# SciPy Tutorial 

Travis E. Oliphant

8th October 2004

## 1 Introduction

SciPy is a collection of mathematical algorithms and convenience functions built on the Numeric extension for Python. It adds significant power to the interactive Python session by exposing the user to high-level commands and classes for the manipulation and visualization of data. With SciPy, an interactive Python session becomes a data-processing and system-prototyping environment rivaling sytems such as Matlab, IDL, Octave, R-Lab, and SciLab.

The additional power of using SciPy within Python, however, is that a powerful programming language is also available for use in developing sophisticated programs and specialized applications. Scientific applications written in SciPy benefit from the development of additional modules in numerous niche's of the software landscape by developers across the world. Everything from parallel programming to web and database subroutines and classes have been made available to the Python programmer. All of this power is available in addition to the mathematical libraries in SciPy.

This document provides a tutorial for the first-time user of SciPy to help get started with some of the features available in this powerful package. It is assumed that the user has already installed the package. Some general Python facility is also assumed such as could be acquired by working through the Tutorial in the Python distribution. Throughout this tutorial it is assumed that the user has imported all of the names defined in the SciPy namespace using the command

```
>>> from scipy import *
```


### 1.1 General Help

Python provides the facility of documentation strings. The functions and classes available in SciPy use this method for on-line documentation. There are two methods for reading these messages and getting help. Python provides the command help in the pydoc module. Entering this command with no arguments (i.e. iii help ) launches an interactive help session that allows searching through the keywords and modules available to all of Python. Running the command help with an object as the argument displays the calling signature, and the documentation string of the object.

The pydoc method of help is sophisticated but uses a pager to display the text. Sometimes this can interfere with the terminal you are running the interactive session within. A scipy-specific help system is also available under the command scipy.info. The signature and documntation string for the object passed to the help command are printed to standard output (or to a writeable object passed as the third argument). The second keyword argument of "scipy.info" defines the maximum width of the line for printing. If a module is passed as the argument to help than a list of the functions and classes defined in that module is printed. For example:

```
>>> info(optimize.fmin)
    fmin(func, x0, args=(), xtol=0.0001, ftol=0.0001, maxiter=None, maxfun=None,
                full_output=0, printmessg=1)
Minimize a function using the simplex algorithm.
Description:
```

```
Uses a Nelder-Mead simplex algorithm to find the minimum of function
of one or more variables.
```


## Inputs:

```
func -- the Python function or method to be minimized.
x0 -- the initial guess.
args -- extra arguments for func.
xtol -- relative tolerance
Outputs: (xopt, {fopt, warnflag})
xopt -- minimizer of function
fopt -- value of function at minimum: fopt = func(xopt)
warnflag -- Integer warning flag:
                                    1 : 'Maximum number of function evaluations.'
                                    2 : 'Maximum number of iterations.'
```

Additional Inputs:

```
xtol -- acceptable relative error in xopt for convergence.
ftol -- acceptable relative error in func(xopt) for convergence.
maxiter -- the maximum number of iterations to perform.
maxfun -- the maximum number of function evaluations.
full_output -- non-zero if fval and warnflag outputs are desired.
printmessg -- non-zero to print convergence messages.
```

Another useful command is source. When given a function written in Python as an argument, it prints out a listing of the source code for that function. This can be helpful in learning about an algorithm or understanding exactly what a function is doing with its arguments. Also don't forget about the Python command dir which can be used to look at the namespace of a module or package.

### 1.2 SciPy Organization

SciPy is organized into subpackages covering different scientific computing domains. Some common functions which several subpackages rely on live under the scipy_base package which is installed at the same directory level as the scipy package itself and could be installed separately. This allows for the possibility of separately distributing the subpackages of scipy as long as scipy_base package is provided as well.

Two other packages are installed at the higher-level: scipy_distutils and weave. These two packages while distributed with main scipy package could see use independently of scipy and so are treated as separate packages and described elsewhere.

The remaining subpackages are summarized in the following table (a * denotes an optional sub-package that requires additional libraries to function or is not available on all platforms).

| Subpackage | Description |
| :---: | :---: |
| cluster | Clustering algorithms |
| cow | Cluster of Workstations code for parallel programming |
| fftpack | FFT based on fftpack - default |
| fftw* $_{\text {* }}$ | FFT based on fftw - requires FFTW libraries (is this still needed?) |
| ga | Genetic algorithms |
| gplt* | Plotting - requires gnuplot |
| integrate | Integration |
| interpolate | Interpolation |
| io | Input and Output |
| linalg | Linear algebra |
| optimize | Optimization and root-finding routines |
| plt* | Plotting - requires wxPython |
| signal | Signal processing |
| special | Special functions |
| stats | Statistical distributions and functions |
| xplt | Plotting with gist |

Because of their ubiquitousness, some of the functions in these subpackages are also made available in the scipy namespace to ease their use in interactive sessions and programs. In addition, many convenience functions are located in the scipy_base package and the in the top-level of the scipy package. Before looking at the sub-packages individually, we will first look at some of these common functions.

## 2 Basic functions in scipy_base and top-level scipy

### 2.1 Interaction with Numeric

To begin with, all of the Numeric functions have been subsumed into the scipy namespace so that all of those functions are available without additionally importing Numeric. In addition, the universal functions (addition, subtraction, division) have been altered to not raise exceptions if floating-point errors are encountered ${ }^{1}$, instead NaN's and Inf's are returned in the arrays. To assist in detection of these events new universal functions (isnan, isfinite, isinf) have been added. In addition, the comparision operators have been changed to allow comparisons and logical operations of complex numbers (only the real part is compared). Also, with the new universal functions in SciPy, the logical operations all return arrays of unsigned bytes (8-bits per element instead of the old 32 -bits, or even 64 -bits) per element ${ }^{2}$.

Finally, some of the basic functions like log, sqrt, and inverse trig functions have been modified to return complex numbers instead of NaN's where appropriate (i.e. scipy.sqrt( -1 ) returns 1 j ).

### 2.2 Scipy_base routines

The purpose of scipy_base is to collect general-purpose routines that the other sub-packages can use. These routines are divided into several files for organizational purposes, but they are all available under the scipy_base namespace (and the scipy namespace). There are routines for type handling and type checking, shape and matrix manipulation, polynomial processing, and other useful functions. Rather than giving a detailed description of each of these functions (which is available using the help, info and source commands), this tutorial will discuss some of the more useful commands which require a little introduction to use to their full potential.

[^0]
### 2.2.1 Type handling

Note the difference between iscomplex (isreal) and iscomplexobj (isrealobj). The former command is array based and returns byte arrays of ones and zeros providing the result of the element-wise test. The latter command is object based and returns a scalar describing the result of the test on the entire object.

Often it is required to get just the real and/or imaginary part of a complex number. While complex numbers and arrays have attributes that return those values, if one is not sure whether or not the object will be complex-valued, it is better to use the functional forms real and imag. These functions succeed for anything that can be turned into a Numeric array. Consider also the function real_if_close which transforms a complex-valued number with tiny imaginary part into a real number.

Occasionally the need to check whether or not a number is a scalar (Python (long)int, Python float, Python complex, or rank-0 array) occurs in coding. This functionality is provided in the convenient function isscalar which returns a 1 or a 0 .

Finally, ensuring that objects are a certain Numeric type occurs often enough that it has been given a convenient interface in SciPy through the use of the cast dictionary. The dictionary is keyed by the type it is desired to cast to and the dictionary stores functions to perform the casting. Thus, >>> a = cast ['f'] (d) returns an array of float32 from d. This function is also useful as an easy way to get a scalar of a certain type: >>> fpi = cast['f'](pi).

### 2.2.2 Index Tricks

Thre are some class instances that make special use of the slicing functionality to provide efficient means for array construction. This part will discuss the operation of $\mathbf{m g r i d}, \mathbf{r}_{-}$, and $\mathbf{c}_{-}$for quickly constructing arrays.

One familiar with Matlab may complain that it is difficult to construct arrays from the interactive session with Python. Suppose, for example that one wants to construct an array that begins with 3 followed by 5 zeros and then contains 10 numbers spanning the range -1 to 1 (inclusive on both ends). Before SciPy, I would need to enter something like the following >>> concatenate ( ([3], [0] $* 5$, arange ( $-1,1.002,2 / 9.0$ ) ). With the $\mathbf{r}_{-}$command one can enter this as $\ggg r_{-}[3,[0] * 5,-1: 1: 10 j]$ which can ease typing in an interactive session. Notice how objects are concatenated, and the slicing syntax is used (abused) to construct ranges. The other term that deserves a little explanation is the use of the complex number 10 j as the step size in the slicing syntax. This non-standard use allows the number to be interpreted as the number of points to produce in the range rather than as a step size (note we would have used the long integer notation, 10L, but this notation may go away in Python as the integers became unified). This non-standard usage may be unsightly to some, but it gives the user the ability to quickly construct complicated vectors in a very readable fashion. When the number of points is specified in this way, the end-point is inclusive.

The "r" stands for row concatenation because if the objects between commas are 2 dimensional arrays, they are stacked by rows (and thus must have commensurate columns). There is an equivalent command $\mathbf{c}_{-}$ that stacks 2 d arrays by columns but works identically to $\mathbf{r}_{-}$for 1 d arrays.

Another very useful class instance which makes use of extended slicing notation is the function mgrid. In the simplest case, this function can be used to construct 1 d ranges as a convenient substitute for arange. It also allows the use of complex-numbers in the step-size to indicate the number of points to place between the (inclusive) end-points. The real purpose of this function however is to produce $\mathrm{N}, \mathrm{N}-\mathrm{d}$ arrays which provide coordinate arrays for an N -dimensional volume. The easiest way to understand this is with an example of its usage:

```
>>> mgrid[0:5,0:5]
array([[[0, 0, 0, 0, 0],
    [1, 1, 1, 1, 1],
    [2, 2, 2, 2, 2],
    [3, 3, 3, 3, 3],
    [4, 4, 4, 4, 4]],
    [[0, 1, 2, 3, 4],
    [0, 1, 2, 3, 4],
    [0, 1, 2, 3, 4],
```

```
    [0, 1, 2, 3, 4],
    [0, 1, 2, 3, 4]]])
>>> mgrid[0:5:4j,0:5:4j]
array([[[ 0. , 0. , 0. , 0. ] ,
    [ 1.6667, 1.6667, 1.6667, 1.6667],
    [ 3.3333, 3.3333, 3.3333, 3.3333],
    [ 5. , 5. , 5. , 5. ]],
    [[ 0. , 1.6667, 3.3333, 5. ],
    [ 0. , 1.6667, 3.3333, 5. ],
    [ 0. , 1.6667, 3.3333, 5. ],
    [0. , 1.6667, 3.3333, 5. ]]])
```

Having meshed arrays like this is sometimes very useful. However, it is not always needed just to evaluate some N-dimensional function over a grid due to the array-broadcasting rules of Numeric and SciPy. If this is the only purpose for generating a meshgrid, you should instead use the function ogrid which generates an "open" grid using NewAxis judiciously to create N, N-d arrays where only one-dimension has length greater than 1 . This will save memory and create the same result if the only purpose for the meshgrid is to generate sample points for evaluation of an $\mathrm{N}-\mathrm{d}$ function.

### 2.2.3 Shape manipulation

In this category of functions are routines for squeezing out length-one dimensions from N -dimensional arrays, ensure that an array is at least 1 -, 2 -, or 3 -dimensional, and stacking (concatenating) arrays by rows, columns, and "pages" (in the third dimension). Routines for splitting arrays (roughly the opposite of stacking arrays).

### 2.2.4 Matrix manipulations

These are functions specifically suited for 2-dimensional arrays that were part of MLab in the Numeric distribution, but have been placed in scipy_base for completeness.

### 2.2.5 Polynomials

There are two (interchangeable) ways to deal with 1-d polynomials in SciPy. The first is to use the poly1d class in scipy_base. This class accepts coefficients or polynomial roots to initialize a polynomial. The polynomial object can then be manipulated in algebraic expressions, integrated, differentiated, and evaluated. It even prints like a polynomial:

```
>>> p = poly1d([3,4,5])
>>> print p
    2
3x + 4x + 5
>>> print p*p
    4 3 2
9x + 24 x + 46 x + 40 x + 25
>>> print p.integ(k=6)
    3 2
x + 2 x + 5 x + 6
>>> print p.deriv()
6 x + 4
>>> p([4,5])
array([ 69, 100])
```

The other way to handle polynomials is as an array of coefficients with the first element of the array giving the coefficient of the highest power. There are explicit functions to add, subtract, multiply, divide, integrate, differentiate, and evaluate polynomials represented as sequences of coefficients.

### 2.2.6 Vectorizing functions (vectorize)

One of the features that SciPy provides is a class vectorize to convert an ordinary Python function which accepts scalars and returns scalars into a "vectorized-function" with the same broadcasting rules as other Numeric functions (i.e. the Universal functions, or ufuncs). For example, suppose you have a Python function named addsubtract defined as:

```
>>> def addsubtract(a,b):
    if a > b:
        return a - b
    else:
        return a + b
```

which defines a function of two scalar variables and returns a scalar result. The class vectorize can be used to "vectorize" this function so that

```
>>> vec_addsubstract = vectorize(addsubtract)
```

returns a function which takes array arguments and returns an array result:

```
>>> vec_addsubtract([0,3,6,9],[1,3,5,7])
```

$\operatorname{array}([1,6,1,2])$

This particular function could have been written in vector form without the use of generalfunction. But, what if the function you have written is the result of some optimization or integration routine. Such functions can likely only be vectorized using vectorize.

