

Lab 2

NumPy and SciPy

Lab Objective: *Create and manipulate NumPy arrays and learn features available in NumPy and SciPy.*

Introduction

NumPy and SciPy¹ are the two Python libraries most used for scientific computing. NumPy is a package for manipulating vectors and arrays, and SciPy is a higher-level library built on NumPy. The basic object in NumPy is the *array*, which is conceptually similar to a matrix. However, unlike a matrix, which has two dimensions, a NumPy `array` can have arbitrarily many dimensions. NumPy is optimized for fast array computations.

The convention is to import NumPy as follows.

```
>>> import numpy as np
```

Learning NumPy

The strategies discussed in the section “Learning Python” of Lab 1 will also help you learn NumPy and SciPy. The following online resources are specific to SciPy:

- Official SciPy Documentation (<http://docs.scipy.org/doc/>)
- Sections 1.3 and 1.5 of the SciPy Lecture Notes (<http://scipy-lectures.github.io/>)

The remainder of this lab is a brief summary of the tools available in NumPy and SciPy, beginning with NumPy arrays.

¹SciPy is also the name of a Python coding environment that includes the NumPy and SciPy libraries, as well as IPython, matplotlib, and other tools.

Arrays

Conceptually, a 1-dimensional array (called a 1-D array) is just a list of numbers. An n -dimensional array (or n -D array) is an array of $(n - 1)$ -dimensional arrays. Thus, any 2-D array is conceptually a matrix, and a 3-D array is a list of matrices, which can be visualized as a cube of numbers. Each dimension is called an *axis*. When a 2-D array is printed to the screen, the 0-axis indexes the rows and the 1-axis indexes the columns.

The NumPy array class is called `ndarray`. The simplest way to create an `ndarray` is to define it explicitly using nested lists.

```
# Create a 1-D array
>>> np.array([0, 3, 8, 6, 3.14])
array([0, 3, 8, 6, 3.14])

# Create a 2-D array
>>> ex1 = np.array([[1, 1, 2], [3, 3, 4]])
>>> ex1
array([[1, 1, 2],
       [3, 3, 4]])
```

You can view the length of each dimension with the `shape` command, and change the shape of an array with the `np.reshape()` function. The number of arguments passed to `reshape` tells NumPy the dimension of the new array, and the arguments specify the length of each dimension. An argument of `-1` tells NumPy to make that dimension as long as necessary.

```
# The 0-axis of ex1 has length 2
>>> ex1.shape
(2, 3)
>>> ex1.reshape(3, 2)
array([[1, 1],
       [2, 3],
       [3, 4]])
>>> ex1.reshape(-1)
array([1, 1, 2, 3, 3, 4])
```

Array objects also support the usual binary operators, including addition `+` and element-wise multiplication `*`. For Python lists, the `+` operator performs concatenation and the `*` operator is not defined.

```
>>> a = [1,2,3]
>>> b = [4,5,6]
>>> a+b
[1, 2, 3, 4, 5, 6]
>>> a*b
TypeError: cannot multiply sequence by non-int of type 'list'

>>> a = np.array(a)
>>> b = np.array(b)
>>> a+b
array([5,7,9])
>>> a*b
array([4, 10, 18])
```

Why Use Arrays?

NumPy arrays are drastically more efficient than nested Python lists for large computations. In this section we will compare matrix multiplication in Python and in NumPy.

Problem 1. A matrix in NumPy is just a 2-D array. Given matrices A and B , there are two different ways we can perform matrix multiplication. We can use `np.dot(A,B)` or `A.dot(B)`.

Perform the matrix multiplication $A * B$ on the following matrices:

$$A = \begin{bmatrix} 2 & 4 & 0 \\ -3 & 1 & -1 \\ 0 & 3 & 2 \end{bmatrix} \quad B = \begin{bmatrix} 3 & -1 & 2 \\ -2 & -3 & 0 \\ 1 & 0 & -2 \end{bmatrix}$$

Remember that to be able to use `np.dot`, we must first define A and B as NumPy arrays.

