

Numerical Python

**David Ascher
Paul F. Dubois
Konrad Hinsen
Jim Hugunin
Travis Oliphant**

March 15, 2001

**Lawrence Livermore National Laboratory, Livermore, CA 94566
UCRL-MA-128569**

Legal Notice

Please see file `Legal.html` in the source distribution.

This open source project has been contributed to by many people, including personnel of the Lawrence Livermore National Laboratory. The following notice covers those contributions including this manual.

Copyright (c) 1999. The Regents of the University of California. All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose without fee is hereby granted, provided that this entire notice is included in all copies of any software which is or includes a copy or modification of this software and in all copies of the supporting documentation for such software.

This work was produced at the University of California, Lawrence Livermore National Laboratory under contract no. W-7405-ENG-48 between the U.S. Department of Energy and The Regents of the University of California for the operation of UC LLNL.

DISCLAIMER

This software was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately-owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Table Of Contents

Numerical Python 1

1 Introduction 2

Where to get information and code 3
Acknowledgments 4

2 Installing NumPy 5

Testing the Python installation 5
Testing the Numeric Python Extension Installation 5
Installing NumPy 5
 Installing on Windows 6
 Installing on Unix 6
At the SourceForge... 6
 The Numeric Discussion List 6
 Bugs and Patches 6
 CVS Repository 6
 FTP Site 6
 Pyfort 6

3 The NumTut package 7

Testing the NumTut package 7
Possible reasons for failure: 7
 Win32 7
 Unix 8

4 High-Level Overview 9

Array Objects 9
Universal Functions 10
Convenience Functions 10

5 Array Basics 12

Basics 12
Creating arrays from scratch 12
 array() and typecodes 12
 Multidimensional Arrays 14
 resize 16
Creating arrays with values specified 'on-the-fly' 17
 zeros() and ones() 17
 arrayrange() 17
 Creating an array from a function: fromfunction() 19
 identity() 20
Coercion and Casting 20
 Automatic Coercions and Binary Operations 20
 Deliberate up-casting: The asarray function 21
 The typecode value table 21
 Consequences of silent upcasting 22
 Saving space 22

Deliberate casts (potentially down): the astype method 22
Operating on Arrays 23
Simple operations 23
Getting and Setting array values 24
Slicing Arrays 25

6 Ufuncs 27

What are Ufuncs? 27
Ufuncs can operate on any Python sequence 28
Ufuncs can take output arguments 28
Ufuncs have special methods 28
 The reduce ufunc method 28
 The accumulate ufunc method 29
 The outer ufunc method 29
 The reduceat ufunc method 29
Ufuncs always return new arrays 30
Which are the Ufuncs? 30
Unary Mathematical Ufuncs (take only one argument) 30
Binary Mathematical Ufuncs 30
Logical Ufuncs 30
Ufunc shorthands 31

7 Pseudo Indices 32

8 Array Functions 34

take(a, indices, axis=0) 34
put (a, indices, values) 35
putmask (a, mask, values) 36
transpose(a, axes=None) 36
repeat(a, repeats, axis=0) 36
choose(a, (b0, ..., bn)) 36
ravel(a) 37
nonzero(a) 37
where(condition, x, y) 37
compress(condition, a, axis=0) 37
diagonal(a, k=0) 37
trace(a, k=0) 38
searchsorted(a, values) 38
sort(a, axis=-1) 38
argsort(a, axis=-1) 39
argmax(a, axis=-1), argmin(a, axis=-1) 39
fromstring(string, typecode) 39
dot(m1, m2) 40
matrixmultiply(m1, m2) 40
clip(m, m_min, m_max) 40
indices(shape, typecode=None) 40
swapaxes(a, axis1, axis2) 41
concatenate((a0, a1, ..., an), axis=0) 41
innerproduct(a, b) 42
outerproduct(a,b) 42
array_repr() 42
array_str() 42

resize(a, new_shape) 42
 diagonal(a, offset=0, axis1=-2, axis2=-1) 42
 repeat..... (a, counts, axis=0) 42
 convolve (a, v, mode=0) 43
 cross_correlate (a, v, mode=0) 43
 where (condition, x, y) 43
 identity(n) 43
 sum(a, index=0) 44
 cumsum(a, index=0) 44
 product(a, index=0) 44
 cumproduct(a, index=0) 44
 alltrue(a, index=0) 44
 sometrue(a, index=0) 44

9 Array Methods 45

itemsize() 45
 iscontiguous() 45
 typecode() 45
 byteswapped() 45
 tostring() 45
 tolist() 46

10 Array Attributes 47

flat 47
real and *imaginary* 47

11 Special Topics 49

Subclassing 49
 Code Organization 49
 Numeric.py and friends 49
 UserArray.py 49
 Matrix.py 49
 Precision.py 49
 ArrayPrinter.py 49
 Mlab.py 49
 bartlett(M) 49
 blackman(M) 50
 corrcoef(x, y=None) 50
 cov(m, y=None) 50
 cumprod(m) 50
 cumsum(m) 50
 diag(v, k=0) 50
 diff(x, n=1) 50
 eig(m) 50
 eye(N, M=N, k=0, typecode=None) 50
 fliplr(m) 50
 flipud(m) 50
 hamming(M) 50
 hanning(M) 50
 kaiser(M, beta) 50
 max(m) 50
 mean(m) 50

- median(m) 50
- min(m) 51
- msort(m) 51
- prod(m) 51
- ptp(m) 51
- rand(d1, ..., dn) 51
- rot90(m,k=1) 51
- sinc(x) 51
- squeeze(a) 51
- std(m) 51
- sum(m) 51
- svd(m) 51
- trapz(y,x=None) 51
- tri(N, M=N, k=0, typecode=None) 51
- tril(m,k=0) 51
- triu(m,k=0) 51
- The multiarray object 52
- Typecodes 52
- Indexing in and out, slicing 53
- Ellipses 54
- NewAxis 54
- Set-indexing and Broadcasting 54
- Axis specifications 55
- Textual representations of arrays 55
- Comparisons 57
- Pickling and Unpickling -- storing arrays on disk 57
- Dealing with floating point exceptions 57

12 Writing a C extension to NumPy 58

- Introduction 58
- Preparing an extension module for NumPy arrays 58
- Accessing NumPy arrays from C 58
 - Types and Internal Structure* 58
 - Element data types* 59
 - Contiguous arrays* 60
 - Zero-dimensional arrays* 60
- A simple example 60
- Accepting input data from any sequence type 61
- Creating NumPy arrays 62
- Returning arrays from C functions 62
- A less simple example 62

13 C API Reference 64

- ArrayObject C Structure and API 64
 - Structures* 64
 - The ArrayObject API* 65
 - Notes* 68
- UfuncObject C Structure and API 68
 - C Structure* 68
 - UfuncObject C API* 70

14 Glossary 73

Optional Packages 75

License and disclaimer for packages MA and RNG 76

15 FFT Reference 77

Python Interface 77

fft(data, n=None, axis=-1) 77
inverse_fft(data, n=None, axis=-1) 77
real_fft(data, n=None, axis=-1) 77
inverse_real_fft(data, n=None, axis=-1) 78
fft2d(data, s=None, axes=(-2,-1)) 78
real_fft2d(data, s=None, axes=(-2,-1)) 78

C API 78

Compilation Notes 79

16 LinearAlgebra Reference 80

Python Interface 80

solve_linear_equations(a, b) 80
inverse(a) 80
eigenvalues(a) 80
eigenvectors(a) 81
singular_value_decomposition(a, full_matrices=0) 81
generalized_inverse(a, rcond=1e-10) 81
determinant(a) 81
linear_least_squares(a, b, rcond=e-10) 81

Compilation Notes 81

17 RandomArray Reference 82

Python Interface 82

seed(x=0, y=0) 82
get_seed() 82
random(shape=ReturnFloat) 82
uniform(minimum, maximum, shape=ReturnFloat) 82
randint(minimum, maximum, shape=ReturnFloat) 82
permutation(n) 82

Floating point random arrays 83

standard_normal(shape=ReturnFloat) 83
normal(mean, stddev, shape=ReturnFloat) 83
multivariate_normal(mean, covariance) or
multivariate_normal(mean, covariance, leadingAxesShape) 83
exponential(mean, shape=ReturnFloat) 83
beta(a, b, shape=ReturnFloat) 83
gamma(a, r, shape=ReturnFloat) 84
chi_square(df, shape=ReturnFloat) 84
noncentral_chi_square(df, nonc, shape=ReturnFloat) 84
F(dfn, dfd, shape=ReturnFloat) 84
noncentral_F(dfn, dfd, nconc, shape=ReturnFloat) 84

Integer random arrays 84

binomial(trials, prob, shape=ReturnInt) 84
negative_binomial(trials, prob, shape=ReturnInt) 84
poisson(mean, shape=ReturnInt) 84
multinomial(trials, probs) or *multinomial(trials, probs, leadingAxesShape)* 84

Examples 85

18 Independent Random Streams 87

Background 87

Usage 87

Module RNG 87

Generator objects 87

Module ranf 88

Examples 88

19 Masked Arrays 89

What is a masked array? 89

Installing and using MA 89

Class MaskedArray 89

Attributes of masked arrays 90

Methods on masked arrays. 90

Constructing masked arrays 92

Working with Masks 93

Copying or not? 94

Behaviors 94

Indexing and Slicing 94

Indexing that produces a scalar result 95

Assignment to elements and slices 95

Module MA: Attributes 95

Module MA: Functions 95

Unary functions 95

Binary functions 95

Comparison operators 96

Logical operators 96

Special array operators 96

Controlling the size of the string representations 98

Helper classes 98

MAError 98

The constant masked 98

Example of masked behavior 98

Class masked_unary_function 99

Class masked_binary_function 99

ActiveAttributes 99

Examples of Using MA 100

Data with a given value representing missing data 100

Filling in the missing data 100

Numerical operations 100

Seeing the mask 100

Filling it your way 101

Ignoring extreme values 101

Averaging an entire multidimensional array 101

Index 103

PART I: Numerical Python

Numerical Python (“Numpy”) adds a fast multidimensional array facility to Python. This part contains all you need to know about “Numpy” arrays and the functions that operate upon them.

1. Introduction

This chapter introduces the Numeric Python extension and outlines the rest of the document.

The Numeric Python extensions (NumPy henceforth) is a set of extensions to the Python programming language which allows Python programmers to efficiently manipulate large sets of objects organized in grid-like fashion. These sets of objects are called arrays, and they can have any number of dimensions: one dimensional arrays are similar to standard Python sequences, two-dimensional arrays are similar to matrices from linear algebra. Note that one-dimensional arrays are also *different* from any other Python sequence, and that two-dimensional matrices are also *different* from the matrices of linear algebra, in ways which we will mention later in this text.

Why are these extensions needed? The core reason is a very prosaic one, and that is that manipulating a set of a million numbers in Python with the standard data structures such as lists, tuples or classes is much too slow and uses too much space. Anything which we can do in NumPy we can do in standard Python – we just may not be alive to see the program finish. A more subtle reason for these extensions however is that the kinds of operations that programmers typically want to do on arrays, while sometimes very complex, can often be decomposed into a set of fairly standard operations. This decomposition has been developed similarly in many array languages. In some ways, NumPy is simply the application of this experience to the Python language – thus many of the operations described in NumPy work the way they do because experience has shown that way to be a good one, in a variety of contexts. The languages which were used to guide the development of NumPy include the infamous APL family of languages, Basis, MATLAB, FORTRAN, S and S+, and others. This heritage will be obvious to users of NumPy who already have experience with these other languages. This tutorial, however, does not assume any such background, and all that is expected of the reader is a reasonable working knowledge of the standard Python language.

This document is the “official” documentation for NumPy. It is both a tutorial and the most authoritative source of information about NumPy with the exception of the source code. The tutorial material will walk you through a set of manipulations of simple, small, arrays of numbers, as well as image files. This choice was made because:

- A concrete data set makes explaining the behavior of some functions much easier to motivate than simply talking about abstract operations on abstract data sets;
- Every reader will have at least an *intuition* as to the meaning of the data and organization of image files, and
- The result of various manipulations can be displayed simply since the data set has a natural graphical representation.

All users of NumPy, whether interested in image processing or not, are encouraged to follow the tutorial with a working NumPy installation at their side, testing the examples, and, more importantly, transferring the understanding gained by working on images to their specific domain. The best way to learn is by doing – the aim of this tutorial is to guide you along this “doing.”

Here is what the rest of this part contains:

- “Installing NumPy” on page 5 provides information on testing Python, NumPy, and compiling and installing NumPy if necessary.
- “The NumTut package” on page 7 provides information on testing and installing the NumTut package, which allows easy visualization of arrays.
- “High-Level Overview” on page 9 gives a high-level overview of the components of the NumPy system as a whole.
- “Array Basics” on page 12 provides a detailed step-by-step introduction to the most important aspect of NumPy, the multidimensional array objects.
- “Ufuncs” on page 27 provides information on universal functions, the mathematical functions which operate on arrays and other sequences elementwise.
- “Pseudo Indices” on page 32 covers syntax for some special indexing operators.
- “Array Functions” on page 34 is a catalog of each of the utility functions which allow easy algorithmic processing of arrays.
- “Array Methods” on page 45 discusses the methods of array objects.
- “Array Attributes” on page 47 presents the attributes of array objects.
- “Special Topics” on page 49 is a collection of special topics, from the organization of the codebase to the mechanisms for customizing printing.
- “Writing a C extension to NumPy” on page 58 is an tutorial on how to write a C extension which uses NumPy arrays.
- “C API Reference” on page 64 is a reference for the C API to NumPy objects (both PyArrayObjects and UFuncObjects).
- “Glossary” on page 73 is a glossary of terms.
- Reference material for the optional packages distributed with Numeric Python are described in the next part, “Optional Packages” on page 75.

Where to get information and code

Numerical Python and its documentation are available at SourceForge. The main web site is:

<http://numpy.sourceforge.net>

Downloads, bug reports, and patch facility, and releases are at the main project page, reachable from the above site or directly at: <http://sourceforge.net/projects/numpy>

The Python web site is www.python.org

Many packages are available from third parties that use Numeric to interface to a variety of mathematical and statistical software.

Acknowledgments

Numerical Python is the outgrowth of a long collaborative design process carried out by the Matrix SIG of the Python Software Activity (PSA). Jim Hugunin, while a graduate student at MIT, wrote most of the code and initial documentation. When Jim joined CNRI and began working on JPython, he didn't have the time to maintain Numerical Python so Paul Dubois at LLNL agreed to become the maintainer of Numerical Python. David Ascher, working as a consultant to LLNL, wrote most of this document, incorporating contributions from Konrad Hinsen and Travis Oliphant, both of whom are major contributors to Numerical Python.

Since the source was moved to SourceForge, the Numeric user community has become a significant part of the process. Numerical Python illustrates the power of the open source software concept.

Please send comments and corrections to this manual to paul@pfdubois.com, or to Paul F. Dubois, L-264, Lawrence Livermore National Laboratory, Livermore, CA 94566, U.S.A.

2. Installing NumPy

This chapter explains how to install and test NumPy, from either the source distribution or from the binary distribution.

Before we start with the actual tutorial, we will describe the steps needed for you to be able to follow along the examples step by step. These steps including installing Python, the NumPy extensions, and some tools and sample files used in the examples of this tutorial.

Testing the Python installation

The first step is to install Python if you haven't already. Python is available from the Python website's download directory at <http://www.python.org/download>. Click on the link corresponding to your platform, and follow the instructions described there. When installed, starting Python by typing `python` at the shell or double-clicking on the Python interpreter should give a prompt such as:

```
Python 1.5.1 (#0, Apr 13 1998, 20:22:04) [MSC 32 bit (Intel)] on win32
Copyright 1991-1995 Stichting Mathematisch Centrum, Amsterdam
>>>
```

If you have problems getting this part to work, consider contacting a local support person or emailing python-help@python.org for help. If neither solution works, consider posting on the `comp.lang.python` newsgroup (details on the newsgroup/ mailing list are available at <http://www.python.org/psa/MailingLists.html#clp>).

Testing the Numeric Python Extension Installation

The standard Python distribution does not come as of this writing with the Numeric Python extensions installed, but your system administrator may have installed them already. To find out if your Python interpreter has NumPy installed, type `import Numeric` at the Python prompt. You'll see one of two behaviors (throughout this document, **bold Courier New** font indicates user input, and `standard Courier New` font indicates output):

```
>>> import Numeric
Traceback (innermost last):
  File "<stdin>", line 1, in ?
ImportError: No module named Numeric
>>>
```

indicating that you don't have NumPy installed, or:

```
>>> import Numeric
>>>
```

indicating that you do. If you do, go on to the next step. If you don't, you have to get the NumPy extensions.

Installing NumPy

The release facility at SourceForge is accessed through the project page, <http://sourceforge.net/projects/numpy>. Click on the "numpy" releases and you will be presented with a list of the available files. The files whose names end in ".tar.gz" are source code releases. The others are "prebuilt" for a given platform. It is possible to get the latest sources directly from our CVS repository using the facilities described at SourceForge. Note that while every effort is made to ensure that the repository is always "good", direct use of the repository is subject to more errors than using a standard release.

Installing on Windows

On Windows, we currently have .zip files that should be unzipped into the top of your Python distribution; there is no “Setup” to run. If you wish to build from source on Windows, the Unix procedure described below can be used, running python in a command-line tool.

In general, there may not be a prebuilt version of a particular kind available in every minor release. If you need a prebuilt version, choose the most recent version available.

Installing on Unix

The source distribution should be uncompressed and unpacked using the the tar program:

```
csh> tar xfz Numeric-n.m.tar.gz
```

Follow the instructions in the top-level directory for compilation and installation. Note that there are options you *must* consider before beginning. Installation is usually as simple as:

```
python setup_all.py install
```

However, please (please!) see the README itself for the latest details.

Important Tip

Just like all Python modules and packages, the Numeric module can be invoked using either the `import Numeric` form, or the `from Numeric import ...` form. Because most of the functions we’ll talk about are in the Numeric module, in this document, all of the code samples will assume that they have been preceded by a statement:

```
from Numeric import *
```

At the SourceForge...

The SourceForge facility is at <http://sourceforge.net/projects/numpy>. Look on SourceForge also for various Numeric-based packages supplied by individuals.

The Numeric Discussion List

You can subscribe to a discussion list about Numeric python using the project page at SourceForge. The list is a good place to ask questions and get help. Send mail to numpy-discussion@lists.sourceforge.net.

Bugs and Patches

Bug tracking and patch-management facilities is provided on the SourceForge project page.

CVS Repository

You can get the latest and greatest (albeit less tested and trustworthy) version of Numeric directly from our CVS repository.

FTP Site

The FTP Site contains this documentation in several formats, plus maybe some other goodies we have lying around.

Pyfort

One tool for connecting Fortran to Numeric and Python is Pyfort, sourceforge.net/projects/pyfortran.

3. The NumTut package

This chapter leads the user through the installation and testing of the NumTut package, which should have been distributed with this document.

Testing the NumTut package

This tutorial assumes that the NumTut package has been installed. This package contains a few sample images and utility functions for displaying arrays and the like. To find out if NumTut has been installed, do:

```
>>> from NumTut import *
>>> view(greece)
```



If a picture of a greek street shows up on your screen, you're all set, and you can go to the next chapter.

Possible reasons for failure:

```
>>> import NumTut
Traceback (innermost last):
  File "<stdin>", line 1, in ?
ImportError: No module named NumTut
```

This message indicates that you do not have the NumTut package installed in your PythonPath. NumTut is distributed along with the Python source in the Demo subdirectory. Copy the NumTut subdirectory somewhere into your Python path, or just execute python from the Demo directory.

On Win32, the NumTut directory can be placed in the main directory of your Python installation. On Unix, it can be placed in the site-packages directory of your installation.

Win32

```
>>> import NumTut
Traceback (innermost last):
[...]
ConfigurationError: view needs Tkinter on Win32, and either threads or
the IDLE editor"
```

or:

```
ConfigurationError: view needs either threads or the IDLE editor to be enabled.
```

On Win32 (Windows 95, 98, NT), the Tk toolkit is needed to view the images. Additionally, either the Python interpreter needs to be compiled with thread support (which is true in the standard win32 distribution) or you need to call the NumTut program from the IDLE interactive development environment.

If you do not wish to modify your Python installation to match these requirements, you can simply ignore the references to the demonstrations which use the `view()` command later in this document. Using NumPy does not require image display tools, they just make some array operations easier to understand.

Unix

On Unix machines, NumTut will work best with a Python interpreter with Tk support (not true in the default configuration), with the Tkinter GUI framework available and optionally with the tkImaging add-on (part of the Python Imaging Library). If this is not the case, it will try to use an external viewer which is able to read PPM files. The default viewer is 'xv', a common image viewer available from <ftp://ftp.cis.upenn.edu/pub/xv>. If xv is not installed, you will get an error message similar to:

```
>>> import NumTut
Traceback (innermost last):
[...]
ConfigurationError: PPM image viewer 'xv' not found
```

You can configure NumTut to use a different image viewer, by typing e.g.:

```
>>> import NumTut
>>> NumTut.view.PPMVIEWER = 'ppmviewer'
>>> from NumTut import *
>>> view(greece)
```

If you do not have a PPM image viewer, you can simply ignore the references to the demonstrations which use the `view()` command later in this document. Using NumPy does not require image display tools, they just make some array operations easier to understand.

4. High-Level Overview

In this chapter, a high-level overview of the extensions is provided, giving the reader the definitions of the key components of the system. This section defines the concepts used by the remaining sections.

Numeric Python consists of a set of modules:

- `Numeric.py` (and its helper modules `multiarray` and `umath`.)

This module defines two new object types, and a set of functions which manipulate these objects, as well as convert between them and other Python types. The objects are the new array object (technically called `multiarray` objects), and universal functions (technically `ufunc` objects).

- Other optional packages shipped with Numeric are discussed in “Optional Packages” on page 75. Among these a packages for linear algebra, random numbers, masked or missing values, and Fast Fourier Transforms.

Array Objects

The array objects are generally homogeneous collections of potentially large numbers of numbers. All numbers in a `multiarray` are the same kind (i.e. number representation, such as double-precision floating point). Array objects must be full (no empty cells are allowed), and their size is immutable. The specific numbers within them can change throughout the life of the array.

Note: In some applications arrays of numbers may contain entries representing invalid or missing values. An optional package “MA” is available to represent such arrays. Attempting to do so by using NaN as a value may lead to disappointment or lack of portability.

Mathematical operations on arrays return new arrays containing the results of these operations performed *elementwise* on the arguments of the operation.

The *size* of an array is the total number of elements therein (it can be 0 or more). It does not change throughout the life of the array.

The *shape* of an array is the number of dimensions of the array and its extent in each of these dimensions (it can be 0, 1 or more). It can change throughout the life of the array. In Python terms, the shape of an array is a tuple of integers, one integer for each dimension that represents the extent in that dimension.

The *rank* of an array is the number of dimensions along which it is defined. It can change throughout the life of the array. Thus, the rank is the length of the shape.

The *typecode* of an array is a single character description of the kind of element it contains (number format, character or Python reference). It determines the *itemsize* of the array.

The *itemsize* of an array is the number of 8-bit bytes used to store a single element in the array. The total memory used by an array tends to its size times its *itemsize*, as the size goes to infinity (there is a fixed overhead per array, as well as a fixed overhead per dimension).

To put this in more familiar mathematical language: A vector is a rank-1 array (it has only one dimension along which it can be indexed). A matrix as used in linear algebra is a rank-2 array (it has two dimensions along which it can be indexed). There are also rank-0 arrays, which can hold single scalars -- they have no dimension along which they can be indexed, but they contain a single number.

Here is an example of Python code using the array objects (bold text refers to user input, non-bold text to computer output):

```

>>> vector1 = array((1,2,3,4,5))
>>> print vector1
[1 2 3 4 5]
>>> matrix1 = array([[0,1],[1,3]])
>>> print matrix1
[[0 1]
 [1 3]]
>>> print vector1.shape, matrix1.shape
(5,) (2,2)
>>> print vector1 + vector1
[ 2  4  6  8 10]
>>> print matrix1 * matrix1
[[0 1]
 [1 9]]

```

note that this is not the matrix
multiplication of linear algebra

If this example does not work for you because it complains of an unknown name “array”, you forgot to begin your session with

```
from Numeric import *
```

See page 6.

Universal Functions

Universal functions (ufuncs) are functions which operate on arrays and other sequences. Most ufuncs perform mathematical operations on their arguments, also elementwise.

Here is an example of Python code using the ufunc objects:

```

>>> print sin([pi/2., pi/4., pi/6.])
[ 1.          ,  0.70710678,  0.5          ]
>>> print greater([1,2,4,5], [5,4,3,2])
[0 0 1 1]
>>> print add([1,2,4,5], [5,4,3,2])
[6 6 7 7]
>>> print add.reduce([1,2,4,5])
12

```

1 + 2 + 3 + 4 + 5

Ufuncs are covered in detail in “Ufuncs” on page27.

Convenience Functions

The Numeric module provides, in addition to the functions which are needed to create the objects above, a set of powerful functions to manipulate arrays, select subsets of arrays based on the contents of other arrays, and other array-processing operations.

```

>>> data = arange(10)
range()
>>> print data
[0 1 2 3 4 5 6 7 8 9]
>>> print where(greater(data, 5), -1, data)
[ 0  1  2  3  4  5 -1 -1 -1 -1] # selection facility
>>> data = resize(array((0,1)), (9, 9))
>>> print data
[[0 1 0 1 0 1 0 1 0]
 [1 0 1 0 1 0 1 0 1]
 [0 1 0 1 0 1 0 1 0]
 [1 0 1 0 1 0 1 0 1]
 [0 1 0 1 0 1 0 1 0]
 [1 0 1 0 1 0 1 0 1]

```

```
[0 1 0 1 0 1 0 1 0]
[1 0 1 0 1 0 1 0 1]
[0 1 0 1 0 1 0 1 0]]
```

All of the functions which operate on NumPy arrays are described in “Array Functions” on page 34.

5. Array Basics

This chapter introduces some of the basic functions which will be used throughout the text.

Basics

Before we explore the world of image manipulation as a case-study in array manipulation, we should first define a few terms which we'll use over and over again. Discussions of arrays and matrices and vectors can get confusing due to disagreements on the nomenclature. Here is a brief definition of the terms used in this tutorial, and more or less consistently in the error messages of NumPy.

The python objects under discussion are formally called “multiarray” objects, but informally we'll just call them “array” objects or just “arrays.” These are different from the array objects defined in the standard Python `array` module (which is an older module designed for processing one-dimensional data such as sound files).

These array objects hold their data in a homogeneous block of elements, i.e. their elements all have the same C type (such as a 64-bit floating-point number). This is quite different from most Python container objects, which can contain heterogeneous collections. (You can, however, have an array of Python objects, as discussed later).

Any given array object has a rank, which is the number of “dimensions” or “axes” it has. For example, a point in 3D space `[1, 2, 1]` is an array of rank 1 – it has one dimension. That dimension has a *length* of 3.

As another example, the array

```
1.0 0.0 0.0
0.0 1.0 2.0
```

is an array of rank 2 (it is 2-dimensional). The first dimension has a length of 2, the second dimension has a length of 3. Because the word “dimension” has many different meanings to different folks, in general the word “axis” will be used instead. Axes are numbered just like Python list indices: they start at 0, and can also be counted from the end, so that axis -1 is the last axis of an array, axis -2 is the penultimate axis, etc.

