

MATLAB[®]

The Language of Technical Computing

Computation

Visualization

Programming

MAT-File Format

Version 5

The MathWorks

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

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

MAT-File Format

© COPYRIGHT 1984 - 1999 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 or for the federal government of the United States. By accepting delivery of the Program, the government hereby agrees that this software qualifies as "commercial" computer software within the meaning of FAR Part 12.212, DFARS Part 227.7202-1, DFARS Part 227.7202-3, DFARS Part 252.227-7013, and DFARS Part 252.227-7014. The terms and conditions of The MathWorks, Inc. Software License Agreement shall pertain to the government's use and disclosure of the Program and Documentation, and shall supersede any conflicting contractual terms or conditions. If this license fails to meet the government's minimum needs or is inconsistent in any respect with federal procurement law, the government agrees to return the Program and Documentation, unused, to MathWorks.

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

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

Printing History: June 1999 Online only New for MATLAB 5.3

MAT-File Format

- Introduction 2**
- MAT-File Formats 2

- Version 5 MAT-File Format 3**
- MAT-File Header Format 5
- Data Element Format 6

- Version 5 MATLAB Array Data Element Formats10**
- Numeric Array and Character Array Data Element Formats . .11
- Sparse Array Data Element Format18
- Cell Array Data Element Format22
- Structure MAT-File Data Element Format25
- MATLAB Object MAT-File Data Element Format29

- Version 4 MAT-File Format32**

Introduction

This document describes the internal format of MATLAB® Version 5 and Version 4 MAT-files. A MAT-file stores data in binary (not human-readable) form. In MATLAB, you create MAT-files by using the `save` command, which writes the arrays currently in memory to a file as a continuous byte stream. By convention, this file has the filename extension `.mat`; thus the name MAT-file. The `load` command reads the arrays from a MAT-file into the MATLAB workspace.

Most MATLAB users do not need to know the internal format of a MAT-file. Even users who must read and write MAT-files from C and Fortran programs do not need to know the MAT-file format if they use the MATLAB Application Program Interface (API). This API shields users from dependence on the details of the MAT-file format.

However, if you need to read or write MAT-files on a system for which the MATLAB API library is not supported, you must write your own read and write routines. The MATLAB API is only available for platforms on which MATLAB is supported. This document provides the details about the MAT-file format you will need to read and write MAT-files on these systems.

Note Whenever possible, The MathWorks strongly advises you to use the MAT-file API routines to read and write MAT-files. Any code you write that depends on the MAT-file format will very likely need to be rewritten when the format changes in future releases.

MAT-File Formats

This document describes both the MATLAB Version 5 and Version 4 MAT-file formats. The MATLAB Version 5 MAT-file format supports all the MATLAB V5 array types, including multidimensional numeric arrays, character arrays, sparse arrays, cell arrays, structures, and objects. “Version 5 MAT-File Format” on page -3 describes this format.

The MATLAB Version 4 MAT-file format is a simpler format but it only supports two-dimensional matrices and character strings. “Version 4 MAT-File Format” on page -32 describes this format.

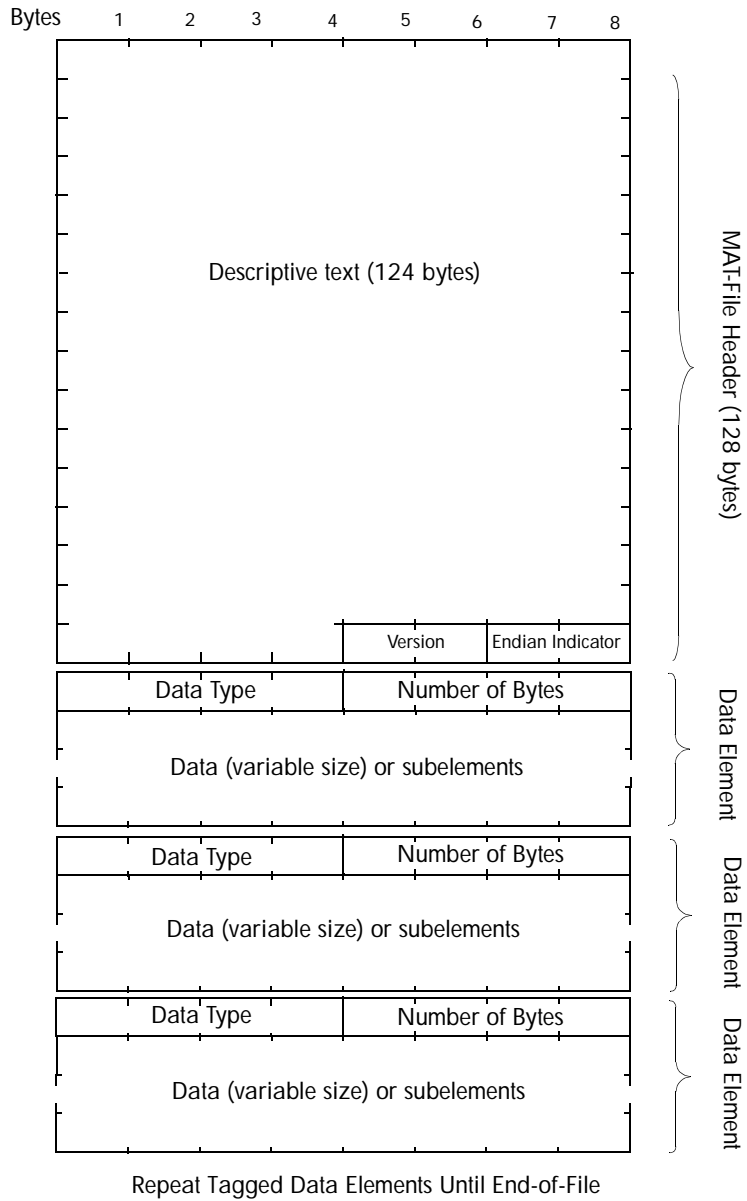
Version 5 MAT-File Format

In MATLAB Version 5, a MAT-file is made up of a 128-byte *header* followed by one or more *data elements*. Each data element is composed of an 8-byte *tag* followed by the data in the element. The tag specifies the number of bytes in the data element and how these bytes should be interpreted; that is, should the bytes be read as 16-bit values, 32-bit values, floating point values or some other data type.

By using tags, the Version 5 MAT-file format provides quick access to individual data elements within a MAT-file. You can move through a MAT-file by finding a tag and then skipping ahead the specified number of bytes until the next tag.