### 2.2.7 Other useful functions

There are several other functions in the scipy_base package including most of the other functions that are also in MLab that comes with the Numeric package. The reason for duplicating these functions is to allow SciPy to potentially alter their original interface and make it easier for users to know how to get access to functions >>> from scipy import *.

New functions which should be mentioned are $\bmod (\mathbf{x}, \mathbf{y})$ which can replace $\mathbf{x} \% \mathbf{y}$ when it is desired that the result take the sign of $\mathbf{y}$ instead of $\mathbf{x}$. Also included is fix which always rounds to the nearest integer towards zero. For doing phase processing, the functions angle, and unwrap are also useful. Also, the linspace and logspace functions return equally spaced samples in a linear or log scale. Finally, mention should be made of the new function select which extends the functionality of where to include multiple conditions and multiple choices. The calling convention is select (condlist, choicelist, default=0). Select is a vectorized form of the multiple if-statement. It allows rapid construction of a function which returns an array of results based on a list of conditions. Each element of the return array is taken from the array in a choicelist corresponding to the first condition in condlist that is true. For example

```
>>> x = r_[-2:3]
>>> x
array([-2, -1, 0, 1, 2])
>>> select([x > 3, x >= 0],[0,x+2])
array([0, 0, 2, 3, 4])
```


### 2.3 Common functions

Some functions depend on sub-packages of SciPy but should be available from the top-level of SciPy due to their common use. These are functions that might have been placed in scipy_base except for their dependence on other sub-packages of SciPy. For example the factorial and comb functions compute $n!$ and $n!/ k!(n-k)$ ! using either exact integer arithmetic (thanks to Python's Long integer object), or by using floating-point precision and the gamma function. The functions rand and randn are used so often that they warranted
a place at the top level. There are convenience functions for the interactive use: disp (similar to print), and who (returns a list of defined variables and memory consumption-upper bounded). Another function returns a common image used in signal processing: lena.

Finally, two functions are provided that are useful for approximating derivatives of functions using discrete-differences. The function central_diff_weights returns weighting coefficients for an equally-spaced $N$-point approximation to the derivative of order $o$. These weights must be multiplied by the function corresponding to these points and the results added to obtain the derivative approximation. This function is intended for use when only samples of the function are avaiable. When the function is an object that can be handed to a routine and evaluated, the function derivative can be used to automatically evaluate the object at the correct points to obtain an N-point approximation to the $o^{\text {th }}$-derivative at a given point.

## 3 Special functions (special)

The main feature of the special package is the definition of numerous special functions of mathematical physics. Available functions include airy, elliptic, bessel, gamma, beta, hypergeometric, parabolic cylinder, mathieu, spheroidal wave, struve, and kelvin. There are also some low-level stats functions that are not intended for general use as an easier interface to these functions is provided by the stats module. Most of these functions can take array arguments and return array results following the same broadcasting rules as other math functions in Numerical Python. Many of these functions also accept complex-numbers as input. For a complete list of the available functions with a one-line description type >>>info(special). Each function also has it's own documentation accessible using help. If you don't see a function you need, consider writing it and contributing it to the library. You can write the function in either C, Fortran, or Python. Look in the source code of the library for examples of each of these kind of functions.

## 4 Integration (integrate)

The integrate sub-package provides several integration techniques including an ordinary differential equation integrator. An overview of the module is provided by the help command:

```
>>> help(integrate)
Methods for Integrating Functions
odeint -- Integrate ordinary differential equations.
quad -- General purpose integration.
dblquad -- General purpose double integration.
tplquad -- General purpose triple integration.
gauss_quad -- Integrate func(x) using Gaussian quadrature of order n.
gauss_quadtol -- Integrate with given tolerance using Gaussian quadrature.
See the orthogonal module (integrate.orthogonal) for Gaussian
    quadrature roots and weights.
```


### 4.1 General integration (integrate.quad)

The function quad is provided to integrate a function of one variable between two points. The points can be $\pm \infty$ ( $\pm$ integrate.inf $)$ to indicate infinite limits. For example, suppose you wish to integrate a bessel function $\mathrm{jv}(2.5, \mathrm{x})$ along the interval $[0,4.5]$.

$$
I=\int_{0}^{4.5} J_{2.5}(x) d x
$$

This could be computed using quad:

```
>>> result = integrate.quad(lambda x: special.jv(2.5,x), 0, 4.5)
>>> print result
(1.1178179380783249, 7.8663172481899801e-09)
>>> I = sqrt(2/pi)*(18.0/27*sqrt(2)*cos(4.5)-4.0/27*sqrt(2)*sin(4.5)+
    sqrt(2*pi)*special.fresnl(3/sqrt(pi))[0])
>>> print I
1.117817938088701
```

>>> print abs(result [0]-I)
$1.03761443881 \mathrm{e}-11$

The first argument to quad is a "callable" Python object (i.e a function, method, or class instance). Notice the use of a lambda-function in this case as the argument. The next two arguments are the limits of integration. The return value is a tuple, with the first element holding the estimated value of the integral and the second element holding an upper bound on the error. Notice, that in this case, the true value of this integral is

$$
I=\sqrt{\frac{2}{\pi}}\left(\frac{18}{27} \sqrt{2} \cos (4.5)-\frac{4}{27} \sqrt{2} \sin (4.5)+\sqrt{2 \pi} \mathrm{Si}\left(\frac{3}{\sqrt{\pi}}\right)\right),
$$

where

$$
\operatorname{Si}(x)=\int_{0}^{x} \sin \left(\frac{\pi}{2} t^{2}\right) d t
$$

is the Fresnel sine integral. Note that the numerically-computed integral is within $1.04 \times 10^{-11}$ of the exact result - well below the reported error bound.

Infinite inputs are also allowed in quad by using $\pm$ integrate.inf (or inf) as one of the arguments. For example, suppose that a numerical value for the exponential integral:

$$
E_{n}(x)=\int_{1}^{\infty} \frac{e^{-x t}}{t^{n}} d t
$$

is desired (and the fact that this integral can be computed as $\operatorname{special} . \operatorname{expn}(\mathrm{n}, \mathrm{x})$ is forgotten). The functionality of the function special.expn can be replicated by defining a new function vec_expint based on the routine quad:

```
>>> from integrate import quad, Inf
>>> def integrand(t,n,x):
    return exp(-x*t) / t**n
>>> def expint(n,x):
    return quad(integrand, 1, Inf, args=(n, x)) [0]
>>> vec_expint = vectorize(expint)
>>> vec_expint(3,arange(1.0,4.0,0.5))
array([ 0.1097, 0.0567, 0.0301, 0.0163, 0.0089, 0.0049])
>>> special.expn(3,arange(1.0,4.0,0.5))
array([ 0.1097, 0.0567, 0.0301, 0.0163, 0.0089, 0.0049])
```

The function which is integrated can even use the quad argument (though the error bound may underestimate the error due to possible numerical error in the integrand from the use of quad). The integral in this case is

$$
I_{n}=\int_{0}^{\infty} \int_{1}^{\infty} \frac{e^{-x t}}{t^{n}} d t d x=\frac{1}{n}
$$

```
>>> result = quad(lambda x: expint(3, x), 0, Inf)
>>> print result
(0.33333333324560266, 2.8548934485373678e-09)
>>> I3 = 1.0/3.0
>>> print I3
0.333333333333
>>> print I3 - result[0]
8.77306560731e-11
```

This last example shows that multiple integration can be handled using repeated calls to quad. The mechanics of this for double and triple integration have been wrapped up into the functions dblquad and tplquad. The function, dblquad performs double integration. Use the help function to be sure that the arguments are defined in the correct order. In addition, the limits on all inner integrals are actually functions which can be constant functions. An example of using double integration to compute several values of $I_{n}$ is shown below:

```
>>> from __future__ import nested_scopes
>>> from integrate import quad, dblquad, Inf
>>> def I(n):
    return dblquad(lambda t, x: exp(-x*t)/t**n, 0, Inf, lambda x: 1, lambda x: Inf)
>>> print I(4)
(0.25000000000435768, 1.0518245707751597e-09)
>>> print I(3)
(0.33333333325010883, 2.8604069919261191e-09)
>>> print I(2)
(0.49999999999857514, 1.8855523253868967e-09)
```


### 4.2 Gaussian quadrature (integrate.gauss_quadtol)

A few functions are also provided in order to perform simple Gaussian quadrature over a fixed interval. The first is fixed_quad which performs fixed-order Gaussian quadrature. The second function is quadrature which performs Gaussian quadrature of multiple orders until the difference in the integral estimate is beneath some tolerance supplied by the user. These functions both use the module special.orthogonal which can calculate the roots and quadrature weights of a large variety of orthogonal polynomials (the polynomials themselves are available as special functions returning instances of the polynomial class - e.g. special.legendre).

### 4.3 Integrating using samples

There are three functions for computing integrals given only samples: trapz, simps, and romb. The first two functions use Newton-Coates formulas of order 1 and 2 respectively to perform integration. These two functions can handle, non-equally-spaced samples. The trapezoidal rule approximates the function as a straight line between adjacent points, while Simpson's rule approximates the function between three adjacent points as a parabola.

If the samples are equally-spaced and the number of samples available is $2^{k}+1$ for some integer $k$, then Romberg integration can be used to obtain high-precision estimates of the integral using the available samples. Romberg integration uses the trapezoid rule at step-sizes related by a power of two and then performs Richardson extrapolation on these estimates to approximate the integral with a higher-degree of accuracy. (A different interface to Romberg integration useful when the function can be provided is also available as integrate.romberg).

### 4.4 Ordinary differential equations (integrate.odeint)

Integrating a set of ordinary differential equations (ODEs) given initial conditions is another useful example. The function odeint is available in SciPy for integrating a first-order vector differential equation:

$$
\frac{d \mathbf{y}}{d t}=\mathbf{f}(\mathbf{y}, t)
$$

given initial conditions $\mathbf{y}(0)=y_{0}$, where $\mathbf{y}$ is a length $N$ vector and $\mathbf{f}$ is a mapping from $\mathcal{R}^{N}$ to $\mathcal{R}^{N}$. A higher-order ordinary differential equation can always be reduced to a differential equation of this type by introducing intermediate derivatives into the $\mathbf{y}$ vector.

For example suppose it is desired to find the solution to the following second-order differential equation:

$$
\frac{d^{2} w}{d z^{2}}-z w(z)=0
$$

with initial conditions $w(0)=\frac{1}{\sqrt[3]{3^{2}} \Gamma\left(\frac{2}{3}\right)}$ and $\left.\frac{d w}{d z}\right|_{z=0}=-\frac{1}{\sqrt[3]{3} \Gamma\left(\frac{1}{3}\right)}$. It is known that the solution to this differential equation with these boundary conditions is the Airy function

$$
w=\operatorname{Ai}(z)
$$

which gives a means to check the integrator using special.airy.
First, convert this ODE into standard form by setting $\mathbf{y}=\left[\frac{d w}{d z}, w\right]$ and $t=z$. Thus, the differential equation becomes

$$
\frac{d \mathbf{y}}{d t}=\left[\begin{array}{c}
t y_{1} \\
y_{0}
\end{array}\right]=\left[\begin{array}{cc}
0 & t \\
1 & 0
\end{array}\right]\left[\begin{array}{l}
y_{0} \\
y_{1}
\end{array}\right]=\left[\begin{array}{cc}
0 & t \\
1 & 0
\end{array}\right] \mathbf{y} .
$$

In other words,

$$
\mathbf{f}(\mathbf{y}, t)=\mathbf{A}(t) \mathbf{y} .
$$

As an interesting reminder, if $\mathbf{A}(t)$ commutes with $\int_{0}^{t} \mathbf{A}(\tau) d \tau$ under matrix multiplication, then this linear differential equation has an exact solution using the matrix exponential:

$$
\mathbf{y}(t)=\exp \left(\int_{0}^{t} \mathbf{A}(\tau) d \tau\right) \mathbf{y}(0)
$$

However, in this case, $\mathbf{A}(t)$ and its integral do not commute.
There are many optional inputs and outputs available when using odeint which can help tune the solver. These additional inputs and outputs are not needed much of the time, however, and the three required input arguments and the output solution suffice. The required inputs are the function defining the derivative, fprime, the initial conditions vector, $y 0$, and the time points to obtain a solution, $t$, (with the initial value point as the first element of this sequence). The output to odeint is a matrix where each row contains the solution vector at each requested time point (thus, the initial conditions are given in the first output row).

