

[SciPy.org](https://scipy.org/) (<https://scipy.org/>) [Docs](https://docs.scipy.org/) (<https://docs.scipy.org/>)
[NumPy v1.17 Manual](https://docs.scipy.org/doc/numpy/index.html) (<https://docs.scipy.org/doc/numpy/index.html>)
[NumPy User Guide](https://docs.scipy.org/doc/numpy/user/index.html) (<https://docs.scipy.org/doc/numpy/user/index.html>)
[index](https://docs.scipy.org/doc/numpy/genindex.html) (<https://docs.scipy.org/doc/numpy/genindex.html>) [next](https://docs.scipy.org/doc/numpy/user/building.html) (<https://docs.scipy.org/doc/numpy/user/building.html>)
[previous](https://docs.scipy.org/doc/numpy/user/misc.html) (<https://docs.scipy.org/doc/numpy/user/misc.html>)

NumPy for Matlab users

Introduction

MATLAB® and NumPy/SciPy have a lot in common. But there are many differences. NumPy and SciPy were created to do numerical and scientific computing in the most natural way with Python, not to be MATLAB® clones. This page is intended to be a place to collect wisdom about the differences, mostly for the purpose of helping proficient MATLAB® users become proficient NumPy and SciPy users.

Some Key Differences

In MATLAB®, the basic data type is a multidimensional array of double precision floating point numbers. Most expressions take such arrays and return such arrays. Operations on the 2-D instances of these arrays are designed to act more or less like matrix operations in linear algebra.	In NumPy the basic type is a multidimensional <code>array</code> . Operations on these arrays in all dimensionalities including 2D are element-wise operations. One needs to use specific functions for linear algebra (though for matrix multiplication, one can use the <code>@</code> operator in python 3.5 and above).
MATLAB® uses 1 (one) based indexing. The initial element of a sequence is found using <code>a(1)</code> . See note INDEXING	Python uses 0 (zero) based indexing. The initial element of a sequence is found using <code>a[0]</code> .
MATLAB®'s scripting language was created for doing linear algebra. The syntax for basic matrix operations is nice and clean, but the API for adding GUIs and making full-fledged applications is more or less an afterthought.	NumPy is based on Python, which was designed from the outset to be an excellent general-purpose programming language. While Matlab's syntax for some array manipulations is more compact than NumPy's, NumPy (by virtue of being an add-on to Python) can do many things that Matlab just cannot, for instance dealing properly with stacks of matrices.
In MATLAB®, arrays have pass-by-value semantics, with a lazy copy-on-write scheme to prevent actually creating copies until they are actually needed. Slice operations copy parts of the array.	In NumPy arrays have pass-by-reference semantics. Slice operations are views into an array.

'array' or 'matrix'? Which should I use?

Historically, NumPy has provided a special matrix type, `np.matrix`, which is a subclass of `ndarray` which makes binary operations linear algebra operations. You may see it used in some existing code instead of `np.array`. So, which one to use?

Short answer

Use arrays.

- They are the standard vector/matrix/tensor type of numpy. Many numpy functions return arrays, not matrices.
- There is a clear distinction between element-wise operations and linear algebra operations.
- You can have standard vectors or row/column vectors if you like.

Until Python 3.5 the only disadvantage of using the array type was that you had to use `dot` instead of `*` to multiply (reduce) two tensors (scalar product, matrix vector multiplication etc.). Since Python 3.5 you can use the matrix multiplication `@` operator.

Given the above, we intend to deprecate `matrix` eventually.

Long answer

NumPy contains both an `array` class and a `matrix` class. The `array` class is intended to be a general-purpose n-dimensional array for many kinds of numerical computing, while `matrix` is intended to facilitate linear algebra computations specifically. In practice there are only a handful of key differences between the two.