“MATLAB Version 5 MAT-File Format” on page -4 graphically illustrates this MAT-file format. The sections that follow provide more details about these MAT-file elements.

Figure 1: MATLAB Version 5 MAT-File Format



MAT-File Header Format

MATLAB Version 5 MAT-files begin with a 128-byte header made up of a 124 byte text field and two, 16-bit flag fields.

Header Text Field

The first 124 bytes of the header can contain text data in human-readable form. This text typically provides information that describes how the MAT-file was created. For example, MAT-files created by MATLAB include the following information in their headers:

- MATLAB version
- Platform on which the file was created
- Date and time the file was created

You can view the text in a MAT-file header using the `cat` command on UNIX systems, or the `type` command on a PC. The output displays the text in this part of the header. (The display of the header is followed by unreadable characters representing the binary data in the file.)

```
cat my_matfile.mat
MATLAB 5.0 MAT-file, Platform: SOL2, Created on: Thu Nov 13
10: 10: 27 1997
```

Programming Note When creating a MAT-file, you *must* write data in the first four bytes of this header. MATLAB uses these bytes to determine if a MAT-file uses a Version 5 format or a Version 4 format. If any of these bytes contain a zero, MATLAB will incorrectly assume the file is a Version 4 MAT-file.

Header Flag Fields

The last four bytes in the header are divided into two, 16-bit flag fields (int 16).


Field	Value
Version	Identifies the version of the MATLAB software used to create the MAT-file format. When creating a MAT-file, set this field to 0x0100.
Endian Indicator	Contains the two characters, M and I, written to the MAT-file in this order, as a 16-bit value. If, when read from the MAT-file as a 16-bit value, the characters appear in reversed order (IM rather than MI), it indicates that the program reading the MAT-file must perform byte-swapping to interpret the data in the MAT-file correctly.

Programming Note Programs that create MAT-files always write data in their native machine format. Programs that read MAT-files are responsible for byte-swapping.

Data Element Format

Each data element begins with an 8-byte tag followed immediately by the data in the element. Figure 2 shows this format. (MATLAB also supports a compressed data element format. See page -9 for more information.)

Figure 2: MAT-File Data Element Format



Tag

The 8-byte data element tag is composed of two, 32-bit fields:

- Data Type
- Number of Bytes

Data Type. The Data Type field specifies how the data in the element should be interpreted, that is, its size and format. The MAT-file format supports many data types including signed and unsigned, 8-bit, 16-bit, 32-bit, and 64-bit data types and a special data type that represents MATLAB arrays. Table 1 lists all these data types with the values used to specify them. The table also includes symbols that are used to represent these data types in the examples in this document.

Table 1: MAT-File Data Types

MAT-File Data Type	Value	Symbol
8 bit, signed	1	mi INT8
8 bit, unsigned	2	mi UINT8
16-bit, signed	3	mi INT16
16-bit, unsigned	4	mi UINT16
32-bit, signed	5	mi INT32
32-bit, unsigned	6	mi UINT32
IEEE 754 single format	7	mi SINGLE
Reserved	8	--
IEEE 754 double format	9	mi DOUBLE
Reserved	10	--
Reserved	11	--
64-bit, signed	12	mi INT64

Table 1: MAT-File Data Types

MAT-File Data Type	Value	Symbol
64-bit, unsigned	13	mi UI NT64
MATLAB array	14	mi MATRIX (For more information about this data type, see “Version 5 MATLAB Array Data Element Formats” on page -10.)

Number of Bytes. The Number of Bytes field is a 32-bit value that specifies the number of bytes of data in the element. This value does not include the eight bytes of the data element’s tag.

Data

The data immediately follows the tag. The data must be aligned on 64-bit boundaries. When writing a MAT-file, if the amount of data in a data element falls short of a 64-bit boundary, you must add bytes of padding to make sure the tag of the next data element falls on a 64-bit boundary. Likewise, when reading data from a MAT-file, be sure to account for these padding bytes.


Programming Note For data elements representing MATLAB arrays, (type `mi MATRIX`), the value of the Number of Bytes field includes padding bytes in the total. For all other MAT-file data types, the value of the Number of Bytes field does *not* include padding bytes.

Example Data Element

For example, Figure 3 illustrates a data element representing an array of six 32-bit, unsigned integers: 1, 2, 3, 4, 5, 6. In the figure, the Data Type field contains the value from Table 1 that specifies unsigned, 32-bit integers (`mi UI NT32`). The Number of Bytes field in the data element tag contains the number of data values multiplied by the number of bytes used to represent

each value. Note that this value does not include the eight bytes in the data element tag.


Figure 3: Example MAT-File Data Element



Compressed Data Element Format

If the data element contains between one and four bytes of data, MATLAB saves storage space by storing the data in a *compressed* data element format. In this format, the Data Type and Number of Bytes fields are stored as 16-bit values, freeing four bytes in the tag in which to store the data. Figure 4 illustrates this format.

Figure 4: Compressed Data Element Format



Programming Note When reading a MAT-file, you can tell if you are processing a compressed data element by comparing the value of the first two bytes of the tag with the value zero (0). If these two bytes are not zero, the tag uses the compressed format. When writing a MAT-file, use of the compressed data element format is optional.

Version 5 MATLAB Array Data Element Formats

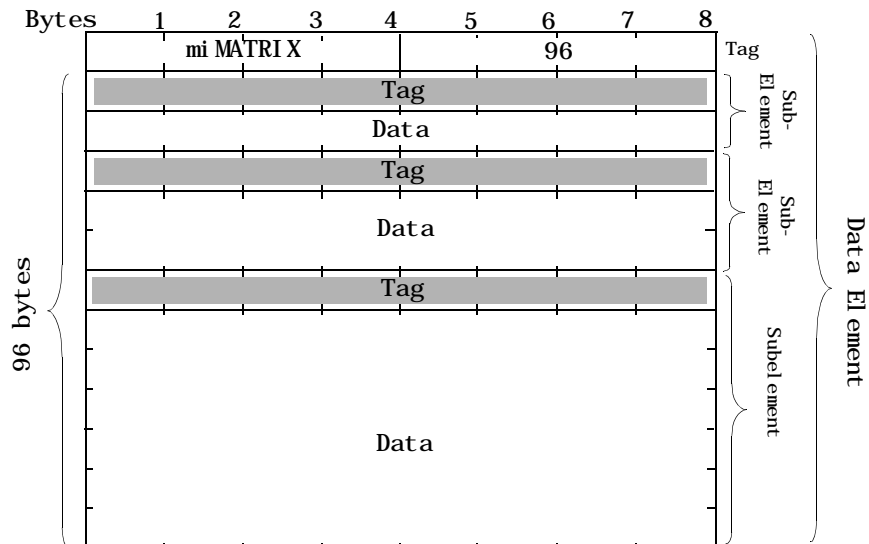
The MAT-file data type `mi MATRIX` (14) is used to represent all types of MATLAB arrays, including:

- Numeric arrays
- Character arrays
- Sparse arrays
- Cell arrays
- Structures
- Objects

The `mi MATRIX` data type is a compound data type. MAT-file data elements of this type are composed of multiple *subelements*. The subelements can be of any other MAT-file data type, including other `mi MATRIX` data types.

