


COMPUTATIONAL METHODS FOR SCIENTISTS

INTRODUCTION TO PYTHON



Michael J. Ware

Department of Physics and Astronomy
Brigham Young University

COMPUTATIONAL METHODS FOR SCIENTISTS

INTRODUCTION TO PYTHON

Michael J. Ware

Department of Physics and Astronomy
Brigham Young University

© 2019 – 2024

Last Revised: April 4, 2024

Acknowledgements

This book is based on the book *Introduction to Scientific Computing in Python*, by Lance J. Nelson and Matthew R. Zachreson at BYU Idaho. With their permission, I've used parts of their material in creating this tutorial. In a sort of virtuous cycle, Ross Spencer and I had previously allowed Lance and Matthew to borrow material from our book *Introduction to Matlab* in the creation of their *Introduction to Scientific Computing in Python*. I express sincere thanks and acknowledgment to Ross, Lance, and Matthew for all of the work that has gone into development this material.

Preface

This book is a tutorial for physics students to get up to speed using Python for scientific computing as quickly as possible. It assumes that the reader is already familiar with the basics of scientific programming in another programming language, and does not spend time systematically going through the fundamentals of programming. This tutorial is designed to work hand-in-hand with the BYU Physics 430 lab manual. Each chapter in this tutorial is designed to give students enough understanding of the Python syntax and ecosystem to tackle a specific scientific computing lab, and in doing so will sample from a range of different topics.

Students in the class for which the book is designed have previously taken a three-credit introductory programming class in C++, a one-credit lab course introducing them to Mathematica, and another one-credit lab course introducing them Matlab. In this book, we build on that foundation, without trying to re-teach (at least not too much) the material already covered in prior classes. We also present material in the order needed to complete the associated labs rather than as a systematic and complete treatment of each topic before moving to the next.

As you find mistakes or have suggestions, send them to me at ware@byu.edu.

Contents

Table of Contents	v
1 Numpy, Scipy, and Plotting	1
1.1 Get Anaconda Installed	1
1.2 Your First Program	1
1.3 The Python Console	2
1.4 Variables and Data Types	3
1.5 Integers	3
1.6 Float variables	4
1.7 Functions and Help	4
1.8 Boolean variables	5
1.9 String Variables	5
1.10 Formatting Printed Values	5
1.11 Functions and Libraries (NumPy)	6
1.12 Numpy Arrays	7
1.13 Making x-y Plots	11
1.14 Customizing the IDE	11
2 Lists, Loops, Logic, and 2D Arrays	13
2.1 Lists	13
2.2 The range Function	14
2.3 For Loops and Shorthand Assignments	14
2.4 Logical Statements	16
2.5 While Loops	17
2.6 Two-Dimensional NumPy Arrays–Matrices	18
2.7 Solving a Set of Linear Equations	20
2.8 NumPy Matrix	20
3 Functions	22
3.1 User-defined functions	22
3.2 Importing User Defined Functions	23
3.3 Writing Readable Code	24
Index	27

Chapter 1

Numpy, Scipy, and Plotting

1.1 Get Anaconda Installed

Python is a popular general-purpose programming language that has become a standard language for many areas of scientific computing. It is open source, and there are many implementations of Python, many development environments for it, and multiple versions of the programming language itself. We will use the Anaconda distribution for Python version 3.7. The Anaconda distribution is geared toward scientific computing, is available as a free download on all major platforms, and comes with an integrated development environment (IDE) called Spyder. You can obtain Anaconda Python for free at their website, anaconda.com.

1.2 Your First Program

Launch the Spyder IDE. You should see a window similar to the one in Fig. 1.1. The Spyder IDE interface is divided into three main windows:

- The code editor on the left is where you edit your Python code files (usually with a .py extension).
- The top right pane has three default tabs: a variable explorer where you can view current values stored in memory, a file explorer where you can browse your files, and a help tab where you can ask questions
- The lower right pane is a console window where your output will be displayed and you can issue Python commands directly to be evaluated.

To write your first program, press the “New File” button on the toolbar, erase any auto-generated text so you have a blank window, and then type

```
print('Hello World')
```

Save your program in a new directory (where you will save all your work for this class) with the name `hello.py`, then click the green arrow above the editor to run your program. Spyder may ask you in which console you’d like to execute the program. Just accept the default values, and then look down at the console window. There you will see some code that Spyder auto-generated to run your program, and under it should be the output from your program:

```
Hello World
```




Figure 1.1 The Spyder IDE window.

Congratulations, you just executed your first Python program. Change the text string `'Hello World'` to something else, and then click the green arrow again. Notice that Spyder saves your code and executes the program again. If you get tired of clicking the green arrow, F5 is the keyboard shortcut to save and execute the code shown in the editor.