- Operators `*` and `@`, functions `dot()`, and `multiply()`:
 - For `array`, `**` means element-wise multiplication, while `@` means matrix multiplication; they have associated functions `multiply()` and `dot()`. (Before python 3.5, `@` did not exist and one had to use `dot()` for matrix multiplication).
 - For `matrix`, `**` means matrix multiplication, and for element-wise multiplication one has to use the `multiply()` function.
- Handling of vectors (one-dimensional arrays)

- For array, the **vector shapes 1xN, Nx1, and N are all different things**. Operations like `A[:, 1]` return a one-dimensional array of shape N, not a two-dimensional array of shape Nx1. Transpose on a one-dimensional array does nothing.
 - For matrix, **one-dimensional arrays are always upconverted to 1xN or Nx1 matrices** (row or column vectors). `A[:, 1]` returns a two-dimensional matrix of shape Nx1.
- Handling of higher-dimensional arrays (`ndim > 2`)
 - array objects **can have number of dimensions > 2**;
 - matrix objects **always have exactly two dimensions**.
- Convenience attributes
 - array **has a .T attribute**, which returns the transpose of the data.
 - matrix **also has .H, .I, and .A attributes**, which return the conjugate transpose, inverse, and `asarray()` of the matrix, respectively.
- Convenience constructor
 - The array constructor **takes (nested) Python sequences as initializers**. As in, `array([[1, 2, 3], [4, 5, 6]])`.
 - The matrix constructor additionally **takes a convenient string initializer**. As in `matrix("1 2 3; 4 5 6")`.

There are pros and cons to using both:

- array
 - `:` Element-wise multiplication is easy: `A*B`.
 - `:` (You have to remember that matrix multiplication has its own operator, `@`.
 - `:` You can treat one-dimensional arrays as *either* row or column vectors. `A @ v` treats `v` as a column vector, while `v @ A` treats `v` as a row vector. This can save you having to type a lot of transposes.
 - `:` array is the “default” NumPy type, so it gets the most testing, and is the type most likely to be returned by 3rd party code that uses NumPy.
 - `:` Is quite at home handling data of any number of dimensions.
 - `:` Closer in semantics to tensor algebra, if you are familiar with that.
 - `:` All operations (`*`, `/`, `+`, `-` etc.) are element-wise.
 - `:` (Sparse matrices from `scipy.sparse` do not interact as well with arrays.
- matrix
 - `:` Behavior is more like that of MATLAB® matrices.
 - `<:` (Maximum of two-dimensional. To hold three-dimensional data you need an array or perhaps a Python list of matrix.
 - `<:` (Minimum of two-dimensional. You cannot have vectors. They must be cast as single-column or single-row matrices.
 - `<:` (Since array is the default in NumPy, some functions may return an array even if you give them a matrix as an argument. This shouldn't happen with NumPy functions (if it does it's a bug), but 3rd party code based on NumPy may not honor type preservation like NumPy does.
 - `:` `A*B` is matrix multiplication, so it looks just like you write it in linear algebra (For Python \geq 3.5 plain arrays have the same convenience with the `@` operator).
 - `<:` (Element-wise multiplication requires calling a function, `multiply(A,B)`.
 - `<:` (The use of operator overloading is a bit illogical: `*` does not work element-wise but `/` does.
 - Interaction with `scipy.sparse` is a bit cleaner.

The array is thus much more advisable to use. Indeed, we intend to deprecate matrix eventually.

Table of Rough MATLAB-NumPy Equivalents

The table below gives rough equivalents for some common MATLAB® expressions. **These are not exact equivalents**, but rather should be taken as hints to get you going in the right direction. For more detail read the built-in documentation on the NumPy functions.

In the table below, it is assumed that you have executed the following commands in Python:

```
from numpy import *
import scipy.linalg
```

Also assume below that if the Notes talk about “matrix” that the arguments are two-dimensional entities.

General Purpose Equivalents

MATLAB	numpy	Notes
<code>help func</code>	<code>info(func)</code> or <code>help(func)</code> or <code>func?</code> (in lpython)	get help on the function <i>func</i>
<code>which func</code>	see note HELP (https://docs.scipy.org/doc/numpy/user/numpy-for-matlab-users.notes)	find out where <i>func</i> is defined
<code>type func</code>	<code>source(func)</code> or <code>func??</code> (in lpython)	print source for <i>func</i> (if not a native function)
<code>a && b</code>	<code>a and b</code>	short-circuiting logical AND operator (Python native operator); scalar arguments only
<code>a b</code>	<code>a or b</code>	short-circuiting logical OR operator (Python native operator); scalar arguments only

MATLAB	numpy	Notes
<code>1*i</code> , <code>1*j</code> , <code>1i</code> , <code>1j</code>	<code>1j</code>	complex numbers
<code>eps</code>	<code>np.spacing(1)</code>	Distance between 1 and the nearest floating point number.
<code>ode45</code>	<code>scipy.integrate.solve_ivp(f)</code>	integrate an ODE with Runge-Kutta 4,5
<code>ode15s</code>	<code>scipy.integrate.solve_ivp(f, method='BDF')</code>	integrate an ODE with BDF method