The following example illustrates the use of odeint including the usage of the Dfun option which allows the user to specify a gradient (with respect to $\mathbf{y}$ ) of the function, $\mathbf{f}(\mathbf{y}, t)$.

```
>>> from integrate import odeint
>>> from special import gamma, airy
>>> y1_0 = 1.0/3**(2.0/3.0)/gamma(2.0/3.0)
>>> y0_0 = -1.0/3**(1.0/3.0)/gamma(1.0/3.0)
>>> y0 = [y0_0, y1_0]
>>> def func(y, t):
    return [t*y[1],y[0]]
>>> def gradient(y,t):
    return [[0,t], [1,0]]
>>> x = arange(0,4.0, 0.01)
```

```
>>> t = x
>>> ychk = airy(x)[0]
>>> y = odeint(func, y0, t)
>>> y2 = odeint(func, y0, t, Dfun=gradient)
>>> import sys
>>> sys.float_output_precision = 6
>>> print ychk[:36:6]
[ 0.355028 0.339511 0.324068 0.308763 0.293658 0.278806]
>>> print y[:36:6,1]
```



```
>>> print y2[:36:6,1]
[ 0.355028 0.339511 0.324067 0.308763 0.293658 0.278806]
```


## 5 Optimization (optimize)

There are several classical optimization algorithms provided by SciPy in the optimize package. An overview of the module is available using help (or pydoc.help):

```
>>> info(optimize)
    Optimization Tools
A collection of general-purpose optimization routines.
    fmin -- Nelder-Mead Simplex algorithm
                                (uses only function calls)
    fmin_powell -- Powell's (modified) level set method (uses only
                                    function calls)
    fmin_bfgs -- Quasi-Newton method (can use function and gradient)
    fmin_ncg -- Line-search Newton Conjugate Gradient (can use
                                    function, gradient and hessian).
    leastsq -- Minimize the sum of squares of M equations in
                                N unknowns given a starting estimate.
Scalar function minimizers
    fminbound -- Bounded minimization of a scalar function.
    brent -- 1-D function minimization using Brent method.
    golden -- 1-D function minimization using Golden Section method
    bracket -- Bracket a minimum (given two starting points)
Also a collection of general_purpose root-finding routines.
    fsolve -- Non-linear multi-variable equation solver.
Scalar function solvers
    brentq -- quadratic interpolation Brent method
    brenth -- Brent method (modified by Harris with
    hyperbolic extrapolation)
    ridder -- Ridder's method
    bisect -- Bisection method
```

```
newton -- Secant method or Newton's method
fixed_point -- Single-variable fixed-point solver.
```

The first four algorithms are unconstrained minimization algorithms (fmin: Nelder-Mead simplex, fmin_bfgs: BFGS, fmin_ncg: Newton Conjugate Gradient, and leastsq: Levenburg-Marquardt). The fourth algorithm only works for functions of a single variable but allows minimization over a specified interval. The last algorithm actually finds the roots of a general function of possibly many variables. It is included in the optimization package because at the (non-boundary) extreme points of a function, the gradient is equal to zero.

### 5.1 Nelder-Mead Simplex algorithm (optimize.fmin)

The simplex algorithm is probably the simplest way to minimize a fairly well-behaved function. The simplex algorithm requires only function evaluations and is a good choice for simple minimization problems. However, because it does not use any gradient evaluations, it may take longer to find the minimum. To demonstrate the minimization function consider the problem of minimizing the Rosenbrock function of $N$ variables:

$$
f(\mathbf{x})=\sum_{i=1}^{N-1} 100\left(x_{i}-x_{i-1}^{2}\right)^{2}+\left(1-x_{i-1}\right)^{2}
$$

The minimum value of this function is 0 which is achieved when $x_{i}=1$. This minimum can be found using the fmin routine as shown in the example below:

```
>>> from scipy.optimize import fmin
>>> def rosen(x): # The Rosenbrock function
        return sum(100.0*(x[1:]-x[:-1]**2.0)**2.0 + (1-x[:-1])**2.0)
>>> x0 = [1.3, 0.7, 0.8, 1.9, 1.2]
>>> xopt = fmin(rosen, x0)
Optimization terminated successfully.
        Current function value: 0.000000
        Iterations: 516
        Function evaluations: 825
```

>>> print xopt
[ 1. 1. 1. 1. 1.]

Another optimization algorithm that needs only function calls to find the minimum is Powell's method available as optimize.fmin_powell.

### 5.2 Broyden-Fletcher-Goldfarb-Shanno algorithm (optimize.fmin_bfgs)

In order to converge more quickly to the solution, this routine uses the gradient of the objective function. If the gradient is not given by the user, then it is estimated using first-differences. The Broyden-Fletcher-Goldfarb-Shanno (BFGS) method requires fewer function calls than the simplex algorithm but unless the gradient is provided by the user, the speed savings won't be significant.

To demonstrate this algorithm, the Rosenbrock function is again used. The gradient of the Rosenbrock function is the vector:

$$
\begin{aligned}
\frac{\partial f}{\partial x_{j}} & =\sum_{i=1}^{N} 200\left(x_{i}-x_{i-1}^{2}\right)\left(\delta_{i, j}-2 x_{i-1} \delta_{i-1, j}\right)-2\left(1-x_{i-1}\right) \delta_{i-1, j} \\
& =200\left(x_{j}-x_{j-1}^{2}\right)-400 x_{j}\left(x_{j+1}-x_{j}^{2}\right)-2\left(1-x_{j}\right)
\end{aligned}
$$

This expression is valid for the interior derivatives. Special cases are

$$
\begin{aligned}
\frac{\partial f}{\partial x_{0}} & =-400 x_{0}\left(x_{1}-x_{0}^{2}\right)-2\left(1-x_{0}\right), \\
\frac{\partial f}{\partial x_{N-1}} & =200\left(x_{N-1}-x_{N-2}^{2}\right) .
\end{aligned}
$$

A Python function which computes this gradient is constructed by the code-segment:

```
>>> def rosen_der(x):
    xm = x[1:-1]
    xm_m1 = x[:-2]
    xm_p1 = x[2:]
    der = zeros(x.shape,x.typecode())
    der[1:-1] = 200*(xm-xm_m1**2) - 400*(xm_p1 - xm**2)*xm - 2*(1-xm)
    der[0] = -400*x[0]*(x[1]-x[0]**2) - 2*(1-x[0])
    der[-1] = 200*(x[-1]-x[-2]**2)
    return der
```

The calling signature for the BFGS minimization algorithm is similar to fmin with the addition of the fprime argument. An example usage of fmin_bfgs is shown in the following example which minimizes the Rosenbrock function.

```
>>> from scipy.optimize import fmin_bfgs
>>> x0 = [1.3, 0.7, 0.8, 1.9, 1.2]
>>> xopt = fmin_bfgs(rosen, x0, fprime=rosen_der)
Optimization terminated successfully.
    Current function value: 0.000000
    Iterations: 109
    Function evaluations: 262
    Gradient evaluations: 110
>>> print xopt
[ 1. 1. 1. 1. 1.]
```


### 5.3 Newton-Conjugate-Gradient (optimize.fmin_ncg)

The method which requires the fewest function calls and is therefore often the fastest method to minimize functions of many variables is fmin_ncg. This method is a modified Newton's method and uses a conjugate gradient algorithm to (approximately) invert the local Hessian. Newton's method is based on fitting the function locally to a quadratic form:

$$
f(\mathbf{x}) \approx f\left(\mathbf{x}_{0}\right)+\nabla f\left(\mathbf{x}_{0}\right) \cdot\left(\mathbf{x}-\mathbf{x}_{0}\right)+\frac{1}{2}\left(\mathbf{x}-\mathbf{x}_{0}\right)^{T} \mathbf{H}\left(\mathbf{x}_{0}\right)\left(\mathbf{x}-\mathbf{x}_{0}\right)
$$

where $\mathbf{H}\left(\mathbf{x}_{0}\right)$ is a matrix of second-derivatives (the Hessian). If the Hessian is positive definite then the local minimum of this function can be found by setting the gradient of the quadratic form to zero, resulting in

$$
\mathbf{x}_{\mathrm{opt}}=\mathbf{x}_{0}-\mathbf{H}^{-1} \nabla f
$$

The inverse of the Hessian is evaluted using the conjugate-gradient method. An example of employing this method to minimizing the Rosenbrock function is given below. To take full advantage of the NewtonCG method, a function which computes the Hessian must be provided. The Hessian matrix itself does not need to be constructed, only a vector which is the product of the Hessian with an arbitrary vector needs to be available to the minimization routine. As a result, the user can provide either a function to compute the Hessian matrix, or a function to compute the product of the Hessian with an arbitrary vector.

### 5.3.1 Full Hessian example:

The Hessian of the Rosenbrock function is

$$
\begin{aligned}
H_{i j}=\frac{\partial^{2} f}{\partial x_{i} \partial x_{j}} & =200\left(\delta_{i, j}-2 x_{i-1} \delta_{i-1, j}\right)-400 x_{i}\left(\delta_{i+1, j}-2 x_{i} \delta_{i, j}\right)-400 \delta_{i, j}\left(x_{i+1}-x_{i}^{2}\right)+2 \delta_{i, j} \\
& =\left(202+1200 x_{i}^{2}-400 x_{i+1}\right) \delta_{i, j}-400 x_{i} \delta_{i+1, j}-400 x_{i-1} \delta_{i-1, j}
\end{aligned}
$$

if $i, j \in[1, N-2]$ with $i, j \in[0, N-1]$ defining the $N \times N$ matrix. Other non-zero entries of the matrix are

$$
\begin{aligned}
\frac{\partial^{2} f}{\partial x_{0}^{2}} & =1200 x_{0}^{2}-400 x_{1}+2, \\
\frac{\partial^{2} f}{\partial x_{0} \partial x_{1}}=\frac{\partial^{2} f}{\partial x_{1} \partial x_{0}} & =-400 x_{0}, \\
\frac{\partial^{2} f}{\partial x_{N-1} \partial x_{N-2}}=\frac{\partial^{2} f}{\partial x_{N-2} \partial x_{N-1}} & =-400 x_{N-2}, \\
\frac{\partial^{2} f}{\partial x_{N-1}^{2}} & =200 .
\end{aligned}
$$

For example, the Hessian when $N=5$ is
$\mathbf{H}=\left[\begin{array}{ccccc}1200 x_{0}^{2}-400 x_{1}+2 & -400 x_{0} & 0 & 0 & 0 \\ -400 x_{0} & 202+1200 x_{1}^{2}-400 x_{2} & -400 x_{1} & 0 & 0 \\ 0 & -400 x_{1} & 202+1200 x_{2}^{2}-400 x_{3} & -400 x_{2} & 0 \\ 0 & & -400 x_{2} & 202+1200 x_{3}^{2}-400 x_{4} & -400 x_{3} \\ 0 & 0 & 0 & -400 x_{3} & 200\end{array}\right]$.
The code which computes this Hessian along with the code to minimize the function using fmin_ncg is shown in the following example:

```
>>> from scipy.optimize import fmin_ncg
>>> def rosen_hess(x):
    x = asarray(x)
    H = diag(-400*x[:-1],1) - diag(400*x[:-1],-1)
    diagonal = zeros(len(x),x.typecode())
    diagonal[0] = 1200*x[0]-400*x[1]+2
    diagonal[-1] = 200
    diagonal[1:-1] = 202 + 1200*x[1:-1]**2 - 400*x[2:]
    H = H + diag(diagonal)
    return H
>>> x0 = [1.3, 0.7, 0.8, 1.9, 1.2]
>>> xopt = fmin_ncg(rosen, x0, rosen_der, fhess=rosen_hess)
Optimization terminated successfully.
    Current function value: 0.000000
    Iterations: 19
    Function evaluations: 40
    Gradient evaluations: 19
    Hessian evaluations: 19
>>> print xopt
[ 0.9999 0.9999 0.9998
```


### 5.3.2 Hessian product example:

For larger minimization problems, storing the entire Hessian matrix can consume considerable time and memory. The Newton-CG algorithm only needs the product of the Hessian times an arbitrary vector. As a
result, the user can supply code to compute this product rather than the full Hessian by setting the fhess_p keyword to the desired function. The fhess_p function should take the minimization vector as the first argument and the arbitrary vector as the second argument. Any extra arguments passed to the function to be minimized will also be passed to this function. If possible, using Newton-CG with the hessian product option is probably the fastest way to minimize the function.

In this case, the product of the Rosenbrock Hessian with an arbitrary vector is not difficult to compute. If $\mathbf{p}$ is the arbitrary vector, then $\mathbf{H}(\mathbf{x}) \mathbf{p}$ has elements:

$$
\mathbf{H}(\mathbf{x}) \mathbf{p}=\left[\begin{array}{c}
\left(1200 x_{0}^{2}-400 x_{1}+2\right) p_{0}-400 x_{0} p_{1} \\
\vdots \\
-400 x_{i-1} p_{i-1}+\left(202+1200 x_{i}^{2}-400 x_{i+1}\right) p_{i}-400 x_{i} p_{i+1} \\
\vdots \\
-400 x_{N-2} p_{N-2}+200 p_{N-1}
\end{array}\right]
$$

Code which makes use of the fhess_p keyword to minimize the Rosenbrock function using fmin_ncg follows:

```
>>> from scipy.optimize import fmin_ncg
>>> def rosen_hess_p(x,p):
    x = asarray(x)
    Hp = zeros(len(x),x.typecode())
    Hp[0] = (1200*x[0]**2 - 400*x[1] + 2)*p[0] - 400*x[0]*p[1]
    Hp[1:-1] = -400*x[:-2]*p[:-2]+(202+1200*x[1:-1]**2-400*x[2:])*p[1:-1] \
                -400*x[1:-1]*p[2:]
    Hp[-1] = -400*x[-2]*p[-2] + 200*p[-1]
    return Hp
>> x0 = [1.3, 0.7, 0.8, 1.9, 1.2]
>>> xopt = fmin_ncg(rosen, x0, rosen_der, fhess_p=rosen_hess_p)
Optimization terminated successfully.
        Current function value: 0.000000
        Iterations: 20
        Function evaluations: 42
        Gradient evaluations: 20
        Hessian evaluations: 44
>>> print xopt
[ 1. 1. 1. 0.9999 0.9999]
```


### 5.4 Least-square fitting (minimize.leastsq)

All of the previously-explained minimization procedures can be used to solve a least-squares problem provided the appropriate objective function is constructed. For example, suppose it is desired to fit a set of data $\left\{\mathbf{x}_{i}, \mathbf{y}_{i}\right\}$ to a known model, $\mathbf{y}=\mathbf{f}(\mathbf{x}, \mathbf{p})$ where $\mathbf{p}$ is a vector of parameters for the model that need to be found. A common method for determining which parameter vector gives the best fit to the data is to minimize the sum of squares of the residuals. The residual is usually defined for each observed data-point as

$$
e_{i}\left(\mathbf{p}, \mathbf{y}_{i}, \mathbf{x}_{i}\right)=\left\|\mathbf{y}_{i}-\mathbf{f}\left(\mathbf{x}_{i}, \mathbf{p}\right)\right\|
$$

An objective function to pass to any of the previous minization algorithms to obtain a least-squares fit is.

$$
J(\mathbf{p})=\sum_{i=0}^{N-1} e_{i}^{2}(\mathbf{p})
$$

The leastsq algorithm performs this squaring and summing of the residuals automatically. It takes as an input argument the vector function $\mathbf{e}(\mathbf{p})$ and returns the value of $\mathbf{p}$ which minimizes $J(\mathbf{p})=\mathbf{e}^{T} \mathbf{e}$
directly. The user is also encouraged to provide the Jacobian matrix of the function (with derivatives down the columns or across the rows). If the Jacobian is not provided, it is estimated.