1.3 The Python Console

Spyder's console window is a powerful environment called Interactive Python (or IPython for short) where you can directly enter Python commands. Put your cursor in the console window at the `In [1] :` prompt, type

```
1+1
```

and press enter. You should see the answer, stored in a variable `Out [1]`. You can use this output in later calculations by typing the following at the `In [2] :` prompt:

```
Out[1]+2
```

Notice that your new result is stored in `Out [2]`. IPython behaves a lot like Matlab's command window, with a series of inputs and variable values accumulating in the workspace. When you run your Python programs, your variables are placed in the console's workspace and can be accessed from the command line afterward. With your cursor in the command window, press the Up-Arrow key, and notice that you can access your command history with the up and down arrow keys.

The console can be useful for quick calculations or for looking at data when debugging, but you should do most of your programming in the editor. Add the following line to your `hello.py` program

```
1+1
```

and run it again. Note that the answer to `1+1` is *not* displayed in the console when you run the program. Now switch your program to read

```
print(1+1)
```

and run it again, and note that the answer displayed in all its glory. Python evaluates each line of code, but will not display the result of a calculation in the console unless you `print` it.

As we go through the remainder of the text, type all the indented example code into a `*.py` file in the editor by hand (don't copy and paste) and execute it using the green arrow (or the F5 keyboard shortcut). This method of interacting with the code will help you better process and understand what each command does. It forces you to read the code like Python will: one line at a time, top to bottom. Also, place each command on a separate line and don't indent any lines of code when you type them in. Python really cares about white space. We'll learn more about that in the next chapter.

1.4 Variables and Data Types

Variables in Python don't need to be declared before being used. You declare a variable and assign it a value in one statement using the assignment operator (=), like this

```
x = 20
```

(Did you type the line of code into your program and execute it? If not, do so now and get in that habit.) This statement creates the variable `x` and assigns it a value of 20. Python didn't print anything since we didn't include a `print` command. To convince yourself that the variable `x` was defined, click on the "Variable explorer" tab in the upper-right pane of Spyder to see that the variable exists, and has a value of 20. Also type `x` in the console window and hit enter, and note that Python displays its value. Now add the line

```
x = x + 1
```

to your program, execute it, and look at the new value of `x` to convince yourself that the program executed correctly. Multiple variables are defined by putting the assignments on separate lines, like this

```
a = 2
b = 4
c = a * b
```

Sometimes you may want your program to prompt the user to enter a value and then save the value that the user inputs in a variable. This can be done like this:

```
a = input('What is your age?')
```

When you run this line of code, you will be prompted to enter your age. When you do, the number you enter will be saved in the variable `a`.

Variable names must start with a letter or an underscore and consist of only letters, numbers, and underscores. Variable names are case sensitive, so watch your capitalization.

1.5 Integers

The simplest type of numerical data is an integer. Python implicitly declares variables as integers when you assign an integer value to the variable, as we did above. You can perform all the common mathematical operations on integer variables. For example:

```
a = 20
b = 15
c = a + b # add two numbers
```



```
d = a/b      # floating point division, single slash
i = a//b     # integer division, double slash
r = a % b    # the remainder of integer division
e = a * b    # multiply two numbers together
f = c**4     # raise number to a power (use **, not ^)
```

Performing an operation on two integers *usually* yields another integer. This can pose an ambiguity for division where the result is rarely an exact integer. Using the regular division operator (as in the line defining `d` above) for two integers yields a float, whereas the double slash (used in the line defining `i`) performs integer division. Look at the variable values in the variable explorer and compare them to your code to convince yourself that you understand the distinction between these two types of division.

1.6 Float variables

Most calculations in physics should be performed using floating point numbers, called floats in Python. Float variables are created and assigned in one of two ways. The first way is to simply include a decimal point in the number, like this

```
a = 20.
```

You can also cast an integer variable to a float variable using the `float` command

```
a = 20
b = float(a)
```

When a float and an integer are used in the same calculation, like this

```
a = 0.1
b = 3 * a #Integer multiplied by a float results in a float.
```

the result is always a float. Only when all of the numbers used in a calculation are integers will the result be an integer. Floats can be entered using scientific notation like this

```
y = 1.23e15
```

1.7 Functions and Help

Table 1.1 shows some common housekeeping functions that you can use with float variables. Practice using one of these functions in your code. Then place your cursor on the function name in your code and press Ctrl-I (think “I is for information”) to display detailed information about this function in the help panel. You can use this method for getting details of calling conventions and function usage for all of the functions we use in this class.

☞ These lines introduce the Python commenting syntax: any text following the symbol `#` on a line is ignored by the Python interpreter.