Linear Algebra Equivalents

MATLAB	NumPy	Notes
<code>ndims(a)</code>	<code>ndim(a)</code> or <code>a.ndim</code>	get the number of dimensions of an array
<code>numel(a)</code>	<code>size(a)</code> or <code>a.size</code>	get the number of elements of an array
<code>size(a)</code>	<code>shape(a)</code> or <code>a.shape</code>	get the "size" of the matrix
<code>size(a,n)</code>	<code>a.shape[n-1]</code>	get the number of elements of the n-th dimension of array <code>a</code> . (Note that MATLAB® uses 1 based indexing while Python uses 0 based indexing, See note INDEXING)
<code>[1 2 3; 4 5 6]</code>	<code>array([[1.,2.,3.], [4.,5.,6.]])</code>	2x3 matrix literal
<code>[a b; c d]</code>	<code>block([[a,b], [c,d]])</code>	construct a matrix from blocks <code>a</code> , <code>b</code> , <code>c</code> , and <code>d</code>
<code>a(end)</code>	<code>a[-1]</code>	access last element in the 1xn matrix <code>a</code>
<code>a(2,5)</code>	<code>a[1,4]</code>	access element in second row, fifth column
<code>a(2,:)</code>	<code>a[1]</code> or <code>a[1,:]</code>	entire second row of <code>a</code>
<code>a(1:5,:)</code>	<code>a[0:5]</code> or <code>a[:5]</code> or <code>a[0:5,:]</code>	the first five rows of <code>a</code>
<code>a(end-4:end,:)</code>	<code>a[-5:]</code>	the last five rows of <code>a</code>
<code>a(1:3,5:9)</code>	<code>a[0:3][:,4:9]</code>	rows one to three and columns five to nine of <code>a</code> . This gives read-only access.
<code>a([2,4,5],[1,3])</code>	<code>a[ix_([1,3,4],[0,2])]</code>	rows 2,4 and 5 and columns 1 and 3. This allows the matrix to be modified, and doesn't require a regular slice.
<code>a(3:2:21,:)</code>	<code>a[2:21:2,:]</code>	every other row of <code>a</code> , starting with the third and going to the twenty-first
<code>a(1:2:end,:)</code>	<code>a[::2,:]</code>	every other row of <code>a</code> , starting with the first
<code>a(end:-1:1,:)</code> or <code>flipud(a)</code>	<code>a[::-1,:]</code>	<code>a</code> with rows in reverse order
<code>a([1:end 1],:)</code>	<code>a[r_[:len(a),0]]</code>	<code>a</code> with copy of the first row appended to the end
<code>a.'</code>	<code>a.transpose()</code> or <code>a.T</code>	transpose of <code>a</code>
<code>a'</code>	<code>a.conj().transpose()</code> or <code>a.conj().T</code>	conjugate transpose of <code>a</code>
<code>a * b</code>	<code>a @ b</code>	matrix multiply
<code>a .* b</code>	<code>a * b</code>	element-wise multiply
<code>a ./ b</code>	<code>a/b</code>	element-wise divide
<code>a.^3</code>	<code>a**3</code>	element-wise exponentiation
<code>(a>0.5)</code>	<code>(a>0.5)</code>	matrix whose i,j th element is $(a_{ij} > 0.5)$. The Matlab result is an array of 0s and 1s. The NumPy result is an array of the boolean values <code>False</code> and <code>True</code> .
<code>find(a>0.5)</code>	<code>nonzero(a>0.5)</code>	find the indices where $(a) > 0.5$
<code>a(:,find(v>0.5))</code>	<code>a[:,nonzero(v>0.5)[0]]</code>	extract the columns of <code>a</code> where vector $v > 0.5$
<code>a(:,find(v>0.5))</code>	<code>a[:,v.T>0.5]</code>	extract the columns of <code>a</code> where column vector $v > 0.5$
<code>a(a<0.5)=0</code>	<code>a[a<0.5]=0</code>	<code>a</code> with elements less than 0.5 zeroed out
<code>a .* (a>0.5)</code>	<code>a * (a>0.5)</code>	<code>a</code> with elements less than 0.5 zeroed out

