
  - and subscripting 3-17
- logical operators 3-15
  - bit-wise 3-18
  - element-wise 3-15
  - short-circuit 3-19
- logical types 2-20
- logical vectors 1-21
- long 6-109
- long integer 6-109
- lookfor 4-8, 4-11
  - and H1 line 4-8
- loops
  - for 3-64
  - while 3-65

## M

- MATLAB
  - data type classes 8-3
  - programming
    - M-files 4-7
    - scripts 4-16
  - structures 8-7
  - version 3-10

**matrices**

- accessing multiple elements 1-19
- accessing single elements 1-17
- concatenating 1-7
- concatenating diagonally 1-40
- constructing a matrix operations
  - constructing 1-4
- data structure query 1-25
- data type query 1-24
- deleting rows and columns 1-30
- diagonal 1-39
- empty 1-41
- expanding 1-26
- flipping 1-32
- functions
  - changing indexing style 1-72
  - creating a matrix 1-69
  - determining data type 1-70
  - finding matrix structure or shape 1-70
  - modifying matrix shape 1-69
  - sorting and shifting 1-71
- functions for creating 1-4
- functions for diagonals 1-71
- getting dimensions of 1-23
- linear indexing 1-18
- logical indexing 1-21
- reshaping 1-31
- rotating 1-32
- scalar 1-44
- shifting 1-35
- sorting column data 1-37
- sorting row data 1-37
- sorting row vectors 1-38
- transposing 1-32
- vectors 1-44

**matrix**

- double-precision 2-4
- for loop index 3-65
- single-precision 2-4
- See also* matrices

**matrix operations**

- concatenating matrices 1-3
- creating matrices 1-3

**mean 1-62****memory**

- function workspace 4-17
- management 10-12
- Out of Memory message 10-16

**message identifiers**

- using with errors 7-4
- using with `lasterr` 7-9
- using with warnings 7-16

**methods 8-2**

- converters 8-20
- determining which is called 4-34
- display 8-12
- end 8-19
- get 8-13
- invoking on objects 8-4
- listing 8-33
- precedence 8-67
- required by MATLAB 8-9
- set 8-13
- subsasgn 8-14
- subsref 8-14

**M-file functions**

- identifying 3-81

**M-files**

- comments 4-12
- contents 4-8
- corresponding to functions 8-22

- creating
  - in MATLAB directory 4-14
- creating with text editor 4-13
- kinds 4-7
- naming 4-7
- operating on structures 2-57
- overview 4-8
- primary function 5-14
- subfunction 5-31
- superseding existing names 5-32

multidimensional arrays

- applying functions 1-62
  - element-by-element functions 1-62
  - matrix functions 1-63
  - vector functions 1-62
- cell arrays 1-66
- computations on 1-62
- creating 1-50
  - at the command line 1-50
  - with functions 1-52
  - with the cat function 1-52
- extending 1-51
- format 1-54
- indexing 1-54
  - avoiding ambiguity 1-58
  - with the colon operator 1-55
- number of dimensions 1-54
- organizing data 1-64
- permuting dimensions 1-60
- removing singleton dimensions 1-59
- reshaping 1-58
- size of 1-54
- storage 1-54
- structure arrays 1-67
  - applying functions 1-68
- subscripts 1-49

multiple conditions for switch 3-63

- multiple inheritance 8-37
- multiplication operators
  - matrix multiplication 3-12
  - multiplication 3-12