☞ In Python 2.x the single slash between two integer variables performs integer division. If you are using this older version of Python, you should be sure to use floats when you want regular division.

<code>abs(x)</code>	The absolute value of <code>x</code>
<code>divmod(x, y)</code>	The quotient and remainder when using long division.
<code>x % y</code>	The remainder of $\frac{x}{y}$
<code>float(x)</code>	Convert <code>x</code> to a float.
<code>int(x)</code>	Convert <code>x</code> to an integer.
<code>round(x)</code>	Round the value of <code>x</code> using standard rounding rules

Table 1.1 A sampling of functions commonly used with integers and floats.

1.8 Boolean variables

Boolean variables store one of two possible values: `True` or `False`. A boolean variable is created and assigned similar to the other variables you've studied so far

```
q = True
```

Boolean variables will be useful when we use loops and logical statements.

1.9 String Variables

String variables contain a sequence of characters, and can be created and assigned using quotes, like this

```
s='This is a string'
```

You may also enclose the characters in double quotes. This syntax is convenient when there are single quotes in your string:

```
t="Don't worry"
```

Some Python functions require options to be passed to them as strings. Make sure you enclose them in quotes, as shown above. Some commonly used functions for strings are shown in Table 1.2

1.10 Formatting Printed Values

The `print` command will take pretty much any variable and dump it to screen in a format of Python's choosing, like this:

```
a = 22
print(a)
```

You can intermingle variable values with text like this:

```
a = 22
b = 3.5
print(f'I am {a} years old and my GPA is: {b}')
```

Notice the letter `f` right before the string constant in the `print` argument. This f-string (think `formatted string`) syntax allows you to place any variable name inside of curly braces within the string to be printed, and Python will replace the variable in the braces with its value in the formatted output.

Often you will want more control over how the variables are displayed. To display the GPA to two digits of precision, you could modify the f-string like this:

```
print(f'I am {a} years old and my GPA is: {b:3.2f}')
```

<code>len(c)</code>	Get the number of characters in the string <code>c</code>
<code>a.count(b)</code>	Count the number of occurrences of string <code>b</code> in string <code>a</code>
<code>a.lower()</code>	Convert upper case letters to lower case
<code>a[x]</code>	Access element <code>x</code> in string <code>a</code>
<code>a[x:y:z]</code>	Slice a string, starting at element <code>x</code> , ending at element <code>y</code> , with a step size of <code>z</code>
<code>a + b</code>	Concatenate (join) strings <code>a</code> and <code>b</code> .

Table 1.2 A sampling of “house-keeping” functions for strings.

The syntax to display a variable with a specific format is a colon after the variable name inside the curly braces, followed by format codes that indicate how you would like the variable formatted when it is printed. The `3.2f` indicates that I'd like the number to be displayed with at least 3 digits and 2 numbers after the decimal. The `:d` indicates an integer variable and `:f` indicates a float. See Table 1.3 for a selection of other format examples.

1.11 Functions and Libraries (NumPy)

The syntax for calling Python functions is fairly standard: you type the name of the function and put parentheses `()` around the arguments, like this

```
x = 3.14
y = round(x)
```

If the function requires more than one argument or returns more than one value, you separate the arguments and outputs with commas, like this

```
x = 3.14
y = 2
b, c = divmod(x, y)
```

Pause a moment after executing this code to compare the values of `b` and `c` with the code above, and convince yourself that you understand how function calls work with multiple arguments and outputs in Python.

Python contains a set of standard functions, called native functions, that are always ready to go whenever you run Python. All the functions we've used so far fall into this category. While useful, these native functions are inadequate for scientific computing. However, user communities have created extensive collections of functions called *libraries* that you can import into Python to extend its native capabilities.

For example, Python doesn't natively include the sine and cosine functions. However, virtually any mathematical function that you will need to perform a scientific calculation can be found in a library called NumPy (say "num pie", not something that rhymes with grumpy). To use this library, we add an `import` statement to the beginning of our program, and then reference the functions in this library like this

```
import numpy

x = numpy.pi           # Get the value of pi
y = numpy.sin(x)       # Find the sine of pi radians
z = numpy.sqrt(x)      # Take the square root of pi
```

After importing the library, the code `"numpy."` before a function tells Python to look in the NumPy library to find the function rather than in the native function list. If you just type

<code>{x:4d}</code>	Display integer with 4 spaces
<code>{x:.4f}</code>	Display float with 4 numbers after the decimal
<code>{x:8.4f}</code>	Display float with at least 8 total spaces and 4 numbers after the decimal
<code>{x:1.2e}</code>	Scientific notation with 2 digits after the decimal

Table 1.3 Formatting options for printing a variable `x` in an f-string.

```
sqrt(x)
```

Python will look for a native function named `sqrt` and will give you an error.