After doing the previous problem, you should know how to implement matrix multiplication in NumPy. On the other hand, a matrix in Python can be represented as a list of lists. We can perform matrix multiplication using lists by using numerical multiplication and addition. The following function will multiply two such matrices in this manner in Python without using NumPy.

```

1 def arr_mult(A,B):
2     new = []
3     # Iterate over the rows of A.
4     for i in range(len(A)):
5         # Create a new row to insert into the product.
6         newrow = []
7         # Iterate over the columns of B.
8         # len(B[0]) returns the length of the first row
9         # (the number of columns).
10        for j in range(len(B[0])):
11            # Initialize an empty total.
12            tot = 0
13            # Multiply the elements of the row of A with
14            # the column of B and sum the products.
15            for k in range(len(B)):
16                tot += A[i][k] * B[k][j]
17            # Insert the value into the new row of the product.
18            newrow.append(tot)
19        # Insert the new row into the product.
20        new.append(newrow)
21    return new

```

arr_mult.py

Table 2.1 documents how long² one computer took to square a $k \times k$ matrix in both Python (using the function `arr_mult`) and NumPy (using the method you

²You can replicate this experiment yourself. In IPython, you can find the execution time of a line of code by prefacing it with `%timeit`. If you aren't using IPython, you will need to use the `timeit` function documented here: <https://docs.python.org/2/library/timeit.html>.

Data Structure	$k \times k$	Time (s)
Python List	10×10	0.0002758503
	100×100	0.1336028576
	1000×1000	200.4009799957
NumPy Array	10×10	0.0000109673
	100×100	0.0009210110
	1000×1000	2.1682999134

Table 2.1: Time for one computer to square a $k \times k$ matrix in Python and NumPy.

Data type	Description
<code>bool</code>	Boolean
<code>int8</code>	8-bit integer
<code>int16</code>	16-bit integer
<code>int32</code>	32-bit integer
<code>int64</code>	64-bit integer
<code>int</code>	Platform integer (depends on platform)
<code>uint8</code>	Unsigned 8-bit integer
<code>uint16</code>	Unsigned 16-bit integer
<code>uint32</code>	Unsigned 32-bit integer
<code>uint64</code>	Unsigned 64-bit integer
<code>float16</code>	Half-precision float
<code>float32</code>	Single-precision float
<code>float64</code>	Double-precision float (also <code>float</code>)
<code>complex64</code>	Complex number represented by two single-precision floats
<code>complex128</code>	Complex number represented by two double-precision floats (also <code>complex</code>)

Table 2.2: Native numerical data types available in NumPy.

found in Problem 1) for various values of k . As you can see, NumPy is much faster. One reason for this is that algorithms in NumPy are usually implemented in C or in Fortran.

Data Types

Unlike Python containers, a NumPy array requires that all of its elements have the same data type. The data types used by NumPy arrays are machine-native and avoid the overhead of Python objects, meaning that they are faster to compute with. A NumPy `int` and a Python `int` are not the same; the former has been optimized to speed up numerical computations. Datatypes supported by NumPy are shown in Table 2.2.

Here is an example of how to manipulate data types in NumPy:

```
# Access the data type of an array
>>> ex2 = np.array(range(5))
>>> ex2.dtype
dtype('int64')

# Specify the data type of an array
```

Function	Description
<code>diag</code>	Extract a diagonal or construct a diagonal array.
<code>empty</code>	Return a new array of given shape and type, without initializing entries.
<code>empty_like</code>	Return a new array with the same shape and type as a given array.
<code>eye</code>	Return a 2-D array with ones on the diagonal and zeros elsewhere.
<code>identity</code>	Return the identity array.
<code>meshgrid</code>	Return coordinate matrices from two coordinate vectors.
<code>ones</code>	Return a new array of given shape and type, filled with ones.
<code>ones_like</code>	Return an array of ones with the same shape and type as a given array.
<code>zeros</code>	Return a new array of given shape and type, filled with zeros.
<code>zeros_like</code>	Return an array of zeros with the same shape and type as a given array.