There are two important and potentially unintuitive behaviors of NumPy arrays which take some getting used to. The first is that by default, operations on arrays are performed element-wise. This means that when adding two arrays, the resulting array has as elements the pairwise sums of the two operand arrays. This is true for all operations, including multiplication. Thus, array multiplication using the `*` operator will default to element-wise multiplication, not matrix multiplication as used in linear algebra. Many people will want to use arrays as linear algebra-type matrices (including their rank-1 versions, vectors). For those users, the `Matrix` class provides a more intuitive interface. We defer discussion of the `Matrix` class until later.

The second behavior which will catch many users by surprise is that functions which return arrays which are simply different views at the same data will in fact *share* their data. This will be discussed at length when we have more concrete examples of what exactly this means.

Now that all of these definitions and warnings are laid out, let's see what we can do with these arrays.

Creating arrays from scratch

array() and *typecodes*

There are many ways to create arrays. The most basic one is the use of the `array()` function:

```
>>> a = array([1.2, 3.5, -1])
to make sure this worked, do:
>>> print a
```

```
[ 1.2  3.5 -1. ]
```

The `array(numbers, typecode=None, savespace=0)` function takes three arguments – the first one is the values, which have to be in a Python sequence object (such as a list or a tuple). The optional second argument is the typecode of the elements. If it is omitted, as in the example above, Python tries to find the one type which can represent all the elements. The third is discussed in “Saving space” on page 22.

Since the elements we gave our example were two floats and one integer, it chose ‘float’ as the type of the resulting array. If one specifies the typecode, one can specify unequivocally the type of the elements – this is especially useful when, for example, one wants to make sure that an array contains floats even though in some cases all of its elements are integers:

```
>>> x,y,z = 1,2,3
>>> a = array([x,y,z])           # integers are enough for 1, 2 and 3
>>> print a
[1 2 3]
>>> a = array([x,y,z], Float)   # not the default type
>>> print a
[ 1.  2.  3.]
```

Important Tip

Pop Quiz: What will be the type of an array defined as follows:

```
>>> mystery = array([1, 2.0, -3j])
```

Hint: `-3j` is an imaginary number.

Answer: try it out!

A very common mistake is to call `array` with a set of numbers as arguments, as in `array(1, 2, 3, 4, 5)`. This doesn’t produce the expected result as soon as at least two numbers are used, because the first argument to `array()` must be the entire data for the array – thus, in most cases, a sequence of numbers. The correct way to write the preceding invocation is most likely `array((1, 2, 3, 4, 5))`.

Possible values for the second argument to the `array` creator function (and indeed to any function which accepts a so-called typecode for arrays) are:

1. One type corresponding to single ASCII characters: `Character`.
2. One unsigned numeric type: `UnsignedInt8`, used to store numbers between 0 and 255.
3. Many signed numeric types:
 - Signed integer choices: `Int`, `Int0`, `Int8`, `Int16`, `Int32`, and on some platforms, `Int64` and `Int128` (their ranges depend on their size).
 - Floating point choices: `Float`, `Float0`, `Float8`, `Float16`, `Float32`, `Float64`, and on some platforms, `Float128`.
 - Complex number choices: `Complex`, `Complex0`, `Complex8`, `Complex16`, `Complex32`, `Complex64`, `Complex128`.

The meaning of these is as follows:

- The versions without any numbers (`Int`, `Float`, `Complex`) correspond to the `int`, `float` and `complex` datatypes in Python. They are thus long integers and double-precision floating point numbers, with a complex number corresponding to two double-precision floats.

- The versions with a number following correspond to whatever words are available on the specific platform you are using which have *at least* that many bits in them. Thus, `Int0` corresponds to the smallest integer word size available, `Int8` corresponds to the smallest integer word size available which has at least 8 bits, etc. The word sizes for the complex numbers refer to the total number of bits used by both the real and imaginary parts (in other words, the data portion of an array of `N Complex128` elements uses up the same amount of memory as the data portions of two arrays of typecode `Float64` with `2N` elements).

4. One non-numeric type, `PyObject`. Arrays of typecode `PyObject` are arrays of Python references, and as such their data area can contain references to any kind of Python objects.

The last typecode deserves a little comment. Indeed, it seems to indicate that arrays can be filled with any Python objects. This appears to violate the notion that arrays are homogeneous. In fact, the typecode `PyObject` *does* allow heterogeneous arrays. However, if you plan to do numerical computation, you're much better off with a homogeneous array with a potentially "large" type than with a heterogeneous array. This is because a heterogeneous array stores references to objects, which incurs a memory cost, and because the speed of computation is much slower with arrays of `PyObject`'s than with uniform number arrays. Why does it exist, then?

A very useful features of arrays is the ability to slice them, dice them, select and choose from them, etc. This feature is so nice that sometimes one wants to do the same operations with, e.g., arrays of class instances. In such cases, computation speed is not as important as convenience. Also, if the array is filled with objects which are instances of classes which define the appropriate methods, then NumPy will let you do math with those objects. For example, if one creates an object class which has an `__add__` method, then arrays (created with the `PyObject` typecode) of instances of such a class can be added together.

Multidimensional Arrays

The following example shows one way of creating multidimensional arrays:

```
>>> ma = array([[1,2,3],[4,5,6]])
>>> print ma
[[1 2 3]
 [4 5 6]]
```

The first argument to `array()` in the code above is a single list containing two lists, each containing three elements. If we wanted floats instead, we could specify, as discussed in the previous section, the optional typecode we wished:

```
>>> ma_floats = array([[1,2,3],[4,5,6]], Float)
>>> print ma_floats
[[ 1.  2.  3.]
 [ 4.  5.  6.]]
```

This array allows us to introduce the notion of 'shape'. The shape of an array is the set of numbers which define its dimensions. The shape of the array `ma` defined above is 2 by 3. More precisely, all arrays have a shape attribute which is a tuple of integers. So, in this case:

```
>>> print ma.shape
(2, 3)
```

Using the earlier definitions, this is a shape of *rank* 2, where the first axis has length 2, and the second axis has length 3. The rank of an array `A` is always equal to `len(A.shape)`.

Note that `shape` is an *attribute* of `array` objects. It is the first of several which we will see throughout this tutorial. If you're not used to object-oriented programming, you can think of attributes as "features" or "qualities" of individual arrays. The relation between an array and its shape is similar to the relation between a person and their hair color. In Python, it's called an object/attribute relation.

What if one wants to change the dimensions of an array? For now, let us consider changing the shape of an array without making it "grow." Say, for example, we want to make the 2x3 array defined above (`ma`) an array of rank 1:

```
>>> flattened_ma = reshape(ma, (6,))
>>> print flattened_ma
```

```
[1 2 3 4 5 6]
```

One can change the shape of arrays to any shape as long as the product of all the lengths of all the axes is kept constant (in other words, as long as the number of elements in the array doesn't change):

```
>>> a = array([1,2,3,4,5,6,7,8])
[1 2 3 4 5 6 7 8]
>>> print a
>>> b = reshape(a, (2,4))          # 2*4 == 8
[[1 2 3 4]
 [5 6 7 8]]
>>> print b
>>> c = reshape(b, (4,2))          # 4*2 == 8
>>> print c
[[1 2]
 [3 4]
 [5 6]
 [7 8]]
```

Notice that we used a new function, `reshape()`. It, like `array()`, is a function defined in the `Numeric` module. It expects an array as its first argument, and a shape as its second argument. The shape has to be a sequence of integers (a list or a tuple). Keep in mind that a tuple with a single element needs a comma at the end; the right shape tuple for a rank-1 array with 5 elements is `(5,)`, not `(5)`.

One nice feature of shape tuples is that one entry in the shape tuple is allowed to be `-1`. The `-1` will be automatically replaced by whatever number is needed to build a shape which does not change the size of the array. Thus:

```
>>> a = reshape(array(range(25)), (5,-1))
>>> print a, a.shape
[[ 0  1  2  3  4]
 [ 5  6  7  8  9]
 [10 11 12 13 14]
 [15 16 17 18 19]
 [20 21 22 23 24]] (5, 5)
```

The shape of an array is a modifiable attribute of the array. You can therefore change the shape of an array simply by assigning a new shape to it:

```
>>> a = array([1,2,3,4,5,6,7,8,9,10])
>>> a.shape
(10,)
>>> a.shape = (2,5)
>>> print a
[[ 1  2  3  4  5]
 [ 6  7  8  9 10]]
>>> a.shape = (10,1)          # second axis has length 1
>>> print a
[[ 1]
 [ 2]
 [ 3]
 [ 4]
 [ 5]
 [ 6]
 [ 7]
 [ 8]
 [ 9]
 [10]]
>>> a.shape = (5,-1)          # note the -1 trick described above
>>> print a
```

```
[[ 1  2]
 [ 3  4]
 [ 5  6]
 [ 7  8]
 [ 9 10]]
```

As in the rest of Python, violating rules (such as the one about which shapes are allowed) results in exceptions:

```
>>> a.shape = (6,-1)
Traceback (innermost last):
  File "<stdin>", line 1, in ?
ValueError: total size of new array must be unchanged
```

For Advanced Users

The default printing routine provided by the Numeric module prints arrays as follows:

- 1 The last axis is always printed left to right
- 2 The next-to-last axis is printed top to bottom

The remaining axes are printed top to bottom with increasing numbers of separators.

This explains why rank-1 arrays are printed from left to right, rank-2 arrays have the first dimension going down the screen and the second dimension going from left to right, etc.

If you want to change the shape of an array so that it has more elements than it started with (i.e. grow it), then you have many options: One solution is to use the `concat()` method discussed later. An alternative is to use the `array()` creator function with existing arrays as arguments:

```
>>> print a
[0 1 2 3 4 5 6 6 7]
>>> b = array([a,a])
>>> print b
[[1 2 3 4 5 6 7 8]
 [1 2 3 4 5 6 7 8]]
>>> print b.shape
(2, 8)
```

resize

A final possibility is the `resize()` function, which takes a “base” array as its first argument and the desired shape as the second argument. Unlike `reshape()`, the shape argument to `resize()` can correspond to a smaller or larger shape than the input array. Smaller shapes will result in arrays with the data at the “beginning” of the input array, and larger shapes result in arrays with data containing as many replications of the input array as are needed to fill the shape. For example, starting with a simple array

```
>>> base = array([0,1])
```

one can quickly build a large array with replicated data:

```
>>> big = resize(base, (9,9))
>>> print big
[[0 1 0 1 0 1 0 1 0]
 [1 0 1 0 1 0 1 0 1]
 [0 1 0 1 0 1 0 1 0]
 [1 0 1 0 1 0 1 0 1]
 [0 1 0 1 0 1 0 1 0]
 [1 0 1 0 1 0 1 0 1]
 [0 1 0 1 0 1 0 1 0]
 [1 0 1 0 1 0 1 0 1]
 [0 1 0 1 0 1 0 1 0]]
```



```
[1 0 1 0 1 0 1 0 1]
[0 1 0 1 0 1 0 1 0]]
```

and if you imported the `view` function from the NumTut package, you can do:

```
>>> view(resize(base, (100,100)))
# grey grid of horizontal lines is shown
>>> view(resize(base, (101,101)))
# grey grid of alternating black and white pixels is shown
```

For Advanced Users

Sections denoted “For Advanced Users” will be used to indicate aspects of the functions which may not be needed for a first introduction at NumPy, but which should be mentioned for the sake of completeness.

The `array` constructor takes a mandatory `data` argument, an optional `typecode`, and optional `savespace` argument, and an optional `copy` argument. If the `data` argument is a sequence, then `array` creates a new object of type `multiarray`, and fills the array with the elements of the `data` object. The shape of the array is determined by the size and nesting arrangement of the elements of `data`.

If `data` is not a sequence, then the array returned is an array of shape `()` (the empty tuple), of typecode `'O'`, containing a single element, which is `data`.

Creating arrays with values specified 'on-the-fly'

`zeros()` and `ones()`

Often, one needs to manipulate arrays filled with numbers which aren't available beforehand. The `Numeric` module provides a few functions which create arrays from scratch:

`zeros()` and `ones()` simply create arrays of a given shape filled with zeros and ones respectively:

```
>>> z = zeros((3,3))
>>> print z
[[0 0 0]
 [0 0 0]
 [0 0 0]]
>>> o = ones([2,3])
>>> print o
[[1 1 1]
 [1 1 1]]
```

Note that the first argument is a shape – it needs to be a list or a tuple of integers. Also note that the default type for the returned arrays is `Int`, which you can feel free to override using something like:

```
>>> o = ones((2,3), Float)
>>> print o
[[ 1.  1.  1.]
 [ 1.  1.  1.]]
```

`arange()`

The `arange()` function is similar to the `range()` function in Python, except that it returns an array as opposed to a list.

```
>>> r = arange(10)
```

```
>>> print r
[0 1 2 3 4 5 6 7 8 9]
```

Combining the `arrayrange()` with the `reshape()` function, we can get:

```
>>> big = reshape(arrayrange(100),(10,10))
>>> print big
[[ 0  1  2  3  4  5  6  7  8  9]
 [10 11 12 13 14 15 16 17 18 19]
 [20 21 22 23 24 25 26 27 28 29]
 [30 31 32 33 34 35 36 37 38 39]
 [40 41 42 43 44 45 46 47 48 49]
 [50 51 52 53 54 55 56 57 58 59]
 [60 61 62 63 64 65 66 67 68 69]
 [70 71 72 73 74 75 76 77 78 79]
 [80 81 82 83 84 85 86 87 88 89]
 [90 91 92 93 94 95 96 97 98 99]]
>>> view(reshape(arrayrange(10000),(100,100)))
# array of increasing lightness from top down (slowly) and from left to
# right (faster) is shown
```

`arange()` is a shorthand for `arrayrange()`.

One can set the start, stop and step arguments, which allows for more varied ranges:

```
>>> print arrayrange(10,-10,-2)
[10  8  6  4  2  0 -2 -4 -6 -8]
```

An important feature of `arrayrange` is that it can be used with non-integer starting points and strides:

```
>>> print arrayrange(5.0)
[ 0.  1.  2.  3.  4.]
>>> print arrayrange(0, 1, .2)
[ 0.  0.2  0.4  0.6  0.8]
```

If you want to create an array with just one value, repeated over and over, you can use the `*` operator applied to lists

```
>>> a = array([[3]*5]*5)
>>> print a
[[3 3 3 3 3]
 [3 3 3 3 3]
 [3 3 3 3 3]
 [3 3 3 3 3]
 [3 3 3 3 3]]
```

but that is relatively slow, since the duplication is done on Python lists. A quicker way would be to start with 0's and add 3:

```
>>> a = zeros([5,5]) + 3
>>> print a
[[3 3 3 3 3]
 [3 3 3 3 3]
 [3 3 3 3 3]
 [3 3 3 3 3]
 [3 3 3 3 3]]
```

The optional typecode argument can force the typecode of the resulting array, which is otherwise the “highest” of the starting and stopping arguments. The starting argument defaults to 0 if not specified. Note that if a typecode is specified which is “lower” than that which `arrayrange` would normally use, the array is the result of a precision-losing cast (a round-down, as that used in the `astype` method for arrays.)

Creating an array from a function: `fromfunction()`

Finally, one may want to create an array with contents which are the result of a function evaluation. This is done using the `fromfunction()` function, which takes two arguments, a shape and a callable object (usually a function). For example:

```
>>> def dist(x,y):
...     return (x-5)**2+(y-5)**2    # distance from point (5,5) squared
...
>>> m = fromfunction(dist, (10,10))
>>> print m
[[50 41 34 29 26 25 26 29 34 41]
 [41 32 25 20 17 16 17 20 25 32]
 [34 25 18 13 10  9 10 13 18 25]
 [29 20 13  8  5  4  5  8 13 20]
 [26 17 10  5  2  1  2  5 10 17]
 [25 16  9  4  1  0  1  4  9 16]
 [26 17 10  5  2  1  2  5 10 17]
 [29 20 13  8  5  4  5  8 13 20]
 [34 25 18 13 10  9 10 13 18 25]
 [41 32 25 20 17 16 17 20 25 32]]
>>> view(fromfunction(dist, (100,100))
# shows image which is dark in topleft corner, and lighter away from it.
>>> m = fromfunction(lambda i,j,k: 100*(i+1)+10*(j+1)+(k+1), (4,2,3))
>>> print m
[[[111 112 113]
  [121 122 123]]
 [[211 212 213]
  [221 222 223]]
 [[311 312 313]
  [321 322 323]]
 [[411 412 413]
  [421 422 423]]]
```

By examining the above examples, one can see that `fromfunction()` creates an array of the shape specified by its second argument, and with the contents corresponding to the value of the function argument (the first argument) evaluated at the indices of the array. Thus the value of `m[3,4]` in the first example above is the value of `dist` when `x=3` and `y=4`. Similarly for the lambda function in the second example, but with a rank-3 array.

The implementation of `fromfunction` consists of:

```
def fromfunction(function, dimensions):
    return apply(function, tuple(indices(dimensions)))
```

which means that the function `function` is called for each element in the sequence `indices(dimensions)`. As described in the definition of `indices`, this consists of arrays of indices which will be of rank one less than that specified by `dimensions`. This means that the function argument must accept the same number of arguments as there are dimensions in `dimensions`, and that each argument will be an array of the same shape as that specified by `dimensions`. Furthermore, the array which is passed as the first argument corresponds to the indices of each element in the resulting array along the first axis, that which is passed as the second argument corresponds to the indices of each element in the resulting array along the second axis, etc. A consequence of this is that the function which is used with `fromfunction` will work as expected only if it performs a separable computation on its arguments, and expects its arguments to be indices along each axis. Thus, no logical operation on the arguments can be performed, or any non-shape preserving operation. The first example below satisfies these requirements, hence works (the `x` and `y` arrays both get 10x10 arrays as input corresponding to the values of the indices along the two dimensions), while the second array attempts to do a comparison test on an array of indices, which fails.

```
>>> def buggy(test):
```

```

...     if test > 4: return 1
...     else: return 0
...
>>> print fromfunction(buggy, (10,))
Traceback (innermost last):
  File "<stdin>", line 1, in ?
  File "C:\PYTHON\LIB\numeric.py", line 157, in fromfunction
    return apply(function, tuple(indices(dimensions)))
  File "<stdin>", line 2, in buggy
TypeError: Comparison of multiarray objects is not implemented.

```

identity()

The simplest array constructor is the `identity(n)` function, which takes a single integer argument and returns a square identity array of that size of integers:

```

>>> print identity(5)
[[1 0 0 0 0]
 [0 1 0 0 0]
 [0 0 1 0 0]
 [0 0 0 1 0]
 [0 0 0 0 1]]
>>> view(identity(100))
# shows black square with a single white diagonal

```

Coercion and Casting

We've mentioned the typecodes of arrays, and how to create arrays with the right typecode, but we haven't covered what happens when arrays with different typecodes interact.

Automatic Coercions and Binary Operations

The rules followed by NumPy when performing binary operations on arrays mirror those used by Python in general. Operations between numeric and non-numeric types are not allowed (e.g. an array of characters can't be added to an array of numbers), and operations between mixed number types (e.g. floats and integers, floats and complex numbers, or in the case of NumPy, operations between any two arrays with different numeric typecodes) first perform a coercion of the 'smaller' numeric type to the type of the 'larger' numeric type. Finally, when scalars and arrays are operated on together, the scalar is converted to a rank-0 array first. Thus, adding a "small" integer to a "large" floating point array is equivalent to first casting the integer "up" to the typecode of the array:

```

>>> arange(0, 1.0, .1) + 12
array([ 12. ,  12.1,  12.2,  12.3,  12.4,  12.5,  12.6,  12.7,  12.8,
        12.9])

```

The automatic coercions are described in Figure 1. Avoiding upcasting is discussed in "Saving space" on page 22.

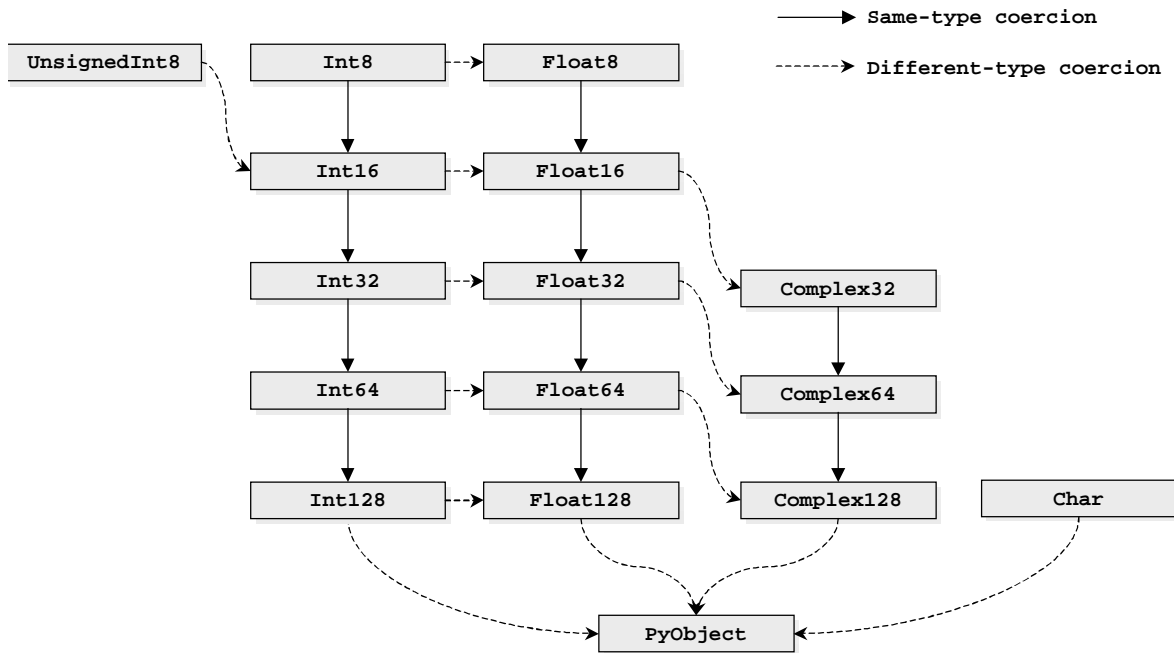


Figure 1 Up-casts are indicated with arrows. Down-casts are allowed by the `astype()` method, but may result in loss of information.

Deliberate up-casting: The `asarray` function

One more array constructor is the `asarray()` function. It is used if you want to have an array of a specific typecode and you don't know what typecode array you have (for example, in a generic function which can operate on all kinds of arrays, but needs them to be converted to complex arrays). If the array it gets as an argument is of the right typecode, it will get sent back unchanged. If the array is not of the right typecode, each element of the new array will be the result of the coercion to the new type of the old elements. `asarray()` will refuse to operate if there might be loss of information -- in other words, `asarray()` only casts 'up'.

`asarray` is also used when you have a function that operates on arrays, but you want to allow people to call it with an arbitrary python sequence object. This gives your function a behavior similar to that of most of the builtin functions that operate on arrays.

The typecode value table

The typecodes identifiers (`Float0`, etc.) have as values single-character strings. The mapping between typecode and character strings is machine dependent. An example of the correspondences between typecode characters and the typecode identifiers for 32-bit architectures are shown in Table 3-X.

Table 1: Typecode character/identifier table on a Pentium computer

Character	# of bytes	# of bits	Identifiers
D	16	128	Complex, Complex64
F	8	64	Complex0, Complex16, Complex32, Complex8
d	8	64	Float, Float64
f	4	32	Float0, Float16, Float32, Float8

Table 1: Typecode character/identifier table on a Pentium computer

Character	# of bytes	# of bits	Identifiers
l	4	32	Int
l	1	8	Int0, Int8
s	2	16	Int16
i	4	32	Int32

Consequences of silent upcasting

When dealing with very large arrays of floats and if precision is not important (or arrays of small integers), then it may be worthwhile to cast the arrays to “small” typecodes, such as `Int8`, `Int16` or `Float32`. As the standard Python integers and floats correspond to the typecodes `Int32` and `Float64`, using them in apparently “innocent” ways will result in up-casting, which may null the benefit of the use of small typecode arrays. For example:

```
>>> mylargearray.typecode()
'f' # a.k.a. Float32 on a Pentium
>>> mylargearray.itemsize()
4
>>> mylargearray = mylargearray + 1 # 1 is an Int64 on a Pentium
>>> mylargearray.typecode() # see Fig. 1 for explanation.
'd'
>>> mylargearray.itemsize()
8
```

Note that the sizes returned by the `itemsize()` method are expressed in bytes.

Saving space

Numeric arrays can be created using an optional, keyworded argument to the constructor, `savespace`. If `savespace` is set to 1, Numeric will attempt to avoid the silent upcasting behavior. The status of an array can be queried with the `spacesaver()` method. If `x.spacesaver()` is true, `x` has its space-saving flag set. The flag can be set with the `savespace` method: `x.savespace(1)` to set it, `x.savespace(0)` to clear it.

Deliberate casts (potentially down): the `astype` method

You may also force NumPy to cast any number array to another number array. For example, to take an array of any numeric type (`IntX` or `FloatX` or `ComplexX` or `UnsignedInt8`) and convert it to a 64-bit float, one can do:

```
>>> floatarray = otherarray.astype(Float64)
```

The typecode can be any of the number typecodes, “larger” or “smaller”. If it is larger, this is a cast-up, as if `asarray()` had been used. If it is smaller, the standard casting rules of the underlying language (C) are used, which means that truncation or loss of precision can occur:

```
>>> print x
[ 0.  0.4  0.8  1.2  1.6]
>>> x.astype(Int)
array([0, 0, 0, 1, 1])
```

If the typecode used with `astype()` is the original array’s typecode, then a copy of the original array is returned.

Operating on Arrays

Simple operations

If you have a keen eye, you have noticed that some of the previous examples did something new. It added a number to an array. Indeed, most Python operations applicable to numbers are directly applicable to arrays:

```
>>> print a
[1 2 3]
>>> print a * 3
[3 6 9]
>>> print a + 3
[4 5 6]
```

Note that the mathematical operators behave differently depending on the types of their operands. When one of the operands is an array and the other is a number, the number is added to all the elements of the array and the resulting array is returned. This is called *broadcasting*. This also occurs for unary mathematical operations such as `sin` and the negative sign

```
>>> print sin(a)
[ 0.84147098  0.90929743  0.14112001]
>>> print -a
[-1 -2 -3]
```

When both elements are arrays with the same shape, then a new array is created, where each element is the sum of the corresponding elements in the original arrays:

```
>>> print a + a
[2 4 6]
```

If the operands of operations such as addition are arrays which have the same rank but different non-integer dimensions, then an exception is generated:

```
>>> print a
[1 2 3]
>>> b = array([4,5,6,7])          # note this has four elements
>>> print a + b
Traceback (innermost last):
  File ``<stdin>``, line 1, in ?
  ArrayError: frames are not aligned
```

This is because there is no reasonable way for NumPy to interpret addition of a (3,) shaped array and a (4,) shaped array.