A library can be imported and then referenced by a different name like this:

```
import numpy as np

x = np.sqrt(5.2) # Take the square root of 5.2
y = np.pi      # Get the value of pi
z = np.sin(34)  # Find the sine of 34 radians
```

This syntax tells Python “I’m going to call the numpy library `np`.” We recommend that you use this syntax for the NumPy unless there is some compelling reason to do otherwise. It is so common to do this, that in this manual we will usually omit writing the code

```
import numpy as np
```

at the beginning of each example and just assume that you’ll include it whenever you see functions with the “`np.`” prefix.

Virtually any mathematical function that you will need to perform a scientific calculation can be found in the library NumPy or in another common library called SciPy (say “sigh pie”, not “skippy”). These two libraries are designed to work together. Here is an example showing how to use the J_0 Bessel function:

```
import numpy as np
import scipy.special as sps

a = 5.5
b = 6.2
c = np.pi/2
d = a * np.sin(b * c)
e = a * sps.j0(3.5 * np.pi/4)
```

If numpy doesn’t have a mathematical function you need, then `scipy.special` probably has it. Table 1.5 provides a small sampling of its functions. See [SciPy’s online help](#) for a more extensive listing.

1.12 Numpy Arrays

In scientific computing we often do calculations involving large data sets. Say you have a large set of numbers, x_i and you want to calculate the summation

$$\sum_{i=1}^N (x_i - 5)^3. \quad (1.1)$$

You could use loops to calculate each term in this sum one-by-one, and then add them up to compute the final sum. However, the Numpy library provides a much

```
sin(x)
cos(x)
tan(x)
arcsin(x)
arccos(x)
arctan(x)
sinh(x)
cosh(x)
tanh(x)
sign(x)
exp(x)
sqrt(x)
log(x)
log10(x)
log2(x)
```

Table 1.4 A very small sampling of functions belonging to the numpy library.

<code>airy(z)</code>	Airy function
<code>j0(z)</code>	0th order Bessel function of 1st kind
<code>j1(z)</code>	1st order Bessel function of 1st kind
<code>jv(v, z)</code>	Bessel function of 1st kind, order v
<code>yv(v, z)</code>	Bessel function of 2nd kind, order v
<code>kv(n, z)</code>	Modified Bessel function of 2nd kind of integer order n
<code>iv(v, z)</code>	Modified Bessel function of 1st kind of real order v
<code>hankel1(v, z)</code>	Hankel function of 1st kind of real order v

Table 1.5 A small sampling of functions belonging to the `scipy.special` library.

slicker method for making these kinds of calculation using something called a NumPy array. For example, if x were a NumPy array containing all of the x_i values, the code to evaluate the sum in Eq. (1.1) is simply:

```
s=sum((x-5)**3)
```

This code subtracts 5 from each element in the array, cubes the result for each element, and then sums the resulting elements all without a loop in sight. Let's explore the details of how to do calculations using NumPy arrays.

Array Creation

Our first task is to create a NumPy array. Here you have several options. You can enter small arrays by providing a list of numbers to the NumPy array function, like this:

```
a = np.array([1,2,3,4,5,6,7,8,9,10])
```

Another common method for creating arrays is to use the function `arange` (think the two words “a range” not the one word “arrange” so you don't misspell this function). This function creates an array of evenly spaced values from a starting value, and ending value, and a step size, like this:

```
b = np.arange(0,10,.1)
```

This code creates the following array:

```
[0., .1, .2, .3, .4, .5, .6, .7, .8, .9, 1.0]
```

Notice that this array does not include the value at the end of the range (i.e. 10). Also, for those of you coming from a Matlab background, notice that the step size (0.1 in this case) is the third argument, not between the two limits as is done in Matlab.

The `linspace` function creates an array by specifying the starting value, ending value, and the number of elements that the array should contain. For example:

```
x = np.linspace(0,10,10)
```

This will create an array that looks like this:

```
[0., 1.111, 2.222, 3.333, 4.444, 5.555, 6.666, 7.777, 8.888, 10.]
```

When you use the `linspace` function, you aren't specifying a step size. You can ask the `linspace` command to tell you what step size results from the array that it created by adding `retstep=True` as an argument to the function, like this:

```
x,dx = np.linspace(0,10,10,retstep = True)
```

Note that since `linspace` is now returning two values (the array and the step

<code>logspace</code>	Returns numbers evenly spaced on a log scale. Same arguments as <code>linspace</code>
<code>empty</code>	Returns an empty array with the specified shape
<code>zeros</code>	Returns an array of zeros with the specified shape
<code>ones</code>	Returns an array of ones with the specified shape.
<code>zeros_like</code>	Returns an array of zeros with the same shape as the provided array.
<code>fromfile</code>	Read in a file and create an array from the data.
<code>copy</code>	Make a copy of another array.