MATLAB	NumPy	Notes
<code>a(:) = 3</code>	<code>a[:] = 3</code>	set all values to the same scalar value
<code>y=x</code>	<code>y = x.copy()</code>	numpy assigns by reference
<code>y=x(2,:)</code>	<code>y = x[1,:].copy()</code>	numpy slices are by reference
<code>y=x(:)</code>	<code>y = x.flatten()</code>	turn array into vector (note that this forces a copy)
<code>1:10</code>	<code>arange(1.,11.)</code> or <code>r_[1.:11.]</code> or <code>r_[1:10:10j]</code>	create an increasing vector (see note RANGES)
<code>0:9</code>	<code>arange(10.)</code> or <code>r_[0:10.]</code> or <code>r_[0:9:10j]</code>	create an increasing vector (see note RANGES)
<code>[1:10]'</code>	<code>arange(1.,11.)[:,newaxis]</code>	create a column vector
<code>zeros(3,4)</code>	<code>zeros((3,4))</code>	3x4 two-dimensional array full of 64-bit floating point zeros
<code>zeros(3,4,5)</code>	<code>zeros((3,4,5))</code>	3x4x5 three-dimensional array full of 64-bit floating point zeros
<code>ones(3,4)</code>	<code>ones((3,4))</code>	3x4 two-dimensional array full of 64-bit floating point ones
<code>eye(3)</code>	<code>eye(3)</code>	3x3 identity matrix
<code>diag(a)</code>	<code>diag(a)</code>	vector of diagonal elements of <code>a</code>
<code>diag(a,0)</code>	<code>diag(a,0)</code>	square diagonal matrix whose nonzero values are the elements of <code>a</code>
<code>rand(3,4)</code>	<code>random.rand(3,4)</code> or <code>random.random_sample((3,4))</code>	random 3x4 matrix
<code>linspace(1,3,4)</code>	<code>linspace(1,3,4)</code>	4 equally spaced samples between 1 and 3, inclusive
<code>[x,y]=meshgrid(0:8,0:5)</code>	<code>mgrid[0:9.,0:6.]</code> or <code>meshgrid(r_[0:9.],r_[0:6.])</code>	two 2D arrays: one of x values, the other of y values
	<code>ogrid[0:9.,0:6.]</code> or <code>ix_(r_[0:9.],r_[0:6.])</code>	the best way to eval functions on a grid
<code>[x,y]=meshgrid([1,2,4],[2,4,5])</code>	<code>meshgrid([1,2,4],[2,4,5])</code>	
	<code>ix_([1,2,4],[2,4,5])</code>	the best way to eval functions on a grid
<code>repmat(a, m, n)</code>	<code>tile(a, (m, n))</code>	create m by n copies of <code>a</code>
<code>[a b]</code>	<code>concatenate((a,b),1)</code> or <code>hstack((a,b))</code> or <code>column_stack((a,b))</code> or <code>c_[a,b]</code>	concatenate columns of <code>a</code> and <code>b</code>
<code>[a; b]</code>	<code>concatenate((a,b))</code> or <code>vstack((a,b))</code> or <code>r_[a,b]</code>	concatenate rows of <code>a</code> and <code>b</code>
<code>max(max(a))</code>	<code>a.max()</code>	maximum element of <code>a</code> (with <code>ndims(a)<=2</code> for matlab)
<code>max(a)</code>	<code>a.max(0)</code>	maximum element of each column of matrix <code>a</code>
<code>max(a,[],2)</code>	<code>a.max(1)</code>	maximum element of each row of matrix <code>a</code>
<code>max(a,b)</code>	<code>maximum(a, b)</code>	compares <code>a</code> and <code>b</code> element-wise, and returns the maximum value from each pair
<code>norm(v)</code>	<code>sqrt(v @ v)</code> or <code>np.linalg.norm(v)</code>	L2 norm of vector <code>v</code>
<code>a & b</code>	<code>logical_and(a,b)</code>	element-by-element AND operator (NumPy ufunc) See note LOGICOPS
<code>a b</code>	<code>logical_or(a,b)</code>	element-by-element OR operator (NumPy ufunc) See note LOGICOPS
<code>bitand(a,b)</code>	<code>a & b</code>	bitwise AND operator (Python native and NumPy ufunc)
<code>bitor(a,b)</code>	<code>a b</code>	bitwise OR operator (Python native and NumPy ufunc)
<code>inv(a)</code>	<code>linalg.inv(a)</code>	inverse of square matrix <code>a</code>
<code>pinv(a)</code>	<code>linalg.pinv(a)</code>	pseudo-inverse of matrix <code>a</code>
<code>rank(a)</code>	<code>linalg.matrix_rank(a)</code>	matrix rank of a 2D array / matrix <code>a</code>
<code>a\b</code>	<code>linalg.solve(a,b)</code> if <code>a</code> is square; <code>linalg.lstsq(a,b)</code> otherwise	solution of <code>a x = b</code> for <code>x</code>
<code>b/a</code>	Solve <code>a.T x.T = b.T</code> instead	solution of <code>x a = b</code> for <code>x</code>
<code>[U,S,V]=svd(a)</code>	<code>U, S, Vh = linalg.svd(a), V = Vh.T</code>	singular value decomposition of <code>a</code>