## N

names

- structure fields 2-51
- superseding 5-32

NaN 2-12

- functions 2-18
- logical operations on 2-13

NaN 3-10

nargin 4-21

nargout 4-21

ndgrid 1-72

ndims 1-54

nested functions 5-15

- creating 5-15
- example — creating a function handle 5-26
- example — function-generating functions 5-27
- separate variable instances 5-24
- using function handles with 5-20
- variable scope in 5-18

nesting

- cell arrays 2-76
- for loops 3-64
- if statements 3-61
- structures 2-63

newlines in string arrays 2-32

not (M-file function equivalent for ~) 3-16

not a number (NaN) 2-12

not equal to operator 3-14

Not-a-Number 3-10

now 2-46

number of arguments 4-21

- numbers
    - date 2-42
    - time 2-42
  - numeric data types 2-4
    - conversion functions 2-40
    - setting display format 2-14
  - numeric to string conversion
    - functions 2-34
- O**
- object-oriented programming
    - features of 8-2
    - inheritance
      - multiple 8-37
      - simple 8-35
    - overloading 8-21
    - subscripting 8-15
    - See also* classes and objects
  - objects
    - accessing data in 8-13
    - as indices into objects 8-19
    - creating 8-4
    - invoking methods on 8-4
    - loading 8-60
    - overview 8-2
    - precedence 8-65
    - saving 8-60
  - offsets for indexing 1-57
  - ones 1-21
  - online help 4-11
  - opening
    - files
      - failing 6-107
      - HDF files 6-90
      - permissions 6-106
      - using low-level functions 6-106
  - operator precedence 3-20
    - overriding 3-21
  - operators
    - addition 3-12
    - applying to cell arrays 2-74
    - applying to structure fields 2-56
    - colon 3-12
    - complex conjugate transpose 3-12
    - deletion 1-30
    - equal to 3-14
    - greater than 3-14
    - greater than or equal to 3-14
    - left division 3-12
    - less than 3-13
    - less than or equal to 3-13
    - logical 3-15
      - bit-wise 3-18
      - element-wise 3-15
      - short-circuit 3-19
    - matrix left division 3-13
    - matrix multiplication 3-12
    - matrix power 3-13
    - matrix right division 3-13
    - multiplication 3-12
    - not equal to 3-14
    - overloading 8-2
    - power 3-12
    - right division 3-12
    - subtraction 3-12
    - table of 8-22
    - transpose 3-12
    - unary minus 3-12
    - unary plus 3-12
  - optimization
    - preallocation, array 10-7, 10-13
    - vectorization 10-4
  - or (M-file function equivalent for |) 3-16

- organizing data
  - cell arrays 2-75
  - multidimensional arrays 1-64
  - structure arrays 2-59
- otherwise 3-62
- Out of Memory message 10-16
- output arguments 4-10
  - order of 4-23
- output formatting functions 2-19
- overloaded functions 3-83
- overloading 8-15
  - arithmetic operators 8-29
  - functions 8-23
  - loadobj 8-61
  - operators 8-2
  - pie3 8-57
  - saveobj 8-61
- P**
- pack 10-12
- page subscripts 1-49
- parentheses
  - for input arguments 4-10
  - overriding operator precedence with 3-21
- Paste Special option 6-7
- path
  - adding directories to 8-6
- pcode 4-15
- percent sign (comments) 4-12
- performance
  - analyzing 10-2
- permission strings 6-106
- permute 1-60
- permuting array dimensions 1-60
  - inverse 1-61
- persistent variables 3-5
- pi 3-10
- pie3 function overloaded 8-57
- plane organization for structures 2-60
- polar 4-16
- polynomials
  - example class 8-24
- power operators
  - matrix power 3-13
  - power 3-12
- preallocation
  - arrays 10-7, 10-13
  - cell array 10-8
- precedence
  - object 8-65
  - operator 3-20
    - overriding 3-21
- precision
  - char 6-109
  - data types 6-109
  - double 6-109
  - float 6-109
  - long 6-109
  - short 6-109
  - single 6-109
  - uchar 6-109
- primary functions 5-14
- private directory 5-33
- private functions 5-33
  - precedence of in classes 8-68
  - precedence of when calling 4-33
- private methods 8-5
- program control
  - break 3-66
  - case 3-62
  - catch 3-67
  - continue 3-66
  - else 3-60