Note what happens when adding arrays with different rank

```
>>> print a
[1 2 3]
>>> print b
[[ 4  8 12]
 [ 5  9 13]
 [ 6 10 14]
 [ 7 11 15]]
>>> print a + b
[[ 5 10 15]
 [ 6 11 16]
 [ 7 12 17]
 [ 8 13 18]]
```

This is another form of broadcasting. To understand this, one needs to look carefully at the shapes of `a` and `b`:

```
>>> a.shape
(3,)
>>> b.shape
```

(4, 3)

Because array `a`'s last dimension had length 3 and array `b`'s last dimension also had length 3, those two dimensions were “matched” and a new dimension was created and automatically “assumed” for array `a`. The data already in `a` was “replicated” as many times as needed (4, in this case) to make the two shapes of the operand arrays conform. This replication (broadcasting) occurs when arrays are operands to binary operations and their shapes differ and when the following conditions are true:

- starting from the last axis, the axis lengths (dimensions) of the operands are compared
- if both arrays have an axis length greater than 1, an exception is raised
- if one array has an axis length greater than 1, then the other array's axis is “stretched” to match the length of the first axis -- if the other array's axis is not present (i.e., if the other array has smaller rank), then a new axis of the same length is created.

This algorithm is complex, but intuitive in practice. For more details, consult the Numeric Reference.

Getting and Setting array values

Just like other Python sequences, array contents are manipulated with the `[]` notation. For rank-1 arrays, there are no differences between list and array notations:

```
>>> a = arrayrange(10)
>>> print a[0]                # get first element
0
>>> print a[1:5]             # get second through fifth element
[1 2 3 4]
>>> print a[:-1]            # get last element
9
```

The first difference with lists comes with multidimensional indexing. If an array is multidimensional (of rank > 1), then specifying a single integer index will return an array of dimension one less than the original array.

```
>>> a = arrayrange(9)
>>> a.shape = (3,3)
>>> print a
[[0 1 2]
 [3 4 5]
 [6 7 8]]
>>> print a[0]                # get first row, not first element!
[0 1 2]
>>> print a[1]                # get second row
[3 4 5]
```

To get to individual elements in a rank-2 array, one specifies both indices separated by commas:

```
>>> print a[0,0]             # get elt at first row, first column
0
>>> print a[0,1]             # get elt at first row, second column
1
>>> print a[1,0]             # get elt at second row, first column
3
>>> print a[2,-1]           # get elt at third row, last column
8
```

Of course, the `[]` notation can be used to *set* values as well:

```
>>> a[0,0] = 123
>>> print a
[[123  1  2]
 [ 3  4  5]
 [ 6  7  8]]
```


Note that when referring to rows, the right hand side of the equal sign needs to be a sequence which “fits” in the referred array subset (in the code sample below, a 3-element row):

```
>>> a[1] = [10,11,12]
>>> print a
[[123  1  2]
 [ 10 11 12]
 [  6  7  8]]
```

Slicing Arrays

The standard rules of Python slicing apply to arrays, on a per-dimension basis. Assuming a 3x3 array:

```
>>> a = reshape(arrayrange(9),(3,3))
>>> print a
[[0 1 2]
 [3 4 5]
 [6 7 8]]
```

The plain [:] operator slices from beginning to end:

```
>>> print a[:,:]
[[0 1 2]
 [3 4 5]
 [6 7 8]]
```

In other words, [:] with no arguments is the same as [:] for lists – it can be read “all indices along this axis. So, to get the second row along the second dimension:

```
>>> print a[:,1]
[1 4 7]
```

Note that what was a “column” vector is now a “row” vector -- any “integer slice” (as in the 1 in the example above) results in a returned array with rank one less than the input array.

If one does not specify as many slices as there are dimensions in an array, then the remaining slices are assumed to be “all”. If A is a rank-3 array, then

```
A[1] == A[1,:] == A[1,:,:]
```

There is one addition to the slice notation for arrays which does not exist for lists, and that is the optional third argument, meaning the “step size” also called stride or increment. Its default value is 1, meaning return every element in the specified range. Alternate values allow one to skip some of the elements in the slice:

```
>>> a = arange(12)
>>> print a
[ 0  1  2  3  4  5  6  7  8  9 10 11]
>>> print a[::2] # return every *other* element
[ 0  2  4  6  8 10]
```

Negative strides are allowed as long as the starting index is greater than the stopping index:

```
>>> a = reshape(arrayrange(9),(3,3))
>>> print a
[[0 1 2]
 [3 4 5]
 [6 7 8]]
>>> print a[:, 0]
[0 3 6]
>>> print a[0:3, 0]
[0 3 6]
>>> print a[2:-1, 0]
[6 3 0]
```

If a negative stride is specified and the starting or stopping indices are omitted, they default to “end of axis” and “beginning of axis” respectively. Thus, the following two statements are equivalent for the array given:

```
>>> print a[2:-1, 0]
[6 3 0]
>>> print a[::-1, 0]
[6 3 0]
>>> print a[::-1]           # this reverses only the first axis
[[6 7 8]
 [3 4 5]
 [0 1 2]]
>>> print a[::-1,::-1]     # this reverses both axes
[[8 7 6]
 [5 4 3]
 [2 1 0]]
```

One final way of slicing arrays is with the keyword `...`. This keyword is somewhat complicated. It stands for “however many” I need depending on the rank of the object I’m indexing, so that the indices I *do* specify are at the end of the index list as opposed to the usual beginning.”

So, if one has a rank-3 array `A`, then `A[... , 0]` is the same thing as `A[:, :, 0]` but if `B` is rank-4, then `B[... , 0]` is the same thing as: `B[:, :, :, 0]`. Only one `...` is expanded in an index expression, so if one has a rank-5 array `C`, then: `C[... , 0, ...]` is the same thing as `C[:, :, :, 0, :]`.

6. Ufuncs

What are Ufuncs?

The operations on arrays that were mentioned in the previous section (element-wise addition, multiplication, etc.) all share some features -- they all follow similar rules for broadcasting, coercion and “element-wise operation”. Just like standard addition is available in Python through the `add` function in the `operator` module, array operations are available through callable objects as well. Thus, the following objects are available in the `Numeric` module:

Table 2: Universal Functions, or `ufuncs`. The operators which invoke them when applied to arrays are indicated in parentheses. The entries in slanted typeface refer to unary `ufuncs`, while the others refer to binary `ufuncs`.

<code>add (+)</code>	<code>subtract (-)</code>	<code>multiply (*)</code>	<code>divide (/)</code>
<code>remainder (%)</code>	<code>power (**)</code>	<i><code>arccos</code></i>	<i><code>arccosh</code></i>
<i><code>arcsin</code></i>	<i><code>arcsinh</code></i>	<i><code>arctan</code></i>	<i><code>arctanh</code></i>
<i><code>cos</code></i>	<i><code>cosh</code></i>	<i><code>exp</code></i>	<i><code>log</code></i>
<i><code>log10</code></i>	<i><code>sin</code></i>	<i><code>sinh</code></i>	<i><code>sqrt</code></i>
<i><code>tan</code></i>	<i><code>tanh</code></i>	<code>maximum</code>	<code>minimum</code>
<i><code>conjugate</code></i>	<code>equal (==)</code>	<code>not_equal (!=)</code>	<code>greater (>)</code>
<code>greater_equal (>=)</code>	<code>less (<)</code>	<code>less_equal (<=)</code>	<code>logical_and (and)</code>
<code>logical_or (or)</code>	<code>logical_xor</code>	<code>logical_not (not)</code>	<code>bitwise_and (&)</code>
<code>bitwise_or ()</code>	<code>bitwise_xor</code>	<code>bitwise_not (~)</code>	

All of these `ufuncs` can be used as functions. For example, to use `add`, which is a binary `ufunc` (i.e. it takes two arguments), one can do either of:

```
>>> a = arange(10)
>>> print add(a,a)
[ 0  2  4  6  8 10 12 14 16 18]
>>> print a + a
[ 0  2  4  6  8 10 12 14 16 18]
```

In other words, the `+` operator on arrays performs exactly the same thing as the `add` `ufunc` when operated on arrays. For a unary `ufunc` such as `sin`, one can do, e.g.:

```
>>> a = arange(10)
>>> print sin(a)
[ 0.          0.84147098  0.90929743  0.14112001 -0.7568025  -0.95892427
 -0.2794155   0.6569866   0.98935825  0.41211849]
```

Unary `ufuncs` return arrays with the same shape as their arguments, but with the contents corresponding to the corresponding mathematical function applied to each element (`sin(0)=0`, `sin(1)=0.84147098`, etc.).

There are three additional features of `ufuncs` which make them different from standard Python functions. They can operate on any Python sequence in addition to arrays; they can take an “output” argument; they have attributes which are themselves callable with arrays and sequences. Each of these will be described in turn.

Ufuncs can operate on any Python sequence

Ufuncs have so far been described as callable objects which take either one or two arrays as arguments (depending on whether they are unary or binary). In fact, any Python sequence which can be the input to the `array()` constructor can be used. The return value from ufuncs is always an array. Thus:

```
>>> add([1,2,3,4], (1,2,3,4))
array([2, 4, 6, 8])
```

Ufuncs can take output arguments

In many computations with large sets of numbers, arrays are often used only once. For example, a computation on a large set of numbers could involve the following step

```
dataset = dataset * 1.20
```

This operation as written needs to create a temporary array to store the results of the computation, and then eventually free the memory used by the original dataset array (provided there are no other references to the data it contains). It is more efficient, both in terms of memory and computation time, to do an “in-place” operation. This can be done by specifying an existing array as the place to store the result of the ufunc. In this example, one can write:

```
multiply(dataset, 1.20, dataset)
```

This is not a step to take lightly, however. For example, the “big and slow” version (`dataset = dataset * 1.20`) and the “small and fast” version above will yield different results in two cases:

- If the typecode of the target array is not that which would normally be computed, the operation will fail and raise a `TypeError` exception.
- If the target array corresponds to a different “view” on the same data as either of the source arrays, inconsistencies will result. For example,

```
>>> a = arange(5, typecode=Float64)
>>> print a[::-1] * 1.2
[ 4.8  3.6  2.4  1.2  0. ]
>>> multiply(a[::-1], 1.2, a)
array([ 4.8 ,  3.6 ,  2.4 ,  4.32,  5.76])
>>> print a
[ 4.8  3.6  2.4  4.32  5.76]
```

This is because the ufunc does not know which arrays share which data, and in this case the overwriting of the data contents follows a different path through the shared data space of the two arrays, thus resulting in strangely distorted data.

Ufuncs have special methods

The reduce ufunc method

If you don't know about the `reduce` command in Python, review section 5.1.1 of the Python Tutorial (<http://www.python.org/doc/tut/functional.html>). Briefly, `reduce` is most often used with two arguments, a callable object (such as a function), and a sequence. It calls the callable object with the first two element of the sequence, then with the result of that operation and the third element, and so on, returning at the end the successive “reduction” of the specified callable object over the sequence elements. Similarly, the `reduce` method of ufuncs is called with a sequence as an argument, and performs the reduction of that ufunc on the sequence. As an example, adding all of the elements in a rank-1 array can be done with:

```
>>> a = array([1,2,3,4])
>>> print add.reduce(a)
10
```

When applied to arrays which are of rank greater than one, the reduction proceeds by default along the first axis:

```
>>> b = array([[1,2,3,4],[6,7,8,9]])
>>> print b
```

```
[[1 2 3 4]
 [6 7 8 9]]
>>> print add.reduce(b)
[ 7  9 11 13]
```

A different axis of reduction can be specified with a second integer argument:

```
>>> print b
[[1 2 3 4]
 [6 7 8 9]]
>>> print add.reduce(b, 1)
[10 30]
```

The accumulate ufunc method

The `accumulate` ufunc method is similar to `reduce`, except that it returns an array containing the intermediate results of the reduction:

```
>>> a = arange(10)
>>> print a
[0 1 2 3 4 5 6 7 8 9]
>>> print add.accumulate(a)
[ 0  1  3  6 10 15 21 28 36 45] # 0, 0+1, 0+1+2, 0+1+2+3, ... 0+...+9
>>> print add.reduce(a)
45 # same as add.accumulate(...)[-1]
```

The outer ufunc method

The third ufunc method is `outer`, which takes two arrays as arguments and returns the “outer ufunc” of the two arguments. Thus the `outer` method of the `multiply` ufunc, results in the outer product. The `outer` method is only supported for binary methods.

```
>>> print a
[0 1 2 3 4]
>>> print b
[0 1 2 3]
>>> print add.outer(a,b)
[[0 1 2 3]
 [1 2 3 4]
 [2 3 4 5]
 [3 4 5 6]
 [4 5 6 7]]
>>> print multiply.outer(b,a)
[[ 0  0  0  0  0]
 [ 0  1  2  3  4]
 [ 0  2  4  6  8]
 [ 0  3  6  9 12]]
>>> print power.outer(a,b)
[[ 1  0  0  0]
 [ 1  1  1  1]
 [ 1  2  4  8]
 [ 1  3  9 27]
 [ 1  4 16 64]]
```

The reduceat ufunc method

The final ufunc method is the `reduceat` method, which I’d love to explain it, but I don’t understand it (XXX).

Ufuncs always return new arrays

Except when the 'output' argument are used as described above, ufuncs always return new arrays which do not share any data with the input array.

Which are the Ufuncs?

Table 1 lists all the ufuncs. We will first discuss the mathematical ufuncs, which perform operations very similar to the functions in the `math` and `cmath` modules, albeit elementwise, on arrays. These come in two forms, unary and binary:

Unary Mathematical Ufuncs (take only one argument)

The following ufuncs apply the predictable functions on their single array arguments, one element at a time: `arccos`, `arccosh`, `arcsin`, `arcsinh`, `arctan`, `arctanh`, `cos`, `cosh`, `exp`, `log`, `log10`, `sin`, `sinh`, `sqrt`, `tan`, `tanh`.

As an example:

```
>>> print x
[0 1 2 3 4]
>>> print cos(x)
[ 1.          0.54030231 -0.41614684 -0.9899925  -0.65364362]
>>> print arccos(cos(x))
[ 0.          1.          2.          3.          2.28318531]
# not a bug, but wraparound: 2*pi%4 is 2.28318531
```

The conjugate ufunc takes an array of complex numbers and returns the array with entries which are the complex conjugates of the entries in the input array. If it is called with real numbers, a copy of the array is returned unchanged.

Binary Mathematical Ufuncs

These ufuncs take two arrays as arguments, and perform the specified mathematical operation on them, one pair of elements at a time: `add`, `subtract`, `multiply`, `divide`, `remainder`, `power`.

Logical Ufuncs

The "logical" ufuncs also perform their operations on arrays in elementwise fashion, just like the "mathematical" ones.

Two are special (`maximum` and `minimum`) in that they return arrays with entries taken from their input arrays:

```
>>> print x
[0 1 2 3 4]
>>> print y
[ 2.  2.5  3.  3.5  4. ]
>>> print maximum(x, y)
[ 2.  2.5  3.  3.5  4. ]
>>> print minimum(x, y)
[ 0.  1.  2.  3.  4.]
```

The others all return arrays of 0's or 1's: `equal`, `not_equal`, `greater`, `greater_equal`, `less`, `less_equal`, `logical_and`, `logical_or`, `logical_xor`, `logical_not`, `bitwise_and`, `bitwise_or`, `bitwise_xor`, `bitwise_not`.

These are fairly self-explanatory, especially with the associated symbols from the standard Python version of the same operations in Table 1 above. The `logical_*` ufuncs perform their operations (and, or, etc.) using the truth value of the elements in the array (equality to 0 for numbers and the standard truth test for PyObject arrays). The `bitwise_*` ufuncs, on the other hand, can be used only with integer arrays (of any word size), and will return integer arrays of the larger bit size of the two input arrays:

```
>>> x
```

```

array([7, 7, 0], 'i')
>>> y
array([4, 5, 6])
>>> bitwise_and(x,y)
array([4, 5, 0], 'i')

```

We've already discussed how to find out about the contents of arrays based on the indices in the arrays – that's what the various slice mechanisms are for. Often, especially when dealing with the result of computations or data analysis, one needs to "pick out" parts of matrices based on the content of those matrices. For example, it might be useful to find out which elements of an array are negative, and which are positive. The comparison ufuncs are designed for just this type of operation. Assume an array with various positive and negative numbers in it (for the sake of the example we'll generate it from scratch):

```

>>> print a
[[ 0  1  2  3  4]
 [ 5  6  7  8  9]
 [10 11 12 13 14]
 [15 16 17 18 19]
 [20 21 22 23 24]]
>>> b = sin(a)
>>> print b
[[ 0.          0.84147098  0.90929743  0.14112001 -0.7568025 ]
 [-0.95892427 -0.2794155   0.6569866   0.98935825  0.41211849]
 [-0.54402111 -0.99999021 -0.53657292  0.42016704  0.99060736]
 [ 0.65028784 -0.28790332 -0.96139749 -0.75098725  0.14987721]
 [ 0.91294525  0.83665564 -0.00885131 -0.8462204  -0.90557836]]
>>> print less_equal(b, 0)
[[1 0 0 0 1]
 [1 1 0 0 0]
 [1 1 1 0 0]
 [0 1 1 1 0]
 [0 0 1 1 1]]

```

This last example has 1's where the corresponding elements are less than or equal to 0, and 0's everywhere else.

```

>>> view(greater(greeceBW, .3))
# shows a binary image with white where the pixel value was greater than
.3

```

Ufunc shorthands

Numeric defines a few functions which correspond to often-used uses of ufuncs: for example, `add.reduce()` is synonymous with the `sum()` utility function:

```

>>> a = arange(5)           # [0 1 2 3 4]
>>> print sum(a)           # 0 + 1 + 2 + 3 + 4
10

```

Similarly, `cumsum` is equivalent to `add.accumulate` (for "cumulative sum"), `product` to `multiply.reduce`, and `cumproduct` to `multiply.accumulate`.

Additional "utility" functions which are often useful are `alltrue` and `sometrue`, which are defined as `logical_and.reduce` and `logical_or.reduce` respectively:

```

>>> a = array([0,1,2,3,4])
>>> print greater(a,0)
[0 1 1 1 1]
>>> alltrue(greater(a,0))
0
>>> sometrue(greater(a,0))
1

```

7. Pseudo Indices

This chapter discusses pseudo-indices, which allow arrays to have their shapes modified by adding axes, sometimes only for the duration of the evaluation of a Python expression.

Consider multiplication of a rank-1 array by a scalar:

```
>>> a = array([1,2,3])
>>> a * 2
[2 4 6]
```

This should be trivial to you by now. We've just multiplied a rank-1 array by a scalar (which is converted to a rank-0 array). In other words, the rank-0 array was broadcast to the next rank. This works for adding some two rank-1 arrays as well:

```
>>> print a
[1 2 3]
>>> a + array([4])
[5 6 7]
```

but it won't work if either of the two rank-1 arrays have non-matching dimensions which aren't 1 – put another way, broadcast only works for dimensions which are either missing (e.g. a lower-rank array) or for dimensions of 1.

With this in mind, consider a classic task, matrix multiplication. Suppose we want to multiply the row vector [10,20] by the column vector [1,2,3].

```
>>> a = array([10,20])
>>> b = array([1,2,3])
>>> a * b
Traceback (innermost last):
  File "<stdin>", line 1, in ?
```

```
ValueError: frames are not aligned example
```

This makes sense – we're trying to multiply a rank-1 array of shape (2,) with a rank-1 array of shape (3,). This violates the laws of broadcast. What we really want to do is make the second vector a vector of shape (3,1), so that the first vector can be broadcast across the second axis of the second vector. One way to do this is to use the reshape function:

```
>>> a.shape
(2,)
>>> b.shape
(3,)
>>> b2 = reshape(b, (3,1))
>>> print b2
[[1]
 [2]
 [3]]
>>> b2.shape
(3, 1)
>>> print a * b2
[[10 20]
 [20 40]
 [30 60]]
```


This is such a common operation that a special feature was added (it turns out to be useful in many other places as well) – the `NewAxis` “pseudo-index”, originally developed in the Yorick language. `NewAxis` is an index, just like integers, so it is used inside of the slice brackets `[]`. It can be thought of as meaning “add a new axis here,” in much the same ways as adding a 1 to an array's shape adds an axis. Again, examples help clarify the situation:

```
>>> print b
[1 2 3]
>>> b.shape
(3,)
>>> c = b[:, NewAxis]
>>> print c
[[1]
 [2]
 [3]]
>>> c.shape
(3,1)
```

Why use such a pseudo-index over the reshape function or shape assignments? Often one doesn't really want a new array with a new axis, one just wants it for an intermediate computation. Witness the array multiplication mentioned above, without and with pseudo-indices:

```
>>> without = a * reshape(b, (3,1))
>>> with = a * b[:,NewAxis]
```

The second is much more readable (once you understand how `NewAxis` works), and it's much closer to the intended meaning. Also, it's independent of the dimensions of the array `b`. You might counter that using something like `reshape(b, (-1,1))` is also dimension-independent, but 1) would you argue that it's as readable? 2) how would you deal with rank-3 or rank-`N` arrays? The `NewAxis`-based idiom also works nicely with higher rank arrays, and with the `...` “rubber index” mentioned earlier. Adding an axis before the last axis in an array can be done simply with:

```
>>> a[... ,NewAxis, :]
```

8. Array Functions

Most of the useful manipulations on arrays are done with functions. This might be surprising given Python's object-oriented framework, and that many of these functions could have been implemented using methods instead. Choosing functions means that the same procedures can be applied to arbitrary python sequences, not just to arrays. For example, while `transpose([[1,2],[3,4]])` works just fine, `[[1,2],[3,4]].transpose()` can't work. This approach also allows uniformity in interface between functions defined in the Numeric Python system, whether implemented in C or in Python, and functions defined in extension modules. The use of array methods is limited to functionality which depends critically on the implementation details of array objects. Array methods are discussed in the next chapter.

We've already covered two functions which operate on arrays, `reshape` and `resize`.

`take(a, indices, axis=0)`

`take` is in some ways like the slice operations. It selects the elements of the array it gets as first argument based on the indices it gets as a second argument. Unlike slicing, however, the array returned by `take` has the same rank as the input array. This is again much easier to understand with an illustration:

```
>>> print a
[[ 0  1  2  3  4]
 [ 5  6  7  8  9]
 [10 11 12 13 14]
 [15 16 17 18 19]]
>>> print take(a, (0,))          # first row
[[ 0  1  2  3  4]]
>>> print take(a, (0,1))        # first and second row
[[0 1 2 3 4]
 [5 6 7 8 9]]
>>> print take(a, (0,-1))       # first and last row
[[ 0  1  2  3  4]
 [15 16 17 18 19]]
```

The optional third argument specifies the axis along which the selection occurs, and the default value (as in the examples above) is 0, the first axis. If you want another axis, then you can specify it:

```
>>> print take(a, (0,), 1)      # first column
[[ 0]
 [ 5]
 [10]
 [15]]
>>> print take(a, (0,1), 1)     # first and second column
[[ 0  1]
 [ 5  6]
 [10 11]
 [15 16]]
>>> print take(a, (0,-1), 1)    # first and last column
[[ 0  4]
 [ 5  9]
 [10 14]
 [15 19]]
```

This is considered to be a “structural” operation, because its result does not depend on the content of the arrays or the result of a computation on those contents but uniquely on the structure of the array. Like all such structural operations, the default axis is 0 (the first rank). I mention it here because later in this tutorial, we will see functions which have a default axis of -1.

Take is often used to create multidimensional arrays with the indices from a rank-1 array. As in the earlier examples, the shape of the array returned by `take()` is a combination of the shape of its first argument and the shape of the array that elements are “taken” from -- when that array is rank-1, the shape of the returned array has the same shape as the index sequence. This, as with many other facets of Numeric, is best understood by experiment.

```
>>> x = arange(10) * 100
>>> print x
[ 0 100 200 300 400 500 600 700 800 900]
>>> print take(x, [[2,4],[1,2]])
[[200 400]
 [100 200]]
```

A typical example of using `take()` is to replace the grey values in an image according to a “translation table”. For example, let’s consider a brightening of a greyscale image. The `view()` function defined in the NumTut package automatically scales the input arrays to use the entire range of grey values, except if the input arrays are of typecode ‘b’ unsigned bytes -- thus to test this brightening function, we’ll first start by converting the greyscale floating point array to a greyscale byte array:

```
>>> BW = (greeceBW*256).astype('b')
>>> view(BW) # shows black and white picture
```

We then create a table mapping the integers 0-255 to integers 0-255 using a “compressive nonlinearity”:

```
>>> table = (255- arange(256)**2 / 256).astype('b')
>>> view(table) # shows the conversion curve
```

To do the “taking” into an array of the right kind, we first create a blank image array with the same shape and typecode as the original array:

```
>>> BW2 = zeros(BW.shape, BW.typecode())
```

and then perform the `take()` operation

```
>>> BW2.flat[:] = take(table, BW.flat)
>>> view(BW2)
```

put (a, indices, values)

`put` is the opposite of `take`. The values of the array `a` at the locations specified in `indices` are set to the corresponding value of `values`. The array `a` must be a contiguous array. The argument `indices` can be any integer sequence object with values suitable for indexing into the flat form of `a`. The argument `values` must be any sequence of values that can be converted to the typecode of `a`.

```
>>> x = arange(6)
>>> put(x, [2,4], [20,40])
>>> print x
[ 0  1 20  3 40  5]
```

Note that the target array `a` is not required to be one-dimensional. Since `a` is contiguous and stored in row-major order, the array `indices` can be treated as indexing `a`’s elements in storage order.

The routine `put` is thus equivalent to the following (although the loop is in C for speed):

```
ind = array(indices, copy=0)
v = array(values, copy=0).astype(a.typecode())
for i in len(ind): a.flat[i] = v[i]
```

putmask(a, mask, values)

putmask sets those elements of a for which mask is true to the corresponding value in values. The array a must be contiguous. The argument mask must be an integer sequence of the same size (but not necessarily the same shape) as a. The argument values will be repeated as necessary; in particular it can be a scalar. The array values must be convertible to the type of a.

```
>>> x=arange(5)
>>> putmask(x, [1,0,1,0,1], [10,20,30,40,50])
>>> print x
[10 1 30 3 50]
>>> putmask(x, [1,0,1,0,1], [-1,-2])
>>> print x
[-1 1 -1 3 -1]
```

Note how in the last example, the third argument was treated as if it was [-1, -2, -1, -2, -1].

transpose(a, axes=None)

transpose takes an array and returns a new array which corresponds to a with the order of axes specified by the second argument. The default corresponds to flipping the order of all the axes (it is equivalent to a.shape[::-1] if a is the input array).

```
>>> print a
[[ 0  1  2  3  4]
 [ 5  6  7  8  9]
 [10 11 12 13 14]
 [15 16 17 18 19]]
>>> print transpose(a)
[[ 0  5 10 15]
 [ 1  6 11 16]
 [ 2  7 12 17]
 [ 3  8 13 18]
 [ 4  9 14 19]]
>>> greece.shape          # it's a 355x242 RGB picture
(355, 242, 3)
>>> view(greece)
# picture of greek street is shown
>>> view(transpose(greece, (1,0,2)))# swap x and y, not color axis!
# picture of greek street is shown sideways
```

repeat(a, repeats, axis=0)

repeat takes an array and returns an array with each element in the input array repeated as often as indicated by the corresponding elements in the second array. It operates along the specified axis. So, to stretch an array evenly, one needs the repeats array to contain as many instances of the integer scaling factor as the size of the specified axis:

```
>>> view(repeat(greece, 2*ones(greece.shape[0]))) # double in X
>>> view(repeat(greece, 2*ones(greece.shape[1]), 1)) # double in Y
```

choose(a, (b0, ..., bn))

a is an array of integers between 0 and n. The resulting array will have the same shape as a, with element selected from b0,...,bn as indicating by the value of the corresponding element in a.