Figure 5 shows a `mi MATRIX` data element composed of three subelements. Note how each subelement is a data element with its own tag. The value of the Number of Bytes field (96 in the figure) in the data element tag includes all the subelements.

Figure 5: MATLAB Array Data Element with Subelements



Each `mi MATRIX` data element representing the different types of MATLAB arrays each has a specific set of subelements. Some of these subelements are common to all MATLAB arrays. Others subelements are unique to a particular type of array. The following sections detail the subelements for each MATLAB array type.

Numeric Array and Character Array Data Element Formats

A MAT-file data element representing a MATLAB numeric array or character array is composed of four subelements and one optional subelement. Table 2 lists the subelements in the order in which they appear in the data element. The table also includes the values of the Data Type and Number of Bytes fields you would use in the tag of each subelement. For an example, see “Examples of Numeric Array Data Elements” on page -15.

Table 2: Numeric and Character Array Subelements with Tag Data


Subelement	Data Type	Number of Bytes
Array Flags	<code>mi UINT32</code>	$2 * \text{size-of-Data-Type}$ (8 bytes)
Dimensions Array	<code>mi INT32</code>	$\text{number-of-dimensions} * \text{size-of-Data-Type}$ (To learn how to determine the number of dimensions, see “Dimensions Array Subelement” on page -13.)
Array Name	<code>mi INT8</code>	$\text{number-of-characters} * \text{size-of-Data-Type}$
Real part (pr)	Any of the numeric data types.	$\text{number-of-values} * \text{size-of-Data-Type}$
Imaginary part (pi) (Optional)	Any of the numeric data types.	$\text{number-of-values} * \text{size-of-Data-Type}$

Array Flags Subelement

This subelement identifies the MATLAB array type (class) represented by the data element and provides other information about the array. The Array Flags subelement is common to all array types.

Figure 6 illustrates the format of the Array Flags subelement. (For sparse matrices, bytes 5 through 8 are used to store the maximum number of nonzero elements in the matrix. See “Sparse Array Data Element Format” on page -18 for more information.)

Figure 6: Array Flags Format



Flags. This field contains three, single-bit flags that indicate whether the numeric data is complex, global, or logical. If the complex bit is set, the data element includes an imaginary part (π). If the global bit is set, MATLAB loads the data element as a global variable in the base workspace. If the logical bit is set, it indicates the array is used for logical indexing.

Class. This field contains a value that identifies the MATLAB array type (class) represented by the data element. Table 3 lists the MATLAB array types with the values you use to specify them. The table also includes symbols that are used to represent the MATLAB array type in the examples in this document.

Note The value of the Class field identifies the MATLAB data type. The value of the Data Type field in the data element tag identifies the data type used to store the data *in the MAT-file*. The MAT-file data types are listed in Table 1. The value of the Class and the Data Type fields do not need to be the same; for more information, see “Data Compression” on page -15.

Table 3: MATLAB Array Types (Classes)

MATLAB Array Type (Class)	Value	Symbol
Cell array	1	mxCELL_CLASS
Structure	2	mxSTRUCT_CLASS
Object	3	mxOBJECT_CLASS
Character array	4	mxCHAR_CLASS
Sparse array	5	mxSPARSE_CLASS
Double precision array	6	mxDOUBLE_CLASS
Single precision array	7	mxSINGLE_CLASS
8-bit, signed integer	8	mxINT8_CLASS
8-bit, unsigned integer	9	mxUINT8_CLASS
16-bit, signed integer	10	mxINT16_CLASS
16-bit, unsigned integer	11	mxUINT16_CLASS
32-bit, signed integer	12	mxINT32_CLASS
32-bit unsigned, integer	13	mxUINT32_CLASS

For numeric arrays, Class can contain any of the numeric array types: mxDOUBLE_CLASS, mxSINGLE_CLASS, mxINT8_CLASS, mxUINT8_CLASS, mxINT16_CLASS, mxUINT16_CLASS, mxINT32_CLASS, or mxUINT32_CLASS.

For character arrays, Class contains mxCHAR_CLASS.

Dimensions Array Subelement

This subelement specifies the size of each dimension of an n -dimensional array in an n -sized array of 32-bit values (mi INT32). All numeric arrays have at least two dimensions. The Dimensions Array subelement is common to all MATLAB array types.

For example, if a data element represents a 2-by-3-by-2 MATLAB array, the Dimensions Array subelement would contain three values: 2, 3, and 2.

Programming Note To calculate the number of dimensions in an array, divide the value stored in the Number of Bytes field in the Dimensions Array subelement tag by 4, the number of bytes in the data type (mi INT32) used in the subelement.

Array Name Subelement

This subelement specifies the name assigned to the array, as an array of signed, 8-bit values (mi INT8). This subelement is common to all array types.

Real Part (pr) Subelement

This subelement contains the numeric data in the MATLAB array. If the array contains complex numbers (the complex bit in the Array Flags is set), this is the real part of the number.

The data type of the values can be any of the numeric data types listed in Table 1, MAT-File Data Types, on page -7.

Imaginary Part (pi) Subelement

This subelement contains the imaginary part of the numeric data in the MATLAB array. This subelement is only present if one or more of the numeric values in the MATLAB array is a complex number (if the complex bit is set in Array Flags). The data type of the values can be any of the numeric data types listed in Table 1, MAT-File Data Types, on page -7.

Programming Note When reading a MAT-file, check the value of the Data Type field in the tag of Real Part and Imaginary Part subelements to identify the data type used to store data. Also note that MATLAB reads and writes these values in column-major order.

Data Compression

MATLAB stores the numeric data in an array in double precision format. When MATLAB writes a numeric (or sparse) array to a MAT-file, it uses the smallest possible data type to store the data, both the real and imaginary parts.