- elseif 3-60
- for 3-64
- if 3-60
- otherwise 3-62
- return 3-68
- switch 3-62
- try 3-67
- while 3-65
- programs
  - running external 3-23
- pseudocode 4-15
  
- Q**
- quit 10-13
  
- R**
- randn 1-52
- reading
  - HDF data 6-69
    - from the command line 6-88
    - selecting HDF data sets 6-71
    - subsets of HDF data 6-72
- realmax 3-10
- realmin 3-10
- reference, subscripted 8-15
- regexp 3-25
- regexpi 3-25
- regexprep 3-25
- regular expression metacharacters
  - character classes
    - match alphanumeric character (\w) 3-28
    - match any character (period) 3-27
    - match any characters but these ([^c1c2c3]) 3-26
    - match any of these characters ([c1c2c3]) 3-28
    - match characters in this range ([c1-c2]) 3-28
    - match digit character (\d) 3-29
    - match nonalphanumeric character (\W) 3-27
    - match nondigit character (\D) 3-27
    - match nonwhitespace character (\S) 3-27
    - match whitespace character (\s) 3-28
  - character representation
    - alarm character (\a) 3-29
    - backspace character (\b) 3-29
    - carriage return character (\r) 3-29
    - escape character (\e) 3-29
    - form feed character (\f) 3-29
    - hexadecimal character (\x) 3-30
    - horizontal tab character (\t) 3-29
    - literal character (\char) 3-30
    - new line character (\n) 3-29
    - octal character (\o) 3-30
    - vertical tab character (\v) 3-29
  - logical operators
    - atomic group ((?>expr)) 3-31
    - comment ((?#expr)) 3-32
    - match exact word (\<expr\>) 3-33
    - match expr1 or expr2 (expr1|expr2) 3-32
    - match if expression begins string (^expr) 3-33
    - match if expression begins word (\<expr) 3-33
    - match if expression ends string (expr\$) 3-33
    - match if expression ends word (expr\>) 3-33
    - noncapturing group ((?:expr)) 3-30
    - standard group ((expr)) 3-31

- lookaround operators
  - match expr1, if followed by expr2  
(expr1(?=expr2)) 3-35
  - match expr1, if not followed by expr2  
(expr1(!expr2)) 3-35
  - match expr2, if not preceded by expr1  
(expr1(?<!expr2)) 3-35
  - match expr2, if preceded by expr1  
(expr1(?<=expr2)) 3-35
- operator summary 3-50
- quantifiers
  - greedy quantifier (qu\_expr\*) 3-39
  - lazy quantifier (qu\_expr?) 3-40
  - match 0 or 1 instance (expr?) 3-37
  - match 0 or more instances (expr?\*) 3-38
  - match 1 or more instances (expr?+) 3-38
  - match n instances (expr{n}) 3-37
  - match n to m instances (expr{n,m}) 3-39
  - possessive quantifier (qu\_expr+) 3-39
- token operators
  - create named token ((?<name>expr)) 3-44
  - create token ((expr)) 3-42
  - if token, match expr ((?(token)expr)) 3-42
  - if token, match expr1, else expr2  
((?(token)expr1 | expr2)) 3-45
  - match named token (\k<name>) 3-42
  - match Nth token (\N) 3-42
  - replace Nth token (\$N) 3-42
- regular expressions
  - [c1 - c2] 3-28
  - [c1c2c3] 3-28
  - character classes 3-26
  - character representation 3-29
  - conditional expressions 3-45
  - functions
    - regex 3-25
    - regexpi 3-25
    - regxprep 3-25
  - introduction 3-25
  - logical operators 3-30
  - lookaround operators 3-33
  - multiple strings
    - finding a single pattern 3-47
    - finding multiple patterns 3-48
    - matching 3-46
    - replacing 3-49
  - quantifiers 3-36
    - greedy 3-39
    - lazy 3-40
    - possessive 3-39
  - tokens 3-40
    - example 1 3-43
    - example 2 3-43
    - named capture 3-44
    - operators 3-42
    - use in replacement string 3-44
    - using token parameter 3-41
- relational operators
  - empty arrays 3-14
  - strings 2-31
- removing
  - cells from cell array 2-72
  - fields from structure arrays 2-56
  - singleton dimensions 1-59
- replacing substring within string 2-33
- repmat 1-52
- reshape 1-58, 2-72
- reshaping
  - cell arrays 2-72
  - multidimensional arrays 1-58
- reshaping matrices 1-31
- return 3-68
- rmfield 2-56
- rotating matrices 1-32