Table 2.3: Some functions for creating arrays in NumPy.

```
>>> ex3 = np.array(range(5), dtype=np.float)
>>> ex3.dtype
dtype('float64')
```

Creating Arrays

In addition to `np.array()`, NumPy provides efficient ways to create special kinds of arrays. The function `np.arange([start], stop, [step])` returns an array of numbers from `start` up to, but not including, `stop`. Like other functions with similar parameters, `start` defaults to 0 and `step` defaults to 1.

```
>>> np.arange(5)
array([0,1,2,3,4])
>>> np.arange(10, 20, 2)
array([10, 12, 14, 16, 18])
```

Use `np.linspace(start, stop, num=50)` to create an array of `num` numbers evenly spaced in the interval from `start` to `stop` inclusive.

```
>>> np.linspace(0, 32, 4)
array([ 0.          , 10.66666667, 21.33333333, 32.          ])
```

We can even create arrays of random values chosen from probability distributions. These probability distributions are stored in the submodule `np.random`.

```
>>> np.random.rand(5) # uniformly distributed values in [0, 1)
array([ 0.21845499,  0.73352537,  0.28064456,  0.66878454,  0.44138609])
```

Some other commonly-used functions are `np.random.randn`, which samples from the normal distribution, `np.random.randint`, which randomly selects integers from a range, and `np.random.random_integers` which returns an array of random integers in a given range. There are many functions for creating arrays besides these, some of which are described in Table 2.3. See <http://docs.scipy.org/doc/numpy/reference/routines.array-creation.html> for more details.

Indexing and Slicing

Indexing for a 1-D NumPy array works exactly like indexing for a Python list. To access a single entry of a multi-dimensional array, say a 3-D array, you should use the syntax `f[i, j, k]`. While the syntax `f[i][j][k]` will also work, it is significantly slower because each bracket returns an array slice. Similarly, slicing an array works just like slicing a list, but with more dimensions.

```
>>> ex4 = np.arange(25).reshape((5,5))
>>> ex4
array([[ 0,  1,  2,  3,  4],
       [ 5,  6,  7,  8,  9],
       [10, 11, 12, 13, 14],
       [15, 16, 17, 18, 19],
       [20, 21, 22, 23, 24]])
>>> ex4[4, -2]
23

# Extract the lower right 2x2 subarray.
>>> ex4[3:, 3:]
array([[18, 19],
       [23, 24]])

# Extract the second column. The returned array is 1-D.
>>> ex4[:, 1]
array([ 1,  6, 11, 16, 21])

# Reverse the order of the columns.
>>> ex4[:, ::-1]
array([[ 4,  3,  2,  1,  0],
       [ 9,  8,  7,  6,  5],
       [14, 13, 12, 11, 10],
       [19, 18, 17, 16, 15],
       [24, 23, 22, 21, 20]])
```

Fancy Indexing

Fancy indexing is a second way to access elements of an array. There are two types of fancy indexing: boolean and integer. Boolean indexing uses an array of `True` or `False` values to determine which elements of the array to take.

```
>>> my_array = np.array([0, 10, 20, 30, 40])
>>> bool_mask = np.array([1,0,1,1,0], dtype=bool)
>>> my_array[bool_mask]
array([0,20,30])

# Return only the values larger than 25
>>> comparison_mask = my_array > 25
>>> comparison_mask
array([False, False, False,  True,  True], dtype=bool)
>>> my_array[comparison_mask]
array([30, 40])

# If we do not need to save the mask for other uses, we can achieve the same ↵
  result by doing the following:
>>> my_array[my_array>25]
```

```
array([30, 40])

# np.argsort returns a list of indices that would sort an array
>>> my_array = np.array([4,1,6,3,7,2,5])
>>> mask = np.argsort(my_array)
>>> my_array[mask]
array([1, 2, 3, 4, 5, 6, 7])
```

Integer indexing uses Python lists to determine what array values to access.

```
>>> my_array = np.array([0,10,20,30,40])
>>> int_mask = [2,3,1]
>>> my_array[int_mask]
array([20,30,10])
```

Fancy indexing can be used for assignment. For example, we can set all values of an array that are less than 25 to 0 in the following way:

```
>>> my_array[my_array<25] = 0
>>> my_array
array([0, 0, 0, 30, 40])
```

Problem 2. Write a function that accepts an array and returns an array with all nonnegative numbers. If an entry in the inputted array is negative, set it to 0. Though one of your first instincts may be to use a `for` loop, it is much more efficient to solve this problem using fancy indexing.