Table 1.6 A sampling of array-building functions in numpy. The arguments to the functions have been omitted for brevity. See online documentation for further details.

size), we need two variables on the left hand side of the equals sign. Table 1.6 gives several other options for creating NumPy arrays.

Math with Arrays

Once a NumPy array object is created, a whole host of mathematical operations become available—you can square the array and Python knows that you want to square each element, you can add two arrays together and Python knows that you want to add the individual elements of the arrays, etc. Here are some examples:

```
x = np.array([2,3,5.2,2,6.7,3])      # Create the array x
y = np.array([4,8,9.8,2.1,8.2,4.5])  # Create the array y

c = x**2    # Square the elements of x
d = x + 3   # Add 3 to every element of x
e = x * 5   # Multiply every element of x by 5
f = x + y   # Add the elements of x to the elements y
g = x * y   # Multiply the elements of x by the elements of y
```

In most respects, math with Numpy arrays behaves like the “dotted” operators with Matlab matrices, where the operations are performed on corresponding elements. However, NumPy arrays are objects, so assigning one array to another, like this

```
x = np.array([2,3,5.2,2,6.7,3])
y = x
```

does not copy the values of `x` into a new variable `y`. Instead, *it creates a new pointer* to the existing object `x`, so that `x` and `y` now refer to the same object. Try changing a value of `y` and note that it also changes the value in `x`. To make a copy of an array, do this

```
x = np.array([2,3,5.2,2,6.7,3])
y = np.copy(x)
```

Switch the values in `x` and `y` and notice that they are now independent variables.

Warning: This array copy business is important and often misunderstood. This results in frustrated students when they use the wrong syntax. Please go back and re-read the last paragraph and make sure you understand the distinctions to save yourself some frustration.

Functions of Arrays

Mathematical functions like $\sin(x)$ and $\sinh(x)$ from the NumPy and SciPy libraries can accept arrays as their arguments and the output will be an array of function values. However, functions from other libraries, such as the `math` library, are not designed to work on arrays. Here is an example of this issue:

```
import numpy as np
import math
x = np.array([2,3,5.2,2,6.7])
c = np.sin(x) # Works just fine, returning array of numbers.
d = math.sin(x) # Returns an error.
```

The math library is much less capable than numpy, so we don't recommend using it for scientific coding. We just introduced it here to emphasize that NumPy arrays usually need to be used with NumPy functions.

Accessing Elements of Numpy Arrays

You can access individual array elements in a NumPy array using square brackets to index the elements like this

```
a = np.array([1,2,3,4,5,6,7,8,9,10])
x = a[1]
```

Look carefully at the value of x after running this code, and convince yourself that Python array indexes are zero-based. That is, the first element has index 0, the second element has index 1, and so forth.

You can extract a contiguous range of values using the colon command in square brackets like this

```
b = a[1:4]
c = a[0:6:2]
```

Note that there can be three numbers inside the brackets, each separated by the : symbol, like this [x:y:z]. The extracted elements start at element x, end just before element y (i.e. the element at y is not in the range), and steps in increments of z. Note that the last number is optional, and if omitted a default value of 1 is used for the step size. Python uses zero-based indexing, so a[1:4] starts at the second element, not the first.

You can also use negative indexes in the colon command to count from the end of the array toward the beginning:

```
b = a[-1]
c = a[1:-2]
```

The index -1 refers to the last element, index -2 refers to the second to last element, etc. You can also omit one of the endpoint arguments to go all the way to the beginning or end of the array, like this:

```
b = a[1:]
c = a[:3]
```

1.13 Making x-y Plots

Plots of y vs. x can be made with a library called `matplotlib`. This library creates plots using a syntax essentially the same as the plotting functions in Matlab using NumPy arrays as data inputs. Here is an example of the basic syntax:

```
import matplotlib.pyplot as plt
import numpy as np
x = np.arange(0,10,0.01)
y = x**2

plt.figure(1)
plt.plot(x,y)
```

If you want to see the actual data points being plotted, you can add the string `'ro'` inside of the plot command

```
plt.plot(x,y,'ro')
```

The `'r'` means make the data points red and the `'o'` means plot circle markers. All of the usual plot commands from Matlab work about the way you'd expect. For instance, here is an example of overlaying two graphs

```
import matplotlib.pyplot as plt
import numpy as np

x = np.linspace(0,2.*np.pi,100)
y1 = np.sin(2*x)
y2 = np.cos(3*x)

plt.figure(1)
plt.plot(x,y1,'r.',x,y2,'b')
plt.legend(['Sine of x','Cosine of x'])
plt.xlabel('x')
plt.title('Two Trig Functions')

plt.figure(2)
plt.plot(y1,y2)
plt.title('Something Fancy')
```

☞ Spyder uses IPython for its console which automatically shows plots. For other consoles, you'll need to add the command `plt.show()` after making a plot to display it on the screen.