Assume a is an array a that you want to "clip" so that no values are greater than 100.0.

```
>>> choose(greater(a, 100.0), (a, 100.0))
```

Everywhere that `greater(a, 100.0)` is false (ie. 0) this will ``choose" the corresponding value in `a`. Everywhere else it will ``choose" 100.0.

This works as well with arrays. Try to figure out what the following does:

```
>>> ret = choose(greater_than(a,b), (c,d))
```

ravel(a)

returns the argument array `a` as a 1d array. It is equivalent to `reshape(a, (-1,))` or `a.flat`. Unlike `a.flat`, however, `ravel` works with non-contiguous arrays.

```
>>> print x
[[ 0  1  2  3]
 [ 5  6  7  8]
 [10 11 12 13]]
>>> x.iscontiguous()
0
>>> x.flat
Traceback (innermost last):
  File "<stdin>", line 1, in ?
ValueError: flattened indexing only available for contiguous array
>>> ravel(x)
array([ 0,  1,  2,  3,  5,  6,  7,  8, 10, 11, 12, 13])
```

nonzero(a)

`nonzero()` returns an array containing the indices of the elements in `a` that are nonzero. These indices only make sense for 1d arrays, so the function refuses to act on anything else. As of 1.0a5 this function does not work for complex arrays.

where(condition, x, y)

`where(condition,x,y)` returns an array shaped like `condition` and has elements of `x` and `y` where `condition` is respectively true or false

compress(condition, a, axis=0)

returns those elements of `a` corresponding to those elements of `condition` that are nonzero. `condition` must be the same size as the given axis of `a`.

```
>>> print x
[0 1 2 3]
>>> print greater(x, 2)
[0 0 0 1]
>>> print compress(greater(x, 2), x)
[3]
```

diagonal(a, k=0)

returns the entries along the `k` th diagonal of `a` (`k` is an offset from the main diagonal). This is designed for 2d arrays. For larger arrays, it will return the diagonal of each 2d sub-array.

```
>>> print x
[[ 0  1  2  3  4]
 [ 5  6  7  8  9]
 [10 11 12 13 14]
 [15 16 17 18 19]
 [20 21 22 23 24]]
```

```

>>> print diagonal(x)
[ 0  6 12 18 24]
>>> print diagonal(x, 1)
[ 1  7 13 19]
>>> print diagonal(x, -1)
[ 5 11 17 23]

```

trace(a, k=0)

returns the sum of the elements in a along the k th diagonal.

```

>>> print x
[[ 0  1  2  3  4]
 [ 5  6  7  8  9]
 [10 11 12 13 14]
 [15 16 17 18 19]
 [20 21 22 23 24]]
>>> print trace(x)           # 0 + 6 + 12 + 18 + 24
60
>>> print trace(x, -1)      # 5 + 11 + 17 + 23
56
>>> print trace(x, 1)       # 1 + 7 + 13 + 19
40

```

searchsorted(a, values)

Called with a rank-1 array sorted in ascending order, `searchsorted()` will return the indices of the positions in a where the corresponding values would fit.

```

>>> print bin_boundaries
[ 0.  0.1  0.2  0.3  0.4  0.5  0.6  0.7  0.8  0.9  1. ]
>>> print data
[ 0.3029573  0.79585496  0.82714031  0.77993884  0.55069605  0.76043182
  0.28511823  0.29987358  0.40286206  0.68617903]
>>> print searchsorted(bin_boundaries, data)
[4 8 9 8 6 8 3 3 5 7]

```

This can be used for example to write a simple histogramming function:

```

>>> def histogram(a, bins):
...     n = searchsorted(sort(a), bins)
...     n = concatenate([n, [len(a)]])
...     return n[1:]-n[:-1]
...
>>> print histogram([0,0,0,0,0,0,0,.33,.33,.33], arange(0,1.0,.1))
[7 0 0 3 0 0 0 0 0 0]
>>> print histogram(sin(arange(0,10,.2)), arange(-1.2, 1.2, .1))
[0 0 4 2 2 2 0 2 1 2 1 3 1 3 1 3 2 3 2 3 4 9 0 0]

```

sort(a, axis=-1)

This function returns an array containing a copy of the data in a, with the same shape as a, but with the order of the elements along the specified axis sorted. The shape of the returned array is the same as a's. Thus, `sort(a, 3)` will be an array of the same shape as a, where the elements of a have been sorted along the fourth axis.

```

>>> print data
[[5 0 1 9 8]
 [2 5 8 3 2]]

```

```

[8 0 3 7 0]
[9 6 9 5 0]
[9 0 9 7 7]]
>>> print sort(data)           # Axis -1 by default
[[0 1 5 8 9]
 [2 2 3 5 8]
 [0 0 3 7 8]
 [0 5 6 9 9]
 [0 7 7 9 9]]
>>> print sort(data, 0)
[[2 0 1 3 0]
 [5 0 3 5 0]
 [8 0 8 7 2]
 [9 5 9 7 7]
 [9 6 9 9 8]]

```

argsort(a, axis=-1)

`argsort` will return the indices of the elements of `a` needed to produce `sort(a)`. In other words, for a rank-1 array, `take(a, argsort(a)) == sort(a)`.

```

>>> print data
[5 0 1 9 8]
>>> print sort(data)
[0 1 5 8 9]
>>> print argsort(data)
[1 2 0 4 3]
>>> print take(data, argsort(data))
[0 1 5 8 9]

```

argmax(a, axis=-1), argmin(a, axis=-1)

The `argmax()` function returns an array with the arguments of the maximum values of its input array `a` along the given axis. The returned array will have one less dimension than `a`. `argmin()` is just like `argmax()`, except that it returns the indices of the minima along the given axis.

```

>>> print data
[[9 6 1 3 0]
 [0 0 8 9 1]
 [7 4 5 4 0]
 [5 2 7 7 1]
 [9 9 7 9 7]]
>>> print argmax(data)
[0 3 0 2 0]
>>> print argmax(data, 0)
[0 4 1 1 4]
>>> print argmin(data)
[4 0 4 4 2]
>>> print argmin(data, 0)
[1 1 0 0 0]

```

fromstring(string, typecode)

Will return the array formed by the binary data given in string of the specified typecode. This is mainly used for reading binary data to and from files, it can also be used to exchange binary data with other modules that use python strings as storage (e.g. PIL). Note that this representation is dependent on the byte order. To find out the byte ordering used, use the `byteswapped()` method described on page 45.

dot(m1, m2)

The `dot()` function returns the dot product of `m1` and `m2`. This is equivalent to matrix multiply for rank-2 arrays (without the transpose). Somebody who does more linear algebra really needs to do this function right some day!

matrixmultiply(m1, m2)

The `matrixmultiply(m1, m2)` multiplies matrices or matrices and vectors as matrices rather than elementwise. Compare:

```
>>> print a
[[0 1 2]
 [3 4 5]]
>>> print b
[1 2 3]
>>> print a*b
[[ 0  2  6]
 [ 3  8 15]]
>>> print matrixmultiply(a,b)
[ 8 26]
```

clip(m, m_min, m_max)

The `clip` function creates an array with the same shape and typecode as `m`, but where every entry in `m` that is less than `m_min` is replaced by `m_min`, and every entry greater than `m_max` is replaced by `m_max`. Entries within the range `[m_min, m_max]` are left unchanged.

```
>>> a = arange(9, Float)
>>> clip(a, 1.5, 7.5)
1.5000 1.5000 2.0000 3.0000 4.0000 5.0000 6.0000 7.0000 7.5000
```

indices(shape, typecode=None)

The `indices` function returns an array corresponding to the shape given. The array returned is an array of a new shape which is based on the specified shape, but has an added dimension of length the number of dimensions in the specified shape. For example, if the shape specified by the `shape` argument is (3,4), then the shape of the array returned will be (2,3,4) since the length of (3,4) is 2. The contents of the returned arrays are such that the `i`th subarray (along index 0, the first dimension) contains the indices for that axis of the elements in the array. An example makes things clearer:

```
>>> i = indices((4,3))
>>> i.shape
(2, 4, 3)
>>> print i[0]
[[0 0 0]
 [1 1 1]
 [2 2 2]
 [3 3 3]]
>>> print i[1]
[[0 1 2]
 [0 1 2]
 [0 1 2]
 [0 1 2]]
```

So, `i[0]` has an array of the specified shape, and each element in that array specifies the index of that position in the subarray for axis 0. Similarly, each element in the subarray in `i[1]` contains the index of that position in the subarray for axis 1.

swapaxes(a, axis1, axis2)

Returns a new array which shares the data of `a`, but which has the two axes specified by `axis1` and `axis2` swapped. If `a` is of rank 0 or 1, `swapaxes` simply returns a new reference to `a`.

```
>>> x = arange(10)
>>> x.shape = (5,2,1)
>>> print x
[[[0]
  [1]
  [2]
  [3]
  [4]
  [5]
  [6]
  [7]
  [8]
  [9]]]
>>> y = swapaxes(x, 0, 2)
>>> print y.shape
(1, 2, 5)
>>> print y
[ [[0 2 4 6 8]
  [1 3 5 7 9]]]
```

concatenate((a0, a1, ... , an), axis=0)

Returns a new array containing copies of the data contained in all arrays `a0 ... an`. The arrays `ai` will be concatenated along the specified axis (0 by default). All arrays `ai` must have the same shape along every axis except for the one given. To concatenate arrays along a newly created axis, you can use `array((a0, ..., an))` as long as all arrays have the same shape.

```
>>> print x
[[ 0  1  2  3]
 [ 5  6  7  8]
 [10 11 12 13]]
>>> print concatenate((x,x))
[[ 0  1  2  3]
 [ 5  6  7  8]
 [10 11 12 13]
 [ 0  1  2  3]
 [ 5  6  7  8]
 [10 11 12 13]]
>>> print concatenate((x,x), 1)
[[ 0  1  2  3  0  1  2  3]
 [ 5  6  7  8  5  6  7  8]
 [10 11 12 13 10 11 12 13]]
>>> print array((x,x))
[[[ 0  1  2  3]
  [ 5  6  7  8]
  [10 11 12 13]]
 [[ 0  1  2  3]
  [ 5  6  7  8]
  [10 11 12 13]]]
```

innerproduct(a, b)

innerproduct produces the inner product of arrays a and b. It is equivalent to `matrixmultiply(a, transpose(b))`.

outerproduct(a,b)

outerproduct(a,b) produces the outer product of vectors a and b, that is `result[i, j] = a[i] * b[j]`

array_repr()

See section on Textual Representations of arrays.

array_str()

See section on Textual Representations of arrays.

resize(a, new_shape)

The `resize` function takes an array and a shape, and returns a new array with the specified shape, and filled with the data in the input array. Unlike the `reshape` function, the new shape does not have to yield the same size as the original array. If the new size is less than that of the input array, the returned array contains the appropriate data from the “beginning” of the old array. If the new size is greater than that of the input array, the data in the input array is repeated as many times as needed to fill the new array.

```
>>> x = arange(10)
>>> y = resize(x, (4,2))           # note that 4*2 < 10
>>> print x
[0 1 2 3 4 5 6 7 8 9]
>>> print y
[[0 1]
 [2 3]
 [4 5]
 [6 7]]
>>> print resize(array((0,1)), (5,5))# note that 5*5 > 2
[[0 1 0 1 0]
 [1 0 1 0 1]
 [0 1 0 1 0]
 [1 0 1 0 1]
 [0 1 0 1 0]]
```

diagonal(a, offset=0, axis1=-2, axis2=-1)

The `diagonal` function takes an array a, and returns an array of rank 1 containing all of the elements of a such that the difference between their indices along the specified axes is equal to the specified offset. With the default values, this corresponds to all of the elements of the diagonal of a along the last two axes. *Currently this is broken for offsets other than -1, 0 and 1, and for non-square arrays.*

repeat (a, counts, axis=0)

The `repeat` function uses repeated copies of a to create a result. The `axis` argument refers to the axis of x which will be replicated. The `counts` argument tells how many copies of each element to make. The length of `counts` must be the `len(shape(a)[axis])`.

In one dimension this is straightforward:

```
>>> y
array([0, 1, 2, 3, 4, 5])
>>> repeat(y, (1,2,0,2,2,3))
```

```
array([0, 1, 1, 3, 3, 4, 4, 5, 5, 5])
```

In more than one dimension it sometimes gets harder to understand. Consider for example this array `x` whose shape is (2,3).

```
>>> x
array([[0, 1, 2],
       [3, 4, 5]])

>>> repeat(x, (2,6))
array([[0, 1, 2],
       [0, 1, 2],
       [3, 4, 5],
       [3, 4, 5],
       [3, 4, 5],
       [3, 4, 5],
       [3, 4, 5],
       [3, 4, 5]])

>>> repeat(x, (6,3), 1)
array([[0, 0, 0, 0, 0, 0, 1, 1, 1],
       [2, 2, 2, 2, 2, 2, 3, 3, 3]])
```

convolve (a, v, mode=0)

The `convolve` function returns the linear convolution of two rank 1 arrays. The output is a rank 1 array whose length depends on the value of `mode` which is zero by default. Linear convolution can be used to find the response of a linear system to an arbitrary input. If the input arrays correspond to the coefficients of a polynomial and `mode=2`, the output of linear convolution corresponds to the coefficients of the product of the polynomials.

The `mode` parameter requires a bit of explanation. True linear convolution is only defined over infinite sequences. As both input arrays must represent finite sequences, the `convolve` operation assumes that the infinite sequences represented by the finite inputs are zero outside of their domain of definition. In other words, the sequences are zero-padded. If `mode` is 2, then the non-zero part of the full linear convolution is returned, so the output has length `len(a)+len(v)-1`. Call this output `f`. If `mode` is 0, then any part of `f` which was affected by the zero-padding is chopped from the result. In other words, let `b` be the input with smallest length and let `c` be the other input. The output when `mode` is 0 is the middle `len(c)-len(b)+1` elements of `f`. When `mode` is 1, the output is the same size as `c` and is equal to the middle `len(c)` elements of `f`.

cross_correlate (a, v, mode=0)

The `cross_correlate` function computes the cross-correlation between two rank 1 arrays. The output is a rank 1 array representing the inner product of `a` with shifted versions of `v`. This is very similar to convolution. The difference is that convolution reverses the axis of one of the input sequences but `cross_correlate` does not. In fact it is easy to verify that `convolve(a, v, mode) = cross_correlate(a, v[::-1], mode)`

where (condition, x, y)

The `where` function creates an array whose values are those of `x` at those indices where `condition` is true, and those of `y` otherwise. The shape of the result is the shape of `condition`. The type of the result is determined by the types of `x` and `y`. Either or both of `x` and `y` can be a scalar, which is then used for any element of `condition` which is true.

identity(n)

The `identity` function returns an `n` by `n` array where the diagonal elements are 1, and the off-diagonal elements are 0.

```
>>> print identity(5)
[[1 0 0 0 0]
 [0 1 0 0 0]
 [0 0 1 0 0]
 [0 0 0 1 0]
 [0 0 0 0 1]]
```

sum(a, index=0)

The `sum` function is a synonym for the `reduce` method of the `add` ufunc. It returns the sum of all of the elements in the sequence given along the specified axis (first axis by default).

```
>>> print x
[[ 0  1  2  3]
 [ 4  5  6  7]
 [ 8  9 10 11]
 [12 13 14 15]
 [16 17 18 19]]
>>> print sum(x)
[40 45 50 55]           # 0+4+8+12+16, 1+5+9+13+17,
2+6+10+14+18, ...
>>> print sum(x, 1)
[ 6 22 38 54 70]      # 0+1+2+3, 4+5+6+7, 8+9+10+11, ...
```

cumsum(a, index=0)

The `cumsum` function is a synonym for the `accumulate` method of the `add` ufunc.

product(a, index=0)

The `product` function is a synonym for the `reduce` method of the `multiply` ufunc.

cumproduct(a, index=0)

The `cumproduct` function is a synonym for the `accumulate` method of the `multiply` ufunc.

alltrue(a, index=0)

The `alltrue` function is a synonym for the `reduce` method of the `logical_and` ufunc.

sometrue(a, index=0)

The `sometrue` function is a synonym for the `reduce` method of the `logical_or` ufunc.

9. Array Methods

As we discussed at the beginning of the last chapter, there are very few array methods for good reasons, and these all depend on the the implementation details. They're worth knowing, though:

itemsize()

The `itemsize()` method applied to an array returns the number of bytes used by any one of its elements.

```
>>> a = arange(10)
>>> a.itemsize()
4
>>> a = array([1.0])
>>> a.itemsize()
8
>>> a = array([1], Complex)
>>> a.itemsize()
16
```

iscontiguous()

Calling an array's `iscontiguous()` method returns true if the memory used by A is contiguous. A non-contiguous array can be converted to a contiguous one by the `copy()` method. This is useful for interfacing to C routines only, as far as I know.

```
>>> XXX example
```

typecode()

The `'typecode()'` method returns the typecode of the array it is applied to. While we've been talking about them as Float, Int, etc., they are represented internally as characters, so this is what you'll get:

```
>>> a = array([1,2,3])
>>> a.typecode()
'1'
>>> a = array([1], Complex)
>>> a.typecode()
'D'
```

byteswapped()

The `byteswapped` method performs a byte swapping operation on all the elements in the array.

```
>>> print a
[1 2 3]
>>> print a.byteswapped()
[16777216 33554432 50331648]
```

tostring()

The `tostring` method returns a string representation of the data portion of the array it is applied to.

```
>>> a = arange(65,100)
>>> print a.tostring()
A B C D E F G H I J K L M N O P Q R S T
```

U V W X Y Z [\] ^ _ ` a b c

tolist()

Calling an array's `tolist()` method returns a hierarchical python list version of the same array:

```
>>> print a
[[65 66 67 68 69 70 71]
 [72 73 74 75 76 77 78]
 [79 80 81 82 83 84 85]
 [86 87 88 89 90 91 92]
 [93 94 95 96 97 98 99]]
>>> print a.tolist()
[[65, 66, 67, 68, 69, 70, 71], [72, 73, 74, 75, 76, 77, 78], [79, 80,
81, 82, 83, 84, 85], [86, 87, 88, 89, 90, 91, 92], [93, 94, 95, 96, 97,
98, 99]]
```

10. Array Attributes

We've already seen a very useful attribute of arrays, the `shape` attribute. There are three more, `flat`, `real` and `imaginary`.

flat

Accessing the `flat` attribute of an array returns the flattened, or `ravel()`'ed version of that array, without having to do a function call. The returned array has the same number of elements as the input array, but is of rank-1. One cannot set the `flat` attribute of an array, but one can use the indexing and slicing notations to modify the contents of the array:

```
>>> print a
[[0 1 2]
 [3 4 5]
 [6 7 8]]
>>> print a.flat
[0 1 2 3 4 5 6 7 8]
>>> a.flat = arange(9,18)
Traceback (innermost last):
  File "<stdin>", line 1, in ?
AttributeError: Attribute does not exist or cannot be set
>>> a.flat[4] = 100
>>> print a
[[ 0  1  2]
 [ 3 100 5]
 [ 6  7  8]]
>>> a.flat[:] = arange(9, 18)
>>> print a
[[ 9 10 11]
 [12 13 14]
 [15 16 17]]
```

real and imaginary

These attributes exist only for complex arrays. They return respectively arrays filled with the real and imaginary parts of their elements. `.imag` is a synonym for `.imaginary`. The arrays returned are not contiguous (except for arrays of length 1, which are always contiguous.). `.real`, `.imag` and `.imaginary` are modifiable:

```
>>> print x
[ 0.          +1.j          0.84147098+0.54030231j  0.90929743-0.41614684j]
>>> print x.real
[ 0.          0.84147098  0.90929743]
>>> print x.imag
[ 1.          0.54030231 -0.41614684]
>>> x.imag = arange(3)
>>> print x
[ 0.          +0.j  0.84147098+1.j  0.90929743+2.j]
>>> x = reshape(arange(10), (2,5)) + 0j# make complex array
>>> print x
[[ 0.+0.j  1.+0.j  2.+0.j  3.+0.j  4.+0.j]
```

```
[ 5.+0.j 6.+0.j 7.+0.j 8.+0.j 9.+0.j]]
>>> print x.real
[[ 0.  1.  2.  3.  4.]
 [ 5.  6.  7.  8.  9.]]
>>> print x.typecode(), x.real.typecode()
D d
>>> print x.itemsize(), x.imag.itemsize()
16 8
```


11. Special Topics

This chapter holds miscellaneous information which did not neatly fit in any of the other chapters.

Subclassing

Subclassing Numeric arrays is not possible due to a limitation of Python. The approach taken in the Masked Array facility (“Masked Arrays” on page 89) is one answer. `UserArray.py`, described below, can be subclassed, but this is often unsatisfactory unless you put in a similar effort to that in `MA`.

Code Organization

Numeric.py and friends

`Numeric.py` is the most commonly used interface to the Numeric extensions. It is a Python module which imports all of the exported functions and attributes from the `multiarray` module, and then defines some utility functions. As some of the functions defined in `Numeric.py` could someday be moved into a supporting C module, the utility functions and the `multiarray` object are documented together, in this section. The `multiarray` objects are the core of Numeric Python – they are extension types written in C which are designed to provide both space- and time-efficiency when manipulating large arrays of homogeneous data types, with special emphasis to numeric data types.

UserArray.py

In the tradition of `UserList.py` and `UserDict.py`, the `UserArray.py` module defines a class whose instances act in many ways like array objects.

Matrix.py

The `Matrix.py` python module defines a class `Matrix` which is a subclass of `UserArray`. The only differences between `Matrix` instances and `UserArray` instances is that the `*` operator on `Matrix` performs a matrix multiplication, as opposed to element-wise multiplication, and that the power operator `**` is disallowed for `Matrix` instances.

Precision.py

The `Precision.py` module contains the code which is used to determine the mapping between typecode names and values, by building small arrays and looking at the number of bytes they use per element.

ArrayPrinter.py

The `ArrayPrinter.py` module defines the functions used for default printing of arrays. See the section on Textual Representations of arrays on page 55,

Mlab.py

The `Mlab.py` module provides some functions which are compatible with the functions of the same name in the MATLAB programming language. These are:

bartlett(M)

returns the M-point Bartlett window.

blackman(M)

returns the M-point Blackman window.

corrcoef(x, y=None)

The correlation coefficient

cov(m,y=None)

returns the covariance

cumprod(m)

returns the cumulative product of the elements along the first dimension of m.

cumsum(m)

returns the cumulative sum of the elements along the first dimension of m.

diag(v, k=0)

returns the k-th diagonal if v is a matrix or returns a matrix with v as the k-th diagonal if v is a vector.

diff(x, n=1)

calculates the first-order, discrete difference approximation to the derivative.

eig(m)

returns the the eigenvalues of m in x and the corresponding eigenvectors in the rows of v.

eye(N, M=N, k=0, typecode=None)

returns a N-by-M matrix where the k-th diagonal is all ones, and everything else is zeros.

fliplr(m)

returns a 2-D matrix m with the rows preserved and columns flipped in the left/right direction. Only works with 2-D arrays.

flipud(m)

returns a 2-D matrix with the columns preserved and rows flipped in the up/down direction. Only works with 2-D arrays.

hamming(M)

returns the M-point Hamming window.

hanning(M)

returns the M-point Hanning window.

kaiser(M, beta)

returns a Kaiser window of length M with shape parameter beta. It depends on the cephes module for the modified bessel function i_0 .

max(m)

returns the maximum along the first dimension of m.

mean(m)

returns the mean along the first dimension of m. Note: if m is an integer array, integer division will occur.

median(m)

returns a mean of m along the first dimension of m.

min(m)

returns the minimum along the first dimension of m.

msort(m)

returns a sort along the first dimension of m as in MATLAB.

prod(m)

returns the product of the elements along the first dimension of m.

ptp(m)

returns the maximum - minimum along the first dimension of m.

rand(d1, ..., dn)

returns a matrix of the given dimensions which is initialized to random numbers from a uniform distribution in the range [0,1).

rot90(m,k=1)

returns the matrix found by rotating m by $k*90$ degrees in the counterclockwise direction.

sinc(x)

returns $\sin(\pi*x)/(\pi*x)$ at all points of array x.

squeeze(a)

removes any ones from the shape of a

std(m)

returns the standard deviation along the first dimension of m. The result is unbiased meaning division by $\text{len}(m)-1$.

sum(m)

returns the sum of the elements along the first dimension of m.

svd(m)

return the singular value decomposition of m [u,x,v]

trapz(y,x=None)

integrates $y = f(x)$ using the trapezoidal rule.

tri(N, M=N, k=0, typecode=None)

returns a N-by-M matrix where all the diagonals starting from lower left corner up to the k-th are all ones.

tril(m,k=0)

returns the elements on and below the k-th diagonal of m. $k=0$ is the main diagonal, $k > 0$ is above and $k < 0$ is below the main diagonal.

triu(m,k=0)

returns the elements on and above the k-th diagonal of m. $k=0$ is the main diagonal, $k > 0$ is above and $k < 0$ is below the main diagonal.

The multiarray object

The array objects which Numeric Python manipulates is actually a multiarray object, given this name to distinguish it from the one-dimensional array object defined in the standard array module. From here on, however, the terms array and multiarray will be used interchangeably to refer to the new object type. multiarray objects are homogeneous multidimensional sequences. Starting from the back, they are sequences. This means that they are container (compound) objects, which contain references to other objects. They are multidimensional, meaning that unlike standard Python sequences which define only a single dimension along which one can iterate through the contents, multiarray objects can have up to 40 dimensions.¹ Finally, they are homogeneous. This means that every object in a multiarray must be of the same type. This is done for efficiency reasons -- storing the type of the contained objects once in the array means that the process of finding the type-specific operation to operate on each element in the array needs to be done only once per array, as opposed to once per element. Furthermore, as the main purpose of these arrays is to process numbers, the numbers can be stored directly, and not as full-fledged Python objects (PyObject *), thus yielding memory savings. It is however possible to make arrays of Python objects, which relinquish both the space and time efficiencies but allow heterogeneous contents (as we shall see, these arrays are still homogeneous from the Numeric perspective, they are just arrays of Python object references).

Typecodes

The kind of number stored in an array is described by its typecode. This code is stored internally as a single-character Python string, but more descriptive names corresponding to the typecodes are made available to the Python programmer in the Precision.py module. The typecodes are defined as follows:

Table 3: Typecode Listing

Variable defined in Typecode module	Typecode character	Description
Char	'c'	Single-character strings
PyObject	'O'	Reference to Python object
UnsignedInt8	'b'	Unsigned integer using a single byte.
Int	'l'	Python standard integers (i.e. C long integers)
Float	'd'	Python standard floating point numbers (i.e. C double-precision floats)
n/a	'f'	Single-precision floating point numbers
Complex	'D'	Complex numbers consisting of two double-precision floats
n/a	'F'	Complex numbers consisting of two single-precision floats
Int0, Int8, Int16, Int32, Int64, Int128	n/a	These correspond to machine-dependent typecodes: Int0 returns the typecode corresponding to the smallest available integer, Int8 that corresponding to the smallest available integer with at least 8 bits, Int16 that with at least 16 bits, etc. If a typecode is not available (e.g. Int64 on a 32-bit machine), the variable is not defined.