Array Views and Copies

NumPy has two ways of returning an array. Slice operations and indexing always return a *view* and fancy indexing always returns a *copy*. Understand that even though they may look the same, views and copies are different.

A view of an array is a distinct object from the original array in Python, but it references the same place in memory. Thus, when you change elements in a view, you also change the array it references.

```
>>> ex5 = np.arange(5)

# Slicing produces a view of ex5.
>>> view_ex5 = ex5[:]
>>> view_ex5

# Check that ex5 and view_ex5 are distinct objects in Python.
>>> id(view_ex5) == id(ex5)
False

# Change the third element of view_ex5 to 500
# Changing view_ex5 also changes ex5.
>>> view_ex5[2] = 500
>>> view_ex5
```

```
array([ 0,  1, 500,  3,  4])
>>> ex5
array([ 0,  1, 500,  3,  4])
```

A copy of an array is a separate array with its own place in memory. Thus, when you change a copy of an array, you do not affect the original array. Because copying an array uses more memory and also more time, it should only be done when necessary. An array can be copied using the `np.copy()` function (also available as a method of the array object, for example `my_array.copy()` returns a copy of `my_array`).

```
>>> copy_ex5 = np.copy(ex5)

# Check that ex5 and copy_ex5 are distinct objects in Python.
>>> id(copy_ex5) == id(ex5)
False

# Change the third element of copy_ex5 to 1000
# Changing copy_ex5 does not affect ex5.
>>> copy_ex5[2] = 1000
>>> copy_ex5
array([ 0,  1, 1000,  3,  4])
>>> ex5
array([ 0,  1, 500,  3,  4])
```

Whenever possible, the function `np.reshape()` returns a view. See the documentation for more information.

More Methods of NumPy Arrays

Some of the more common methods of NumPy arrays are described in Table 2.4. A more comprehensive list can be found at <http://docs.scipy.org/doc/numpy/reference/generated/numpy.ndarray.html>.

Many of these methods have the option to operate *along an axis*. When called in this way on an n -D array, these methods return an $(n-1)$ -D array (the specified axis is collapsed in the evaluation process).

```
>>> ex6 = np.arange(9).reshape(3, 3)
>>> ex6
array([[0, 1, 2],
       [3, 4, 5],
       [6, 7, 8]])

# Return the maximum value in the array
>>> ex6.max()
8

# Return the maximum values evaluated along the 0-axis
>>> ex6.max(axis=0)
array([6, 7, 8])

# Return the maximum values evaluated along the 1-axis
>>> ex6.max(axis=1)
array([2, 5, 8])
```


Function	Description
<code>all</code>	returns True if all elements evaluate to True
<code>any</code>	returns True if any elements evaluate to True
<code>argmax</code>	indices of maximum value(s)
<code>argmin</code>	indices of minimum value(s)
<code>argsort</code>	indices that would sort the array
<code>astype</code>	casts a copy of an array to a different data type
<code>clip</code>	restrict values in an array to fit within a given range
<code>conj</code>	return the complex conjugate of the array
<code>copy</code>	return a copy of the array
<code>diagonal</code>	return a given diagonal of the array
<code>dot</code>	matrix multiplication
<code>max</code>	max element of the array
<code>mean</code>	average of the array
<code>min</code>	minimum element of the array
<code>prod</code>	product of elements of the array
<code>ravel</code>	make a flattened version of an array, return a view if possible
<code>reshape</code>	return a view of the array with a changed shape
<code>round</code>	return a rounded version of the array
<code>sort</code>	sort the array in place
<code>std</code>	compute the standard deviation
<code>sum</code>	sum the elements of the array
<code>swapaxes</code>	return a view with the given axes swapped
<code>tolist</code>	return the array represented as a list or nested list
<code>trace</code>	return the sum of the elements along the main diagonal
<code>var</code>	return the variance of the array
<code>vstack</code>	stack arrays in sequence vertically

Table 2.4: A few of the methods of NumPy arrays.