1.14 Customizing the IDE

The Spyder development environment can be customized to suit your preferences. We recommend at least the following changes, and you can also look around at some of the other options while making these changes.

1. By default plots are displayed inline in the console. To have them pop out in their own window go to the Tools menu in Spyder and select Prefer-

ences → IPython Console → Graphics Tab and set the Graphics backend to “Automatic.” The popup windows will allow you to zoom in on different parts of the graph. The separate window will also allow us to make simple animations by repeatedly redrawing plots in the same window. If you don’t make this change, your animations will scroll by as separate plots in the console window rather than displaying on the same axis.

2. By default, all the variables that you define by direct interaction with the console and during the execution of your code remain in memory and available for subsequent use, both in the console and subsequent executions of programs. Sometimes this can be convenient, but it can also cause debugging problems. For example, if you define a variable using the console and then write a program that depends on that variable, then the mode of execution of your program depends on what you’ve manually entered in the console. When learning to program, we recommend that you have Spyder clear all variables before executing a program. To do this go to the Tools menu in Spyder and select Preferences → Run and check the “Remove all variables before execution” option.

Chapter 2

Lists, Loops, Logic, and 2D Arrays

2.1 Lists

You'll encounter lists quite a bit in Python. In fact we already used some in the previous chapter without telling you. Lists are ordered collections of values. To create a list, put the elements inside of square brackets like this:

```
x = [5.6, 2.1, 3.4, 2.9]
```

You can make a list of any variable type. Here is a list of strings:

```
a = ['electron', 'proton', 'neutron']
```

List elements don't need to be the same data type, so this list is perfectly valid

```
a = ['Ben', 90, 'Chad', 75, 'Andrew', 22]
```

You can even define a list of lists:

```
a = [[4, 3, 2], [1, 2.5, 9], [4.2, 2.9, 10.5], [39.4, 1.4], [2.7, 98, 42, 16.2]]
```

When looking at the numerical lists above, you might be tempted to think of lists as mathematical matrices. Please don't. Lists are not designed for numerical modeling. To convince yourself of this, add the statement

```
a[1][2] = 'george'
```

after the list of lists above and then examine a. You can't calculate on elements that aren't numbers! As another example, execute this code

```
import numpy as np
x = np.array([1, 2, 3]) # x is a NumPy array
print(x+x)             # adding two NumPy arrays adds the values

y = [1, 2, 3]          # y is a Python list
print(y+y)             # adding lists concatenates (joins) the lists
```

Study the output from this code and convince yourself that NumPy arrays are the right object for doing math, not lists. However, lists are otherwise quite useful in Python. Table 2.1 gives some common list functions. Note that lists are objects with methods (e.g. `a.sort()`), but can also be used as arguments to other functions (e.g. `len(a)`). Like many things in Python, the reasons for the different syntaxes are known only to those who developed Python. For example,

<code>a[x]</code>	Access element <code>x</code> in list <code>a</code>
<code>a[x:y:z]</code>	Extract a slice of list <code>a</code>
<code>a.append(x)</code>	Append <code>x</code> to list <code>a</code>
<code>a.pop()</code>	Remove the last element of list <code>a</code>
<code>len(a)</code>	The number of elements in <code>a</code>
<code>a.insert(x, y)</code>	Insert <code>y</code> at location <code>x</code> in list <code>a</code>
<code>a.sort()</code>	Sort list <code>a</code> from least to greatest
<code>filter(f, x)</code>	Filter list <code>x</code> using the criteria function <code>f</code> (see lambda functions)
<code>a.reverse()</code>	Reverse the order of list <code>a</code>
<code>a.index(x)</code>	Find the index where element <code>x</code> resides
<code>a + b</code>	Join list <code>a</code> to list <code>b</code> to form one list
<code>max(a)</code>	The largest element of <code>a</code>
<code>min(a)</code>	The smallest element of <code>a</code>
<code>sum(a)</code>	The sum of the elements of <code>a</code>

Table 2.1 A sampling of “house-keeping” functions for lists.

why does `len(a)` work, but `a.len()` return an error? The organic nature of the development of Python syntax results in some inconsistencies in syntax that just need to be learned.