An example should clarify the usage. Suppose it is believed some measured data follow a sinusoidal pattern

$$
y_{i}=A \sin \left(2 \pi k x_{i}+\theta\right)
$$

where the parameters $A, k$, and $\theta$ are unknown. The residual vector is

$$
e_{i}=\left|y_{i}-A \sin \left(2 \pi k x_{i}+\theta\right)\right| .
$$

By defining a function to compute the residuals and (selecting an appropriate starting position), the leastsquares fit routine can be used to find the best-fit parameters $\hat{A}, \hat{k}, \hat{\theta}$. This is shown in the following example and a plot of the results is shown in Figure 1.

```
>>> x = arange(0,6e-2,6e-2/30)
>>> A,k,theta = 10, 1.0/3e-2, pi/6
>> y_true = A*sin( }2*\mathrm{ pi*k*x+theta)
>>> y_meas = y_true + 2*randn(len(x))
>>> def residuals(p, y, x):
    A,k,theta = p
    err = y-A*sin(2*pi*k*x+theta)
    return err
>>> def peval(x, p):
    return p[0]*sin(2*pi*p[1]*x+p[2])
>>> p0 = [8, 1/2.3e-2, pi/3]
>>> print array(p0)
[ 8. 43.4783 1.0472]
>>> from optimize import leastsq
>>> plsq = leastsq(residuals, p0, args=(y_meas, x))
>>> print plsq[0]
[ 10.9437 33.3605 0.5834]
>>> print array([A, k, theta])
[ 10. 33.3333 0.5236]
>>> from xplt import * # Only on X-windows systems
>>> plot(x,peval(x,plsq[0]),x,y_meas,'o',x,y_true)
>>> title('Least-squares fit to noisy data')
>>> legend(['Fit', 'Noisy', 'True'])
>>> gist.eps('leastsqfit') # Make epsi file.
```


### 5.5 Scalar function minimizers

Often only the minimum of a scalar function is needed (a scalar function is one that takes a scalar as input and returns a scalar output). In these circumstances, other optimization techniques have been developed that can work faster.

### 5.5.1 Unconstrained minimization (optimize.brent)

There are actually two methods that can be used to minimize a scalar function (brent and golden), but golden is included only for academic purposes and should rarely be used. The brent method uses Brent's


Figure 1: Least-square fitting to noisy data using scipy.optimize.leastsq
algorithm for locating a minimum. Optimally a bracket should be given which contains the minimum desired. A bracket is a triple $(a, b, c)$ such that $f(a)>f(b)<f(c)$ and $a<b<c$. If this is not given, then alternatively two starting points can be chosen and a bracket will be found from these points using a simple marching algorithm. If these two starting points are not provided 0 and 1 will be used (this may not be the right choice for your function and result in an unexpected minimum being returned).

### 5.5.2 Bounded minimization (optimize.fminbound)

Thus far all of the minimization routines described have been unconstrained minimization routines. Very often, however, there are constraints that can be placed on the solution space before minimization occurs. The fminbound function is an example of a constrained minimization procedure that provides a rudimentary interval constraint for scalar functions. The interval constraint allows the minimization to occur only between two fixed endpoints.

For example, to find the minimum of $J_{1}(x)$ near $x=5$, fminbound can be called using the interval $[4,7]$ as a constraint. The result is $x_{\min }=5.3314$ :

```
>>> from scipy.special import j1
>>> from scipy.optimize import fminbound
>>> xmin = fminbound(j1, 4, 7)
>>> print xmin
5.33144184241
```


### 5.6 Root finding

### 5.6.1 Sets of equations

To find the roots of a polynomial, the command roots is useful. To find a root of a set of non-linear equations, the command optimize.fsolve is needed. For example, the following example finds the roots of the single-variable transcendental equation

$$
x+2 \cos (x)=0,
$$

and the set of non-linear equations

$$
\begin{aligned}
x_{0} \cos \left(x_{1}\right) & =4 \\
x_{0} x_{1}-x_{1} & =5
\end{aligned}
$$

The results are $x=-1.0299$ and $x_{0}=6.5041, x_{1}=0.9084$.

```
>>> def func(x):
    return x + 2*\operatorname{cos(x)}
>>> def func2(x):
    out = [x[0]*\operatorname{cos(x[1]) - 4]}
    out.append(x[1]*x[0] - x[1] - 5)
    return out
>>> from optimize import fsolve
>>> x0 = fsolve(func, 0.3)
>>> print x0
-1.02986652932
>>> x02 = fsolve(func2, [1, 1])
>>> print x02
[ 6.5041 0.9084]
```


### 5.6.2 Scalar function root finding

If one has a single-variable equation, there are four different root finder algorithms that can be tried. Each of these root finding algorithms requires the endpoints of an interval where a root is suspected (because the function changes signs). In general brentq is the best choice, but the other methods may be useful in certain circumstances or for academic purposes.

### 5.6.3 Fixed-point solving

A problem closely related to finding the zeros of a function is the problem of finding a fixed-point of a function. A fixed point of a function is the point at which evaluation of the function returns the point: $g(x)=x$. Clearly the fixed point of $g$ is the root of $f(x)=g(x)-x$. Equivalently, the root of $f$ is the fixed_point of $g(x)=f(x)+x$. The routine fixed_point provides a simple iterative method using Aitkens sequence acceleration to estimate the fixed point of $g$ given a starting point.

## 6 Interpolation (interpolate)

There are two general interpolation facilities available in SciPy . The first facility is an interpolation class which performs linear 1-dimensional interpolation. The second facility is based on the FORTRAN library FITPACK and provides functions for 1- and 2-dimensional (smoothed) cubic-spline interpolation.

### 6.1 Linear 1-d interpolation (interpolate.linear_1d)

The linear_1d class in scipy.interpolate is a convenient method to create a function based on fixed data points which can be evaluated anywhere within the domain defined by the given data using linear interpolation. An instance of this class is created by passing the 1-d vectors comprising the data. The instance of this class defines a _call__ method and can therefore by treated like a function which interpolates between known data values to obtain unknown values (it even has a docstring for help). Behavior at the boundary can be specified at instantiation time. The following example demonstrates it's use.

```
>>> x = arange(0,10)
>>> y = exp(-x/3.0)
>>> f = interpolate.linear_1d(x,y)
>>> help(f)
Instance of class: linear_1d
    <name>(x_new)
Find linearly interpolated y_new = <name>(x_new).
Inputs:
    x_new -- New independent variables.
Outputs:
    y_new -- Linearly interpolated values corresponding to x_new.
>>> xnew = arange(0,9,0.1)
>>> xplt.plot(x,y,'x',xnew,f(xnew),'.')
```

Figure shows the result:

### 6.2 Spline interpolation in 1-d (interpolate.splXXX)

Spline interpolation requires two essential steps: (1) a spline representation of the curve is computed, and (2) the spline is evaluated at the desired points. In order to find the spline representation, there are two


Figure 2: One-dimensional interpolation using the class interpolate.linear_1d.
different was to represent a curve and obtain (smoothing) spline coefficients: directly and parametrically. The direct method finds the spline representation of a curve in a two-dimensional plane using the function interpolate.splrep. The first two arguments are the only ones required, and these provide the $x$ and $y$ components of the curve. The normal output is a 3 -tuple, $(t, c, k)$, containing the knot-points, $t$, the coefficients $c$ and the order $k$ of the spline. The default spline order is cubic, but this can be changed with the input keyword, $k$.

For curves in $N$-dimensional space the function interpolate.splprep allows defining the curve parametrically. For this function only 1 input argument is required. This input is a list of $N$-arrays representing the curve in $N$-dimensional space. The length of each array is the number of curve points, and each array provides one component of the $N$-dimensional data point. The parameter variable is given with the keword argument, $u$, which defaults to an equally-spaced monotonic sequence between 0 and 1 . The default output consists of two objects: a 3-tuple, $(t, c, k)$, containing the spline representation and the parameter variable $u$.

The keyword argument, $s$, is used to specify the amount of smoothing to perform during the spline fit. The default value of $s$ is $s=m-\sqrt{2 m}$ where $m$ is the number of data-points being fit. Therefore, if no smoothing is desired a value of $s=0$ should be passed to the routines.

Once the spline representation of the data has been determined, functions are available for evaluating the spline (interpolate.splev) and its derivatives (interpolate.splev, interpolate.splade) at any point and the integral of the spline between any two points (interpolate.splint). In addition, for cubic splines $(k=3)$ with 8 or more knots, the roots of the spline can be estimated (interpolate.sproot). These functions are demonstrated in the example that follows (see also Figure 3).

```
>>> # Cubic-spline
>>> x = arange(0,2*pi+pi/4,2*pi/8)
>>> y = sin(x)
>>> tck = interpolate.splrep(x,y,s=0)
>>> xnew = arange(0,2*pi,pi/50)
>>> ynew = interpolate.splev(xnew,tck,der=0)
>>> xplt.plot(x,y,'x',xnew,ynew,xnew,sin(xnew),x,y,'b')
>>> xplt.legend(['Linear','Cubic Spline', 'True'],['b-x','m','r'])
>>> xplt.limits(-0.05,6.33,-1.05,1.05)
>>> xplt.title('Cubic-spline interpolation')
>>> xplt.eps('interp_cubic')
>>> # Derivative of spline
>>> yder = interpolate.splev(xnew,tck,der=1)
>>> xplt.plot(xnew,yder,xnew, cos(xnew),'|')
>>> xplt.legend(['Cubic Spline', 'True'])
>>> xplt.limits(-0.05,6.33,-1.05,1.05)
>>> xplt.title('Derivative estimation from spline')
>>> xplt.eps('interp_cubic_der')
>>> # Integral of spline
>>> def integ(x,tck,constant=-1):
>>> x = asarray_1d(x)
>>> out = zeros(x.shape, x.typecode())
>>> for n in xrange(len(out)):
>>> out[n] = interpolate.splint(0,x[n],tck)
>>> out += constant
>>> return out
>>>
>>> yint = integ(xnew,tck)
>>> xplt.plot(xnew,yint,xnew,-cos(xnew),'|')
>>> xplt.legend(['Cubic Spline', 'True'])
```



Figure 3: Examples of using cubic-spline interpolation.

```
>>> xplt.limits(-0.05,6.33,-1.05,1.05)
>>> xplt.title('Integral estimation from spline')
>>> xplt.eps('interp_cubic_int')
>>> # Roots of spline
>>> print interpolate.sproot(tck)
[ 0. 3.1416]
>>> # Parametric spline
>>> t = arange(0,1.1,.1)
>>> x = sin(2*pi*t)
>>> y = cos(2*pi*t)
>>> tck,u = interpolate.splprep([x,y],s=0)
>>> unew = arange(0,1.01,0.01)
>>> out = interpolate.splev(unew,tck)
>>> xplt.plot(x,y,'x',out[0],out[1],sin(2*pi*unew), cos(2*pi*unew),x,y,'b')
>>> xplt.legend(['Linear','Cubic Spline', 'True'],['b-x','m','r'])
>>> xplt.limits(-1.05,1.05,-1.05,1.05)
>>> xplt.title('Spline of parametrically-defined curve')
>>> xplt.eps('interp_cubic_param')
```


### 6.3 Two-dimensionsal spline representation (interpolate.bisplrep)

For (smooth) spline-fitting to a two dimensional surface, the function interpolate.bisplrep is available. This function takes as required inputs the $\mathbf{1 - D}$ arrays $x, y$, and $z$ which represent points on the surface $z=f(x, y)$. The default output is a list $[t x, t y, c, k x, k y]$ whose entries represent respectively, the components of the knot positions, the coefficients of the spline, and the order of the spline in each coordinate. It is convenient to hold this list in a single object, $t c k$, so that it can be passed easily to the function interpolate.bisplev. The keyword, $s$, can be used to change the amount of smoothing performed on the data while determining the appropriate spline. The default value is $s=m-\sqrt{2 m}$ where $m$ is the number of data points in the $x, y$, and $z$ vectors. As a result, if no smoothing is desired, then $s=0$ should be passed to interpolate.bisplrep.