Problem 3. Write a function which accepts an integer n as input and does the following:

1. Creates an $n \times n$ array of `floats` randomly chosen from a normal distribution.
2. Computes the mean of each row (use a built-in command).
3. Computes the variance of these means (use a built-in command).

As you increase n , what happens to the output of your function? This illustrates one version of the Law of Large Numbers, about which you will learn more later on.

Problem 4. One good application of array slicing is the Jacobi method for solving Laplace's equation, which is used to model steady-state heat flow on a square. This problem will help you implement the Jacobi method.

Make a function that accepts an array and a float (which represents a tolerance) as input and does the following:

1. Makes a copy of the array.
2. Creates a variable to track the difference between the arrays. Initialize it as the tolerance parameter your function accepts.
3. While the difference is greater than or equal to the tolerance
 - (a) Sets all points that are not on an edge of the new array equal to the average of their 4 immediate neighbors. Use the values from the old array for this computation. This should only take one line and should be based entirely on array slicing, NOT iterating through the array. (Hint: given a 2D array `A`, the slice `A[1:-1,1:-1]` references all non-edge entries, `A[:-2,1:-1]` references the upper neighbors, and `A[1:-1,2:]` references the right neighbors.)
 - (b) Updates the difference to be the maximum of the absolute value of the new array minus the old one.
 - (c) Copies the values from the new array into the old one (without creating a new array).

Now use the following code to generate a plot of your results.


```

1 from matplotlib import pyplot as plt
2 from mpl_toolkits.mplot3d import Axes3D
   n = 100
4 tol = .0001
   U = np.ones ((n, n))
6 U[:,0] = 100 # sets west boundary condition
   U[:, -1] = 100 # sets east boundary condition
8 U[0] = 0 # sets north boundary condition
   U[-1] = 0 # sets south boundary condition
10 # U has been changed in place (note that laplace is the name of
   # the function in this case).
12 laplace(U, tol)
   x = np.linspace(0, 1, n)
14 y = np.linspace(0, 1, n)
   X, Y = np.meshgrid(x, y)
16 fig = plt.figure()
   ax = fig.gca(projection = '3d')
18 ax.plot_surface(X, Y, U, rstride=5)
   plt.show()

```

laplace_plot.py

It should resemble the following figure.



Iterating Through Arrays

Iterating through an array negates most speed advantages of NumPy. You should avoid doing this whenever possible. You can often avoid iterating through arrays by using *array broadcasting* and *universal functions*, discussed in the next section.

It is occasionally valid to iterate through an array. The function `np.nditer()` will create an object that iterates through an array as quickly as possible.

NumPy and SciPy

We now introduce some additional features of NumPy and SciPy.

Array Broadcasting

Many matrix operations make sense only when the two operands have the same shape. Two examples are addition and element-wise multiplication. Broadcasting is NumPy's way of extending such operations to accept some (not all) operands with different shapes. Broadcasting happens automatically whenever it is necessary.

To understand broadcasting, let us look at an example.

```
>>> A = np.ones(3);
>>> A
array([ 1.,  1.,  1.])
>>> B = np.vstack([1, 2, 3])
>>> B
array([[1],
       [2],
       [3]])
>>> A+B
array([[ 2.,  2.,  2.],
       [ 3.,  3.,  3.],
       [ 4.,  4.,  4.]])
```