1. This limit is modifiable in the source code if higher dimensionality is needed.

Table 3: Typecode Listing

Variable defined in Typecode module	Typecode character	Description
Float0, Float8, Float16, Float32, Float64, Float128	n/a	Same as Int0, Int8 etc. except for floating point numbers.
Complex0, Complex8, Complex16, Complex32, Complex64, Complex128	n/a	Same as Float0, etc., except that the number of bits refers to the precision of each of the two (real and imaginary) parts of the complex number.

Note on number format: the binary format used by Python is that of the underlying C library. [notes about IEEE formats, etc?]

Indexing in and out, slicing

Indexing arrays works like indexing of other Python sequences, but supports some extensions which are as of yet not implemented for other sequence types¹. The standard [start:stop] notation is supported, with start defaulting to 0 (the first index position) and stop defaulting to the length of the sequence, as for lists and tuples. In addition, there is an optional stride argument, which specifies the stride size between successive indices in the slice. It is expressed by a integer following a second `:` immediately after the usual start:stop slice. Thus `[0:11:2]` will slice the array at indices 0, 2, 4, .. 10. The start and stop indices are optional, but the first `:` must be specified for the stride interpretation to occur. Therefore, `[: : 2]` means slice from beginning to end, with a stride of 2 (i.e. skip an index for each stride). If the start index is omitted and the stride is negative, the indexing starts from the end of the sequence and works towards the beginning of the sequence. If the stop index is omitted and the stride is negative, the indexing stops at the beginning of the sequence.

```
>>> print x
[ 0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19]
>>> print x[10]
10
>>> print x[:10]
[0 1 2 3 4 5 6 7 8 9]
>>> print x[5:15:3]
[ 5  8 11 14]
>>> print x[:10:2]
[0 2 4 6 8]
>>> print x[10::-2]
[10 8 6 4 2 0]
>>> print x[::-1]
[19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0]
```

It is important to note that the out-of-bounds conditions follow the same rules as standard Python indexing, so that slices out of bounds are trimmed to the sequence boundaries, but element indexing with out-of-bound indices yields an `IndexError`:

```
>>> print x[:100]
[ 0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19]
>>> print x[-200:4]
[0 1 2 3]
>>> x[100]
```

1. The Python syntax can allow other Python datatypes to use both the stride notation and multidimensional indexing, and it is relatively simple to write Python classes which support these operations. See the Python Reference manual for details.

```
Traceback (innermost last):
  File "<stdin>", line 1, in ?
IndexError: index out of bounds
```

The second difference between array indexing and other sequences is that arrays provide multidimensional indexing. An array of rank N can be indexed with up to N indices or slices (or combinations thereof. Indices should be integers (with negative integers indicating offsets from the end of the dimension, as for other Python sequences), and slices can have, as explained above, one or two ':'s separating integer arguments. These indices and slices must be separated by commas, and correspond to sequential dimensions starting from the leftmost (first) index on. Thus `a[3]` means index 3 along dimension 0. `a[3, :, -4]` means the slice of a along three dimensions: index 3 along the first dimension, the entire range of indices along the second dimension, and the 4th from the end index along the third dimension. If the array being indexed has more dimensions than are specified in the multidimensional slice, those dimensions are assumed to be sliced from beginning to end. Thus, if a is a rank 3 array,

```
a[0] == a[0,:] == a[0,:,:]
```

Ellipses

A special slice element called Ellipses (and written `...`) is used to refer to a variable number of slices from beginning to end along the current dimension. It is a shorthand for a set of such slices, specifically the number of dimensions of the array being indexed minus those which are already specified. Only the first (leftmost) Ellipses in an multidimensional slice is expanded, while the others are single dimensional slices from beginning to end.

Thus, if a is a rank-6 array,

```
a[3,:,:,:,-1,:] == a[3,...,-1,:] == a[3,...,-1,...].
```

NewAxis

There is another special symbol which can be used inside indexing operations to create new dimensions in the returned array. The reference `NewAxis`, used as one of the comma-separated slice elements, does not change the selection of the subset of the array being indexed, but changes the shape of the array returned by the indexing operation, so that an additional dimension (of length 1) is created, at the dimension position corresponding to the location of `NewAxis` within the indexing sequence. Thus, `a[:,3,NewAxis,-3]` will perform the indexing of a corresponding to the slice `[a[:,3,-3]`, but will also modify the shape of a so that the new shape of a is `(a.shape[0], a.shape[1], 1, a.shape[2])`. This operation is especially useful in conjunction with the broadcasting feature described next, as it replaces a lengthy but common operation with a simple notation (in the example above, the same effect can be had with

```
reshape(a[:,3,-1], (a.shape[0], a.shape[1], 1, a.shape[2])).
```

Set-indexing and Broadcasting

The indexing rules described so far specify exactly the behavior of get-indexing. For set-indexing, the rules are exactly the same, and describe the slice of the array on the left hand side of the assignment operator which is the target of the assignment. The only point left to mention is the process of assigning from the source (on the right hand side of the assignment) to the target (on the left hand side).

If both source and target have the same shape, then the assignment is done element by element. The typecode of the target specifies the casting which can be applied in the case of a typecode mismatch between source and target. If the typecode of the source is "lower" than that of the target, then an 'up-cast' is performed and no loss in precision results. If the typecode of the source is "higher" than that of the target, then a downcast is performed, which may lose precision (as discussed in the description of the array call, these casts are truncating casts, not rounding casts). Complex numbers cannot be cast to non-complex numbers.

If the source and the target have different shapes, Numeric Python attempts to broadcast the contents of the source over the range of the target. This broadcasting occurs for all dimensions where the source has dimension 1 or 0 (i.e., is absent). If there exists a dimension for which the two arrays have differing lengths, and the length of that dimension in the source is not 1, then the assignment fails and an exception (`ValueError`) is raised, notifying the user that the arrays are not aligned.

Axis specifications

In many of the functions defined in this document, indices are used to refer to axes. The numbering scheme is the same as that used by indexing in Python: the first (leftmost) axis is axis 0, the second axis is axis 1, etc. Axis -1 refers to the last axis, -2 refers to the next-to-last axis, etc.

Textual representations of arrays

The algorithm used to display arrays as text strings is defined in the file `ArrayPrinter.py`, which defines a function `array2string` (imported into Numeric's namespace) which offers considerable control over how arrays are output. The range of options to the `array2string` function will be described first, followed by a description of which options are used by default by `str` and `repr`.

Note that the optional package `MA`, if imported, modifies this process so that very long arrays are not printed; rather, a summary of their shape and type are shown. You may wish to import `MA` even if you do not use it otherwise, to get this effect, because without it accidentally attempting to print a very long array can take a very long time to convert, giving the appearance that the program has hung.

```
array2string(a, max_line_width = None, precision = None,
             suppress_small = None, separator=' ', array_output=0):
```

The `array2string` function takes an array and returns a textual representation of it. Each dimension is indicated by a pair of matching square brackets (`[]`), within which each subset of the array is output. The orientation of the dimensions is as follows: the last (rightmost) dimension is always horizontal, so that the frequent rank-1 arrays use a minimum of screen real-estate. The next-to-last dimension is displayed vertically if present, and any earlier dimension is displayed with additional bracket divisions. For example:

```
>>> a = arange(24)
>>> print array2string(a)
[ 0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23]
>>> a.shape = (2,10)
>>> print array2string(a)
[[ 0  1  2  3  4  5  6  7  8  9 10 11]
 [12 13 14 15 16 17 18 19 20 21 22 23]]
>>> a.shape = (2,3,4)
>>> print array2string(a)
[[[ 0  1  2  3]
 [ 4  5  6  7]
 [ 8  9 10 11]]
 [[12 13 14 15]
 [16 17 18 19]
 [20 21 22 23]]]
```

The `max_line_width` argument specifies the maximum number of characters which the `array2string` routine uses in a single line. If it is set to `None`, then the value of the `sys.output_line_width` attribute is looked up. If it exists, it is used. If not, the default of 77 characters is used.

```
>>> print array2string(x)
[ 0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25
 26 27 28 29]
>>> sys.output_line_width = 30
>>> print array2string(x)
[ 0  1  2  3  4  5  6  7  8  9
```

```

10 11 12 13 14 15 16 17
18 19 20 21 22 23 24 25
26 27 28 29]

```

The `precision` argument specifies the number of digits after the decimal point which are used. If a value of `None` is used, the value of the `sys.float_output_precision` is looked up. If it exists, it is used. If not, the default of 8 digits is used.

```

>>> x = array((10.11111111111123123111, pi))
>>> print array2string(x)
[ 10.11111111  3.14159265]
>>> print array2string(x, precision=3)
[ 10.111  3.142]
>>> sys.float_output_precision = 2
>>> print array2string(x)
[ 10.11  3.14]

```

The `suppress_small` argument specifies whether small values should be suppressed (and output as 0). If a value of `None` is used, the value of the `sys.float_output_suppress_small` is looked up. If it exists, it is used (all that matters is whether it evaluates to true or false). If not, the default of 0 (false) is used. This variable also interacts with the precision parameters, as it can be used to suppress the use of exponential notation.

```

>>> print x
[ 1.00000000e-005  3.14159265e+000]
>>> print array2string(x)
[ 1.00000000e-005  3.14159265e+000]
>>> print array2string(x, suppress_small=1)
[ 0.00001  3.14159265]
>>> print array2string(x, precision=3)
[ 1.000e-005  3.142e+000]
>>> print array2string(x, precision=3, suppress_small=1)
[ 0.  3.142]

```

The `separator` argument is used to specify what character string should be placed between two numbers which do not straddle a dimension. The default is a single space.

```

>>> print array2string(x)
[ 0 100 200 300 400 500 600 700 800 900 100]
>>> print array2string(x, separator = ', ')
[ 0, 100, 200, 300, 400, 500, 600, 700, 800, 900, 100]

```

Finally, the last attribute, `array_output`, specifies whether to prepend the string "array(" and append either the string ")" or ", 'X'" where X is a typecode for non-default typecodes (in other words, the typecode will only be displayed if it is not that corresponding to Float, Complex or Int, which are the standard typecodes associated with floating point numbers, complex numbers and integers respectively). The `array()` is so that an eval of the returned string will return an array object (provided a comma separator is also used).

```

>>> array2string(arange(3))
[0 1 2]
>>> eval(array2string(arange(3), array_output=1))
Traceback (innermost last):
  File "<stdin>", line 1, in ?
  File "<string>", line 1
    array([0 1 2])
           ^
SyntaxError: invalid syntax
>>> type(eval(array2string(arange(3), array_output=1, separator=', ')))
<type 'array'>
>>> array2string(arange(3), array_output=1)
'array([0, 1, 2])'
>>> array2string(zeros((3,),'i') + arange(3), array_output=1)

```



```
"array([0, 1, 2], 'i')"
```

The `str` and `repr` operations on arrays call `array2string` with the `max_line_width`, `precision` and `suppress_small` all set to `None`, meaning that the defaults are used, but that modifying the attributes in the `sys` module will affect array printing. `str` uses the default separator and does not use the `array()` text, while `repr` uses a comma as a separator and does use the `array(...)` text.

```
>>> x = arange(3)
>>> print x
[0 1 2]
>>> str(x)
'[0 1 2]'
>>> repr(x)
'array([0, 1, 2])'          # note the array(...) and ','s
>>> x = arange(0,.01,.001)
>>> print x
[ 0.      0.001  0.002  0.003  0.004  0.005  0.006  0.007  0.008  0.009]
>>> import sys
>>> sys.float_output_precision = 2
>>> print x
[ 0.    0.    0.    0.    0.    0.01  0.01  0.01  0.01  0.01]
```

Comparisons

Comparisons of multiarray objects results using the normal comparison operators (such as `==` or `>`) result in exceptions. Python requires that the result of a comparison be a scalar, not the vector we would want as the result of elementwise comparison.

Therefore, for comparisons you must use the routines for comparison describe in “Logical Ufuncs” on page 30.

Pickling and Unpickling -- storing arrays on disk

This documentation has not yet been written, but pickling of Numeric arrays is possible.

Dealing with floating point exceptions

Attempts to use NaN's as missing values have proven frustrating and not very portable. Consider “Masked Arrays” on page 89 instead.

`fpectl` should be documented here ...

12. Writing a C extension to NumPy

Introduction

There are two applications that require using the NumPy array type in C extension modules:

- Access to numerical libraries: Extension modules can be used to make numerical libraries written in C (or languages linkable to C, such as Fortran) accessible to Python programs. The NumPy array type has the advantage of using the same data layout as arrays in C and Fortran.
- Mixed-language numerical code: In most numerical applications, only a small part of the total code is CPU time intensive. Only this part should thus be written in C, the rest can be written in Python. NumPy arrays are important for the interface between these two parts, because they provide equally simple access to their contents from Python and from C.

This document is a tutorial for using NumPy arrays in C extensions.

Preparing an extension module for NumPy arrays

To make NumPy arrays available to an extension module, it must include the header file `arrayobject.h`, after the header file `Python.h` that is obligatory for all extension modules. The file `arrayobject.h` comes with the NumPy distribution; depending on where it was installed on your system you might have to tell your compiler how to find it. By default Distutils installed in a subdirectory `Numeric` in your Python include path, and so you should include it this way:

```
#include "Numeric/arrayobject.h"
```

Important Tip

Is your C extension using Numeric blowing up? Maybe you didn't call `import_array()`.

In addition to including `arrayobject.h`, the extension must call `import_array()` in its initialization function, after the call to `Py_InitModule()`. This call makes sure that the module which implements the array type has been imported, and initializes a pointer array through which the NumPy functions are called. If you forget this call, your extension module will crash on the first call to a NumPy function! If you will be manipulating ufunc objects, you should also include the file `ufuncobject.h`, also available as part of the NumPy distribution in the `Include` directory and usually installed in subdirectory `Numeric`.

All of the rules related to writing extension modules for Python apply. The reader unfamiliar with these rules is encouraged to read the standard text on the topic, "Extending and Embedding the Python Interpreter," available as part of the standard Python documentation distribution.

Accessing NumPy arrays from C

Types and Internal Structure

NumPy arrays are defined by the structure `PyArrayObject`, which is an extension of the structure `PyObject`. Pointers to `PyArrayObject` can thus safely be cast to `PyObject` pointers, whereas the inverse is safe only if the object is known to be an array. The type structure corresponding to array objects is `PyArray_Type`. The structure `PyArrayObject` has four elements that are needed in order to access the array's data from C code:

`int nd`

The number of dimensions in the array.

`int *dimensions`

A pointer to an array of `nd` integers, describing the number of elements along each dimension. The sizes are in the conventional order, so that for any array `a`,

`a.shape==(dimensions[0], dimensions[1], ..., dimensions[nd]).`

`int *strides`

A pointer to an array of `nd` integers, describing the address offset between two successive data elements along each dimension. Note that strides can also be negative! Each number gives the number of bytes to add to a pointer to get to the next element in that dimension. For example, if `myptr` currently points to element of a rank-5 array at indices `1, 0, 5, 3, 2` and you want it to point to element `1, 0, 5, 4, 2` then you should add `strides[3]` to the pointer: `myptr += strides[3]`. This works even if (and is especially useful when) the array is not contiguous in memory.

`char *data`

A pointer to the first data element of the array.

The address of a data element can be calculated from its indices and the data and strides pointers. For example, element `[i, j]` of a two-dimensional array has the address `data + i*array->strides[0] + j*array->strides[1]`. Note that the stride offsets are in bytes, not in storage units of the array elements. Therefore address calculations must be made in bytes as well, starting from the data pointer, which is always a char pointer. To access the element, the result of the address calculation must be cast to a pointer of the required type. The advantage of this arrangement is that purely structural array operations (indexing, extraction of sub-arrays, etc.) do not have to know the type of the array elements.

Element data types

The type of the array elements is indicated by a type number, whose possible values are defined as constants in `arrayobject.h`, as given in Table 3.

Table 4: C constants corresponding to storage types

Constant	element data type
<code>PyArray_CHAR</code>	<code>char</code>
<code>PyArray_UBYTE</code>	<code>unsigned char</code>
<code>PyArray_SBYTE</code>	<code>signed char</code>
<code>PyArray_SHORT</code>	<code>short</code>
<code>PyArray_INT</code>	<code>int</code>
<code>PyArray_LONG</code>	<code>long</code>
<code>PyArray_FLOAT</code>	<code>float</code>
<code>PyArray_DOUBLE</code>	<code>double</code>
<code>PyArray_CFLOAT</code>	<code>float[2]</code>
<code>PyArray_CDOUBLE</code>	<code>double[2]</code>
<code>PyArray_OBJECT</code>	<code>PyObject *</code>

The type number is stored in `array->descr->type_num`. Note that the names of the element type constants refer to the C data types, not the Python data types. A Python `int` is equivalent to a C `long`, and a Python `float` corresponds to a C `double`. Many of the element types listed above do not have corresponding Python scalar types (e.g. `PyArray_INT`).

Contiguous arrays

An important special case of a NumPy array is the contiguous array. This is an array whose elements occupy a single contiguous block of memory and have the same order as a standard C array. In a contiguous array, the value of `array->strides[i]` is equal to the size of a single array element times the product of `array->dimensions[j]` for `j` up to `i-1`. Arrays that are created from scratch are always contiguous; non-contiguous arrays are the result of indexing and other structural array operations. The main advantage of contiguous arrays is easier handling in C; the pointer `array->data` is cast to the required type and then used like a C array, without any reference to the stride values. This is particularly important when interfacing to existing libraries in C or Fortran, which typically require this standard data layout. A function that requires input arrays to be contiguous must call the conversion function `PyArray_ContiguousFromObject()`, described in the section "Accepting input data from any sequence type".

Zero-dimensional arrays

NumPy permits the creation and use of zero-dimensional arrays, which can be useful to treat scalars and higher-dimensional arrays in the same way. However, library routines for general use should not return zero-dimensional arrays, because most Python code is not prepared to handle them. Moreover, zero-dimensional arrays can create confusion because they behave like ordinary Python scalars in many circumstances but are of a different type. A comparison between a Python scalar and a zero-dimensional array will always fail, for example, even if the values are the same. NumPy provides a conversion function from zero-dimensional arrays to Python scalars, which is described in the section "Returning arrays from C functions".

A simple example

The following function calculates the sum of the diagonal elements of a two-dimensional array, verifying that the array is in fact two-dimensional and of type `PyArray_DOUBLE`.

```
static PyObject *
trace(PyObject *self, PyObject *args)
{
    PyArrayObject *array;
    double sum;
    int i, n;

    if (!PyArg_ParseTuple(args, "O!",
                           &PyArray_Type, &array))
        return NULL;
    if (array->nd != 2 || array->descr->type_num != PyArray_DOUBLE) {
        PyErr_SetString(PyExc_ValueError,
                        "array must be two-dimensional and of type float");
        return NULL;
    }

    n = array->dimensions[0];
    if (n > array->dimensions[1])
        n = array->dimensions[1];
    sum = 0.;
    for (i = 0; i < n; i++)
        sum += *(double *) (array->data + i*array->strides[0] + i*array->strides[1]);
}
```

```
    return PyFloat_FromDouble(sum);
}
```

Accepting input data from any sequence type

The example in the last section requires its input to be an array of type double. In many circumstances this is sufficient, but often, especially in the case of library routines for general use, it would be preferable to accept input data from any sequence (lists, tuples, etc.) and to convert the element type to double automatically where possible. NumPy provides a function that accepts arbitrary sequence objects as input and returns an equivalent array of specified type (this is in fact exactly what the array constructor `Numeric.array()` does in Python code):

```
PyObject *
PyArray_ContiguousFromObject(PyObject *object,
                             int type_num,
                             int min_dimensions,
                             int max_dimensions);
```

The first argument, `object`, is the sequence object from which the data is taken. The second argument, `type_num`, specifies the array element type (see the table in the section “Element data types”. If you want the function to select the “smallest” type that is sufficient to store the data, you can pass the special value `PyArray_NOTYPE`. The remaining two arguments let you specify the number of dimensions of the resulting array, which is guaranteed to be no smaller than `min_dimensions` and no larger than `max_dimensions`, except for the case `max_dimensions == 0`, which means that no upper limit is imposed.

If the input data is not compatible with the type or dimension restrictions, an exception is raised. Since the array returned by `PyArray_ContiguousFromObject()` is guaranteed to be contiguous, this function also provides a method of converting a non-contiguous array to a contiguous one. If the input object is already a contiguous array of the specified type, it is passed on directly; there is thus no performance or memory penalty for calling the conversion function when it is not required. Using this function, the example from the last section becomes

```
static PyObject *
trace(PyObject *self, PyObject *args)
{
    PyObject *input;
    PyArrayObject *array;
    double sum;
    int i, n;

    if (!PyArg_ParseTuple(args, "O", &input))
        return NULL;
    array = (PyArrayObject *)
        PyArray_ContiguousFromObject(input, PyArray_DOUBLE, 2, 2);
    if (array == NULL)
        return NULL;

    n = array->dimensions[0];
    if (n > array->dimensions[1])
        n = array->dimensions[1];
    sum = 0.;
    for (i = 0; i < n; i++)
        sum += *(double *) (array->data + i*array->strides[0] + i*array->strides[1]);

    Py_DECREF(array);
    return PyFloat_FromDouble(sum);
}
```

Note that no explicit error checking is necessary in this version, and that the array reference that is returned by `PyArray_ContiguousFromObject()` must be destroyed by calling `Py_DECREF()`.

Creating NumPy arrays

NumPy arrays can be created by calling the function

```
PyObject *
PyArray_FromDims(int n_dimensions,
                 int dimensions[n_dimensions],
                 int type_num);
```

The first argument specifies the number of dimensions, the second one the length of each dimension, and the third one the element data type (see the table in the section “Element data types”. The array that is returned is contiguous, but the contents of its data space are undefined. There is a second function which permits the creation of an array object that uses a given memory block for its data space:

```
PyObject *
PyArray_FromDimsAndData(int n_dimensions,
                        int dimensions[n_dimensions]
                        int item_type
                        char *data);
```

The first three arguments are the same as for `PyArray_FromDims()`. The fourth argument is a pointer to the memory block that is to be used as the array's data space. It is the caller's responsibility to ensure that this memory block is not freed before the array object is destroyed. With few exceptions (such as the creation of a temporary array object to which no reference is passed to other functions), this means that the memory block may never be freed, because the lifetime of Python objects are difficult to predict. Nevertheless, this function can be useful in special cases, for example for providing Python access to arrays in Fortran common blocks.

Returning arrays from C functions

Array objects can of course be passed out of a C function just like any other object. However, as has been mentioned before, care should be taken not to return zero-dimensional arrays unless the receiver is known to be prepared to handle them. An equivalent Python scalar object should be returned instead. To facilitate this step, NumPy provides a special function

```
PyObject *
PyArray_Return(PyArrayObject *array);
```

which returns the array unchanged if it has one or more dimensions, or the appropriate Python scalar object in case of a zero-dimensional array.

A less simple example

The function shown below performs a matrix-vector multiplication by calling the BLAS function `DGEMV`. It takes three arguments: a scalar prefactor, the matrix (a two-dimensional array), and the vector (a one-dimensional array). The return value is a one-dimensional array. The input values are checked for consistency. In addition to providing an illustration of the functions explained above, this example also demonstrates how a Fortran routine can be integrated into Python. Unfortunately, mixing Fortran and C code involves machine-specific peculiarities. In this example, two assumptions have been made:

- The Fortran function `DGEMV` must be called from C as `dgemv_`. Many Fortran compilers apply this rule, but the C name could also be `dgemv` or `DGEMV` (or in principle anything else; there is no fixed standard).
- Fortran integers are equivalent to C longs, and Fortran double precision numbers are equivalent to C doubles. This works for all systems that I have personally used, but again there is no standard.

Also note that the libraries that this function must be linked to are system-dependent; on my Linux system (using `gcc/g77`), the libraries are `blas` and `f2c`. So here is the code:

```
static PyObject *
matrix_vector(PyObject *self, PyObject *args)
```

```

{
    PyObject *input1, *input2;
    PyArrayObject *matrix, *vector, *result;
    int dimensions[1];
    double factor[1];
    double real_zero[1] = {0.};
    long int_one[1] = {1};
    long dim0[1], dim1[1];

    extern dgemv_(char *trans, long *m, long *n,
                 double *alpha, double *a, long *lda,
                 double *x, long *incx,
                 double *beta, double *Y, long *incy);

    if (!PyArg_ParseTuple(args, "dOO", factor, &input1, &input2))
        return NULL;
    matrix = (PyArrayObject *)
        PyArray_ContiguousFromObject(input1, PyArray_DOUBLE, 2, 2);
    if (matrix == NULL)
        return NULL;
    vector = (PyArrayObject *)
        PyArray_ContiguousFromObject(input2, PyArray_DOUBLE, 1, 1);
    if (vector == NULL)
        return NULL;
    if (matrix->dimensions[1] != vector->dimensions[0]) {
        PyErr_SetString(PyExc_ValueError,
            "array dimensions are not compatible");
        return NULL;
    }

    dimensions[0] = matrix->dimensions[0];
    result = (PyArrayObject *)PyArray_FromDims(1, dimensions,
        PyArray_DOUBLE);
    if (result == NULL)
        return NULL;

    dim0[0] = (long)matrix->dimensions[0];
    dim1[0] = (long)matrix->dimensions[1];
    dgemv_("T", dim1, dim0, factor, (double *)matrix->data, dim1,
        (double *)vector->data, int_one,
        real_zero, (double *)result->data, int_one);

    return PyArray_Return(result);
}

```

Note that `PyArray_Return()` is not really necessary in this case, since we know that the array being returned is one-dimensional. Nevertheless, it is a good habit to always use this function; its performance cost is practically zero.

13. C API Reference

This chapter describes the API for ArrayObjects and Ufuncs.

ArrayObject C Structure and API

Structures

The PyArrayObject is, like all Python types, a kind of PyObject. Its definition is:

```
typedef struct {
    PyObject_HEAD
    char *data;
    int nd;
    int *dimensions, *strides;
    PyObject *base;
    PyArray_Descr *descr;
    int flags;
} PyArrayObject;
```

Where PyObject_HEAD is the standard PyObject header, and the other fields are:

`char *data`

A pointer to the first data element of the array.

`int nd`

The number of dimensions in the array.

`int *dimensions`

A pointer to an array of `nd` integers, describing the number of elements along each dimension. The sizes are in the conventional order, so that for any array `a`, `a.shape==(dimensions[0], dimensions[1], ..., dimensions[nd])`.

`int *strides`

A pointer to an array of `nd` integers, describing the address offset between two successive data elements along each dimension. Note that strides can also be negative! Each number gives the number of bytes to add to a pointer to get to the next element in that dimension. For example, if `myptr` currently points to an element in a rank-5 array at indices `1, 0, 5, 3, 2` and you want it to point to element `1, 0, 5, 4, 2` then you should add `strides[3]` to the pointer: `myptr += strides[3]`. This works even if (and is especially useful when) the array is not contiguous in memory.

`PyObject *base`

Used internally in arrays that are created as slices of other arrays. Since the new array shares its data area with the old one, the original array's reference count is incremented. When the subarray is garbage collected, the base array's reference count is decremented.