MATLAB	NumPy	Notes
<code>chol(a)</code>	<code>linalg.cholesky(a).T</code>	cholesky factorization of a matrix (<code>chol(a)</code> in matlab returns an upper triangular matrix, but <code>linalg.cholesky(a)</code> returns a lower triangular matrix)
<code>[V,D]=eig(a)</code>	<code>D,V = linalg.eig(a)</code>	eigenvalues and eigenvectors of <code>a</code>
<code>[V,D]=eig(a,b)</code>	<code>D,V = scipy.linalg.eig(a,b)</code>	eigenvalues and eigenvectors of <code>a</code> , <code>b</code>
<code>[V,D]=eigs(a,k)</code>		find the <code>k</code> largest eigenvalues and eigenvectors of <code>a</code>
<code>[Q,R,P]=qr(a,θ)</code>	<code>Q,R = scipy.linalg.qr(a)</code>	QR decomposition
<code>[L,U,P]=lu(a)</code>	<code>L,U = scipy.linalg.lu(a)</code> or <code>LU,P=scipy.linalg.lu_factor(a)</code>	LU decomposition (note: <code>P(Matlab) == transpose(P(numpy))</code>)
<code>conjgrad</code>	<code>scipy.sparse.linalg.cg</code>	Conjugate gradients solver
<code>fft(a)</code>	<code>fft(a)</code>	Fourier transform of <code>a</code>
<code>ifft(a)</code>	<code>ifft(a)</code>	inverse Fourier transform of <code>a</code>
<code>sort(a)</code>	<code>sort(a)</code> or <code>a.sort()</code>	sort the matrix
<code>[b,I] = sortrows(a,i)</code>	<code>I = argsort(a[:,i])</code> , <code>b=a[I,:]</code>	sort the rows of the matrix
<code>regress(y,X)</code>	<code>linalg.lstsq(X,y)</code>	multilinear regression
<code>decimate(x, q)</code>	<code>scipy.signal.resample(x, len(x)/q)</code>	downsample with low-pass filtering
<code>unique(a)</code>	<code>unique(a)</code>	
<code>squeeze(a)</code>	<code>a.squeeze()</code>	

Notes

Submatrix: Assignment to a submatrix can be done with lists of indexes using the `ix_` command. E.g., for 2d array `a`, one might do: `ind=[1,3]; a[np.ix_(ind,ind)]+=100`.

HELP: There is no direct equivalent of MATLAB's `which` command, but the commands `help` and `source` will usually list the filename where the function is located. Python also has an `inspect` module (do `import inspect`) which provides a `getfile` that often works.

INDEXING: MATLAB® uses one based indexing, so the initial element of a sequence has index 1. Python uses zero based indexing, so the initial element of a sequence has index 0. Confusion and flamewars arise because each has advantages and disadvantages. One based indexing is consistent with common human language usage, where the "first" element of a sequence has index 1. Zero based indexing simplifies indexing

(<https://groups.google.com/group/comp.lang.python/msg/1bf4d925dfbf368?q=g:thl3498076713d&hl=en>). See also a text by prof.dr. Edsger W. Dijkstra (<https://www.cs.utexas.edu/users/EWD/transcriptions/EWD08xx/EWD831.html>).

RANGES: In MATLAB®, `0:5` can be used as both a range literal and a 'slice' index (inside parentheses); however, in Python, constructs like `0:5` can *only* be used as a slice index (inside square brackets). Thus the somewhat quirky `r_` object was created to allow numpy to have a similarly terse range construction mechanism. Note that `r_` is not called like a function or a constructor, but rather *indexed* using square brackets, which allows the use of Python's slice syntax in the arguments.