2.2 The range Function

Python's `range` function represents a sequence of integer values that behaves a lot like a list. The syntax of this function is

```
r = range(10,20,2)
```

This statement represents a sequence of integers that starts at 10, ends at 18, and steps in increments of 2. The `range` function does not include the item at the endpoint of the range. In Python 2.x, the `range` function returned an actual list object, but in Python 3.x (which we are using) it returns an object called an iterator. Rather than store a huge list of numbers, an iterator just stores the rule for generating the number at any index. So, if you print the range variable

```
print(r)
```

Python will just tell you that the variable `r` represents a range, but if you execute this code

```
r[2]
```

you can get an actual value.

The `range` function can also be called with one or two arguments and default values will be assigned to the missing ones.

```
b = range(3,9) # starts at 3, ends at 8, steps of 1
c = range(10)  # starts at 0, ends at 9, steps of 1
```

2.3 For Loops and Shorthand Assignments

The primary use for the `range` function is writing `for` loops. A `for` loop iterates over each element in a list or range. Here is a `for` loop used to sum a list of values

```
s = 0
for x in [3,42,1,9.9]:
    s = s + x
print(s)
```

This code iterates over the list `[3, 42, 1, 9.9]`, assigning the variable `x` each value in the list, one by one, until it reaches the end of the list. Notice the required colon at the end of the `for` statement that marks the beginning of the loop, and also that there is no code indicating the end of the `for` loop. **Python uses indentation to**

☞ Usually you would just use a `sum` function, but we are learning about loops here.

denote the lines of code that will be executed during each iteration. To see this, indent the `print(s)` statement to the same level as the statement above it and execute the code. Note that the `print` function now executes on each iteration of the loop instead of just displaying the final value of `s` after the loop is finished.

You can also iterate over the values from the `range` function, like this

```
for y in range(5,50,3):
    print(y)
```

The `range` function only produces integers, so the most common use case is to use it with a `for` loop to index a separate array, like this

```
x = np.arange(5,10,.1)

for n in range(len(x)):
    print(x[n])
```

You can also iterate directly over the values of an array, like this:

```
x = np.arange(5,10,.1)

for y in x:
    print(y)
```

To practice `for` loops, let's build a loop to calculate the sum

$$\sum_{n=1}^{1000} \frac{1}{n^2} \quad (2.1)$$

```
s = 0
for n in range(1,1001): # use 1001 to go to 1000. Last value omitted
    s += 1/n**2
print(s)
```

The code `s += 1/n**2` is equivalent to `s = s + 1/n**2`, but runs a little faster and is a little bit easier to read (once you get used to it). See Table 2.2 for more shorthand notations. Here's another example of a loop used to calculate the value of $p = 20!$ (20 factorial) using another shorthand notation

```
p = 1
for n in range(1,21): # use 21 to go to 20
    p *= n # Multiply p by n
print(p)
```

Remember that the `range` function starts at the first argument (i.e. 1), and goes up to, but not including the second argument, so our product will only go up to $n = 20$.

Operation	Shorthand
<code>a = a + 1</code>	<code>a += 1</code>
<code>a = a - 2</code>	<code>a -= 2</code>
<code>a = 5*a</code>	<code>a *= 5</code>
<code>a = a/c</code>	<code>a /= c</code>
<code>a = a % 10</code>	<code>a %= 10</code>
<code>a = a**3</code>	<code>a **= 3</code>
<code>a = a // 12</code>	<code>a //= 12</code>

Table 2.2 A list of shorthand variable reassignment notation.

2.4 Logical Statements

Often we want to run a section of code only when some condition is satisfied. This requires the use of logic. The simplest logic is the `if ... elif ... else` construct, which works like this:

```
a = 1
b = 3
if a > 0:
    c = 1
else:
    c = 0

if a >= 0 or b >= 0: # If condition 1 met
    c = a + b
elif a > 0 and b >= 0: # If condition 2 met
    c = a - b
else: # If neither condition is met.
    c = a * b
```

Note the locations of colons in these statements, and also note that indentation again defines the regions of code that execute as a group. Table 2.3 lists the standard elements for constructing logical statements.

A word of caution about comparing Python floats is in order here. Because of the way floats are represented in a computer, the number that is stored in a float is often not *exactly* the number that you think it is. For instance, execute the following, and study the output:

```
a = 0.1
b = 3 * a
c = 0.3
print(b==c) # Are they the same number?
print(b)    # It sure looks like they are the same.
print(c)    # It sure looks like they are the same.
print(f' {b:.45f} ') #b, to 45 decimal places
print(f' {c:.45f} ') #c, to 45 decimal places
```

This can cause problems when you are comparing two numbers that you think should be equal but actually aren't equal in the computer. The take home here is that *comparing two floats directly to see if they are equal is a bad idea*. A better way to check to see if two floats are equal (or close enough that we can say they are equal) is to check if the absolute value of their difference is very small, like this:

```
a = 0.1
b = 3 * a
c = 0.3
print(abs(b - c) < 1e-10)
```

<code>==</code>	Equal
<code>>=</code>	Greater than or equal
<code>></code>	Greater than
<code><</code>	Less than
<code><=</code>	Less than or equal
<code>!=</code>	Not equal
<code>and</code>	True if both conditions joined by <code>and</code> are true
<code>or</code>	True if either of the conditions joined by <code>or</code> are true
<code>not</code>	True if the following condition is false.