`PyArray_Descr *descr`

See below.

`int flags`

A bitfield indicating whether the array:

- is contiguous (rightmost bit)
- owns the dimensions (next bit to the left) (???)
- owns the strides (next bit to the left) (???)
- owns the data area

The ownership bits are used by NumPy internally to manage memory allocation and deallocation. They can be false if the array is the result of e.g. a slicing operation on an existing array.

`PyArray_Descr *descr`

a pointer to a data structure that describes the array and has some handy functions. The slots in this structure are:

`PyArray_VectorUnaryFunc *cast[]`

an array of function pointers which will cast this arraytype to each of the other data types.

`PyArray_GetItemFunc *getitem`

a pointer to a function which returns a PyObject of the appropriate type given a (char) pointer to the data to get.

`PyArray_SetItemFunc *setitem`

a pointer to a function which sets the element pointed to by the second argument to converted Python Object given as the first argument.

`int type_num`

A number indicating the datatype of the array (i.e. a `PyArray_XXXX`)

`char *one`

A pointer to a representation of one for this datatype.

`char *zero`

A pointer to a representation of zero for this datatype (especially useful for `PyArray_OBJECT` types)

`char type`

A character representing the array's typecode (one of 'cb1sildfDO').

The ArrayObject API

In the following `op` is a pointer to a PyObject and `arp` is a pointer to a PyArrayObject. Routines which return PyObject * return NULL to indicate failure (and follow the standard exception-setting mechanism). Functions followed by a dagger (†) are functions which return PyObjects whose reference count has been increased by one (new references). See the Python Extending/Embedding manual for details on reference-count management.

`int PyArray_Check(op)`

returns 1 if `op` is a PyArrayObject or 0 if it is not.

`int PyArray_SetNumericOps(d)`

internally used by `umath` to setup some of its functions.

`int PyArray_INCREF(op)`

Used for arrays of python objects (`PyArray_OBJECT`) to increment the reference count of every

python object in the array `op`. User code does not typically need to call this.

```
int PyArray_XDECREF(op)
```

Used for arrays of python objects (`PyArray_OBJECT`) to decrement the reference count of every python object in the array `op`.

```
PyArrayError
```

Exports the array error object. I don't know its use.

```
void PyArray_SetStringFunction(op, repr)
```

Sets the function for representation of all arrays to `op` which should be a callable `PyObject`. If `repr` is non-zero then the function corresponding to the `repr` string representation is set, otherwise, that for the `str` string representation is set.

```
PyArray_Descr PyArray_DescrFromType(type)
```

returns a `PyArray_Descr` structure for the datatype given by `type`. The input type can be either the enumerated types (`PyArray_Float`, etc.) or a character (`'cblsilfdFDO'`).

```
PyObject *PyArray_Cast(arp, type) †
```

returns a pointer to a `PyArrayObject` that is `arp` cast to the array type specified by `type`. It is just a wrapper around the function defined in `arp->descr->cast` that handles non-contiguous arrays and arrays of Python objects appropriately.

```
int PyArray_CanCastSafely(fromtype, totype)
```

returns 1 if the array with type `fromtype` can be cast to an array of type `totype` without loss of accuracy, otherwise it returns 0. It allows conversion of longs to ints which is not safe on 64-bit machines. The inputs `fromtype` and `totype` are the enumerated array types (e.g. `PyArray_SBYTE`).

```
int PyArray_ObjectType(op, min_type)
```

returns the typecode to use for a call to an array creation function given an input python sequence object `op` and a minimum type value, `min_type`. It looks at the datatypes used in `op`, compares this with `min_type` and returns a consistent type value that can be used to store all of the data in `op` and satisfying at the minimum the precision of `min_type`.

```
int _PyArray_multiply_list(list, n)
```

is a utility routine to multiply an array of `n` integers pointed to by `list`.

```
int PyArray_Size(op)
```

is a useful function for returning the total number of elements in `op` if `op` is a `PyArrayObject`, 0 otherwise.

```
PyObject *PyArray_FromDims(nd, dims, type) †
```

returns a pointer to a newly constructed `PyArrayObject` (returned as a `PyObject`) given the number of dimensions in `nd`, an array `dims` of `nd` integers specifying the size of the array, and the enumerated type of the array in `type`.

```
PyObject *PyArray_FromDimsAndData(nd, dims, type, data) †
```

This function should only be used to access global data that will never be freed (like FORTRAN common blocks). It builds a `PyArrayObject` in the same way as `PyArray_FromDims` but instead of allocating new memory for the array elements it uses the bytes pointed to by `data` (a `char *`).

```
PyObject *PyArray_ContiguousFromObject(op, type, min_dim, max_dim) †
```

returns a contiguous array of type `type` from the (possibly nested) sequence object `op`. If `op` is a contiguous `PyArrayObject` then a reference is made; if `op` is a non-contiguous then a copy is performed to get a contiguous array; if `op` is not a `PyArrayObject` then a new `PyArrayObject` is created from the sequence object and returned. The two parameters `min_dim` and `max_dim` let you specify the expected rank of the input sequence. An error will result if the resulting `PyArrayObject` does not have rank bounded by these limits. To specify an exact rank requirement set `min_dim = max_dim`. To allow for an arbitrary number of dimensions specify `min_dim = max_dim = 0`.

```
PyObject *PyArray_CopyFromObject(op, type, min_dim, max_dim) †
```

returns a contiguous array similar to `PyArray_ContiguousFromObject` except that a copy of `op` is performed even if a shared array could have been used.

```
PyObject *PyArray_FromObject(op, type, min_dim, max_dim) †
```

returns a reference to `op` if `op` is a `PyArrayObject` and a newly constructed `PyArrayObject` if `op` is any other (nested) sequence object. You must use strides to access the elements of this possibly discontinuous array correctly.

```
PyObject *PyArray_Return(apr)
```

returns a pointer to `apr` with some extra code to check for errors and be sure that zero-dimensional arrays are returned as scalars. If a scalar is returned instead of `apr` then `apr`'s reference count is decremented, so it is safe to use this function in the form :

```
return PyArray_Return (apr);
```

```
PyObject *PyArray_Reshape(apr, op) †
```

returns a reference to `apr` with a new shape specified by `op` which must be a one dimensional sequence object. One dimension may be specified as unknown by giving a value less than zero, its value will be calculated from the size of `apr`.

```
PyObject *PyArray_Copy(apr) †
```

returns an element-for-element copy of `apr`

```
PyObject *PyArray_Take(a, indices, axis) †
```

the equivalent of `take(a, indices, axis)` which is a method defined in the `Numeric` module that just calls this function.

```
int PyArray_As1D(*op, char **ptr, int *n, int type)
```

This function replaces `op` with a pointer to a contiguous 1-D `PyArrayObject` (using `PyArray_ContiguousFromObject`) and sets as output parameters a pointer to the first byte of the array in `ptr` and the number of elements in the array in `n`. It returns `-1` on failure (`op` is not a 1-D array or sequence object that can be cast to type `type`) and `0` on success.

```
int PyArray_As2D(*op, char **ptr, int *m, int *n, int type)
```

This function replaces `op` with a pointer to a contiguous 2-D `PyArrayObject` (using `PyArray_ContiguousFromObject`). It returns `-1` on failure (`op` is not a 2-D array or nested sequence object that can be cast to type `type`) and `0` on success. It also sets as output parameters: an array of pointers in `ptr` which can be used to access the data as a 2-D array so that `ptr[i][j]` is a pointer to the first byte of element `[i,j]` in the array; `m` and `n` are set to respectively the number of rows and columns of the array.

```
int PyArray_Free(op, ptr)
```

is supposed to free the allocated data structures and decrease object references when using `PyArray_As1D` and `PyArray_As2D` but there are suspicions that this code is buggy.

Notes

Number formats, overflow issues, NaN/Inf representations, fpectl module, how to deal with 'missing' values.

UfuncObject C Structure and API

C Structure

The `ufuncobject` is a generic function object that can be used to perform fast operations over Numeric Arrays with very useful broadcasting rules and type conversions performed automatically. The `ufuncobject` and its API make it easy and graceful to add arbitrary functions to Python which operate over Numeric arrays. All of the unary and binary operators currently available in the Numerical extensions (like `sin`, `cos`, `+`, `logical_or`, etc.) are implemented using this object. The hooks are all in place to make it very easy to add any function that takes one or two (double) arguments and returns a single (double) argument. It is not difficult to add support routines in order to handle arbitrary functions whose total number of input/output arguments is less than some maximum number (currently 10).

```
typedef struct {
    PyObject_HEAD
    int *ranks, *canonical_ranks;
    int nin, nout, nargs;
    int identity;
    PyUFuncGenericFunction *functions;
    void **data;
    int ntypes, nranks, attributes;
    char *name, *types;
    int check_return;
} PyUFuncObject;
```

where:

```
int *ranks
```

unused.

```
int *canonical_ranks
```

unused

```
int nin
```

the number of input arguments to function

```
int nout
```

the number of output arguments for the function

```
int nargs
```

the total number of arguments = `nin` + `nout`

```
int identity
```

a flag telling whether the identity for this function is 0 or 1 for use in the `reduce` method for a zero size array input.

```
PyUFuncGenericFunction *functions
```

an array of functions that perform the innermost looping over the input and output arrays (I think this is over a single axis). These functions call the underlying math function with the data from the input arguments along this axis and return the outputs of the function into the correct place in the output arrayobject (with appropriate typecasting). These functions are called by the general looping code. There is one function for each of the supported datatypes. Function pointers to do this looping for types 'f', 'd', 'F', and 'D', are provided in the C-API for functions that take one or two argu-

ments and return one argument. Each `PyUFuncGenericFunction` returns `void` and has the following argument list (in order):

`args`

an array of pointers to the data for each of the input and output arguments with input arguments first and output arguments immediately following. Each element of `args` is a `char *` to the first byte in the corresponding input or output array.

`dimensions`

a pointer to a single `int` giving the size of the axis being looped over.

`steps`

an array of `ints` giving the number of bytes to skip to go to the next element of the array for this loop. There is an entry in the array for each of the input and output arguments, with input arguments first and output arguments immediately following.

`func`

a pointer to the underlying math function to be called at each point in this inner loop. This is a `void *` and must be recast to the required type before actually calling the function e.g. to a pointer to a function that takes two `doubles` and returns a `double`). If you need to write your own `PyUFuncGenericFunction`, it is most readable to also have a `typedef` statement that defines your specific underlying function type so the function pointer cast is somewhat readable.

`void **data`

a pointer to an array of functions (each cast to `void *`) that compute the actual mathematical function for each set of inputs and outputs. There should be a function in the array for each supported data type. This function will be called from the `PyUFuncGenericFunction` for the corresponding type.

`int ntypes`

the number of datatypes supported by this function. For datatypes that are not directly supported, a coercion will be performed if possible safely, otherwise an error will be reported.

`int nrank`

unused.

`int attributes`

unused.

`char *name`

the name of this function (not the same as the dictionary label for this function object, but it is usually set to the same string). It is printed when `__repr__` is called for this object, defaults to "?" if set to `NULL`.

`char *types`

an array of supported types for this function object. I'm not sure why but each supported datatype (`PyArray_FLOAT`, etc.) is entered as many times as there are arguments for this function. (`nargs`)

`int check_return`

Usually best to set to 1. If this is non-zero then returned matrices will be cleaned up so that rank-0 arrays will be returned as python scalars. Also, if non-zero, then any math error that sets the `errno` global variable will cause an appropriate Python exception to be raised.

UfuncObject C API

There are currently 15 pointers in the C-API array for the ufuncobject which is loaded by `import_ufunc()`. The macros implemented by this API, available by including the file `ufuncobject.h`, are given below. The only function normally called by user code is the ufuncobject creation function `PyUFunc_FromFuncAndData`. Some of the other functions can be used as elements of an array to be passed to this creation function.

```
int PyUFunc_Check(op)
```

returns 1 if `op` is a ufunc object otherwise returns 0.

```
PyObject *PyUFunc_FromFuncAndData(functions, data, types, ntypes, nin,  
nout, identity, name, check_return)
```

returns the ufunc object given its parameters. This is the most important function call. It requires defining three arrays to be passed as parameters: `functions`, `data`, and `types`. The arguments to be passed are:

`functions`

an array of functions of type `PyUFuncGenericFunction`, there should be one function for each supported datatype. The functions should be in order so that datatypes listed toward the beginning of the array could be cast as datatypes listed toward the end.

`data`

an array of pointers to `void*` the same size as the functions array and in the same datatype order. Each element of this array is the actual underlying math function (recast to a `void *`) that will be called from one of the `PyUFuncGenericFunctions`. It will operate on each element of the input NumPy arrayobject(s) and return its element-by-element result in the output NumPy arrayobject(s). There is one function call for each datatype supported, (though functions can be repeated if you handle the typecasting appropriately with the `PyUFuncGenericFunction`).

`types`

an array of `PyArray_Types`. The size of this array should be $(nin+nout)$ times the size of one of the previous two arrays. There should be $nin+nout$ copies of `PyArray_XXXXX` for each datatype explicitly supported. (Remember datatypes not explicitly supported will still be accepted as input arguments to the ufunc if they can be cast safely to a supported type.)

`ntypes`

the number of supported types for this ufunc.

`nin`

the number of input arguments

`nout`

the number of output arguments

`identity`

`PyUFunc_One`, `PyUFunc_Zero`, or `PyUFunc_None`, depending on the desired value for the identity. This is only relevant for functions that take two input arguments and return one output argument. If not relevant use `PyUFunc_None`.

`name`

the name of this ufuncobject for use in the `__repr__` method.

`check_return`

the desired value for `check_return` for this ufuncobject.

```
int PyUFunc_GenericFunction(self, args, mps)
```

allows calling the ufunc from user C routine. It returns 0 on success and -1 on any failures. This is the core of what happens when a ufunc is called from Python. Its arguments are:

`self`

the ufunc object to be called. INPUT

`args`

a Python tuple object containing the input arguments to the ufunc (should be Python sequence objects). INPUT

`mps`

an array of pointers to PyArrayObjects for the input and output arguments to this function. The input NumPy arrays are elements `mps[0]...mps[self->nin-1]`. The output NumPy arrays are elements `mps[self->nin]...mps[self->nargs-1]`. OUTPUT

The following are all functions of type `PyUFuncGenericFunction` and are suitable for use in the `functions` argument passed to `PyUFunc_FromFuncAndData`:

```
PyUFunc_f_f_As_d_d
```

for a unary function that takes a double input and returns a double output as a ufunc that takes `PyArray_FLOAT` input and returns `PyArray_FLOAT` output.

```
PyUFunc_d_d
```

for a using a unary function that takes a double input and returns a double output as a ufunc that takes `PyArray_DOUBLE` input and returns `PyArray_DOUBLE` output.

```
PyUFunc_F_F_As_D_D
```

for a unary function that takes a `Py_complex` input and returns a `Py_complex` output as a ufunc that takes `PyArray_CFLOAT` input and returns `PyArray_CFLOAT` output.

```
PyUFunc_D_D
```

for a unary function that takes a `Py_complex` input and returns a `Py_complex` output as a ufunc that takes `PyArray_CFLOAT` input and returns `PyArray_CFLOAT` output.

```
PyUFunc_O_O
```

for a unary function that takes a `Py_Object *` input and returns a `Py_Object *` output as a ufunc that takes `PyArray_OBJECT` input and returns `PyArray_OBJECT` output

```
PyUFunc_ff_f_As_dd_d
```

for a binary function that takes two double inputs and returns one double output as a ufunc that takes `PyArray_FLOAT` input and returns `PyArray_FLOAT` output.

```
PyUFunc_dd_d
```

for a binary function that takes two double inputs and returns one double output as a ufunc that takes `PyArray_DOUBLE` input and returns `PyArray_DOUBLE` output.

```
PyUFunc_FF_F_As_DD_D
```

for a binary function that takes two `Py_complex` inputs and returns a `Py_complex` output as a ufunc that takes `PyArray_CFLOAT` input and returns `PyArray_CFLOAT` output.

```
PyUFunc_DD_D
```

for a binary function that takes two `Py_complex` inputs and returns a `Py_complex` output as a ufunc that takes `PyArray_CFLOAT` input and returns `PyArray_CFLOAT` output

PyUFunc_OO_O

for a unary function that takes two `Py_Object *` input and returns a `Py_Object *` output as a ufunc that takes `PyArray_OBJECT` input and returns `PyArray_OBJECT` output

PyUFunc_O_O_method

for a unary function that takes a `Py_Object *` input and returns a `Py_Object *` output and is pointed to by a Python method as a ufunc that takes `PyArray_OBJECT` input and returns `PyArray_OBJECT` output

PyArrayMap

an exported API that was apparently considered but never implemented probably because the functionality is already available with Python's `map` function.

14. Glossary

This section will define a few of the technical words used throughout this document. [Please let us know of any additions to this list which you feel would be helpful -- the authors]

typecode: a single character describing the format of the data stored in an array. For example, 'b' refers to unsigned byte-sized integers (0-255).

ufunc / universal function: a ufunc is a callable object which performs operations on all of the elements of its arguments, which can be lists, tuples, or arrays. Many ufuncs are defined in the `umath` module.

array / multiarray: an array refers to the Python object type defined by the NumPy extensions to store and manipulate numbers efficiently.

UserArray: The UserArray module defines a UserArray class which should be subclassed by users wishing to have classes which behave similarly to the array object type.

Matrix: The Matrix module defines a subclass Matrix of the UserArray class which is specialized for linear algebra matrices. Most notably, it overrides the multiplication operator on Matrix instances to perform matrix multiplication instead of element-wise multiplication.

rank: the rank of an array is the number of dimensions it has, or the number of integers in its shape tuple.

shape: array objects have an attribute called shape which is necessarily a tuple. An array with an empty tuple shape is treated like a scalar (it holds one element).

PART II: Optional Packages

This part contains descriptions of the packages that are included with the distribution but which are not necessary for using Numeric arrays. The packages are for the most part in the Packages subdirectory of the source distribution, and can be installed anywhere in the Python module search path. Each has its own “setup.py” to use to build and install the package.

For historical reasons, some of these packages are currently installed inside the Numeric package rather than on their own. We hope to remedy this in the future.

The subdirectory Packages contains directories, each of which contains its own installation script setup.py. As with the main directory, these packages are generally compiled and installed using the command

```
python setup.py install
```

The Makefile in the main directory will do this for all the packages provided.

In addition, many people make available libraries that use Numeric. At the moment a centralized reference for these does not exist, but they are usually announced on the discussion list; also check the project web page.

Table 5: Descriptions of the Optional Packages

Package Name	Description	Reference
FFT	Fast Fourier Transforms	“FFT Reference” on pag e77
LinearAlgebra	Basic linear algebra	“LinearAlgebra Reference” on page 80
RandomArray	Arrays of random numbers.	“RandomArray Reference” on page 82
RNG	Generators for independent streams of random numbers from various distributions and arrays of same.	“Independent Random Streams” on page 87
MA	Masked arrays, that is, arrays that have missing or invalid entries.	“Masked Arrays” on page 89

License and disclaimer for packages MA and RNG

Package MA was written by Paul Dubois, LLNL. Package RNG was written by Konrad Hinsen after modifying an earlier package UNRG by Paul Dubois and Fred Fritsch.

Copyright (c) 1999, 2000. The Regents of the University of California. All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose without fee is hereby granted, provided that this entire notice is included in all copies of any software which is or includes a copy or modification of this software and in all copies of the supporting documentation for such software.

This work was produced at the University of California, Lawrence Livermore National Laboratory under contract no. W-7405-ENG-48 between the U.S. Department of Energy and The Regents of the University of California for the operation of UC LLNL.

DISCLAIMER

This software was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately-owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

15. FFT Reference

The `FFT.py` module provides a simple interface to the `FFTPACK` FORTRAN library, which is a powerful standard library for doing fast Fourier transforms of real and complex data sets, or the C `fftpack` library, which is algorithmically based on `FFTPACK` and provides a compatible interface. On some platforms, optimized version of one of these libraries may be available, and can be used to provide optimal performance (see “Compilation Notes” on page 79).

Python Interface

The Python user imports the `FFT` module, which provides a set of utility functions which provide access to the most commonly used FFT routines, and allows the specification of which axes (dimensions) of the input arrays are to be used for the FFT's. These routines are:

`fft(data, n=None, axis=-1)`

Performs a *n*-point discrete Fourier transform of the array `data`. *n* defaults to the size of `data`. It is most efficient for *n* a power of two. If *n* is larger than `data`, then `data` will be zero-padded to make up the difference. If *n* is smaller than `data`, then `data` will be aliased to reduce its size. This also stores a cache of working memory for different sizes of `fft`'s, so you could theoretically run into memory problems if you call this too many times with too many different *n*'s.

The FFT is performed along the axis indicated by the `axis` argument, which defaults to be the last dimension of `data`.

The format of the returned array is a complex array of the same shape as `data`, where the first element in the result array contains the DC (steady-state) value of the FFT, and where each successive ...XXX

Example of use:

```
>>> print fft(array((1,0,1,0,1,0,1,0))+ 10).real
[ 84.  0.  0.  0.  4.  0.  0.  0.]
>>> print fft(array((0,1,0,1,0,1,0,1))+ 10).real
[ 84.  0.  0.  0. -4.  0.  0.  0.]
>>> print fft(array((0,1,0,0,0,1,0,0))+ 10).real
[ 82.  0.  0.  0. -2.  0.  0.  0.]
```

`inverse_fft(data, n=None, axis=-1)`

Will return the *n* point inverse discrete Fourier transform of `data`. *n* defaults to the length of `data`. This is most efficient for *n* a power of two. If *n* is larger than `data`, then `data` will be zero-padded to make up the difference. If *n* is smaller than `data`, then `data` will be aliased to reduce its size. This also stores a cache of working memory for different sizes of FFT's, so you could theoretically run into memory problems if you call this too many times with too many different *n*'s.

`real_fft(data, n=None, axis=-1)`

Will return the *n* point discrete Fourier transform of the real valued array `data`. *n* defaults to the length of `data`. This is most efficient for *n* a power of two. The returned array will be one half of the symmetric complex transform of the real array.

```

>>> x = cos(arange(30.0)/30.0*2*pi)
>>> print real_fft(x)
[ -1.          +0.j          13.69406641+2.91076367j
  -0.91354546-0.40673664j  -0.80901699-0.58778525j
  -0.66913061-0.74314483j  -0.5          -0.8660254j
  -0.30901699-0.95105652j  -0.10452846-0.9945219j
   0.10452846-0.9945219j    0.30901699-0.95105652j
   0.5          -0.8660254j    0.66913061-0.74314483j
   0.80901699-0.58778525j  0.91354546-0.40673664j
   0.9781476  -0.20791169j   1.          +0.j          ]

```

inverse_real_fft(data, n=None, axis=-1)

Will return the inverse FFT of the real valued array data.

fft2d(data, s=None, axes=(-2,-1))

Will return the 2-dimensional FFT of the array data.

real_fft2d(data, s=None, axes=(-2,-1))

Will return the 2d FFT of the real valued array data.

C API

The interface to the FFTPACK library is performed via the `fftpackmodule` module, which is responsible for making sure that the arrays sent to the FFTPACK routines are in the right format (contiguous memory locations, right numerical storage format, etc). It provides interfaces to the following FFTPACK routines, which are also the names of the Python functions:

- `cffti(i)`
- `cfftf(data, savearea)`
- `cfftb(data, savearea)`
- `rffti(i)`
- `rfftf(data, savearea)`
- `rfftb(data, savearea)`

The routines which start with `c` expect arrays of complex numbers, the routines which start with `r` expect real numbers only. The routines which end with `i` are the initialization functions, those which end with `f` perform the forward FFTs and those which end with `b` perform the backwards FFTs.

The initialization functions require a single integer argument corresponding to the size of the dataset, and returns a work array. The forward and backwards FFTs require two array arguments -- the first is the data array, the second is the work array returned by the initialization function. They return arrays corresponding to the coefficients of the FFT, with the first element in the returned array corresponding to the DC component, the second one to the first fundamental, etc. The length of the returned array is 1 + half the length of the input array in the case of real FFTs, and the same size as the input array in the case of complex data.

```

>>> x = cos(arange(30.0)/30.0*2*pi)
>>> w = rffti(30)
>>> f = rfftf(x, w)
>>> f[0]
(-1+0j)
>>> f[1]
(13.6940664103+2.91076367145j)
>>> f[2]

```

(-0.913545457643-0.406736643076j)

Compilation Notes

On some platforms, precompiled optimized versions of the FFTPACK library are preinstalled on the operating system, and the compilation procedure needs to be modified to force the fftpackmodule file to be linked against those rather than the fftpacklite.c file which is shipped with NumPy.

16. LinearAlgebra Reference

The LinearAlgebra.py module provides a simple interface to the low-level linear algebra routines provided by either the LAPACK FORTRAN library or the compatible lapack_lite C library.

Python Interface

solve_linear_equations(a, b)

This function solves a system of linear equations with a square non-singular matrix *a* and a right-hand-side vector *b*. Several right-hand-side vectors can be treated simultaneously by making *b* a two-dimensional array (i.e. a sequence of vectors). The function *inverse(a)* calculates the inverse of the square non-singular matrix *a* by calling *solve_linear_equations(a, b)* with a suitable *b*.

inverse(a)

This function returns the inverse of the specified matrix *a* which must be square and non-singular. To within floating point precision, it should always be true that:

```
matrixmultiply(a, inverse(a)) == identity(len(a))
```

To test this claim, one can do e.g.:

```
>>> a = reshape(arange(25.0), (5,5)) + identity(5)
>>> print a
[[ 1.  1.  2.  3.  4.]
 [ 5.  7.  7.  8.  9.]
 [10. 11. 13. 13. 14.]
 [15. 16. 17. 19. 19.]
 [20. 21. 22. 23. 25.]]
>>> inv_a = inverse(a)
>>> print inv_a
[[ 0.20634921 -0.52380952 -0.25396825  0.01587302  0.28571429]
 [-0.5026455  0.63492063 -0.22751323 -0.08994709  0.04761905]
 [-0.21164021 -0.20634921  0.7989418  -0.1957672  -0.19047619]
 [ 0.07936508 -0.04761905 -0.17460317  0.6984127  -0.42857143]
 [ 0.37037037  0.11111111 -0.14814815 -0.40740741  0.33333333]]
>>> # Verify the inverse by printing the largest absolute element
... # of a * a^{-1} - identity(5)
... print "Inversion error:", \
... maximum.reduce(fabs(ravel(dot(a, inv_a)-identity(5))))
Inversion error: 2.6645352591e-015
```

eigenvalues(a)

This function returns the eigenvalues of the square matrix *a*.

```
>>> print a
[[ 1.  0.  0.  0.  0.]
 [ 0.  2.  0.  0.  1.]
```



```

[ 0.  0.  3.  0.  0.]
[ 0.  0.  0.  4.  0.]
[ 0.  0.  0.  0.  1.]]
>>> print eigenvalues(a)
[ 1.  2.  3.  4.  1.]