LOGICOPS: `&` or `|` in NumPy is bitwise AND/OR, while in Matlab `&` and `|` are logical AND/OR. The difference should be clear to anyone with significant programming experience. The two can appear to work the same, but there are important differences. If you would have used Matlab's `&` or `|` operators, you should use the NumPy `ufuncs` `logical_and`/`logical_or`. The notable differences between Matlab's and NumPy's `&` and `|` operators are:

- Non-logical {0,1} inputs: NumPy's output is the bitwise AND of the inputs. Matlab treats any non-zero value as 1 and returns the logical AND. For example `(3 & 4)` in NumPy is 0, while in Matlab both 3 and 4 are considered logical true and `(3 & 4)` returns 1.
- Precedence: NumPy's `&` operator is higher precedence than logical operators like `<` and `>`; Matlab's is the reverse.

If you know you have boolean arguments, you can get away with using NumPy's bitwise operators, but be careful with parentheses, like this: `z = (x > 1) & (x < 2)`. The absence of NumPy operator forms of `logical_and` and `logical_or` is an unfortunate consequence of Python's design.

RESHAPE and LINEAR INDEXING: Matlab always allows multi-dimensional arrays to be accessed using scalar or linear indices, NumPy does not. Linear indices are common in Matlab programs, e.g. `find()` on a matrix returns them, whereas NumPy's `find` behaves differently. When converting Matlab code it might be necessary to first reshape a matrix to a linear sequence, perform some indexing operations and then reshape back. As `reshape` (usually) produces views onto the same storage, it should be possible to do this fairly efficiently. Note that the scan order used by `reshape` in NumPy defaults to the 'C' order, whereas Matlab uses the Fortran order. If you are simply converting to a linear sequence and back this doesn't matter. But if you are converting reshapes from Matlab code which relies on the scan order, then this Matlab code: `z = reshape(x,3,4)`; should become `z = x.reshape(3,4,order='F').copy()` in NumPy.

Customizing Your Environment

In MATLAB® the main tool available to you for customizing the environment is to modify the search path with the locations of your favorite functions. You can put such customizations into a startup script that MATLAB will run on startup.

NumPy, or rather Python, has similar facilities.

- To modify your Python search path to include the locations of your own modules, define the `PYTHONPATH` environment variable.
- To have a particular script file executed when the interactive Python interpreter is started, define the `PYTHONSTARTUP` environment variable to contain the name of your startup script.

Unlike MATLAB®, where anything on your path can be called immediately, with Python you need to first do an 'import' statement to make functions in a particular file accessible.

For example you might make a startup script that looks like this (Note: this is just an example, not a statement of "best practices"):

```
# Make all numpy available via shorter 'np' prefix
import numpy as np
# Make all matlib functions accessible at the top level via M.func()
import numpy.matlib as M
# Make some matlib functions accessible directly at the top level via, e.g. rand(3,3)
from numpy.matlib import rand,zeros,ones,empty,eye
# Define a Hermitian function
def hermitian(A, **kwargs):
    return np.transpose(A,**kwargs).conj()
# Make some shortcuts for transpose,hermitian:
#   np.transpose(A) --> T(A)
#   hermitian(A) --> H(A)
T = np.transpose
H = hermitian
```

Links

See <http://mathesaurus.sf.net/> (<http://mathesaurus.sf.net/>) for another MATLAB®/NumPy cross-reference.

An extensive list of tools for scientific work with python can be found in the topical software page (<https://scipy.org/topical-software.html>).

MATLAB® and SimuLink® are registered trademarks of The MathWorks.

Table of Contents (<https://docs.scipy.org/doc/numpy/contents.html>)

- NumPy for Matlab users
 - Introduction
 - Some Key Differences
 - 'array' or 'matrix'? Which should I use?
 - Short answer
 - Long answer
 - Table of Rough MATLAB-NumPy Equivalents
 - General Purpose Equivalents
 - Linear Algebra Equivalents
 - Notes
 - Customizing Your Environment
 - Links

Previous topic

Miscellaneous (<https://docs.scipy.org/doc/numpy/user/misc.html>)

Next topic

Building from source (<https://docs.scipy.org/doc/numpy/user/building.html>)

Quick search