For example, if MATLAB determines that the data stored in double precision format can actually be stored in an 8-bit format, it will use an 8-bit data type to store it in a MAT-file. Note, however, that if any of the numeric values in the array requires a 64-bit representation, MATLAB stores all of the data in a 64-bit data type. See “Compressed Data Element” on page -17 for an example.

When you create a MAT-file, compressing data is optional.

Programming Note When MATLAB uses a smaller data type to store data in a MAT-file, the value of the Class field in the Array Flags subelement identifies the original MATLAB data type.

Examples of Numeric Array Data Elements

This section uses examples to illustrate both the compressed and uncompressed numeric array data element formats.

Uncompressed Data Element. Figure 7 shows how this 2-by-2 numeric array, `my_array`, is represented in a MAT-file.

```
my_array = [ 1. 1+1.1i 2 ; 3 4 ]
```

```
my_array =
```


```
1. 1000 + 1. 1000i    2. 0000
3. 0000              4. 0000
```

In the figure, note:

- The data element includes five subelements. Because one of the numeric values in the array is a complex number, the complex bit flag in the Array Flags subelement is set and the Imaginary Part (pi) subelement is included.

- The value of the Number of Bytes field in the data element tag includes all the subelements, but not the eight bytes of the tag itself.

Figure 7: Example Numeric Array MAT-File Data Element



Compressed Data Element. Figure 8 shows how the three-dimensional numeric array in this example, `arr`, is represented in a MAT-file when compression is used to conserve storage space.

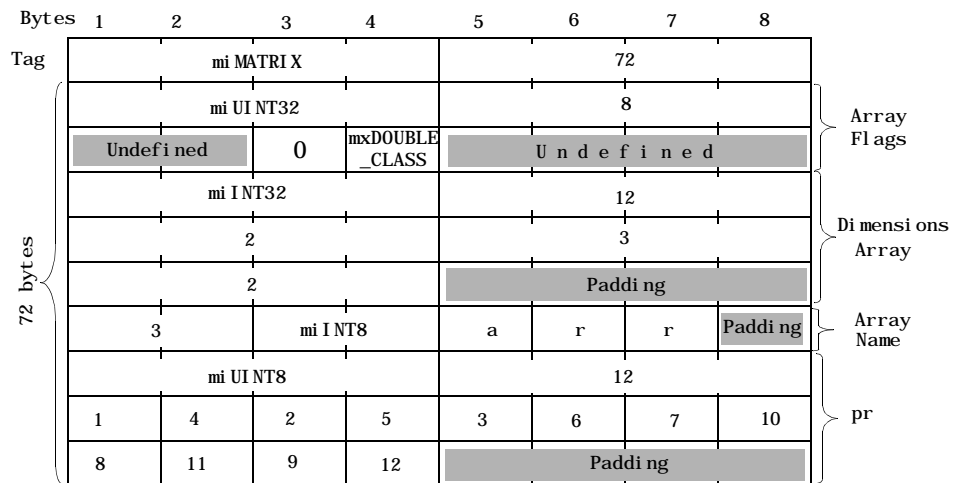
```
A = [ 1 2 3 ; 4 5 6 ];
B = [ 7 8 9 ; 10 11 12];
arr = cat(3,A,B)
arr(:,:,1) =
     1     2     3
     4     5     6

arr(:,:,2) =
     7     8     9
    10    11    12
```

In the figure, note:

- The Array Name subelement uses the compressed data element format.
- The numeric data in the array, stored in double precision format in MATLAB, is stored as 8-bit, unsigned values in the `pr` subelement. The Class field in the Array Flags subelement identifies the original MATLAB data type.

Figure 8: Example Numeric Array MAT-file Data Element (Compressed)



Sparse Array Data Element Format

A MAT-file data element representing a MATLAB sparse array is composed of six subelements and one optional subelement. Table 4 lists the subelements in the order in which they appear in the data element. The table lists the values of the Data Type and Number of Bytes fields of the tag for each subelement.

Table 4: Sparse Array Subelements with Tag Data


Subelement	Data Type	Number of Bytes
Array Flags	mi UINT32	2*size-of-Data-Type (8 bytes)
Dimensions Array	mi INT32	number-of-dimensions*size-of-Data-Type where number-of-dimensions can be 0, 1 or 2.
Array Name	mi INT8	number-of-characters*size-of-Data-Type
Row Index (i r)	mi INT32	nzmax*size-of-Data-Type (The nzmax value is stored in Array Flags.)
Column Index (j c)	mi INT32	(N+1)*sizeof(int32) where N is the second element of the Dimensions array subelement.
Real part (pr)	Any numeric data type	number-of-nonzero-values*size-of-Data-Type
Imaginary part (pi) (Optional)	Any numeric data type	number-of-nonzero-values*size-of-Data-Type

Array Flags Subelement

This subelement identifies the MATLAB array type (class) represented by the data element and provides other information about the array. The Array Flags subelement is common to all array types.

Figure 9 shows the Array Flags format. For sparse arrays, this value also contains the maximum number of non-zero elements in the array (nzmax).

Figure 9: Array Flags Format for Sparse Arrays



Flag. For more information, see “Flags” on page -12.

Class. This field contains a value that identifies the MATLAB data type represented by the data element. For sparse arrays, Class contains the value 5 (mxSPARSE_CLASS). See “Class” on page -12 for more information.

Dimensions Array Subelement

This subelement specifies the size of each dimension of the array. This subelement is common to all array types. For more information, see “Dimensions Array Subelement” on page -13.

Note that MATLAB only supports two-dimensional sparse arrays.

Array Name Subelement

This subelement specifies the name assigned to the array. This subelement is common to all array types. For more information, see “Array Name Subelement” on page -14.

Row Index for Non-zero Values (ir) Subelement

This subelement specifies the row indices of the non-zero elements in the real part (pr) of the matrix data and the imaginary part (pi) of the matrix data, if present. This subelement is a series of 32-bit (mi INT32) values.

Column Index for Non-Zero Values (jc) Subelement

This subelement contains column index information as a series of 32-bit (mi INT32) values. For more information about what this subelement contains, see the *MATLAB Application Program Interface Guide*.

Real Part (pr) Subelement

This subelement contains the numeric data in the MATLAB array. If the array contains complex numbers (the complex bit in the Array Flags is set), this is the real part of the number.

Because MATLAB uses data compression to save storage space, the data type of the values can be any of the numeric data types listed in Table 1, MAT-File Data Types, on page -7. For more information, see “Data Compression” on page -15.