```

eigenvectors(a)

This function returns both the eigenvalues and the eigenvectors, the latter as a two-dimensional array (i.e. a sequence of vectors).

```

>>> print a
[[ 1.  0.  0.  0.  0.]
 [ 0.  2.  0.  0.  1.]
 [ 0.  0.  3.  0.  0.]
 [ 0.  0.  0.  4.  0.]
 [ 0.  0.  0.  0.  1.]]
>>> evalues, evectors = eigenvectors(a)
>>> print evalues
[ 1.  2.  3.  4.  1.]
>>> print evectors
[[ 1.          0.          0.          0.          0.          ]
 [ 0.          1.          0.          0.          0.          ]
 [ 0.          0.          1.          0.          0.          ]
 [ 0.          0.          0.          1.          0.          ]
 [ 0.          -0.70710678  0.          0.          0.70710678]]

```

singular_value_decomposition(a, full_matrices=0)

This function returns three arrays V , S , and W^T whose matrix product is the original matrix a . V and W^T are unitary matrices (rank-2 arrays), whereas S is the vector (rank-1 array) of diagonal elements of the singular-value matrix. This function is mainly used to check whether (and in what way) a matrix is ill-conditioned.

generalized_inverse(a, rcond=1e-10)

This function returns the generalized inverse (also known as pseudo-inverse or Moore-Penrose-inverse) of the matrix a . It has numerous applications related to linear equations and least-squares problems.

determinant(a)

This function returns the determinant of the square matrix a .

linear_least_squares(a, b, rcond=e-10)

This function returns the least-squares solution of an overdetermined system of linear equations. An optional third argument indicates the cutoff for the range of singular values (defaults to 10^{-10}). There are four return values: the least-squares solution itself, the sum of the squared residuals (i.e. the quantity minimized by the solution), the rank of the matrix a , and the singular values of a in descending order.

Compilation Notes

On some platforms, precompiled optimized versions of the LAPACK library are preinstalled on the operating system, and the setup procedure needs to be modified to force the `lapackmodule.c` file to be linked against those rather than the `lapack_lite` library.

17. RandomArray Reference

The `RandomArray.py` module (in conjunction with the `ranlibmodule.c` file) provides a high-level interface to the `ranlib` module, which provides a good quality C implementation of a random-number generator.

Python Interface

seed(x=0, y=0)

The `seed()` function takes two integers and sets the two seeds of the random number generator to those values. If the default values of 0 are used for both `x` and `y`, then a seed is generated from the current time, providing a pseudo-random seed.

get_seed()

The `get_seed()` function returns the two seeds used by the current random-number generator. It is most often used to find out what seeds the `seed()` function chose at the last iteration. [thread-safety issue?]

random(shape=ReturnFloat)

The `random()` function takes a shape, and returns an array of double-precision floating point numbers between 0.0 and 1.0. Neither 0.0 nor 1.0 is ever returned by this function. If no argument is specified, the function returns a single floating point number (not an array). The array is filled from the generator following the canonical array organization (see discussion of the `.flat` attribute)

uniform(minimum, maximum, shape=ReturnFloat)

The `uniform()` function returns an array of the specified shape and containing double-precision floating point random numbers strictly between minimum and maximum. If no shape is specified, a single number is returned.

randint(minimum, maximum, shape=ReturnFloat)

The `randint()` function returns an array of the specified shape and containing random (standard) integers greater than or equal to minimum and strictly less than maximum. If no shape is specified, a single number is returned.

permutation(n)

The `permutation()` function returns an array of the integers between 0 and `n-1`, in an array of shape `(n,)`, and with its elements randomly permuted.

An example use of the `RandomArray` module (exact output will be different each time!):

```
>>> from RandomArray import *
>>> seed()                               # Set seed based on current time
>>> print get_seed()                     # Find out what seeds were used
(897800491, 192000)
>>> print random()
0.0528018975065
>>> print random((5,2))
[[ 0.14833829  0.99031458]
```

```

[ 0.7526806  0.09601787]
[ 0.1895229  0.97674777]
[ 0.46134511 0.25420982]
[ 0.66132009 0.24864472]]
>>> print uniform(-1,1,(10,))
[ 0.72168852 -0.75374185 -0.73590945  0.50488248 -0.74462822  0.09293685
 -0.65898308  0.9718067  -0.03252475  0.99611011]
>>> print randint(0,100, (12,))
[28  5 96 19  1 32 69 40 56 69 53 44]
>>> print permutation(10)
[4 2 8 9 1 7 3 6 5 0]
>>> seed(897800491, 192000)      # resetting the same seeds
>>> print random()              # yields the same numbers
0.0528018975065

```

Floating point random arrays

standard_normal (shape=ReturnFloat)

The `standard_normal ()` function returns an array of the specified shape that contains double precision floating point numbers normally (Gaussian) distributed with mean zero and variance and standard deviation one. If no shape is specified, a single number is returned.

normal (mean, stddev, shape=ReturnFloat)

The `normal ()` function returns an array of the specified shape that contains double precision floating point numbers normally distributed with the specified mean and standard deviation. If no shape is specified, a single number is returned.

multivariate_normal (mean, covariance) or multivariate_normal (mean, covariance, leadingAxesShape)

The `multivariate_normal ()` function takes a one dimensional array argument `mean` and a two dimensional array argument `covariance`. Suppose the shape of `mean` is $(n,)$. Then the shape of `covariance` must be (n,n) . The `multivariate_normal ()` function returns a double precision floating point array. The effect of the `leadingAxesShape` parameter is:

- If no `leadingAxesShape` is specified, then an array with shape $(n,)$ is returned containing a vector of numbers with a multivariate normal distribution with the specified mean and covariance.
- If `leadingAxesShape` is specified, then an array of such vectors is returned. The shape of the output is `leadingAxesShape.append ((n,))`. The leading indices into the output array select a multivariate normal from the array. The final index selects one number from within the multivariate normal.

In either case, the behavior of `multivariate_normal ()` is undefined if `covariance` is not symmetric and positive definite.

exponential (mean, shape=ReturnFloat)

The `exponential ()` function returns an array of the specified shape that contains double precision floating point numbers exponentially distributed with the specified mean. If no shape is specified, a single number is returned.

beta (a, b, shape=ReturnFloat)

The `beta ()` function returns an array of the specified shape that contains double precision floating point numbers beta distributed with alpha parameter `a` and beta parameter `b`. If no shape is specified, a single number is returned.

gamma (a, r, shape=ReturnFloat)

The `gamma ()` function returns an array of the specified shape that contains double precision floating point numbers beta distributed with location parameter `a` and distribution shape parameter `r`. If no shape is specified, a single number is returned.

chi_square (df, shape=ReturnFloat)

The `chi_square()` function returns an array of the specified shape that contains double precision floating point numbers with the chi square distribution with `df` degrees of freedom. If no shape is specified, a single number is returned.

noncentral_chi_square (df, nonc, shape=ReturnFloat)

The `noncentral_chi_square()` function returns an array of the specified shape that contains double precision floating point numbers with the chi square distribution with `df` degrees of freedom and noncentrality parameter `nonc`. If no shape is specified, a single number is returned.

F (dfn, dfd, shape=ReturnFloat)

The `F ()` function returns an array of the specified shape that contains double precision floating point numbers with the F distribution with `dfn` degrees of freedom in the numerator and `dfd` degrees of freedom in the denominator. If no shape is specified, a single number is returned.

noncentral_F (dfn, dfd, nonc, shape=ReturnFloat)

The `noncentral_F ()` function returns an array of the specified shape that contains double precision floating point numbers with the F distribution with `dfn` degrees of freedom in the numerator, `dfd` degrees of freedom in the denominator, and noncentrality parameter `nonc`. If no shape is specified, a single number is returned.

Integer random arrays

binomial (trials, prob, shape=ReturnInt)

The `binomial ()` function returns an array with the specified shape that contains integer numbers with the binomial distribution with trials `trials` and event probability `prob`. In other words, each value in the returned array is the number of times an event with probability `prob` occurred within `trials` repeated trials. If no shape is specified, a single number is returned.

negative_binomial (trials, prob, shape=ReturnInt)

The `negative_binomial ()` function returns an array with the specified shape that contains integer numbers with the negative binomial distribution with trials `trials` and event probability `prob`. If no shape is specified, a single number is returned.

poisson (mean, shape=ReturnInt)

The `poisson ()` function returns an array with the specified shape that contains integer numbers with the Poisson distribution with the specified mean. If no shape is specified, a single number is returned.

multinomial (trials, probs) or multinomial (trials, probs, leadingAxesShape)

The `multinomial ()` function returns an array with that contains integer numbers with the multinomial distribution with trials `trials` and event probabilities given in `probs`. `probs` must be a one dimensional array. There are `len(probs)+1` events. `probs[i]` is the probability of the `i`-th event for $0 \leq i < \text{len}(\text{probs})$. The probability of event `len(probs)` is `1.-Numeric.sum(prob)`.

The first form returns an integer array of shape $(\text{len}(\text{probs})+1,)$ containing one multinomially distributed vector. The second form returns an array of shape $(m, n, \dots, \text{len}(\text{probs})+1)$ where (m, n, \dots) is `leadingAxesShape`. In this case, each output `[i,j,...:]` is an integer array of shape $(\text{len}(\text{prob})+1,)$ containing one multinomially distributed vector..

Examples

Most of the functions in this package take zero or more distribution specific parameters plus an optional shape parameter. The shape parameter gives the shape of the output array:

```
>>> from RandomArray import *
>>> print standard_normal()
-0.435568600893
>>> print standard_normal(5)
[-1.36134553  0.78617644 -0.45038718  0.18508556  0.05941355]
>>> print standard_normal((5,2))
[[ 1.33448863 -0.10125473]
 [ 0.66838062  0.24691346]
 [-0.95092064  0.94168913]
 [-0.23919107  1.89288616]
 [ 0.87651485  0.96400219]]
>>> print normal(7., 4., (5,2)) #mean=7, std. dev.=4
[[ 2.66997623 11.65832615]
 [ 6.73916003  6.58162862]
 [ 8.47180378  4.30354905]
 [ 1.35531998 -2.80886841]
 [ 7.07408469 11.39024973]]
>>> print exponential(10., 5) #mean=10
[ 18.03347754  7.11702306  9.8587961  32.49231603  28.55408891]
>>> print beta(3.1, 9.1, 5) # alpha=3.1, beta=9.1
[ 0.1175056  0.17504358  0.3517828  0.06965593  0.43898219]
>>> print chi_square(7, 5) # 7 degrees of freedom (dfs)
[ 11.99046516  3.00741053  4.72235727  6.17056274  8.50756836]
>>> print noncentral_chi_square(7, 3, 5) # 7 dfs, noncentrality 3
[ 18.28332138  4.07550335  16.0425396  9.51192093  9.80156231]
>>> F(5, 7, 5) # 5 and 7 dfs
array([ 0.24693671,  3.76726145,  0.66883826,  0.59169068,  1.90763224])
>>> noncentral_F(5, 7, 3., 5) # 5 and 7 dfs, noncentrality 3
array([ 1.17992553,  0.7500126 ,  0.77389943,  9.26798989,  1.35719634])
>>> binomial(32, .5, 5) # 32 trials, prob of an event = .5
array([12, 20, 21, 19, 17])
>>> negative_binomial(32, .5, 5) # 32 trials: prob of an event = .5
array([21, 38, 29, 32, 36])
```

Two functions that return generate multivariate random numbers (that is, random vectors with some known relationship between the elements of each vector, defined by the distribution). They are `multivariate_normal()` and `multinomial()`. For these two functions, the lengths of the leading axes of the output may be specified. The length of the last axis is determined by the length of some other parameter.

```
>>> multivariate_normal([1,2], [[1,2],[2,1]], [2,3])
array([[[ 0.14157988,  1.46232224],
        [-1.11820295, -0.82796288],
        [ 1.35251635, -0.2575901 ]],
       [[-0.61142141,  1.0230465 ],
        [-1.08280948, -0.55567217],
        [ 2.49873002,  3.28136372]]])
```

```

>>> x = multivariate_normal([10,100], [[1,2],[2,1]], 10000)
>>> x_mean = sum(x)/10000
>>> print x_mean
[ 9.98599893 100.00032416]
>>> x_minus_mean = x - x_mean
>>> cov = matrixmultiply(transpose(x_minus_mean), x_minus_mean) / 9999.
>>> cov
array([[ 2.01737122,  1.00474408],
       [ 1.00474408,  2.0009806 ]])

```

The a priori probabilities for a multinomial distribution must sum to one. The prior probability argument to `multinomial()` doesn't give the prior probability of the last event: it is computed to be one minus the sum of the others.

```

>>> multinomial(16, [.1, .4, .2]) # prior probabilities [.1, .4, .2, .3]
array([2, 7, 1, 6])
>>> multinomial(16, [.1, .4, .2], [2,3]) # output shape [2,3,4]
array([[[ 1, 9, 1, 5],
        [ 0, 10, 3, 3],
        [ 4, 9, 3, 0]],
       [[ 1, 6, 1, 8],
        [ 3, 4, 5, 4],
        [ 1, 5, 2, 8]]])

```

18. Independent Random Streams

The RNG package provides any number of independent random number generators tied to a distribution. Distributions include exponential, normal, and log-normal distributions, but adding others is not difficult. Contributions of code for other distributions are welcome!

Background

RNG was written by Konrad Hinsien based on the package URNG by Paul Dubois and Fred Fritsch of LLNL. This package has been released for unlimited redistribution. Please see “License and disclaimer for packages MA and RNG” on page 76.

Usage

Package RNG installs two modules: RNG.RNG, and RNG.ranf. The former is a C extension that does the generation. The latter is an easy-to-use interface for a single uniform distribution.

Module RNG

Module RNG defines the function:

```
CreateGenerator(s, distribution=UniformDistribution(0., 1.))
```

creates a new random number generator with a distribution. The random numbers produced by the generator sample the distribution and are independent of other generators created earlier or later. Its first argument, an integer, determines the initial state:

- 0 ; Use the default initial seed value.
- <0 ; Set a random value for the seed from the system clock.
- >0 ; Set seed directly (32 bits only).

The default distribution is a uniform distribution on [0., 1.); other distributions are obtained by supplying a second argument which must be a distribution. Currently RNG defines the following distribution types:

- UniformDistribution(a, b) -- a uniform distribution of numbers in the interval [a, b)
- NormalDistribution(mu, sigma) -- a normal distribution with mean mu and standard deviation sigma
- ExponentialDistribution(l) -- an exponential distribution of positive numbers with decay constant l.
- LogNormalDistribution(mean, std) -- a log normal distribution with given mean and standard deviation.

Generator objects

Once a generator is created, it contains these methods:

- sample(n) will return an array of n samples from the generator.
- ranf() will return one sample from the generator.

Module ranf

Module `ranf`, whose main function `ranf()` is equivalent to the old `ranf` generator on Cray 1 computers, defines these facilities.

Attribute `standard_generator` is an instance of `RNG.UniformDistribution(0., 1.)`.

`ranf()`: returns a random number from the `standard_generator`.

`random_sample(*n)` returns a Numeric array of samples from the `standard_generator`.

```
random_sample(n) = array of n random numbers;  
random_sample(n1, n2, ...)= array of shape (n1, n2, ..)
```

Examples

The test routine `Demo/RNGdemo.py` illustrates some common usage of both `RNG` and `Numeric`.

The test routine `RNGtest2.py` combines `RNG` with Konrad Hinsén's `Statistics` package to do a test of the log normal distribution.

Here is one function from `RNGdemo.py`, showing a test of a normal distribution.

```
from Numeric import *  
import RNG  
  
def test_normal (mean, std, n=10000):  
    dist = RNG.NormalDistribution(mean, std)  
    rng = RNG.CreateGenerator(0, dist)  
    values = rng.sample(n)  
    m = sum(values)/n  
    s = sqrt(sum((values-m)**2)/n)  
    return (m, s)
```


19. Masked Arrays

Masked arrays are arrays that may have missing or invalid entries. Module MA provides a work-alike replacement for Numeric that supports data arrays with masks.

What is a masked array?

Masked arrays are arrays that may have missing or invalid entries. Module MA provides a work-alike replacement for Numeric that supports data arrays with masks. A mask is either None or an array of ones and zeros, that determines for each element of the masked array whether or not it contains an invalid entry. The package assures that invalid entries are not used in calculations.

A particular element is said to be masked (invalid) if the mask is not None and the corresponding element of the mask is 1; otherwise it is unmasked (valid).

This package was written by Paul F. Dubois at Lawrence Livermore National Laboratory. Please see the legal notice in the software and on “License and disclaimer for packages MA and RNG” on page 76.

Installing and using MA

MA is one of the optional Packages and installing it requires a separate step as explained in the Numeric README. To install just the MA package using Distutils, in the MA top directory enter:

```
python setup.py install
```

Use MA as a replacement for Numeric:

```
from MA import *
x = array([1, 2, 3])
```

To create an array with the second element invalid, we would do:

```
y = array([1, 2, 3], mask = [0, 1, 0])
```

To create a masked array where all values “near” 1.e20 are invalid, we can do:

```
z = masked_values ([1.0, 1.e20, 3.0, 4.0], 1.e20)
```

For a complete discussion of creation methods for masked arrays please see “Constructing masked arrays” on page 92.

The Numeric module is an attribute in MA, so to execute a method *foo* from Numeric, you can reference it as `Numeric.foo(...)`.

Usually people use both MA and Numeric this way, but of course you can always fully-qualify the names:

```
import MA
x = MA.array([1, 2, 3])
```

The principal feature of module MA is class `MaskedArray`, the class whose instances are returned by the array constructors and most functions in module MA. We will discuss this class first, and later cover the attributes and functions in module MA. For now suffice it to say that among the attributes of the module are the constants from module Numeric including those for declaring typecodes, `NewAxis`, and the mathematical constants such as `pi` and `e`. An additional typecode, `MaskType`, is the typecode used for masks.

Class `MaskedArray`

In Module MA, an array is an instance of class `MaskedArray`, which is defined in the module MA. An instance of class `MaskedArray` can be thought of as containing the following parts:

- An array of data, of any shape;
- A mask of ones and zeros of the same shape as the data; and,
- A “fill value” -- this is a value that may be used to replace the invalid entries in order to return a plain Numeric array. The chief method that does this is the method **filled** discussed below.

We will use the terms “invalid value” and “invalid entry” to refer to the data value at a place corresponding to a mask value of 1. It should be emphasized that the invalid values are *never* used in any computation, and that the fill value is not used for *any* computational purpose. When an instance *x* of class MaskedArray is converted to its string representation, it is the result returned by **filled (x)** that is converted to a string.

Attributes of masked arrays

flat: (deprecated) returns the masked array as one-dimensional. This is provided for compatibility with Numeric. **ravel (x)** is preferred.

real: returns the real part of the array if complex.

imaginary: returns the imaginary part of the array if complex.

shape: The shape of a masked array can be accessed or changed by using the special attribute **shape**, as with Numerical arrays.

shared_data: This read-only flag if true indicates that the masked array shared a reference with the original data used to construct it at the time of construction. Changes to the original array will affect the masked array. (This is not the default behavior; see “Copying or not?” on page 94.) This flag is informational only.

shared_mask: This read-only flag if true indicates that the masked array *currently* shares a reference to the mask used to create it. Unlike **shared_data**, this flag may change as the result of modifying the array contents, as the mask uses copy on write semantics if it is shared.

Methods on masked arrays.

Table 6: Methods on masked arrays; attributes, constructors and operations discussed separately.

Method	Description	Sample syntax
astype (typecode)	return self as array of given type.	y = x.astype (Float32)
compressed ()	return an array of the valid elements. Result is one-dimensional.	y = x.compressed()
filled (fill_value=None)	filled(self, self.fill_value()); see description of module method filled.	y = x.filled()
fill_value ()	Get the current fill value.	v = x.fill_value ()
filled (value = None)	Same as filled(self, value); see “filled is very important. It converts its argument to a plain Numeric array.” on page 93.	numar = x.filled ()
get_shape ()	Return the tuple giving the current shape. Same as shape attribute.	s = x.get_shape () s = x.shape
ids ()	Return the ids of the data and mask areas	id1, id2 = x.ids ()
iscontiguous ()	Is the data area contiguous? See Numeric manual.	if x.iscontiguous ()
mask ()	Return the data mask, or None.	m = x.mask ()

Table 6: Methods on masked arrays; attributes, constructors and operations discussed separately.

Method	Description	Sample syntax
<code>put (values)</code>	Set the value at each non-masked entry to the corresponding entry in <i>values</i> . The mask is unchanged. See also module method <code>put</code> .	<code>x.put (values)</code>
<code>putmask (values)</code>	Eliminate any masked values by setting the value at each masked entry to the corresponding entry in <i>values</i> . Set the mask to <code>None</code> .	<code>x.putmask(values)</code> <code>assert getmask(x) is None</code>
<code>raw_data ()</code>	A reference to the non-filled data; portions may be meaningless. Expert use only.	<code>d = x.raw_data ()</code>
<code>savespace (v)</code>	Set the spacesaver attribute to <i>v</i> .	<code>x.savespace (1)</code>
<code>set_fill_value ()</code>	Set the fill value to <i>v</i> . Omit <i>v</i> to restore default.	<code>x.set_fill_value (1.e21)</code>
<code>set_shape (args...)</code>	<code>shape (n, m, ...)</code> sets the shape.	<code>x.set_shape (3, 12)</code>
<code>size (axis)</code>	Number of elements in array, or in a particular axis.	<code>totalsize = x.size ()</code> <code>col_len = x.size (1)</code>
<code>spacesaver()</code>	Query the spacesave flag.	<code>flag = x.spacesaver()</code>
<code>unshare_mask()</code>	If <code>shared_mask</code> is currently true, replaces the reference to it with a copy.	<code>x.unshare_mask()</code>
<code>typecode ()</code>	Return the type of the data. See module <code>Precision</code> .	<code>z = x.typecode()</code>

Constructing masked arrays

1. **array** (*data*, *typecode* = None, *copy* = 1, *savespace* = 0, *mask* = None, *fill_value* = None)
Creates a masked array with the given *data* and *mask*. The name **array** is simply an alias for the class name, **MaskedArray**. The fill value is set to *fill_value*, and the *savespace* flag is applied. If *data* is a **MaskedArray**, its *mask*, *typecode*, *savespace* flag, and fill value will be used unless specifically specified by one of the remaining arguments. In particular, if *d* is a masked array, `array(d, copy=0)` is *d*.
2. **masked_array** (*data*, *mask* = None, *fill_value* = None)
This is an easier-to-use version of **array**, for the common case of *typecode* = None, *copy* = 0. When *data* is newly-created this function can be used to make it a masked array without copying the data if *data* is already a Numeric array.
3. **masked_values** (*data*, *value*, *rtol*=1.e-5, *atol*=1.e-8, *typecode* = None, *copy* = 1, *savespace* = 0)
Constructs a masked array whose mask is set at those places where $\text{abs}(data - value) < atol + rtol * \text{abs}(data)$.
That is a careful way of saying that those elements of the *data* that have value = *value* (to within a tolerance) are to be treated as invalid. If *data* is not of a floating point type, calls `masked_object` instead.
4. **masked_object** (*data*, *value*, *copy*=1, *savespace*=0)
Creates a masked array with those entries marked invalid that are equal to *value*. Again, *copy* and *savespace* are passed on to the Numeric array constructor.
5. **asarray**(*data*, *typecode*=None)
This is the same as `array(data, typecode, copy=0)`. This is a short way of ensuring that something is an instance of **MaskedArray** of a given type before proceeding, as in `data = asarray(data)`.

If *data* already is a masked array and *typecode* is None then the return value is *data*; nothing is copied in that case.

6. **masked_where** (*condition*, *data*, *copy*=1)
Creates a masked array whose shape is that of *condition*, whose values are those of *data*, and which is masked where elements of *condition* are true.

The following additional constructors are provided for convenience.

- **masked_greater** (*data*, *value*, *copy*=1) is equivalent to **masked_where** (**greater**(*data*, *value*), *data*). Similarly, **masked_greater_equal**, **masked_equal**, **masked_not_equal**, **masked_less**, **masked_less_equal** are called in the same way with the obvious meanings. Note that for floating point data, **masked_values** is preferable to **masked_equal** in most cases.
- **masked_inside** (*data*, *v1*, *v2*, *copy*=1) creates an array with values in the closed interval [*v1*, *v2*] masked. *v1* and *v2* may be in either order.
- **masked_outside** (*data*, *v1*, *v2*, *copy*=1) creates an array with values outside the closed interval [*v1*, *v2*] masked. *v1* and *v2* may be in either order.

On entry to any of these constructors, *data* must be any object which the Numeric package can accept to create an array (with the desired *typecode*, if specified). The *mask* if given must be None or any object that can be turned into a Numeric array of integer type (it will be converted to *typecode* `MaskType`, if necessary), have the same shape as *data*, and contain only values of 0 or 1.

If the *mask* is not None but its shape does not match that of *data*, an exception will be thrown, unless one of the two is of length 1, in which case the scalar will be resized (using `Numeric.resize`) to match the other.

See “Copying or not?” on page 94 for a discussion of whether or not the resulting array shares its data or its mask with the arguments given to these constructors.

Important Tip

filled is very important. It converts its argument to a plain Numeric array.

filled (*x*, *value* = *None*) returns *x* with any invalid locations replaced by a fill value. **filled** is guaranteed to return a plain Numeric array. The argument *x* does not have to be a masked array or even an array, just something that Numeric can turn into one.

- If *x* is not a masked array, and not a Numeric array, `Numeric.array(x)` is returned.
- If *x* is a contiguous Numeric array then *x* is returned. (A Numeric array is contiguous if its data storage region is laid out in column-major order; Numeric allows non-contiguous arrays to exist but they are not allowed in certain operations).
- If *x* is a masked array, but the mask is *None*, and *x*'s data array is contiguous, then it is returned. If the data array is not contiguous, a (contiguous) copy of it is returned.
- If *x* is a masked array with an actual mask, then an array formed by replacing the invalid entries with *value*, or **fill_value** (*x*) if *value* is *None*, is returned. If the fill value used is of a different type or precision than *x*, the result may be of a different type or precision than *x*.

Note that a new array is created only if necessary to create a correctly filled contiguous Numeric array.

The function **filled** plays a central role in our design. It is the “exit” back to Numeric, and is used whenever the invalid values must be replaced before an operation. For example, adding two masked arrays *a* and *b* is roughly:

```
masked_array(filled(a, 0)+filled(b, 0), mask_or(getmask(a), getmask(b)))
```

That is, fill the invalid entries *a* and *b* with zeros, add them up, and declare any entry of the result invalid if either *a* or *b* was invalid at that spot. The functions **getmask** and **mask_or** are discussed later.

filled (*x*) also can be used to simply be certain that some expression is a contiguous Numerical array at little cost. If its argument is a Numeric array already, it is returned without copying.

fill_value (*x*), and the method *x.fill_value*() of the same name on masked arrays, returns a value suitable for filling *x* based on its type. If *x* is a masked array, then *x.fill_value*() results. The returned value for a given type can be changed by assigning to these names in module MA: They should be set to scalars or one element arrays.

```
default_real_fill_value = Numeric.array([1.0e20], Float32)
default_complex_fill_value = Numeric.array([1.0e20 + 0.0j], Complex32)
default_character_fill_value = masked
default_integer_fill_value = Numeric.array([0]).astype(UnsignedInt8)
default_object_fill_value = masked
```

The variable **masked** is a module variable of MA and is discussed in “Working with Masks” on page 93. Calling **filled** with a *fill_value* of **masked** sometimes produces a useful printed representation of a masked array. The function **fill_value** works on any kind of object.