We will now describe the algorithm used to obtain the result for $\mathbf{A}+\mathbf{B}$ above. First, the shapes of \mathbf{A} and \mathbf{B} are lined up, starting at the far right, and 1's are prepended to the shorter tuple. So

```
A (1-D array): 3
B (2-D array): 3 x 1
```

becomes

```
A (1-D array): 1 x 3
B (2-D array): 3 x 1
```

For broadcasting to work, the dimensions must be compatible; that is, in a given axis, the lengths are equal, or one of the lengths is 1. Second, the arrays \mathbf{A} and \mathbf{B} are “stretched” one axis at a time until the lengths of their axes are the same. In each axis, if the lengths are different, the smaller array is copied along that axis (or “stretched”), until it is the size of the larger array. Conceptually, we are creating new 3×3 arrays \mathbf{A}' and \mathbf{B}' where

$$\mathbf{A}' = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} \quad \text{and} \quad \mathbf{B}' = \begin{bmatrix} 1 & 1 & 1 \\ 2 & 2 & 2 \\ 3 & 3 & 3 \end{bmatrix}.$$

Finally, NumPy returns the sum $\mathbf{A}' + \mathbf{B}'$.

We emphasize that the “stretching” in this example is only conceptual, and no new array \mathbf{A}' or \mathbf{B}' is created. However, you should still be careful when broadcasting large arrays because you can fill the RAM on your computer, which can sometimes freeze the system. For a more detailed description of array broadcasting rules, see <http://docs.scipy.org/doc/numpy/user/basics.broadcasting.html>.

Problem 5. Create a $100 \times 100 \times 3$ array of random integers taking values in the range $[0, 255]$. Such an array can represent an RGB image of 100×100 pixels, where each pixel is associated with an array of three integers indicating the amounts of red, green, and blue color present in that pixel. Use array broadcasting to multiply the red and green values by 0.5. (Such an operation would tone down the red and green colors and make the image appear more blue.)

To visualize your solutions, you may use the following function.

```
1 original, blue = blue_shift()
2 original = 255 - original
  blue = 255 - blue
4 plt.subplot(1,2,1)
  plt.imshow(original)
6 plt.subplot(1,2,2)
  plt.imshow(blue)
8 plt.show()
```

blue_shift_plot.py

Universal Functions

A universal function, or `ufunc`, operates on an array element-wise. It outputs an array of the same shape and datatype as the input array. Using a universal function is usually much faster than iterating through the array yourself.

Many scalar functions from the Python standard library have a universal analog that operates on arrays. For example, `math.sin()` operates on scalars, and `numpy.sin()` operates on arrays. If you have a simple operation that you want to perform element-wise on an array, you should see if SciPy has a universal function that will do it (it probably does). For a list of available `ufuncs`, see <http://docs.scipy.org/doc/numpy/reference/ufuncs.html#available-ufuncs>.

Most universal functions also allow you to specify an output array, which must have the same shape as the input array. Doing so can reduce memory allocation.

```
>>> ex7 = np.arange(3, dtype=float)

# Take exp(ex7) and store the result in ex7.
>>> np.exp(ex7, out=ex7)
>>> ex7
array([ 1.          ,  2.71828183,  7.3890561 ])
```

Although universal functions also accept scalar inputs, they can be much slower than the corresponding standard library function. Thus, use standard library functions on scalars and universal functions on arrays.

Linear Algebra

Both NumPy and SciPy have a linear algebra library, but the SciPy library is larger. The SciPy linear algebra library is typically imported as follows:

```
from scipy import linalg as la
```

The linear algebra library contains several functions to construct special matrices, located in `linalg.special_matrices`. There are also functions that will invert matrices, find determinants and norms, solve linear systems and least squares problems, and find special matrix decompositions. You can read more about the linear algebra capabilities of SciPy in the documentation for the `linalg` module found at <http://docs.scipy.org/doc/scipy/reference/linalg.html>.

Finally, the `scipy.linalg` library has a `matrix` class that is very similar to a 2-D NumPy array. The `matrix` class can be convenient when doing matrix operations. However, in such situations we still recommend using NumPy arrays, which have many of the same features and are also compatible with all other SciPy operations.

The `scipy.linalg` library will be essential for the remainder of the labs in this manual. We will address the details of this package at length in future labs.