To evaluate the two-dimensional spline and it's partial derivatives (up to the order of the spline), the function interpolate.bisplev is required. This function takes as the first two arguments two 1-D arrays whose cross-product specifies the domain over which to evaluate the spline. The third argument is the tck list returned from interpolate.bisplrep. If desired, the fourth and fifth arguments provide the orders of the partial derivative in the $x$ and $y$ direction respectively.

It is important to note that two dimensional interpolation should not be used to find the spline representation of images. The algorithm used is not amenable to large numbers of input points. The signal processing toolbox contains (soon) more appropriate algorithms for finding the spline representation of an image. The two dimensional interpolation commands are intended for use when interpolating a two dimensional function as shown in the example that follows (See also Figure 4). This example uses the mgrid command in SciPy which is useful for defining a "mesh-grid" in many dimensions. (See also the ogrid command if the full-mesh is not needed). The number of output arguments and the number of dimensions of each argument is determined by the number of indexing objects passed in mgrid[].

```
>>> # Define function over sparse 20x20 grid
>>> x,y = mgrid[-1:1:20j,-1:1:20j]
>>> z = (x+y)*exp(-6.0*(x*x+y*y))
>>> xplt.surf(z,x,y,shade=1,palette='rainbow')
>>> xplt.title3("Sparsely sampled function.")
>>> xplt.eps("2d_func")
>>> # Interpolate function over new 70x70 grid
>>> xnew,ynew = mgrid[-1:1:70j,-1:1:70j]
>>> tck = interpolate.bisplrep(x,y,z,s=0)
```

```
>>> znew = interpolate.bisplev(xnew[:,0],ynew[0,:],tck)
```

>>> xplt.surf(znew, xnew, ynew,shade=1,palette='rainbow')
>>> xplt.title3("Interpolated function.")
>>> xplt.eps("2d_interp")

# Interpolated function. 

Sparsely sampled function.



Figure 4: Example of two-dimensional spline interpolation.

## 7 Signal Processing (signal)

The signal processing toolbox currently contains some filtering functions, a limited set of filter design tools, and a few B-spline interpolation algorithms for one- and two-dimensional data. While the B-spline algorithms could technically be placed under the interpolation category, they are included here because they only work with equally-spaced data and make heavy use of filter-theory and transfer-function formalism to provide a fast B-spline transform. To understand this section you will need to understand that a signal in SciPy is an array of real or complex numbers.

### 7.1 B-splines

A B-spline is an approximation of a continuous function over a finite-domain in terms of B-spline coefficients and knot points. If the knot-points are equally spaced with spacing $\Delta x$, then the B -spline approximation to a 1-dimensional function is the finite-basis expansion.

$$
y(x) \approx \sum_{j} c_{j} \beta^{o}\left(\frac{x}{\Delta x}-j\right)
$$

In two dimensions with knot-spacing $\Delta x$ and $\Delta y$, the function representation is

$$
z(x, y) \approx \sum_{j} \sum_{k} c_{j k} \beta^{o}\left(\frac{x}{\Delta x}-j\right) \beta^{o}\left(\frac{y}{\Delta y}-k\right) .
$$

In these expressions, $\beta^{o}(\cdot)$ is the space-limited B -spline basis function of order, o. The requirement of equally-spaced knot-points and equally-spaced data points, allows the development of fast (inverse-filtering)
algorithms for determining the coefficients, $c_{j}$, from sample-values, $y_{n}$. Unlike the general spline interpolation algorithms, these algorithms can quickly find the spline coefficients for large images.

The advantage of representing a set of samples via B-spline basis functions is that continuous-domain operators (derivatives, re-sampling, integral, etc.) which assume that the data samples are drawn from an underlying continuous function can be computed with relative ease from the spline coefficients. For example, the second-derivative of a spline is

$$
y^{\prime \prime}(x)=\frac{1}{\Delta x^{2}} \sum_{j} c_{j} \beta^{o \prime \prime}\left(\frac{x}{\Delta x}-j\right)
$$

Using the property of B-splines that

$$
\frac{d^{2} \beta^{o}(w)}{d w^{2}}=\beta^{o-2}(w+1)-2 \beta^{o-2}(w)+\beta^{o-2}(w-1)
$$

it can be seen that

$$
y^{\prime \prime}(x)=\frac{1}{\Delta x^{2}} \sum_{j} c_{j}\left[\beta^{o-2}\left(\frac{x}{\Delta x}-j+1\right)-2 \beta^{o-2}\left(\frac{x}{\Delta x}-j\right)+\beta^{o-2}\left(\frac{x}{\Delta x}-j-1\right)\right] .
$$

If $o=3$, then at the sample points,

$$
\begin{aligned}
\left.\Delta x^{2} y^{\prime}(x)\right|_{x=n \Delta x} & =\sum_{j} c_{j} \delta_{n-j+1}-2 c_{j} \delta_{n-j}+c_{j} \delta_{n-j-1} \\
& =c_{n+1}-2 c_{n}+c_{n-1}
\end{aligned}
$$

Thus, the second-derivative signal can be easily calculated from the spline fit. if desired, smoothing splines can be found to make the second-derivative less sensitive to random-errors.

The savvy reader will have already noticed that the data samples are related to the knot coefficients via a convolution operator, so that simple convolution with the sampled B-spline function recovers the original data from the spline coefficients. The output of convolutions can change depending on how boundaries are handled (this becomes increasingly more important as the number of dimensions in the data-set increases). The algorithms relating to B-splines in the signal-processing sub package assume mirror-symmetric boundary conditions. Thus, spline coefficients are computed based on that assumption, and data-samples can be recovered exactly from the spline coefficients by assuming them to be mirror-symmetric also.

Currently the package provides functions for determining seond- and third-order cubic spline coefficients from equally spaced samples in one- and two-dimensions (signal.qspline1d, signal.qspline2d, signal.cspline1d, signal.cspline2d). The package also supplies a function (signal.bspline) for evaluating the bspline basis function, $\beta^{o}(x)$ for arbitrary order and $x$. For large $o$, the B -spline basis function can be approximated well by a zero-mean Gaussian function with standard-deviation equal to $\sigma_{o}=(o+1) / 12$ :

$$
\beta^{o}(x) \approx \frac{1}{\sqrt{2 \pi \sigma_{o}^{2}}} \exp \left(-\frac{x^{2}}{2 \sigma_{o}}\right) .
$$

A function to compute this Gaussian for arbitrary $x$ and $o$ is also available (signal.gauss_spline). The following code and Figure uses spline-filtering to compute an edge-image (the second-derivative of a smoothed spline) of Lena's face which is an array returned by the command lena(). The command signal.sepfir2d was used to apply a separable two-dimensional FIR filter with mirror-symmetric boundary conditions to the spline coefficients. This function is ideally suited for reconstructing samples from spline coefficients and is faster than signal.convolve2d which convolves arbitrary two-dimensional filters and allows for choosing mirror-symmetric boundary conditions.

```
>>> image = lena().astype(Float32)
>>> derfilt = array([1.0,-2,1.0],Float32)
>>> ck = signal.cspline2d(image,8.0)
>>> deriv = signal.sepfir2d(ck, derfilt, [1]) + \
```

```
>>> signal.sepfir2d(ck, [1], derfilt)
>>>
>>> ## Alternatively we could have done:
>>> ## laplacian = array([[0,1,0],[1,-4,1],[0,1,0]],Float32)
>>> ## deriv2 = signal.convolve2d(ck,laplacian,mode='same',boundary='symm')
>>>
>>> xplt.imagesc(image[::-1]) # flip image so it looks right-side up.
>>> xplt.title('Original image')
>>> xplt.eps('lena_image')
>>> xplt.imagesc(deriv[::-1])
>>> xplt.title('Output of spline edge filter')
>>> xplt.eps('lena_edge')
```

Original image


Output of spline edge filter


Figure 5: Example of using smoothing splines to filter images.

### 7.2 Filtering

Filtering is a generic name for any system that modifies an input signal in some way. In SciPy a signal can be thought of as a Numeric array. There are different kinds of filters for different kinds of operations. There are two broad kinds of filtering operations: linear and non-linear. Linear filters can always be reduced to multiplication of the flattened Numeric array by an appropriate matrix resulting in another flattened Numeric array. Of course, this is not usually the best way to compute the filter as the matrices and vectors involved may be huge. For example filtering a $512 \times 512$ image with this method would require multiplication of a $512^{2} x 512^{2}$ matrix with a $512^{2}$ vector. Just trying to store the $512^{2} \times 512^{2}$ matrix using a standard Numeric array would require $68,719,476,736$ elements. At 4 bytes per element this would require 256 GB of memory. In most applications most of the elements of this matrix are zero and a different method for computing the output of the filter is employed.

### 7.2.1 Convolution/Correlation

Many linear filters also have the property of shift-invariance. This means that the filtering operation is the same at different locations in the signal and it implies that the filtering matrix can be constructed from knowledge of one row (or column) of the matrix alone. In this case, the matrix multiplication can be accomplished using Fourier transforms.

Let $x[n]$ define a one-dimensional signal indexed by the integer $n$. Full convolution of two one-dimensional signals can be expressed as

$$
y[n]=\sum_{k=-\infty}^{\infty} x[k] h[n-k] .
$$

This equation can only be implemented directly if we limit the sequences to finite support sequences that can be stored in a computer, choose $n=0$ to be the starting point of both sequences, let $K+1$ be that value for which $y[n]=0$ for all $n>K+1$ and $M+1$ be that value for which $x[n]=0$ for all $n>M+1$, then the discrete convolution expression is

$$
y[n]=\sum_{k=\max (n-M, 0)}^{\min (n, K)} x[k] h[n-k] .
$$

For convenience assume $K \geq M$. Then, more explicitly the output of this operation is

$$
\begin{aligned}
y[0] & =x[0] h[0] \\
y[1] & =x[0] h[1]+x[1] h[0] \\
y[2] & =x[0] h[2]+x[1] h[1]+x[2] h[0] \\
\vdots & \vdots \\
y[M] & \vdots \\
y[M+1] & =x[1] h[M]+x[2] h[M-1]+\cdots+x[M+1] h[0] \\
\vdots & \vdots \\
y[K] & =x[K-M] h[M]+\cdots+x[K] h[0] \\
y[K+1] & =x[K+1-M] h[M]+\cdots+x[K] h[1] \\
\vdots & \vdots \\
y[K+M-1] & =x[K-1] h[M]+x[K] h[M-1] \\
y[K+M] & =x[K] h[M] .
\end{aligned}
$$

Thus, the full discrete convolution of two finite sequences of lengths $K+1$ and $M+1$ respectively results in a finite sequence of length $K+M+1=(K+1)+(M+1)-1$.

One dimensional convolution is implemented in SciPy with the function signal. convolve. This function takes as inputs the signals $x, h$, and an optional flag and returns the signal $y$. The optional flag allows for specification of which part of the output signal to return. The default value of 'full' returns the entire signal. If the flag has a value of 'same' then only the middle $K$ values are returned starting at $y\left[\left\lfloor\frac{M-1}{2}\right\rfloor\right]$ so that the output has the same length as the largest input. If the flag has a value of 'valid' then only the middle $K-M+1=(K+1)-(M+1)+1$ output values are returned where $z$ depends on all of the values of the smallest input from $h[0]$ to $h[M]$. In other words only the values $y[M]$ to $y[K]$ inclusive are returned.

This same function signal. convolve can actually take $N$-dimensional arrays as inputs and will return the $N$-dimensional convolution of the two arrays. The same input flags are available for that case as well.

Correlation is very similar to convolution except for the minus sign becomes a plus sign. Thus

$$
w[n]=\sum_{k=-\infty}^{\infty} y[k] x[n+k]
$$

is the (cross) correlation of the signals $y$ and $x$. For finite-length signals with $y[n]=0$ outside of the range $[0, K]$ and $x[n]=0$ outside of the range $[0, M]$, the summation can simplify to

$$
w[n]=\sum_{k=\max (0,-n)}^{\min (K, M-n)} y[k] x[n+k] .
$$

Assuming again that $K \geq M$ this is

$$
\begin{aligned}
w[-K] & =y[K] x[0] \\
w[-K+1] & =y[K-1] x[0]+y[K] x[1] \\
\vdots & \vdots \\
w[M-K] & =y[K-M] x[0]+y[K-M+1] x[1]+\cdots+y[K] x[M] \\
w[M-K+1] & =y[K-M-1] x[0]+\cdots+y[K-1] x[M] \\
\vdots & \vdots \\
w[-1] & =y[1] x[0]+y[2] x[1]+\cdots+y[M+1] x[M] \\
w[0] & =y[0] x[0]+y[1] x[1]+\cdots+y[M] x[M] \\
w[1] & =y[0] x[1]+y[1] x[2]+\cdots+y[M-1] x[M] \\
w[2] & =y[0] x[2]+y[1] x[3]+\cdots+y[M-2] x[M] \\
\vdots & \vdots \\
w[M-1] & =y[0] x[M-1]+y[1] x[M] \\
w[M] & =y[0] x[M] .
\end{aligned}
$$

The SciPy function signal. correlate implements this operation. Equivalent flags are available for this operation to return the full $K+M+1$ length sequence ('full') or a sequence with the same size as the largest sequence starting at $w\left[-K+\left\lfloor\frac{M-1}{2}\right\rfloor\right]$ ('same') or a sequence where the values depend on all the values of the smallest sequence ('valid'). This final option returns the $K-M+1$ values $w[M-K]$ to $w[0]$ inclusive.