Working with Masks

create_mask(*ashape*) returns an array suitable for use as a mask, having the given shape and initialized to zeros.

is_mask (*m*) is true if *m* is of a type and precision that would be allowed as the mask field of a masked array (that is, it is an array of integers with Numeric's typecode `MaskType`, or it is *None*). To be a legal mask, *m* should contain only zeros or ones, but this is not checked.

make_mask (*m*, *copy*=0, *flag*=0) returns an object whose entries are equal to *m* and for which **is_mask** would return true. If *m* is already a mask or *None*, it returns *m* or a copy of it. Otherwise it will attempt to make a mask, so it will accept any sequence of integers of for *m*. If *flag* is true, **make_mask** returns *None* if its return value otherwise would contain no true elements. To make a legal mask, *m* should contain only zeros or ones, but this is not checked.

getmask (*x*) returns *x.mask()*, the mask of *x*, if *x* is a masked array, and None otherwise. Note that **getmask** may return None if *x* is a masked array but has a mask of None.

getmaskarray (*x*) returns *x.mask()* if *x* is a masked array and has a mask that is not None; otherwise it returns a zero mask array of the same shape as *x*. Unlike **getmask**, **getmaskarray** always returns an Numeric array of typecode MaskType.

mask_or (*m1*, *m2*) returns an object which when used as a mask behaves like the element-wise “logical or” of *m1* and *m2*, where *m1* and *m2* are either masks or None (e.g., they are the results of calling **getmask**). A None is treated as everywhere false. If both *m1* and *m2* are None, it returns None. If just one of them is None, it returns the other. If *m1* and *m2* refer to the same object, a reference to that object is returned.

masked is a module constant equal to an instance of a class that prints as the word ‘masked’ and which will throw an exception of type MAError if any attempt is made to do arithmetic upon it. This constant is returned when an indexing operation results in a scalar result at a masked location.

set_fill_value (*a*, *fill_value*) is the same as *a.set_fill_value* (*fill_value*) if *a* is a masked array; otherwise it does nothing.

Copying or not?

Depending on the arguments results of constructors may or may not contain a separate copy of the data or mask arguments. The easiest way to think about this is as follows: the given field, be it data or a mask, is required to be a Numerical array, possibly with a given typecode, and a mask’s shape must match that of the data. If the copy argument is zero, and the candidate array otherwise qualifies, a reference will be made instead of a copy. If for any reason the data is unsuitable as is, an attempt will be made to make a copy that is suitable. Should that fail, an exception will be thrown. Thus, a *copy=0* argument is more of a hope than a command.

If the basic array constructor is given a masked array as the first argument, its mask, typecode, spacesaver flag, and fill value will be used unless specifically specified by one of the remaining arguments. In particular, if *d* is a masked array, *array(d, copy=0)* is *d*.

Since the default behavior for masks is to use a reference if possible, rather than a copy, which produces a sizeable time and space savings, it is especially important not to modify something you used as a mask argument to a masked array creation routine, if it was a Numeric array of typecode MaskType.

Behaviors

A masked array defines the conversion operators *str* (*x*), *repr* (*x*), *float* (*x*), and *int* (*x*) by applying the corresponding operator to the Numeric array **filled** (*x*)

Indexing and Slicing

Indexing and slicing differ from Numeric: while generally the same, they return a copy, not a reference, when used in an expression that produces a non-scalar result. Consider this example:

```
from Numeric import *
x = array([1., 2., 3.])
y = x[1:]
y[0] = 9.
print x
```

This will print [1., 9., 3.] since *x[1:]* returns a reference to a portion of *x*. Doing the same operation using MA,

```
from MA import *
x = array([1., 2., 3.])
y = x[1:]
y[0] = 9.
print x
```

will print [1., 2., 3.], while *y* will be a separate array whose present value would be [9., 3.]. While sentiment on the correct semantics here is divided amongst the Numeric community as a whole, it is not divided amongst the author’s community, on whose behalf this package is written.

Indexing that produces a scalar result

If indexing into a masked array with one or more indices produces a scalar result, then a scalar value is returned rather than a one-element masked array. This raises the issue of what to return if that location is masked. The answer is that the module constant **masked**, discussed above, is returned.

Assignment to elements and slices

Assignment of a normal value to a single element or slice of a masked array has the effect of clearing the mask in those locations. In this way previously invalid elements become valid. The value being assigned is filled first, so that you are guaranteed that all the elements on the left-hand side are now valid.

Assignment of None to a single element or slice of a masked array has the effect of setting the mask in those locations, and the locations become invalid.

Since these operations change the mask, the result afterwards will no longer share a mask, since masks have copy-on-write semantics.

Module MA: Attributes

Constants **e**, **pi**, **NewAxis** from Numeric, and the constants from module Precision that define nice names for the typecodes.

The special variable **masked** is discussed in “The constant masked” on page 98.

The module **Numeric** is an element of MA, so after from MA import *, you can refer to the functions in Numeric such as Numeric.ones.

Module MA: Functions

Each of the operations discussed below returns an instance of **class MaskedArray**, having performed the desired operation element-wise. In most cases the array arguments can be masked arrays or Numeric arrays or something that Numeric can turn into a Numeric array, such as a list of real numbers.

Where Numeric has a function of the same name, the behavior of the one in MA is the same, except that it “respects” the mask.

Unary functions

The result of a unary operation will be masked wherever the original operand was masked. It may also be masked if the argument is not in the domain of the function. Functions available are:

sqrt, **log**, **log10**, **exp**, **conjugate**, **sin**, **cos**, **tan**, **arcsin**, **arccos**, **arctan**, **sinh**, **cosh**, **tanh**, **absolute**, **fabs**, **negative** (also as operator -x), **nonzero**, **around**, **floor**

fabs (x) is the absolute value of x as a Float32 array. The other functions have their standard meaning.

Binary functions

Binary functions return a result that is masked wherever either of the operands were masked; it may also be masked where the arguments are not in the domain of the function.

add (also as operator +), **subtract** (also as operator -), **multiply** (also as operator *), **divide** (also as operator /), **power** (also as operator **), **remainder**, **fmod**, **hypot**, **arctan2**, **bitwise_and**, **bitwise_or**, **bitwise_xor**.

Comparison operators

Important Tip

Due to limitations in Python, it is not meaningful to compare arrays using the symbolic comparison operators such as “<”. Unfortunately, you can do it; the result just won’t mean anything.

To compare arrays, use the following binary functions. Each of them returns a masked array of 1’s and 0’s.

equal, **not_equal**, **less_equal**, **greater_equal**, **less**, **greater**

Note that as in Numeric, you can use a scalar for one argument and an array for the other.

Logical operators

Arrays of logical values can be manipulated with:

logical_not (unary), **logical_or**, **logical_and**, **logical_xor**.

alltrue (x) returns 1 if all elements of x are true. Masked elements are treated as true.

sometrue (x) returns 1 if any element of x is true. Masked elements are treated as false.

Special array operators

isarray (x), **isMA** (x) return true if x is a masked array.

rank (x) is the number of dimensions in x.

shape (x) returns the shape of x, a tuple of array extents.

resize (x, new_shape) returns a new array with specified shape.

reshape (x, new_shape) returns a copy of x with the given new shape.

ravel (x) returns x as one-dimensional.

concatenate (arrays, axis=0) concatenates the arrays along the specified axis.

repeat (array, repeats, axis = 0) repeat elements of a repeats times along axis. repeats is a sequence of length a.shape[axis] telling how many times to repeat each element.

identity (n) returns the identity matrix of shape n by n.

indices (dimensions, typecode = None) returns an array representing a grid of indices with row-only and column-only variation.

len (x) is defined to be the length of the first dimension of x. This definition, peculiar from the array point of view, is required by the way Python implements slicing. Use **size** (x) for the total length of x.

size (x, axis = None) is the total size of x, or the length of a particular dimension axis whose index is given. When axis is given the dimension of the result is one less than the dimension of x.

count (x, axis = None) counts the number of (non-masked) elements in the array, or in the array along a certain axis. When axis is given the dimension of the result is one less than the dimension of x.

arange, **arrayrange**, **ones**, and **zeros** are the same as in Numeric, but return masked arrays.

sum, and **product** are called the same way as count; the difference is that the result is the sum, product, or average respectively of the unmasked element.

average (x, axis=0, weights=None) computes the average value of the non-masked elements of x along the selected axis. If *weights* is given, it must match the size and shape of x, and the value returned is:

$$\frac{\sum (weights_i \cdot x_i)}{\sum weights_i}$$

In computing these sums, elements that correspond to those that are masked in x or $weights$ are ignored.

allclose (x , y , $fill_value = 1$, $rtol = 1.e-5$, $atol = 1.e-8$) tests whether or not arrays x and y are equal subject to the given relative and absolute tolerances. If $fill_value$ is 1, masked values are considered equal, otherwise they are considered different. The formula used for elements where both x and y have a valid value is:

$$|x - y| < atol + rtol * |y|$$

This means essentially that both elements are small compared to $atol$ or their difference divided by their value is small compared to $rtol$.

allequal (x , y , $fill_value = 1$) is similar to **allclose**, except that exact equality is demanded.

take (a , $indices$, $axis=0$) returns a selection of items from a . See the documentation in the Numeric manual.

transpose (a , $axes=None$) performs a reordering of the axes depending on the tuple of indices $axes$; the default is to reverse the order of the axes.

put (a , $indices$, $values$) is the opposite of **take**. The values of the array a at the locations specified in $indices$ are set to the corresponding value of $values$. The array a must be a contiguous array. The argument $indices$ can be any integer sequence object with values suitable for indexing into the flat form of a . The argument v must be any sequence of values that can be converted to the typecode of a .

```
>>> x = arange(6)
>>> put(x, [2,4], [20,40])
>>> print x
[ 0  1 20  3 40  5 ]
```

Note that the target array a is not required to be one-dimensional. Since it is contiguous and stored in row-major order, the array $indices$ can be treated as indexing a 's elements in storage order.

The wrinkle on this for masked arrays is that if the locations being set by **put** are masked, the mask is cleared in those locations.

choose ($condition$, t) has a result shaped like $condition$. t must be a tuple of two arrays $t1$ and $t2$. Each element of the result is the corresponding element of $t1$ where $condition$ is true, and the corresponding element of $t2$ where $condition$ is false. The result is masked where $condition$ is masked or where the selected element is masked.

If one element of t is the special element **masked** (See “The constant masked” on pag e98.), the type of the result will be the type of the other array. Otherwise, the type of the result is computed using the standard coercion rules.

where ($condition$, x , y) returns an array that is filled (x) where $condition$ is true, filled (y) where the condition is false, and masked where any of the three arguments is masked. It is implemented using **choose**.

innerproduct (a , b) and **dot** (a , b) work as in Numeric, but missing values don't contribute. The result is always a masked array, possibly of length one, because of the possibility that one or more entries in it may be invalid since all the data contributing to that entry was invalid.

outerproduct (a , b) produces a masked array such that $result[i, j] = a[i] * b[j]$. The result will be masked where $a[i]$ or $b[j]$ is masked.

compress ($condition$, x , $dimension=-1$) compresses out only those *valid* values where $condition$ is true.

maximum (x , $y = None$) and **minimum** (x , $y = None$) compute the minimum and maximum valid values of x if y is $None$; with two arguments, they return the element-wise larger or smaller of valid values, and mask the result where either x or y is masked.

sort (x , $axis=-1$, $value = None$) returns the array x sorted along the given axis, with masked values treated as if they have a sort value of $value$ but locations containing $value$ are masked in the result if x had a mask to start with. Thus if x contains $value$ at a non-masked spot, but has other spots masked, the result may not be what you want.

argsort (*x*, *axis* = -1, *fill_value* = None) is unusual in that it returns a Numeric array, equal to Numeric.argsort (filled (*x*, *fill_value*), *axis*); this is an array of indices for sorting along a given axis.

Controlling the size of the string representations

The functions **get_print_limit** () and **set_print_limit** (*n*=0) query and set the limit for converting arrays using **str**() or **repr** (). If an array is printed that is larger than this, the values are not printed; rather you are informed of the type and size of the array. If *n* is zero, the standard Numeric conversion functions are used.

When imported, MA sets this limit to 300, and the limit is also made to apply to standard Numeric arrays as well.

Helper classes

This section discusses other classes defined in module MA.

MAError

Class MAError inherits from Exception, used to raise exceptions in the MA module. Other exceptions are possible, such as errors from the underlying Numeric module.

The constant masked

A constant named **masked**, in Module MA, serves several purposes.

1. When a indexing operation on an MaskedArray instance returns a scalar result, but the location indexed was masked, then **masked** is returned. For example, given a one-dimensional array **x** such that **x.mask(3)** is 1, then **x[3]** is **masked**.
2. When **masked** is assigned to elements of an array via indexing or slicing, those elements become masked. So after **x[3] = masked**, **x[3]** is masked.
3. **masked.display()** is a string that may be used to indicate those elements of an array that are masked when the array is converted to a string, as happens with the print statement.
4. **masked.set_display** (string) can be used to change the value; the default is '--'.
5. **masked.enable**(flag) can be used to enable (flag = 1, default) the use of the display string. If disabled (flag=0), the conversion to string becomes equivalent to **str(self.filled())**.
6. **masked.enabled()** returns the state of the display-enabling flag.
7. Most operations on masked will result in an exception being raised.

Example of masked behavior

```
>>> from MA import *
>>> x=arange(10)
>>> x[3] = masked
>>> print x
[0 ,1 ,2 ,-- ,4 ,5 ,6 ,7 ,8 ,9 ,]
>>> print repr(x)
*** Masked array, mask present ***
Data:
[0 ,1 ,2 ,-- ,4 ,5 ,6 ,7 ,8 ,9 ,]
Mask (fill value [0,])
[0,0,0,1,0,0,0,0,0,0,]

>>> print x[3]
--
>>> print x[3] + 1.0
Traceback (innermost last):
  File "<stdin>", line 1, in ?
```

```
File "/pcmdi/dubois/prerelease/linux/lib/python1.5/site-packages/MA/
__init__.py", line 62, in nope
    raise MAError, 'Cannot do requested operation with a masked value.'
MA.MAError: Cannot do requested operation with a masked value.
```

Class *masked_unary_function*

Given a unary array function $f(x)$, **masked_unary_function** (f , $fill = 0$, $domain = None$) is a function which when applied to an argument x returns f applied to the array filled (x , $fill$), with a mask equal to $mask_or$ ($getmask(x)$, $domain(x)$).

The argument $domain$ therefore should be a callable object that returns true where x is not in the domain of f . The following domains are also supplied as members of module MA:

- `domain_check_interval(a, b)(x) = true` where $x < a$ or $y > b$.
- `domain_tan(eps)(x)` is true where $abs(\cos(x)) < eps$, that is, a domain suitable for the tangent function.
- `domain_greater(v)(x)` is true where $x \leq v$.
- `domain_greater_equal(v)(x)` is true where $x < v$.

Class *masked_binary_function*

Given a binary array function $f(x, y)$, **masked_binary_function** (f , $fillx=0$, $filly=0$, $domain=None$) defines a function whose value at x is f (filled (x , $fillx$), filled (y , $filly$)) with a resulting mask of $mask_or$ ($getmask(x)$, $getmask(y)$, $mask_or$)'d again with those locations where $domain(x, y)$ is true. The values $fillx$ and $filly$ must be chosen so that ($fillx$, $filly$) is in the domain of f .

In addition, an instance of **masked_binary_function** has two methods defined upon it:

- **reduce** ($target$, $axis = 0$)
- **accumulate** ($target$, $axis = 0$)

These methods perform reduction and accumulation as discussed in the section “Ufuncs have special methods” on page 28.

The following domains are available for use as the domain argument:

- `domain_safe_divide()(x, y)` is true where $absolute(x) * 1.e-35 > absolute(y)$. As the comments in the code say, “better ideas welcome”. This domain is used for the divide operator.

ActiveAttributes

MA contains a subpackage, MA.activeattr, which defines the class ActiveAttributes. Class MaskedArray inherits from ActiveAttributes.

An active attribute is a name, say *active*, that appears to be an attribute of a class instance but which in fact is implemented by a triplet of functions, one each corresponding to the operations $x.active$, $x.active = value$, and $del x.active$. To create such an attribute, you inherit from ActiveAttributes and in your classes’ initialization routine you do:

```
ActiveAttributes.__init__(self) # safe for multiple inheritance
self.add_active_attribute_handler ("active", self.actg,
                                  self.acts, self.actd)
```

Here `actg`, `acts`, and `actd` are the three handlers, which should be methods of this class with signatures `actg(self)`, `acts(self, value)`, and `actd(self)`. The last two arguments to `add_active_attribute_handler` can be None, in which case the “active” attribute will behave as if it is read-only.

The “attributes” `shape`, `flat`, `real`, and `imag` in class MaskedArray are actually “active” attributes.

ActiveAttributes also contains methods:

```
def get_active_attribute_handler (self, name):
    "Get current attribute handler associated with a name."
```

```

def get_active_attributes (self):
    "Return the list of attributes that have handlers."
def get_attribute_mode (self, name):
    "Get the mode of an attribute readonly ('r') or writeable ('w')."
def get_basic_attribute_handler (self):
    "Returns the underlying methods that
    handle the three events."
def remove_active_attribute_handler (self, name):
    "Remove the 'active' behavior for
    name."

```

Examples of Using MA

Data with a given value representing missing data

Suppose we have read a one-dimensional list of elements named *x*. We also know that if any of the values are 1.e20, they represent missing data. We want to compute the average value of the data and the vector of deviations from average.

```

>>> from MA import *
>>> x = arange(5)
>>> x[2] = 1.e20
>>> y = masked_values (x, 1.e20)
>>> print average(y)
2.0
>>> print y-average(y)
[-2.00000000e+00, -1.00000000e+00,  --,  1.00000000e+00,
 2.00000000e+00,]

```

Filling in the missing data

Suppose now that we wish to print that same data, but with the missing values replaced by the average value.

```

>>> print filled (y, average(y))

```

Numerical operations

We can do numerical operations without worrying about missing values, dividing by zero, square roots of negative numbers, etc.

```

>>> from MA import *
>>> x=array([1., -1., 3., 4., 5., 6.], mask=[0,0,0,0,1,0])
>>> y=array([1., 2., 0., 4., 5., 6.], mask=[0,0,0,0,0,1])
>>> print sqrt(x/y)
[ 1.00000000e+00,  --,  --,  1.00000000e+00,  --,  --,]

```

Note that four values in the result are invalid: one from a negative square root, one from a divide by zero, and two more where the two arrays *x* and *y* had invalid data. Since the result was of a real type, the print command printed str (filled (sqrt (x/y))).

Seeing the mask

There are various ways to see the mask. One is to print it directly, the other is to convert to the **repr** representation, and a third is get the mask itself. Use of **getmask(x)** is more robust than *x.mask()*, since it will work (returning None) if *x* is a Numeric array or list.

```

>>> x=arange(10)
>>> x[3:5] = masked
>>> print x
[0 ,1 ,2 ,-- ,-- ,5 ,6 ,7 ,8 ,9 ,]

```

```

>>> print repr(x)
*** Masked array, mask present ***
Data:
[0 ,1 ,2 ,-- ,-- ,5 ,6 ,7 ,8 ,9 ,]
Mask (fill value [0,])
[0,0,0,1,1,0,0,0,0,0,]

>>> print getmask(x)
[0,0,0,1,1,0,0,0,0,0,]

```

Filling it your way

If we want to print the data with -1's where the elements are masked, we use **filled**.

```

>>> print filled(z, -1)
[ 1.,-1.,-1., 1.,-1.,-1.,]

```

Ignoring extreme values

Suppose we have an array `d` and we wish to compute the average of the values in `d` but ignore any data outside the range -100. to 100.

```

v = masked_outside(d, -100., 100.)
print average(v)

```

Averaging an entire multidimensional array

The problem with averaging over an entire array is that the average function only reduces one dimension at a time. So to average the entire array, **ravel** it first.

```

>>> x
*** Masked array, no mask ***
Data:
[[ 0, 1, 2,]
 [ 3, 4, 5,]
 [ 6, 7, 8,]
 [ 9,10,11,]]

>>> average(x)
*** Masked array, no mask ***
Data:
[ 4.5, 5.5, 6.5,]

>>> average(ravel(x))
5.5

```


Index

Symbols

- 97
- ... 26
- : 25
- :: 25, 53
- `_PyArray_multiply_list` 66

A

- `alltrue()` 44
- `argmax()` 39
- `argsort()` 39
- `array` 73
- Array Attributes 47
- Array Functions 34
- Array Methods 45
- `array()` 12
- `array_repr()` 42
- `array_str()` 42
- `array2string` 55
- `arrayobject.h` 58
- `arrayrange()` 17
- arrays with missing values 89
- `astype` 22
- Automatic Coercions 20
- axes 12

B

- Broadcasting 54
- Bugs 6
- `byteswapped()` 45

C

- Casting 20
- Character 13
- `choose` 36
- `clip()` 40
- Code Organization 49
- Coercion 20
- Complex 13
- complex numbers 13
- `Complex0` 13
- `Complex128` 13
- `Complex16` 13
- `Complex32` 13

Complex64 13
Complex8 13
compress() 37
concatenate() 41
Contiguous arrays 60
Convenience 10
convolve 43
correlation 43
cross_correlate 43
cumproduct() 44
cumsum() 44
CVS 6

D

determinant() 81
diagonal() 37, 42
dimensions 12
Discussion list 6
Distutils 58
Dubois, Paul F. 89

E

eigenvalues() 80
eigenvectors() 81
element-wise 12
Ellipses 54

F

FFT 77
fft() 77
flat (attribute) 47
Float 13
Float0 13
Float128 13
Float16 13
Float32 13
Float64 13
Float8 13
floating point exceptions 57
Fortran and Python 6
fpectl 57
fromfunction() 19
fromstring() 39, 43
FTP Site 6

G

gather

see put() 35
generalized_inverse() 81
get_seed() 82
Getting array values 24
greece 7

H

header files 58
homogeneous 12

I

identity() 20, 43
IDLE 8
imaginary (attribute) 47
imaginary numbers 13
import_array() 58
Indexing 53
indices() 40
innerproduct() 42
Installing NumPy 5
Int 13
Int0 13
Int128 13
Int16 13
Int32 13
Int64 13
Int8 13
inverse() 80
inverse_fft() 77
inverse_real_fft() 78
iscontiguous() 45
itemsize() 45

L

limiting printing of large arrays 98
linear_least_squares() 81
Logical Ufuncs 30

M

MA
absolute 95
add 95
allclose 97
allequal 97
alltrue 96
arange 96
arccos 95

arcsin 95
 arctan 95
 arctan2 95
 argsort 98
 arithmetic operators 95
 around 95
 array (constructor) 92
 arrayrange 96
 average 96
 bitwise_and 95
 bitwise_or 95
 bitwise_xor 95
 choose 97
 class MA 95
 comparisons 96
 compress 97
 concatenate 96
 conjugate 95
 cos 95
 cosh 95
 count 96
 default_character_fill_value 93
 default_complex_fill_value 93
 default_integer_fill_value 93
 default_object_fill_value 93
 default_real_fill_value 93
 divide 95
 equal 96
 exp 95
 fabs 95
 fill_value 93
 floor 95
 fmod 95
 functions and operators 94
 get_print_limit 98
 getmask 94
 getmaskarray 94
 greater 96
 greater_equal 96
 hypot 95
 identity 96
 indices 96
 innerproduct 97
 installing 89
 invalid, defined 89
 is_mask 93
 isarray 96
 isMA 96
 len 96
 less 96
 less_equal 96
 log 95
 log10 95
 logical operators 96
 logical_and 96
 logical_not 96
 logical_or 96
 logical_xor 96
 MAError 98
 make_mask 93
 mask method 94
 mask_object (constructor) 92
 mask_or 94
 mask_values 89
 masked (constant) 98
 masked (constant), role in
 printing 98
 masked (module constant) 93
 masked (special constant) 95
 masked, defined 89
 masked, special role in choose 97
 masked_array 92
 masked_array (constructor) 92
 masked_equal 92
 masked_greater 92
 masked_greater_equal 92
 masked_less 92
 masked_less_equal 92
 masked_not_equal 92
 masked_unary_function 99
 masked_values 92
 masked_where 92
 masks 93
 maximum 97
 maximum with two arguments 97
 minimum 97
 multiply 95
 negative 95
 nonzero 95
 not_equal 96
 ones 96
 outerproduct 97
 power 95
 product 96

- put 97
- rank 96
- ravel 96
- remainder 95
- reshape 96
- resize 96
- set_fill_value 94
- set_print_limit 98
- shape 96
- sin 95
- sinh 95
- size 96
- sometrue 96
- sort 97
- sqrt 95
- subtract 95
- sum 96
- take 97
- tan 95
- tanh 95
- transpose 97
- where 97
- zeros 96
- MA (masked arrays) 9
- mailing list 6
- masked arrays 89
- Matrix 73
- matrixmultiply() 40
- missing values 9, 89
- multiarray 9, 12, 52, 73

N

- NaNs 57
- NewAxis 33, 54
- nonzero() 37
- Numeric.py 9
- NumTut 7

O

- observations, dealing with missing 89
- ones() 17
- outerproduct 42

P

- Patches 6
- permutation() 82
- Pickling 57

- printing
 - using MA to limit printing of large arrays 98
- product() 44
- Pseudo Indices 32
- put() 35
- PyArray_As1D 67
- PyArray_As2D 67
- PyArray_CanCastSafely 66
- PyArray_Cast 66
- PyArray_Check 65
- PyArray_ContiguousFromObject 67
- PyArray_Copy 67
- PyArray_CopyFromObject 67
- PyArray_DescrFromType 66
- PyArray_Free 67
- PyArray_FromDims 66
- PyArray_FromDimsAndData 66
- PyArray_FromObject 67
- PyArray_INCREf 65
- PyArray_ObjectType 66
- PyArray_Reshape 67
- PyArray_Return 67
- PyArray_SetNumericOps 65
- PyArray_SetStringFunction 66
- PyArray_Size 66
- PyArray_Take 67
- PyArray_XDECREF 66
- PyArrayError 66
- PyArrayObject 58, 64
- Pyfort 6
- PyObject 14
- PyUFunc_Check 70
- PyUFunc_FromFuncAndData 70
- PyUFunc_GenericFunction 71

R

- randint() 82
- random() 82
- RandomArray 82
- rank 12, 73
- ravel 37
- real (attribute) 47
- real_fft() 77
- repeat() 36, 42
- reshape 15
- resize 16

resize() 42

S

seed() 82

Setting array values 24

shape 73

singular_value_decomposition() 81

Slicing Arrays 25

solve_linear_equations() 80

sometrue() 44

sort() 38

SourceForge 6

sum() 44

swapaxes() 41

T

take() 34

Textual representations 55

thread 8

Tk 8

tolist() 46

tostring() 45

trace() 38

transpose() 36

typecode 12, 21, 73

typecode() 45

Typecodes 52

U

ufunc 9, 73

Ufunc shorthands 31

UfuncObject 68

ufuncobject.h 58

Ufuncs 27

Unary Mathematical Ufuncs 30

uniform() 82

universal function 73

universal functions 9

Unix 6

Unpickling 57

UnsignedInt8 13

UserArray 73

V

view 7

W

where() 37, 43

Z

Zero-dimensional arrays 60

zeros() 17