Imaginary Part (pi) Subelement

This subelement contains the imaginary data in the array, if one or more of the numeric values in the MATLAB array is a complex number (if the complex bit is set in Array Flags).

Because MATLAB uses data compression to save storage space, the data type of the values can be any of the numeric data types listed in Table 1, MAT-File Data Types, on page -7. For more information, see “Data Compression” on page -15.

Programming Note You must check the value of the Data Type field in the tag of Real Part and Imaginary Part subelements to identify the type of the data. Also note that MATLAB reads and writes these values in column-major order.

Example Sparse Array

Figure 10 illustrates the MAT-file data element format of this 3-by-3 sparse matrix:

```
a = [ 1 2 3 ];
S = sparse(a, a, a+. 5)

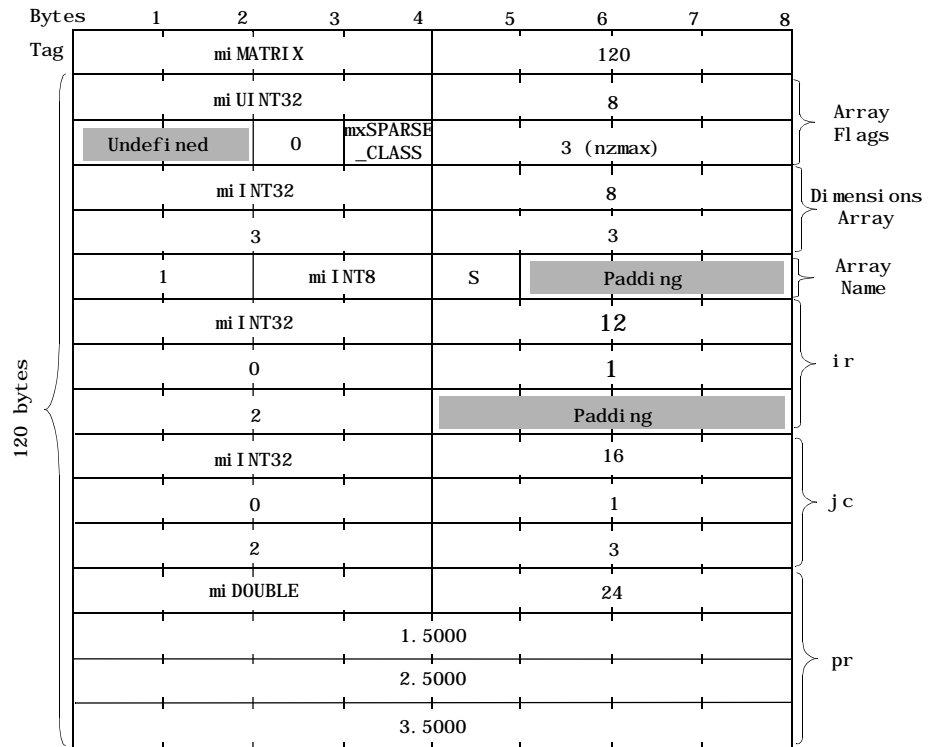
S =

    (1, 1)    1.5000
    (2, 2)    2.5000
    (3, 3)    3.5000
```

In the figure, note:

- The data element contains six subelements.
- The value of the Number of Bytes field in the data element tag includes all the subelements, but not the eight bytes of the tag itself.
- Bytes 5 through 8 of the Array Flags subelement contain the maximum number of non-zero elements (nzmax) in the sparse array.
- The Array Name subelement uses the compressed data element format.

Figure 10: Example Sparse Array MAT-file Data Element



Cell Array Data Element Format

A MAT-file data element representing a MATLAB cell array is composed of four subelements. Table 5 lists the subelements in the order in which they appear in the data element. The table lists the values of the Data Type and Number of Bytes fields of the tag for each subelement.


Table 5: Cell Array Subelements with Tag Data

Subelement	Data Type	Number of Bytes
Array Flags	mi UINT32	2*size-of-Data-Type (8 bytes)
Dimensions Array	mi INT32	number-of-dimensions*size-of-Data-Type
Array Name	mi INT8	number-of-characters*size-of-Data-Type
Cells	Each cell is written in place as an mi MATRIX element.	

Array Flags Subelement

This subelement identifies the MATLAB array type (class) represented by the data element and provides other information about the array. Figure 11 shows the Array Flags format. The Array Flags subelement is common to all array types.

Figure 11: Array Flags Format



Flags. See “Flags” on page -12 for more information.

Class. This field contains a value that identifies the MATLAB data type represented by the data element. For cell arrays, Class contains the value 1 (mxCELL_CLASS). For more information, see “Class” on page -12.

Dimensions Array Subelement

This subelement specifies the size of each dimension of the array. This subelement is common to all array types. For more information, see “Dimensions Array Subelement” on page -13.

Array Name Subelement

This subelement specifies the name assigned to the array. This subelement is common to all array types. For more information, see “Array Name Subelement” on page -14.

Cells Subelement

This subelement contains the value stored in a cell. These values are MATLAB arrays, represented using the `mi MATRIX` format specific to the array type: numeric array, sparse array, structure, object or other cell array. See the appropriate section in this document for details about the MAT-file representation of each of these array types. Cells are written in column-major order.

Example Cell Array

Figure 12 illustrates the MAT-file data element format of this cell array:

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

B = [ 7 8 9 ; 10 11 12 ]
B =
     7     8     9
    10    11    12

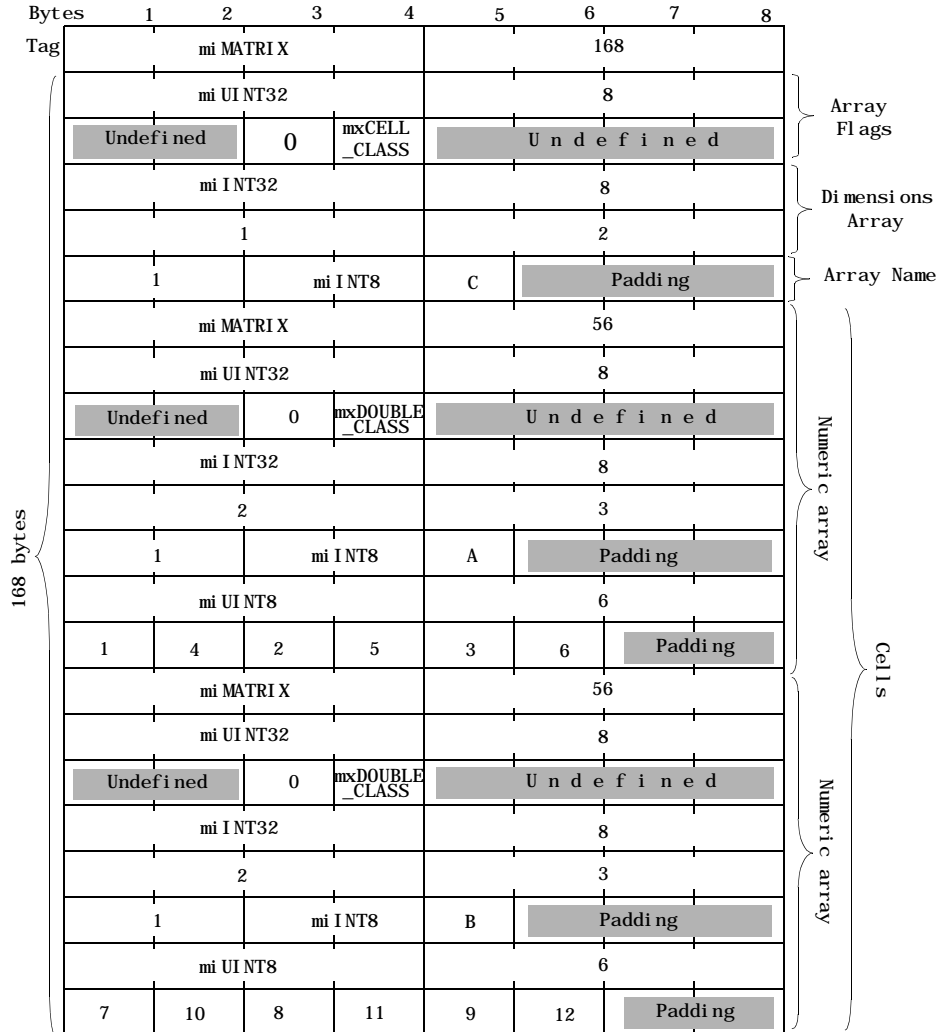
C = { A, B }
C =
    [2x3 double]    [2x3 double]
```

In the figure, note:

- The data element contains five subelements, the three common subelements; Array Flags, Dimensions and Array Name; and two cell subelements.
- The value of the Number of Bytes field in the data element tag includes all the subelements, but not the eight bytes of the tag itself.
- Each cell subelement is an `mi MATRIX` type. In the example, each cell contains a numeric array. For more information about the format of these elements,

see “Numeric Array and Character Array Data Element Formats” on page -11.

Figure 12: Example Cell Array Data Element



Structure MAT-File Data Element Format

A MAT-file data element representing a MATLAB structure is composed of six subelements. Table 6 lists the subelements in the order in which they appear in the data element. The table lists the values of the Data Type and Number of Bytes fields of the tag for each subelement.


Table 6: Structure Subelements with Tag Data

Subelements	Data Type	Number of Bytes
Array Flags	mi UI NT32	2*si ze- of- Data- Type (8 bytes)
Dimensions Array	mi I NT32	number- of- di mensi ons*si ze- of- Dat a- Type
Array Name	mi I NT8	number- of- characters*si ze- of- Dat a- Type
Field Name Length	mi I NT32	si ze- of- Dat a- Type (4 bytes)
Field Names	mi I NT8	number- of- fi el ds*Fi el d- Name- Length
Fields	Each field is written in place as an array. Fields are written in column order.	

Array Flags Subelement

This subelement identifies the MATLAB array type (class) represented by the data element and provides other information about the array. Figure 13 shows the Array Flags format. The Array Flags subelement is common to all array types.

Figure 13: Array Flags Format



Flags. See “Flags” on page -12 for more information.

Class. This field contains a value that identifies the MATLAB data type represented by the data element. For structures, Class contains the value 2 (mxSTRUCT_CLASS). For more information, see “Class” on page -12.

Dimensions Array Subelement

This subelement specifies the size of each dimension of the array. This subelement is common to all array types. For more information, see “Dimensions Array Subelement” on page -13.

Array Name Subelement

This subelement specifies the name assigned to the structure. This subelement is common to all array types. For more information, see “Array Name Subelement” on page -14.

Field Name Length Subelement

This subelement specifies the maximum length of a Field Name. MATLAB sets this limit to 32 (31 characters and a NULL terminator). In a MAT-file created by MATLAB, this subelement always uses the compressed data element format.

Field Names Subelement

This subelement specifies the name of each field in the structure as a series of 8-bit (mi INT8) character arrays. The value of the Field Name Length subelement determines the length of each field name array (32 bytes). Field names must be NULL-terminated.

Fields Subelement

This subelement contains the value stored in a field. These values are MATLAB arrays, represented using the mi MATRIX format specific to the array type: numeric array, sparse array, cell, object or other structure. See the appropriate section of this document for details about the MAT-file format of each of these array type. MATLAB reads and writes these fields in column-major order.

Example

Figure 14 illustrates the MAT-file data element format for this MATLAB structure:

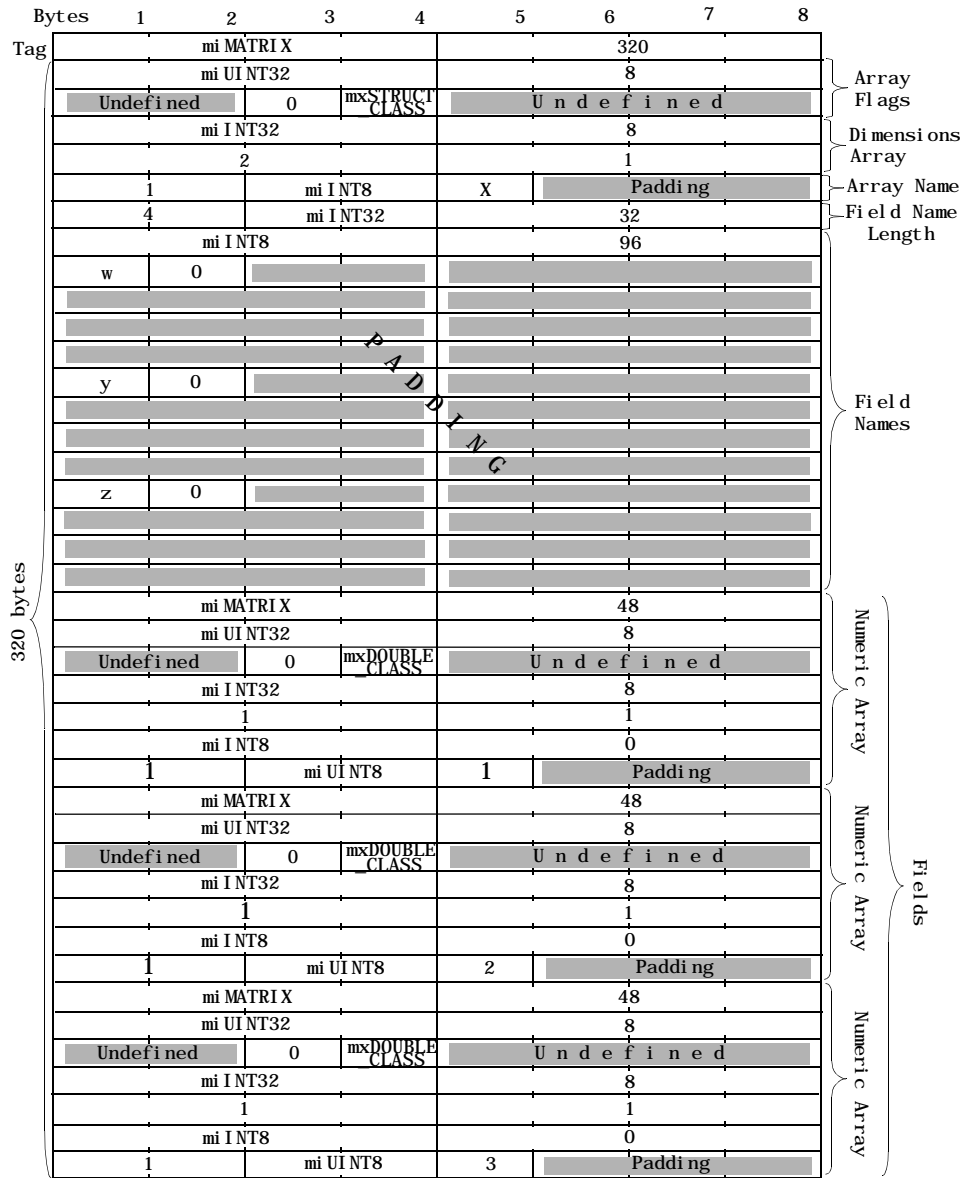
```
X.w = [1];  
X.y = [2];  
X.z = [3];  
X  
X =
```

w: 1
y: 2
z: 3

In the figure, note:

- The data element contains eight subelements: the three common subelements (Array Flags, Dimensions and Array Name) and five structure-specific subelements (Field Name Length, Field Names, and three Field subelements).
- The value of the Number of Bytes field in the data element tag includes all the subelements, but not the eight bytes of the tag itself.
- The Field Names subelement allocates 32 bytes of storage for each field name. A NULL terminator indicates the end of each field name.
- Each Field subelement is an `mi MATRIX` data type. In the example, each field contains a numeric array. For more information about the format of these elements, see “Numeric Array and Character Array Data Element Formats” on page -11.
- Each of the numeric arrays contain zero-length Array Name subelements. The Field Names subelement contains the names of the numeric arrays.

Figure 14: Example Structure MAT-File Data Element



MATLAB Object MAT-File Data Element Format

A MAT-file data element representing a MATLAB object is composed of seven subelements. Table 7 lists the subelements in the order in which they appear in the data element. An object data element has the same subelements as a structure with the addition of the Class Name subelement. The table lists the values of the Data Type and Number of Bytes fields of the tag for each subelement.


Table 7: MATLAB Object Subelements with Tag Data

Subelement	Data Type	Number of Bytes
Array Flags	mi UI NT32	2*si ze-of-Data-Type (8 bytes)
Dimensions Array	mi I NT32	number-of-di mensi ons*si ze-of-Data-Type
Array Name	mi I NT8	number-of-characters*si ze-of-Data-Type
Class Name	mi I NT8	number-of-characters*si ze-of-Data-Type
Field Name Length	mi I NT32	si ze-of-Data-Type (4 bytes)
Field Names	mi I NT8	number-of-fi el ds*Fi el d-Name-Length
Fields	Each field is written in place as an array.	

Array Flags Subelement

This subelement identifies the MATLAB array type (class) represented by the data element and provides other information about the array. Figure 15 shows the Array Flags format. The Array Flags subelement is common to all array types.

Figure 15: Array Flags Format



Flags. See “Flags” on page -12 for more information.

Class. This field contains a value that identifies the MATLAB data type represented by the data element. For objects, the Class byte has the value 3 (mxOBJECT_CLASS). For more information, see “Class” on page -12.

Dimensions Array Subelement

This subelement specifies the size of each dimension of the array. This subelement is common to all array types. For more information, see “Dimensions Array Subelement” on page -13.

Array Name Subelement

This subelement specifies the name assigned to the array. This subelement is common to all array types. For more information, see “Array Name Subelement” on page -14.

Class Name Subelement

This subelement specifies the name assigned to the object class. This subelement is an array of 8-bit characters (mi INT8).

Field Name Length Subelement

This subelement specifies the maximum length of a Field Name. See “Field Name Length Subelement” on page -26 for more information.

Field Names Subelement

This subelement specifies the name of each field in the structure. See “Field Names Subelement” on page -26 for more information.

Fields Subelement

This subelement contains the value stored in a field. See “Fields Subelement” on page -26 for more information.

Example

Figure 16 illustrates how the MATLAB object in this example is represented in a MAT-file.