The function signal.correlate can also take arbitrary $N$-dimensional arrays as input and return the $N$-dimensional convolution of the two arrays on output.

When $N=2$, signal.correlate and/or signal.convolve can be used to construct arbitrary image filters to perform actions such as blurring, enhancing, and edge-detection for an image.

Convolution is mainly used for filtering when one of the signals is much smaller than the other $(K \gg M)$, otherwise linear filtering is more easily accomplished in the frequency domain (see Fourier Transforms).

### 7.2.2 Difference-equation filtering

A general class of linear one-dimensional filters (that includes convolution filters) are filters described by the difference equation

$$
\sum_{k=0}^{N} a_{k} y[n-k]=\sum_{k=0}^{M} b_{k} x[n-k]
$$

where $x[n]$ is the input sequence and $y[n]$ is the output sequence. If we assume initial rest so that $y[n]=0$ for $n<0$, then this kind of filter can be implemented using convolution. However, the convolution filter sequence $h[n]$ could be infinite if $a_{k} \neq 0$ for $k \geq 1$. In addition, this general class of linear filter allows initial conditions to be placed on $y[n]$ for $n<0$ resulting in a filter that cannot be expressed using convolution.

The difference equation filter can be thought of as finding $y[n]$ recursively in terms of it's previous values

$$
a_{0} y[n]=-a_{1} y[n-1]-\cdots-a_{N} y[n-N]+\cdots+b_{0} x[n]+\cdots+b_{M} x[n-M] .
$$

Often $a_{0}=1$ is chosen for normalization. The implementation in SciPy of this general difference equation filter is a little more complicated then would be implied by the previous equation. It is implemented so that
only one signal needs to be delayed. The actual implementation equations are (assuming $a_{0}=1$ ).

$$
\begin{aligned}
y[n] & =b_{0} x[n]+z_{0}[n-1] \\
z_{0}[n] & =b_{1} x[n]+z_{1}[n-1]-a_{1} y[n] \\
z_{1}[n] & =b_{2} x[n]+z_{2}[n-1]-a_{2} y[n] \\
\vdots & \vdots \\
z_{K-2}[n] & =b_{K-1} x[n]+z_{K-1}[n-1]-a_{K-1} y[n] \\
z_{K-1}[n] & =b_{K} x[n]-a_{K} y[n],
\end{aligned}
$$

where $K=\max (N, M)$. Note that $b_{K}=0$ if $K>M$ and $a_{K}=0$ if $K>N$. In this way, the output at time $n$ depends only on the input at time $n$ and the value of $z_{0}$ at the previous time. This can always be calculated as long as the $K$ values $z_{0}[n-1] \ldots z_{K-1}[n-1]$ are computed and stored at each time step.

The difference-equation filter is called using the command signal.lfilter in SciPy. This command takes as inputs the vector $b$, the vector, $a$, a signal $x$ and returns the vector $y$ (the same length as $x$ ) computed using the equation given above. If $x$ is $N$-dimensional, then the filter is computed along the axis provided. If, desired, initial conditions providing the values of $z_{0}[-1]$ to $z_{K-1}[-1]$ can be provided or else it will be assumed that they are all zero. If initial conditions are provided, then the final conditions on the intermediate variables are also returned. These could be used, for example, to restart the calculation in the same state.

Sometimes it is more convenient to express the initial conditions in terms of the signals $x[n]$ and $y[n]$. In other words, perhaps you have the values of $x[-M]$ to $x[-1]$ and the values of $y[-N]$ to $y[-1]$ and would like to determine what values of $z_{m}[-1]$ should be delivered as initial conditions to the difference-equation filter. It is not difficult to show that for $0 \leq m<K$,

$$
z_{m}[n]=\sum_{p=0}^{K-m-1}\left(b_{m+p+1} x[n-p]-a_{m+p+1} y[n-p]\right) .
$$

Using this formula we can find the intial condition vector $z_{0}[-1]$ to $z_{K-1}[-1]$ given initial conditions on $y$ (and $x$ ). The command signal.lfiltic performs this function.

### 7.2.3 Other filters

The signal processing package affords many more filters as well.
Median Filter A median filter is commonly applied when noise is markedly non-Gaussian or when it is desired to preserve edges. The median filter works by sorting all of the array pixel values in a rectangular region surrounding the point of interest. The sample median of this list of neighborhood pixel values is used as the value for the output array. The sample median is the middle array value in a sorted list of neighborhood values. If there are an even number of elements in the neighborhood, then the average of the middle two values is used as the median. A general purpose median filter that works on N -dimensional arrays is signal.medfilt. A specialized version that works only for two-dimensional arrays is available as signal.medfilt2d.

Order Filter A median filter is a specific example of a more general class of filters called order filters. To compute the output at a particular pixel, all order filters use the array values in a region surrounding that pixel. These array values are sorted and then one of them is selected as the output value. For the median filter, the sample median of the list of array values is used as the output. A general order filter allows the user to select which of the sorted values will be used as the output. So, for example one could choose to pick the maximum in the list or the minimum. The order filter takes an additional argument besides the input array and the region mask that specifies which of the elements in the sorted list of neighbor array values should be used as the output. The command to perform an order filter is signal.order_filter.

Wiener filter The Wiener filter is a simple deblurring filter for denoising images. This is not the Wiener filter commonly described in image reconstruction problems but instead it is a simple, local-mean filter. Let $x$ be the input signal, then the output is

$$
y=\left\{\begin{array}{cl}
\frac{\sigma^{2}}{\sigma_{x}^{2}} m_{x}+\left(1-\frac{\sigma^{2}}{\sigma_{x}^{2}}\right) x & \sigma_{x}^{2} \geq \sigma^{2} \\
m_{x} & \sigma_{x}^{2}<\sigma^{2}
\end{array}\right.
$$

Where $m_{x}$ is the local estimate of the mean and $\sigma_{x}^{2}$ is the local estimate of the variance. The window for these estimates is an optional input parameter (default is $3 \times 3$ ). The parameter $\sigma^{2}$ is a threshold noise parameter. If $\sigma$ is not given then it is estimated as the average of the local variances.

Hilbert filter The Hilbert transform constructs the complex-valued analytic signal from a real signal. For example if $x=\cos \omega n$ then $y=\operatorname{hilbert}(x)$ would return (except near the edges) $y=\exp (j \omega n)$. In the frequency domain, the hilbert transform performs

$$
Y=X \cdot H
$$

where $H$ is 2 for positive frequencies, 0 for negative frequencies and 1 for zero-frequencies.

## Detrend

### 7.3 Filter design

7.3.1 Finite-impulse response design
7.3.2 Inifinite-impulse response design
7.3.3 Analog filter frequency response
7.3.4 Digital filter frequency response

### 7.4 Linear Time-Invariant Systems

### 7.4.1 LTI Object

7.4.2 Continuous-Time Simulation
7.4.3 Step response
7.4.4 Impulse response

## 8 Input/Output

### 8.1 Binary

8.1.1 Arbitrary binary input and output (fopen)
8.1.2 Read and write Matlab .mat files
8.1.3 Saving workspace

### 8.2 Text-file

8.2.1 Read text-files (read_array)
8.2.2 Write a text-file (write_array)

## 9 Fourier Transforms

### 9.1 One-dimensional

### 9.2 Two-dimensional

### 9.3 N -dimensional

9.4 Shifting

### 9.5 Sample frequencies

### 9.6 Hilbert transform

### 9.7 Tilbert transform

## 10 Linear Algebra

When SciPy is built using the optimized ATLAS LAPACK and BLAS libraries, it has very fast linear algebra capabilities. If you dig deep enough, all of the raw lapack and blas libraries are available for your use for even more speed. In this section, some easier-to-use interfaces to these routines are described.

All of these linear algebra routines expect an object that can be converted into a 2-dimensional array. The output of these routines is also a two-dimensional array. There is a matrix class defined in Numeric that scipy inherits and extends. You can initialize this class with an appropriate Numeric array in order to get objects for which multiplication is matrix-multiplication instead of the default, element-by-element multiplication.

### 10.1 Matrix Class

The matrix class is initialized with the SciPy command mat which is just convenient short-hand for Matrix. Matrix. If you are going to be doing a lot of matrix-math, it is convenient to convert arrays into matrices using this command. One convencience of using the mat command is that you can enter two-dimensional matrices in using MATLAB-like syntax with commas or spaces separating columns and semicolons separting rows as long as the matrix is placed in a string passed to mat.

### 10.2 Basic routines

### 10.2.1 Finding Inverse

The inverse of a matrix $\mathbf{A}$ is the matrix $\mathbf{B}$ such that $\mathbf{A B}=\mathbf{I}$ where $\mathbf{I}$ is the identity matrix consisting of ones down the main diagonal. Usually $\mathbf{B}$ is denoted $\mathbf{B}=\mathbf{A}^{-1}$. In SciPy, the matrix inverse of the Numeric array, A, is obtained using linalg.inv(A), or using A.I if A is a Matrix. For example, let

$$
\mathbf{A}=\left[\begin{array}{lll}
1 & 3 & 5 \\
2 & 5 & 1 \\
2 & 3 & 8
\end{array}\right]
$$

then

$$
\mathbf{A}^{-\mathbf{1}}=\frac{\mathbf{1}}{\mathbf{2 5}}\left[\begin{array}{ccc}
-37 & 9 & 22 \\
14 & 2 & -9 \\
4 & -3 & 1
\end{array}\right]=\left[\begin{array}{ccc}
-1.48 & 0.36 & 0.88 \\
0.56 & 0.08 & -0.36 \\
0.16 & -0.12 & 0.04
\end{array}\right]
$$

The following example demonstrates this computation in SciPy

```
>>> A = mat('[1 3 5; 2 5 1; 2 3 8]')
>>> A
Matrix([[1, 3, 5],
    [2, 5, 1],
    [2, 3, 8]])
>>> A.I
Matrix([[-1.48, 0.36, 0.88],
    [ 0.56, 0.08, -0.36],
    [ 0.16, -0.12, 0.04]])
>>> linalg.inv(A)
array([[-1.48, 0.36, 0.88],
    [ 0.56, 0.08, -0.36],
    [ 0.16, -0.12, 0.04]])
```


### 10.2.2 Solving linear system

Solving linear systems of equations is straightforward using the scipy command linalg.solve. This command expects an input matrix and a right-hand-side vector. The solution vector is then computed. An option for entering a symmetrix matrix is offered which can speed up the processing when applicable. As an example, suppose it is desired to solve the following simultaneous equations:

$$
\begin{aligned}
x+3 y+5 z & =10 \\
2 x+5 y+z & =8 \\
2 x+3 y+8 z & =3
\end{aligned}
$$

We could find the solution vector using a matrix inverse:

$$
\left[\begin{array}{l}
x \\
y \\
z
\end{array}\right]=\left[\begin{array}{lll}
1 & 3 & 5 \\
2 & 5 & 1 \\
2 & 3 & 8
\end{array}\right]^{-1}\left[\begin{array}{c}
10 \\
8 \\
3
\end{array}\right]=\frac{1}{25}\left[\begin{array}{c}
-232 \\
129 \\
19
\end{array}\right]=\left[\begin{array}{c}
-9.28 \\
5.16 \\
0.76
\end{array}\right]
$$

However, it is better to use the linalg.solve command which can be faster and more numerically stable. In this case it gives the same answer as shown in the following example:

```
>>> A = mat('[1 3 5; 2 5 1; 2 3 8]')
>>> b = mat('[10;8;3]')
>>> A.I*b
Matrix([[-9.28],
    [ 5.16],
    [ 0.76]])
>>> linalg.solve(A,b)
array([[-9.28],
    [ 5.16],
    [ 0.76]])
```


### 10.2.3 Finding Determinant

The determinant of a square matrix $\mathbf{A}$ is often denoted $|\mathbf{A}|$ and is a quantity often used in linear algebra. Suppose $a_{i j}$ are the elements of the matrix $\mathbf{A}$ and let $M_{i j}=\left|\mathbf{A}_{i j}\right|$ be the determinant of the matrix left by removing the $i^{\text {th }}$ row and $j^{\text {th }}$ column from $\mathbf{A}$. Then for any row $i$,

$$
|\mathbf{A}|=\sum_{j}(-1)^{i+j} a_{i j} M_{i j}
$$

This is a recursive way to define the determinant where the base case is defined by accepting that the determinant of a $1 \times 1$ matrix is the only matrix element. In SciPy the determinant can be calculated with linalg.det. For example, the determinant of

$$
\mathbf{A}=\left[\begin{array}{lll}
1 & 3 & 5 \\
2 & 5 & 1 \\
2 & 3 & 8
\end{array}\right]
$$

is

$$
\begin{aligned}
|\mathbf{A}| & =1\left|\begin{array}{cc}
5 & 1 \\
3 & 8
\end{array}\right|-3\left|\begin{array}{cc}
2 & 1 \\
2 & 8
\end{array}\right|+5\left|\begin{array}{cc}
2 & 5 \\
2 & 3
\end{array}\right| \\
& =1(5 \cdot 8-3 \cdot 1)-3(2 \cdot 8-2 \cdot 1)+5(2 \cdot 3-2 \cdot 5)=-25 .
\end{aligned}
$$

In SciPy this is computed as shown in this example:

```
>>> A = mat('[11 3 5; 2 5 1; 2 3 8]')
>>> linalg.det(A)
-25.000000000000004
```


### 10.2.4 Computing norms

Matrix and vector norms can also be computed with SciPy. A wide range of norm definitions are available using different parameters to the order argument of linalg.norm. This function takes a rank-1 (vectors) or a rank- 2 (matrices) array and an optional order argument (default is 2). Based on these inputs a vector or matrix norm of the requested order is computed.