**S**

- save 10-12
- saveobj example 8-62
- saving
  - objects 8-60
- scalar
  - and relational operators 2-31
  - expansion 3-13
  - string 2-31
- scalars 1-44
- scheduling program execution
  - using timers 9-2
- Scientific Data API
  - programming model 6-89
- scripts 4-7
  - example 4-16
  - executing 4-17
- search path
  - M-files on 5-32
- set method 8-13
- shell escape functions 3-23
- shiftdim 1-72
- shifting matrix elements 1-35
- short 6-109
- short integer 6-109
- simple inheritance 8-35
- sin 1-62
- single precision 6-109
- single-precision matrix 2-4
- size
  - structure arrays 2-55
  - structure fields 2-55
- size 2-55
- smallest value system can represent 3-10
- sorting matrix column data 1-37
- sorting matrix row data 1-37
- sorting matrix row vectors 1-38
- sparse matrix functions 1-46
- sprintf 6-116
- square brackets
  - for output arguments 4-10
- squeeze 1-59
  - with multidimensional arguments 1-63
- sscanf 6-115
- starting
  - timers 9-9
- statements
  - conditional 4-21
- stopping
  - timers 9-9
- strcmp 2-30
- string to numeric conversion
  - functions 2-36
- strings 2-25
  - functions to create 2-38
- strings, cell arrays of 2-27
- structs 2-51
  - for nested structures 2-63
- structure arrays 2-49
  - accessing data 2-53
  - adding fields to 2-56
  - applying functions to 2-56
  - building 2-50
    - using structs 2-51
  - data organization 2-59
  - deleting fields 2-56
  - dynamic field names 2-54
  - element-by-element organization 2-61
  - expanding 2-50
  - fields 2-49
    - assigning data to 2-50
  - growing 1-27, 1-29

- indexing
    - nested structures 2-64
    - within fields 2-54
  - multidimensional 1-67
    - applying functions 1-68
  - nesting 2-63
  - obtaining field names 2-51
  - organizing data 2-59
    - example 2-62
  - plane organization 2-60
  - size 2-55
  - subarrays, accessing 2-53
  - subscripting 2-50
  - used with classes 8-7
  - within cell arrays 2-78
  - writing M-files for 2-57
    - example 2-57
  - structures
    - field names
      - dynamic 2-54
    - functions 2-65
  - subfunctions 5-31
    - accessing 5-32
    - creating 5-31
    - debugging 5-32
    - definition line 5-31
    - precedence of 4-32
  - subsasgn
    - for index reference 8-14
    - for subscripted assignment 8-17
  - subscripted assignment 8-17
  - subscripting
    - how MATLAB calculates indices 1-57
    - multidimensional arrays 1-49
    - overloading 8-15
    - page 1-49
    - structure arrays 2-50
      - with logical expression 3-17
      - with logical vectors 1-21
      - with the find function 3-17
  - subsref 8-15
  - subsref method 8-14
  - substring within a string 2-33
  - subtraction operator 3-12
  - sum 1-62
  - superiorto 8-66
  - superseding existing M-file names 5-32
  - switch 3-62
    - case groupings 3-62
    - example 3-63
    - multiple conditions 3-63
  - symbols 3-69
    - asterisk \* 3-70
    - at sign @ 3-70
    - colon : 3-71
    - comma , 3-72
    - curly braces { } 3-73
    - dot . 3-73
    - dot-dot .. 3-74
    - dot-dot-dot ... 3-74
    - dot-parentheses . ( ) 3-75
    - exclamation point ! 3-75
    - parentheses ( ) 3-76
    - percent % 3-76
    - percent-brace %{ and %} 3-77
    - semicolon ; 3-77
    - single quotes ' 3-78
    - space character 3-78
    - square brackets [ ] 3-79
- T**
- tabs in string arrays 2-32
  - tempdir 6-108