```
X = inline('t^2');
```

The figure only shows the first four subelements of the object. For an example that shows the remaining subelements, see “Example” on page -26.

In the figure, note:

- The Array Flag Class byte is set to `mxOBJECT_CLASS`.
- The data element includes the Class Name subelement.

Figure 16: Example Object MAT-file Data Element



Version 4 MAT-File Format

Note This section is taken from the MATLAB V4.2 *External Interface Guide*, which is no longer available in printed form.

This section presents the internal structure of Level 1.0 MAT-files. This information is provided to enable users to read and write MAT-files on machines for which the MAT-file access routine library is not available. It is not needed when using the MAT-file subroutine library to read and write MAT-files, and we strongly advise that you do use the External Interface Library if it is available for all of the machines that you are working with.

A MAT-file may contain one or more matrices. The matrices are written sequentially on disk, with the bytes forming a continuous stream. Each matrix starts with a fixed-length 20-byte header that contains information describing certain attributes of the Matrix. The 20-byte header consists of five long (4-byte) integers:

Table 8: MATLAB Version 4 MAT-File Matrix Header Format

Field	Description
type	The type flag contains an integer whose decimal digits encode storage information. If the integer is represented as MOPT where M is the thousands digit, O is the hundreds digit, P is the tens digit, and T is the ones digit, then:
	M indicates the numeric format of binary numbers on the machine that wrote the file. Use this table to determine the number to use for your machine:
0	IEEE Little Endian (PC, 386, 486, DEC Risc)
1	IEEE Big Endian (Macintosh, SPARC, Apollo, SGI, HP 9000/300, other Motorola)
2	VAX D-float
3	VAX G-float
4	Cray

Table 8: MATLAB Version 4 MAT-File Matrix Header Format

	0 is always 0 (zero) and is reserved for future use.	
	P indicates which format the data is stored in according to the following table:	
0	double-precision (64-bit) floating point numbers	
1	single-precision (32-bit) floating point numbers	
2	32-bit signed integers	
3	16-bit signed integers	
4	16-bit unsigned integers	
5	8-bit unsigned integers	
	The precision used by the save command depends on the size and type of each matrix. Matrices with any noninteger entries and matrices with 10,000 or fewer elements are saved in floating point formats requiring 8 bytes per real element. Matrices with all integer entries and more than 10,000 elements are saved in the following formats, requiring fewer bytes per element.	
	Element range	Bytes per element
	[0:255]	1
	[0:65535]	2
	[-32767:32767]	2
	$[-2^{31}+1:2^{31}-1]$	4
	other	8
	T indicates the matrix type according to the following table:	
0	Numeric (Full) matrix	
1	Text matrix	
2	Sparse matrix	
	Note that the elements of a text matrix are stored as floating point numbers between 0 and 255 representing ASCII-encoded characters.	

Table 8: MATLAB Version 4 MAT-File Matrix Header Format

<code>mrows</code>	The row dimension contains an integer with the number of rows in the matrix.
<code>ncols</code>	The column dimension contains an integer with the number of columns in the matrix.
<code>imagf</code>	The imaginary flag is an integer whose value is either 0 or 1. If 1, then the matrix has an imaginary part. If 0, there is only real data.
<code>namlen</code>	The name length contains an integer with 1 plus the length of the matrix name.

Immediately following the fixed length header is the data whose length is dependent on the variables in the fixed length header:

Table 9: MATLAB Version 4 MAT-File Matrix Data Format

Field	Description
<code>name</code>	The matrix name consists of <code>namlen</code> ASCII bytes, the last one of which must be a null character (' <code>\0</code> ').
<code>real</code>	Real part of the matrix consists of <code>mrows * ncols</code> numbers in the format specified by the <code>P</code> element of the type flag. The data is stored column-wise such that the second column follows the first column, etc.
<code>imag</code>	Imaginary part of the matrix, if any. If the imaginary flag <code>imagf</code> is nonzero, the imaginary part of a matrix is placed here. It is stored in the same manner as the real data.

This structure is repeated for each matrix stored in the file.

The following C language code demonstrates how to write a single matrix to disk in Level 1.0 MAT-file format.