For vector $\mathbf{x}$, the order parameter can be any real number including inf or -inf. The computed norm is

$$
\|\mathbf{x}\|=\left\{\begin{array}{cl}
\max \left|x_{i}\right| & \begin{array}{l}
\text { ord }=\inf \\
\min \left|x_{i}\right| \\
\text { ord }=-\inf
\end{array} \\
\left(\sum_{i}\left|x_{i}\right|^{\text {ord }}\right)^{1 / \text { ord }} & \text { ord } \mid<\infty
\end{array}\right.
$$

For matrix A the only valid values for norm are $\pm 2, \pm 1, \pm i n f$, and 'fro' (or 'f') Thus,

$$
\|\mathbf{A}\|=\left\{\begin{array}{cc}
\max _{i} \sum_{j}\left|a_{i j}\right| & \text { ord }=\inf \\
\min _{i} \sum_{j}\left|a_{i j}\right| & \text { ord }=-\inf \\
\max _{j} \sum_{i}\left|a_{i j}\right| & \text { ord }=1 \\
\min _{j} \sum_{i}\left|a_{i j}\right| & \text { ord }=-1 \\
\max \sigma_{i} & \text { ord }=2 \\
\min \sigma_{i} & \text { ord }=-2 \\
\sqrt{\operatorname{trace}\left(\mathbf{A}^{H} \mathbf{A}\right)} & \text { ord }=\text { 'fro }
\end{array}\right.
$$

where $\sigma_{i}$ are the singular values of $\mathbf{A}$.

### 10.2.5 Solving linear least-squares problems and pseudo-inverses

Linear least-squares problems occur in many branches of applied mathematics. In this problem a set of linear scaling coefficients is sought that allow a model to fit data. In particular it is assumed that data $y_{i}$ is related to data $\mathbf{x}_{i}$ through a set of coefficients $c_{j}$ and model functions $f_{j}\left(\mathbf{x}_{i}\right)$ via the model

$$
y_{i}=\sum_{j} c_{j} f_{j}\left(\mathbf{x}_{i}\right)+\epsilon_{i}
$$

where $\epsilon_{i}$ represents uncertainty in the data. The strategy of least squares is to pick the coefficients $c_{j}$ to minimize

$$
J(\mathbf{c})=\sum_{i}\left|y_{i}-\sum_{j} c_{j} f_{j}\left(x_{i}\right)\right|^{2}
$$

Theoretically, a global minimum will occur when

$$
\frac{\partial J}{\partial c_{n}^{*}}=0=\sum_{i}\left(y_{i}-\sum_{j} c_{j} f_{j}\left(x_{i}\right)\right)\left(-f_{n}^{*}\left(x_{i}\right)\right)
$$

or

$$
\begin{aligned}
\sum_{j} c_{j} \sum_{i} f_{j}\left(x_{i}\right) f_{n}^{*}\left(x_{i}\right) & =\sum_{i} y_{i} f_{n}^{*}\left(x_{i}\right) \\
\mathbf{A}^{H} \mathbf{A c} & =\mathbf{A}^{H} \mathbf{y}
\end{aligned}
$$

where

$$
\{\mathbf{A}\}_{i j}=f_{j}\left(x_{i}\right) .
$$

When $\mathbf{A}^{\mathbf{H}} \mathbf{A}$ is invertible, then

$$
\mathbf{c}=\left(\mathbf{A}^{H} \mathbf{A}\right)^{-1} \mathbf{A}^{H} \mathbf{y}=\mathbf{A}^{\dagger} \mathbf{y}
$$

where $\mathbf{A}^{\dagger}$ is called the pseudo-inverse of $\mathbf{A}$. Notice that using this definition of $\mathbf{A}$ the model can be written

$$
\mathbf{y}=\mathbf{A c}+\boldsymbol{\epsilon}
$$

The command linalg.lstsq will solve the linear least squares problem for $\mathbf{c}$ given $\mathbf{A}$ and $\mathbf{y}$. In addition linalg.pinv or linalg.pinv2 (uses a different method based on singular value decomposition) will find $\mathbf{A}^{\dagger}$ given $\mathbf{A}$.

The following example and figure demonstrate the use of linalg.lstsq and linalg.pinv for solving a data-fitting problem. The data shown below were generated using the model:

$$
y_{i}=c_{1} e^{-x_{i}}+c_{2} x_{i}
$$

where $x_{i}=0.1 i$ for $i=1 \ldots 10, c_{1}=5$, and $c_{2}=4$. Noise is added to $y_{i}$ and the coefficients $c_{1}$ and $c_{2}$ are estimated using linear least squares.

```
c1,c2= 5.0,2.0
i = r_[1:11]
xi = 0.1*i
yi = c1*exp(-xi)+c2*xi
zi = yi + 0.05*max(yi)*randn(len(yi))
A = c_[exp(-xi)[:,NewAxis],xi[:,NewAxis]]
c,resid,rank,sigma = linalg.lstsq(A,zi)
xi2 = r_[0.1:1.0:100j]
yi2 = c[0]*exp(-xi2) + c[1]*xi2
xplt.plot(xi,zi,'x',xi2,yi2)
xplt.limits(0,1.1,3.0,5.5)
xplt.xlabel('x_i')
xplt.title('Data fitting with linalg.lstsq')
xplt.eps('lstsq_fit')
```

Data fitting with linalg.Istsq


### 10.2.6 Generalized inverse

The generalized inverse is calculated using the command linalg.pinv or linalg.pinv2. These two commands differ in how they compute the generalized inverse. The first uses the linalg.lstsq algorithm while the second uses singular value decomposition. Let $\mathbf{A}$ be an $M \times N$ matrix, then if $M>N$ the generalized inverse is

$$
\mathbf{A}^{\dagger}=\left(\mathbf{A}^{H} \mathbf{A}\right)^{-1} \mathbf{A}^{H}
$$

while if $M<N$ matrix the generalized inverse is

$$
\mathbf{A}^{\#}=\mathbf{A}^{H}\left(\mathbf{A} \mathbf{A}^{H}\right)^{-1}
$$

In both cases for $M=N$, then

$$
\mathbf{A}^{\dagger}=\mathbf{A}^{\#}=\mathbf{A}^{-1}
$$

as long as $\mathbf{A}$ is invertible.

### 10.3 Decompositions

In many applications it is useful to decompose a matrix using other representations. There are several decompositions supported SciPy.

### 10.3.1 Eigenvalues and eigenvectors

The eigenvalue-eigenvector problem is one of the most commonly employed linear algebra operations. In one popular form, the eigenvalue-eigenvector problem is to find for some square matrix $\mathbf{A}$ scalars $\lambda$ and corresponding vectors $\mathbf{v}$ such that

$$
\mathbf{A} \mathbf{v}=\lambda \mathbf{v}
$$

For an $N \times N$ matrix, there are $N$ (not necessarily distinct) eigenvalues - roots of the (characteristic) polynomial

$$
|\mathbf{A}-\lambda \mathbf{I}|=0 .
$$

The eigenvectors, $\mathbf{v}$, are also sometimes called right eigenvectors to distinguish them from another set of left eigenvectors that satisfy

$$
\mathbf{v}_{L}^{H} \mathbf{A}=\lambda \mathbf{v}_{L}^{H}
$$

or

$$
\mathbf{A}^{H} \mathbf{v}_{L}=\lambda^{*} \mathbf{v}_{L}
$$

With it's default optional arguments, the command linalg.eig returns $\lambda$ and $\mathbf{v}$. However, it can also return $\mathbf{v}_{L}$ and just $\lambda$ by itself (linalg.eigvals returns just $\lambda$ as well).

In addtion, linalg.eig can also solve the more general eigenvalue problem

$$
\begin{aligned}
\mathbf{A v} & =\lambda \mathbf{B} \mathbf{v} \\
\mathbf{A}^{H} \mathbf{v}_{L} & =\lambda^{*} \mathbf{B}^{H} \mathbf{v}_{L}
\end{aligned}
$$

for square matrices $\mathbf{A}$ and $\mathbf{B}$. The standard eigenvalue problem is an example of the general eigenvalue problem for $\mathbf{B}=\mathbf{I}$. When a generalized eigenvalue problem can be solved, then it provides a decomposition of $\mathbf{A}$ as

$$
\mathbf{A}=\mathbf{B} \mathbf{V} \boldsymbol{\Lambda} \mathbf{V}^{-1}
$$

where $\mathbf{V}$ is the collection of eigenvectors into columns and $\boldsymbol{\Lambda}$ is a diagonal matrix of eigenvalues.
By definition, eigenvectors are only defined up to a constant scale factor. In SciPy, the scaling factor for the eigenvectors is chosen so that $\|\mathbf{v}\|^{2}=\sum_{i} v_{i}^{2}=1$.

As an example, consider finding the eigenvalues and eigenvectors of the matrix

$$
\mathbf{A}=\left[\begin{array}{lll}
1 & 5 & 2 \\
2 & 4 & 1 \\
3 & 6 & 2
\end{array}\right]
$$

The characteristic polynomial is

$$
\begin{aligned}
|\mathbf{A}-\lambda \mathbf{I}|= & (1-\lambda)[(4-\lambda)(2-\lambda)-6]- \\
& 5[2(2-\lambda)-3]+2[12-3(4-\lambda)] \\
= & -\lambda^{3}+7 \lambda^{2}+8 \lambda-3
\end{aligned}
$$

The roots of this polynomial are the eigenvalues of $\mathbf{A}$ :

$$
\begin{aligned}
& \lambda_{1}=7.9579 \\
& \lambda_{2}=-1.2577 \\
& \lambda_{3}=0.2997
\end{aligned}
$$

The eigenvectors corresponding to each eigenvalue can be found using the original equation. The eigenvectors associated with these eigenvalues can then be found.

```
>>> A = mat('[1 5 2; 2 4 1; 3 6 2]')
>>> la,v = linalg.eig(A)
>>> 11,12,13 = la
>>> print l1, 12, 13
(7.95791620491+0j) (-1.25766470568+0j) (0.299748500767+0j)
>>> print v[:,0]
array([-0.5297, -0.4494, -0.7193])
>>> print v[:,1]
[-0.9073 0.2866 0.3076]
>>> print v[:,2]
[ 0.2838 -0.3901 0.8759]
>>> print sum(abs(v**2),axis=0)
[ 1. 1. 1.]
>>> v1 = mat(v[:,0]).T
>>> print max(ravel(abs(A*v1-l1*v1)))
4.4408920985e-16
```


### 10.3.2 Singular value decomposition

Singular Value Decompostion (SVD) can be thought of as an extension of the eigenvalue problem to matrices that are not square. Let $\mathbf{A}$ be an $M \times N$ matrix with $M$ and $N$ arbitrary. The matrices $\mathbf{A}^{H} \mathbf{A}$ and $\mathbf{A A}^{H}$ are square hermitian matrices ${ }^{3}$ of size $N \times N$ and $M \times M$ respectively. It is known that the eigenvalues of square hermitian matrices are real and non-negative. In addtion, there are at most $\min (M, N)$ identical non-zero eigenvalues of $\mathbf{A}^{H} \mathbf{A}$ and $\mathbf{A} \mathbf{A}^{H}$. Define these positive eigenvalues as $\sigma_{i}^{2}$. The square-root of these are called singular values of $\mathbf{A}$. The eigenvectors of $\mathbf{A}^{H} \mathbf{A}$ are collected by columns into an $N \times N$ unitary ${ }^{4}$ matrix $\mathbf{V}$ while the eigenvectors of $\mathbf{A} \mathbf{A}^{H}$ are collected by columns in the unitary matrix $\mathbf{U}$, the singular values are collected in an $M \times N$ zero matrix $\boldsymbol{\Sigma}$ with main diagonal entries set to the singular values. Then

$$
\mathbf{A}=\mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^{H}
$$

is the singular-value decomposition of A. Every matrix has a singular value decomposition. Sometimes, the singular values are called the spectrum of $\mathbf{A}$. The command linalg.svd will return $\mathbf{U}, \mathbf{V}^{H}$, and $\sigma_{i}$ as an array of the singular values. To obtain the matrix $\boldsymbol{\Sigma}$ use linalg.diagsvd. The following example illustrates the use of linalg.svd.

```
>>> A = mat('[1 3 2; 1 2 3]')
>>> M,N = A.shape
>>> U,s,Vh = linalg.svd(A)
>>> Sig = mat(diagsvd(s,M,N))
>>> U, Vh = mat(U), mat(Vh)
>>> print U
Matrix([[-0.7071, -0.7071],
    [-0.7071, 0.7071]])
>>> print Sig
Matrix([[ 5.1962, 0. , 0. ],
    [0. , 1. , 0. ]])
>>> print Vh
Matrix([[-0.2722, -0.6804, -0.6804],
    [-0. , -0.7071, 0.7071],
    [-0.9623, 0.1925, 0.1925]])
```

```
>>> print A
Matrix([[1, 3, 2],
    [1, 2, 3]])
>>> print U*Sig*Vh
Matrix([[ 1., 3., 2.],
    [ 1., 2., 3.]])
```


### 10.3.3 LU decomposition

The LU decompostion finds a representation for the $M \times N$ matrix $\mathbf{A}$ as

$$
\mathbf{A}=\mathbf{P L U}
$$

where $\mathbf{P}$ is an $M \times M$ permutation matrix (a permutation of the rows of the identity matrix), $\mathbf{L}$ is in $M \times K$ lower triangular or trapezoidal matrix $(K=\min (M, N))$ with unit-diagonal, and $\mathbf{U}$ is an upper triangular or trapezoidal matrix. The SciPy command for this decomposition is linalg.lu.