- tempname 6-108
  - temporary files
    - creating 6-108
  - text files
    - importing 6-7
    - reading 6-112
  - time
    - numbers 2-42
  - time and date functions 2-47
  - timer objects
    - blocking the command line 9-11
    - callback functions 9-13
    - creating 9-4
    - deleting 9-4
    - execution modes 9-18
    - finding all existing timers 9-23
    - naming convention 9-5
    - overview 9-2
    - properties 9-6
    - starting 9-9
    - stopping 9-9
  - timers
    - starting and stopping 9-9
    - using 9-2
  - times and dates 2-41
  - tips, programming
    - additional information 11-53
    - command and function syntax 11-3
    - debugging 11-19
    - demos 11-52
    - development environment 11-10
    - evaluating expressions 11-30
    - files and filenames 11-43
    - function arguments 11-14
    - help 11-6
    - input/output 11-46
    - MATLAB path 11-32
    - M-file functions 11-12
    - operating system compatibility 11-50
    - program control 11-36
    - program development 11-16
    - save and load 11-40
    - starting MATLAB 11-49
    - strings 11-27
    - variables 11-23
  - token in string 2-33
  - tokens
    - regular expressions 3-40
  - tolerance 3-10
  - transpose 1-61
  - transpose operator 3-12
  - transposing matrices 1-32
  - trigonometric functions 1-62
  - try 3-67
  - try-catch 7-2
  - type identification functions 2-18
- U**
- uchar data type 6-109
  - unary minus operator 3-12
  - unary plus operator 3-12
  - user classes, designing 8-9
- V**
- value
    - data type 6-109
    - largest system can represent 3-10
  - varargin 2-74, 4-24
    - in argument list 4-25
    - unpacking contents 4-24

- varargout 4-24
  - in argument list 4-25
  - packing contents 4-24
- variables
  - global 3-3
    - alternatives 3-4
    - creating 3-3
    - displaying 3-4
    - suggestions for use 3-4
  - local 3-2
  - naming 3-6
  - persistent 3-5
  - replacing list with a cell array 2-72
  - usage guidelines 3-7
- vector
  - logical 1-21
  - of dates 2-44
  - preallocation 10-7, 10-13
- vectorization 10-4
  - example 10-4
  - replacing for
    - vectorization 3-60
- vectors 1-44
- version
  - obtaining 3-10
- version 3-10
- visualizing
  - cell array 2-68

## W

- warning control 7-14
  - saving and restoring state 7-20
- warning control statements
  - message identifiers 7-16
  - output from 7-17
  - output structure array 7-18

- warnings
  - debugging 7-23
  - identifying 7-13
  - syntax 7-14
  - warning control statements 7-15
  - warning states 7-16
- Web content access 6-118
- which 4-34
  - used with methods 8-70
- while 3-65
  - empty arrays 3-66
  - example 3-65
  - syntax 3-65
- white space
  - finding in string 2-32
- whos 1-54
  - interpreting memory use 10-12
- wildcards, in filenames 3-70
- workspace
  - context 4-17
  - of individual functions 4-17
- writing
  - ASCII data 6-36
  - HDF data 6-98
  - in HDF4 format 6-96
  - in HDF5 format 6-65

## Z

- zeros 1-52