```
#include <stdio.h>

main()
{
    typedef struct {
        long type;
        long mrows;
        long ncols;
    }
```

```
long imagf;
long namelen;
} Fmatrix;

char *pname;
double *pr;
double *pi;
Fmatrix x;
int mn;

FILE *fp;

double real_data = 1.0;
double imag_data = 2.0;

fp=fopen("mymatfile.mat", "wb");

if(fp==NULL)
    printf("File could not be opened.\n");
else
{
    pname = "x";
    x.type = 1000;
    x.mrows = 1;
    x.ncols = 1;
    x.imagf = 1;
    x.namelen = 2;

    pr = &real_data;
    pi = &imag_data;

    fwrite(&x, sizeof(Fmatrix), 1, fp);

    fwrite(pname, sizeof(char), x.namelen, fp);

    mn = x.mrows *x.ncols;

    fwrite(pr, sizeof(double), mn, fp);

    if(x.imagf)
```

```
        fwrite(pi, sizeof(double), mn, fp);  
    }  
  
    fclose(fp);  
}
```

Again, we strongly advise against using this approach, and recommend that you instead use the MAT-file access routines provided in the External Interface Library. You will need to write your own C code as shown above only if you do not have the MAT-file access routines for the particular platform on which you need to read and write MAT-files.

A

- array flags subelement 11
 - in sparse arrays 18
- array name subelement 14

B

- byte swapping 6

C

- cell arrays
 - example 23
 - MAT-file format 22
- character arrays
 - MAT-file format 11
- classes
 - MATLAB arrays 13
- complex numbers
 - in MAT-files 12
- compression
 - data element format 9
 - example in data element 17
 - in data element 15

D

- data elements
 - alignment 8
 - cell arrays 22
 - character array format 11
 - compressed format 9
 - defined 3
 - format 6
 - MATLAB arrays 10
 - numeric array format 11
 - objects 29

- padding bytes 8
- sparse array format 18
- structures 25
- subelements 10

- data types
 - changed by compression 15
 - MAT-file vs. MATLAB 12
 - used in MAT-files 7
- dimensions
 - determining number of 14
- dimensions array subelement 13

E

- Endian indicator 6

F

- field name length
 - in structure data elements 26
- field names
 - in structure data elements 26
- flags
 - format 12

G

- global variables
 - in MAT-files 12

H

- header
 - defined 3
 - flag fields 6
 - format 5

text field 5

I

IEEE 754 double format 7

imaginary data
in data element 14

L

logical arrays
in MAT-files 12

M

MAT-files (V4)
distinguishing from V5 format 5
format 32

MAT-files (V5)
data types 7
distinguishing from V4 format 5
header format 5
header text field 5
numeric array data elements 11
overall format 3
version field 6

MATLAB array types 12, 13
data element formats 10

mi MATRIX format 10

N

Number of Bytes field
tag 8

numeric array
compressed example 17
example 15

numeric arrays
MAT-file format 11

O

objects
MAT-file format 29

P

padding bytes
data elements 8
including in Number of Bytes total 8

pi 14

pr 14

R

real data
in data element 14

S

sparse arrays
example 20
in MAT-file 18

structures
example 26
MAT-file format 25

subelements
defined 10

T

tags
defined 3
format 7
number of bytes field 8

V

version field

 MAT-file V5 format 6