Such a decomposition is often useful for solving many simultaneous equations where the left-hand-side does not change but the right hand side does. For example, suppose we are going to solve

$$
\mathbf{A} \mathbf{x}_{i}=\mathbf{b}_{i}
$$

[^1]for many different $\mathbf{b}_{i}$. The LU decomposition allows this to be written as
$$
\mathbf{P L U x}_{i}=\mathbf{b}_{i} .
$$

Because $\mathbf{L}$ is lower-triangular, the equation can be solved for $\mathbf{U} \mathbf{x}_{i}$ and finally $\mathbf{x}_{i}$ very rapidly using forwardand back-substitution. An initial time spent factoring A allows for very rapid solution of similar systems of equations in the future. If the intent for performing LU decomposition is for solving linear systems then the command linalg.lu_factor should be used followed by repeated applications of the command linalg.lu_solve to solve the system for each new right-hand-side.

### 10.3.4 Cholesky decomposition

Cholesky decomposition is a special case of LU decomposition applicable to Hermitian positive definite matrices. When $\mathbf{A}=\mathbf{A}^{H}$ and $\mathbf{x}^{H} \mathbf{A} \mathbf{x} \geq 0$ for all $\mathbf{x}$, then decompositions of $\mathbf{A}$ can be found so that

$$
\begin{aligned}
\mathbf{A} & =\mathbf{U}^{H} \mathbf{U} \\
\mathbf{A} & =\mathbf{L L}^{H}
\end{aligned}
$$

where $\mathbf{L}$ is lower-triangular and $\mathbf{U}$ is upper triangular. Notice that $\mathbf{L}=\mathbf{U}^{H}$. The command linagl.cholesky computes the cholesky factorization. For using cholesky factorization to solve systems of equations there are also linalg.cho_factor and linalg.cho_solve routines that work similarly to their LU decomposition counterparts.

### 10.3.5 QR decomposition

The QR decomposition (sometimes called a polar decomposition) works for any $M \times N$ array and finds an $M \times M$ unitary matrix $\mathbf{Q}$ and an $M \times N$ upper-trapezoidal matrix $\mathbf{R}$ such that

$$
\mathbf{A}=\mathbf{Q R}
$$

Notice that if the SVD of $\mathbf{A}$ is known then the QR decomposition can be found

$$
\mathbf{A}=\mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^{H}=\mathbf{Q} \mathbf{R}
$$

implies that $\mathbf{Q}=\mathbf{U}$ and $\mathbf{R}=\boldsymbol{\Sigma} \mathbf{V}^{H}$. Note, however, that in SciPy independent algorithms are used to find QR and SVD decompositions. The command for QR decomposition is linalg.qr.

### 10.3.6 Schur decomposition

For a square $N \times N$ matrix, $\mathbf{A}$, the Schur decomposition finds (not-necessarily unique) matrices $\mathbf{T}$ and $\mathbf{Z}$ such that

$$
\mathbf{A}=\mathbf{Z} \mathbf{T} \mathbf{Z}^{H}
$$

where $\mathbf{Z}$ is a unitary matrix and $\mathbf{T}$ is either upper-triangular or quasi-upper triangular depending on whether or not a real schur form or complex schur form is requested. For a real schur form both $\mathbf{T}$ and $\mathbf{Z}$ are real-valued when $\mathbf{A}$ is real-valued. When $\mathbf{A}$ is a real-valued matrix the real schur form is only quasi-upper triangular because $2 \times 2$ blocks extrude from the main diagonal corresponding to any complex-valued eigenvalues. The command linalg.schur finds the Schur decomposition while the command linalg.rsf2csf converts $\mathbf{T}$ and $\mathbf{Z}$ from a real Schur form to a complex Schur form. The Schur form is especially useful in calculating functions of matrices.

The following example illustrates the schur decomposition:

```
>>> A = mat('[1 3 2; 1 4 5; 2 3 6]')
>>> T,Z = linalg.schur(A)
>>> T1,Z1 = linalg.schur(A,'complex')
>>> T2,Z2 = linalg.rsf2csf(T,Z)
>>> print T
```

```
Matrix([[ 9.9001, 1.7895, -0.655 ],
    [ 0. , 0.5499, -1.5775],
    [ 0. , 0.5126, 0.5499]])
>>> print T2
Matrix([[ 9.9001+0.j , -0.3244+1.5546j, -0.8862+0.569j ],
    [0. +0.j , 0.5499+0.8993j, 1.0649-0.j ],
    [ 0. +0.j , 0. +0.j , 0.5499-0.8993j]])
>>> print abs(T1-T2) # different
[[ 0. 2.1184 0.1949]
    [ 0. 0. 1.2676]
    [ 0. 0. 0. ]}
>>> print abs(Z1-Z2) # different
[[[ 0.0683 1.1175 0.1973]
    [ 0.1186 0.5644 0.247 ]
    [ 0.1262 0.7645 0.1916]]
>>> T,Z,T1,Z1,T2,Z2 = map(mat, (T,Z,T1,Z1,T2,Z2))
>>> print abs(A-Z*T*Z.H)
Matrix([[ 0., 0., 0.],
    [ 0., 0., 0.],
    [ 0., 0., 0.]])
>>> print abs(A-Z1*T1*Z1.H)
Matrix([[ 0., 0., 0.],
    [ 0., 0., 0.],
    [ 0., 0., 0.]])
>>> print abs(A-Z2*T2*Z2.H)
Matrix([[ 0., 0., 0.],
    [0., 0., 0.],
    [0., 0., 0.]])
```


### 10.4 Matrix Functions

Consider the function $f(x)$ with Taylor series expansion

$$
f(x)=\sum_{k=0}^{\infty} \frac{f^{(k)}(0)}{k!} x^{k} .
$$

A matrix function can be defined using this Taylor series for the square matrix $\mathbf{A}$ as

$$
f(\mathbf{A})=\sum_{k=0}^{\infty} \frac{f^{(k)}(0)}{k!} \mathbf{A}^{k}
$$

While, this serves as a useful representation of a matrix function, it is rarely the best way to calculate a matrix function.

### 10.4.1 Exponential and logarithm functions

The matrix exponential is one of the more common matrix functions. It can be defined for square matrices as

$$
e^{\mathbf{A}}=\sum_{k=0}^{\infty} \frac{1}{k!} \mathbf{A}^{k} .
$$

The command linalg.expm3 uses this Taylor series definition to compute the matrix exponential. Due to poor convergence properties it is not often used.

Another method to compute the matrix exponential is to find an eigenvalue decomposition of $\mathbf{A}$ :

$$
\mathbf{A}=\mathbf{V} \boldsymbol{\Lambda} \mathbf{V}^{-1}
$$

and note that

$$
e^{\mathbf{A}}=\mathbf{V} e^{\boldsymbol{\Lambda}} \mathbf{V}^{-1}
$$

where the matrix exponential of the diagonal matrix $\boldsymbol{\Lambda}$ is just the exponential of its elements. This method is implemented in linalg.expm2.

The preferred method for implementing the matrix exponential is to use scaling and a Padé approximation for $e^{x}$. This algorithm is implemented as linalg.expm.

The inverse of the matrix exponential is the matrix logarithm defined as the inverse of the matrix exponential.

$$
\mathbf{A} \equiv \exp (\log (\mathbf{A}))
$$

The matrix logarithm can be obtained with linalg.logm.

### 10.4.2 Trigonometric functions

The trigonometric functions $\sin , \cos$, and tan are implemented for matrices in linalg.sinm, linalg.cosm, and linalg.tanm respectively. The matrix sin and cosine can be defined using Euler's identity as

$$
\begin{aligned}
\sin (\mathbf{A}) & =\frac{e^{j \mathbf{A}}-e^{-j \mathbf{A}}}{2 j} \\
\cos (\mathbf{A}) & =\frac{e^{j \mathbf{A}}+e^{-j \mathbf{A}}}{2}
\end{aligned}
$$

The tangent is

$$
\tan (x)=\frac{\sin (x)}{\cos (x)}=[\cos (x)]^{-1} \sin (x)
$$

and so the matrix tangent is defined as

$$
[\cos (\mathbf{A})]^{-1} \sin (\mathbf{A}) .
$$

### 10.4.3 Hyperbolic trigonometric functions

The hyperbolic trigonemetric functions sinh, cosh, and tanh can also be defined for matrices using the familiar definitions:

$$
\begin{aligned}
\sinh (\mathbf{A}) & =\frac{e^{\mathbf{A}}-e^{-\mathbf{A}}}{2} \\
\cosh (\mathbf{A}) & =\frac{e^{\mathbf{A}}+e^{-\mathbf{A}}}{2} \\
\tanh (\mathbf{A}) & =[\cosh (\mathbf{A})]^{-1} \sinh (\mathbf{A}) .
\end{aligned}
$$

These matrix functions can be found using linalg.sinhm, linalg.coshm, and linalg.tanhm.

### 10.4.4 Arbitrary function

Finally, any arbitrary function that takes one complex number and returns a complex number can be called as a matrix function using the command linalg.funm. This command takes the matrix and an arbitrary Python function. It then implements an algorithm from Golub and Van Loan's book "Matrix Computations" to compute function applied to the matrix using a Schur decomposition. Note that the function needs to accept complex numbers as input in order to work with this algorithm. For example the following code computes the zeroth-order Bessel function applied to a matrix.

```
>>> A = rand(3,3)
>>> B = linalg.funm(A,lambda x: special.jv(0,real(x)))
>>> print A
[[ [0.0593 0.5612 0.4403]
    [ 0.8797 0.2556 0.1452]
    [ 0.964 0.9666 0.1243]}
>>> print B
[[ 0.8206 -0.1212 -0.0612]
    [-0.1323 0.8256 -0.0627]
    [-0.2073 -0.1946 0.8516]}
```


## 11 Statistics

SciPy has a tremendous number of basic statistics routines with more easily added by the end user (if you create one please contribute it). All of the statistics functions are located in the sub-package stats and a fairly complete listing of these functions can be had using info(stats).

### 11.1 Random Variables

There are two general distribution classes that have been implemented for encapsulating continuous random variables and discrete random variables. Over 80 continuous random variables and 10 discrete random variables have been implemented using these classes. The list of the random variables available is in the docstring for the stats sub-package. A detailed description of each of them is also located in the files continuous.lyx and discrete.lyx in the stats sub-directories.

## 12 Interfacing with Python Imaging Library

If you have the Python Imaging Library installed, SciPy provides some convenient functions that make use of it's facilities particularly for reading, writing, displaying, and rotating images. In SciPy an image is always a two- or three-dimensional array. Gray-scale, and colormap images are always two-dimensional arrays while RGB images are three-dimensional with the third dimension specifying the channel.

Commands available include

- fromimage - convert a PIL image to a Numeric array
- toimage - convert Numeric array to PIL image
- imsave - save Numeric array to an image file
- imread - read an image file into a Numeric array
- imrotate - rotate an image (a Numeric array) counter-clockwise
- imresize - resize an image using the PIL
- imshow - external viewer of a Numeric array (using PIL)
- imfilter - fast, simple built-in filters provided by PIL
- radon - simple radon transform based on imrotate


## 13 Plotting with xplt

### 13.1 Gist

The underlying graphics library for xplt is the pygist library. All of the commands of pygist are available under xplt as well. For more information on the pygist commands you can read the documentation of that package in html here http://bonsai.ims.u-tokyo.ac.jp/~mdehoon/software/python/pygist_ html/pygist.html or in pdf at this location http://bonsai.ims.u-tokyo.ac.jp/~mdehoon/software/ python/pygist.pdf.

### 13.2 Basic commands

### 13.3 Adding a legend

### 13.4 Drawing a histogram

### 13.5 Drawing a barplot

13.6 Plotting color arrays

### 13.7 Contour plots

13.8 Three-dimensional plots
13.9 Placing text and arrows
13.10 Special plots


[^0]:    ${ }^{1}$ These changes are all made in a new module (fastumath) that is part of the scipy_base package. The old functionality is still available in umath (part of Numeric) if you need it (note: importing umath or fastumath resets the behavior of the infix operators to use the umath or fastumath ufuncs respectively).
    ${ }^{2}$ Be careful when treating logical expressions as integers as the 8 -bit integers may silently overflow at 256.

[^1]:    ${ }^{3}$ A hermition matrix $\mathbf{D}$ satisfies $\mathbf{D}^{H}=\mathbf{D}$.
    ${ }^{4} \mathrm{~A}$ unitary matrix $\mathbf{D}$ satisfies $\mathbf{D}^{H} \mathbf{D}=\mathbf{I}=\mathbf{D D}^{H}$ so that $\mathbf{D}^{-1}=\mathbf{D}^{H}$.