Table 2.3 Python's logic elements.

2.5 While Loops

A `while` loop iterates until a certain condition is met. This loop is a good choice when you don't know beforehand exactly how many iterations of the loop will be executed, but rather what condition you want to be met before finishing. As an example, let's compute the sum

$$\sum \frac{1}{n^2} \quad (2.2)$$

by looping until the terms become smaller than 1.00×10^{-10} .

```
term = 1 # Load the first term in the sum
s = term # Initialize the sum
n = 1 # Set a counter
while term > 1e-10: # Loop while term is bigger than 1e-10
    n += 1 # Add 1 to n so that it will count: 2,3,4,5
    term = 1./n**2 # Calculate the next term to add
    s += term # Add 1/n^2 to the running total
```

This loop will continue to execute until `term` is less than 10^{-10} . Note that indentation again defines the extent of the code that will execute each iteration. Unlike the `for` loop, you have to do your own counting in a `while` loop. Be careful about what value `n` starts at and when it is incremented (`n+=1`). Also notice that the variable `term` must be assigned prior to the start of the loop. If it isn't, the loop's first logical test will fail and the loop won't execute at all.

The `break` statement

Sometimes `while` loops are awkward to use because you can get stuck in an infinite loop if your check condition is never true. The `break` command is designed to help you here. When `break` is executed in a loop the script jumps to just after the end at the bottom of the loop. The `break` command also works with `for` loops. Here is our sum loop rewritten with `break`

```
term = 1 # Load the first term in the sum
s = term # Initialize the sum
n = 1 # Set a counter
while term > 1e-10: # Loop while term is bigger than 1e-10
    n += 1 # Add 1 to n so that it will count: 2,3,4,5
    term = 1./n**2 # Calculate the next term to add
    s += term # Add 1/n^2 to the running total
    if n > 1000:
        print("This is taking too long. I'm outta here...")
        break
```

☞ If you are stuck in a loop, you can force your program to stop by pressing Ctrl+C or Cmd+C

The `continue` statement

Another statement that is used with loops is the `continue` statement. When `continue` is used, the remainder of the code for the current iteration is skipped and the next iteration of the loop begins. If you wanted to do the sum in Eq. (2.2) but only include those terms for which n is a multiple of 3, you could do this

```
term = 1 # Load the first term in the sum
s = term # Initialize the sum
n = 1 # Set a counter
while term > 1e-10: # Loop while term is bigger than 1e-10
    n += 1 # Add 1 to n so that it will count: 2,3,4,5
    if n % 3 != 0:
        continue # Skip the rest of the code, start the next iteration
    term = 1./n**2 # Calculate the next term to add
    s += term # Add 1/n^2 to the running total
```

Now, when the value of n is not a multiple of 3, the last two lines of code in the loop will be skipped.

2.6 Two-Dimensional NumPy Arrays–Matrices

We introduced one-dimensional NumPy arrays in the last chapter. The syntax for two-dimensional arrays (matrices) builds on this foundation. Small matrices can be built by giving the array function an argument that is a list of lists, like this

```
a = np.array([[1,2,3],[4,5,6],[7,8,9]])
```

which could be interpreted as this matrix:

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \quad (2.3)$$

Once you have a 2D array, single elements can be extracted with this syntax:

```
b = a[0,2]
```

This code extracts the number 3, the element in the first row and third column in the array `a`. If you wanted to slice out the following 2×2 sub-matrix:

$$\begin{pmatrix} 5 & 6 \\ 8 & 9 \end{pmatrix} \quad (2.4)$$

you could do it like this:

```
b[1:3,1:3] # Slice out a sub-array
```

If you want all of the elements in a given dimension, use the `:` alone with no numbers surrounding it. For example, the following:

```
b[:,1:3]
```

would extract all of the rows on columns 1 and 2:

$$\begin{pmatrix} 2 & 3 \\ 5 & 6 \\ 8 & 9 \end{pmatrix} \quad (2.5)$$

This kind of slicing can't be done with lists.

Creating Larger Matrices

You can create larger arrays using several functions from NumPy (see Table 1.6). For example, to create an array full of zeros, you use this syntax:

```
A = np.zeros((100,200))
```

This creates an array with 100 rows and 200 columns, completely full of zeros. The argument that you provide to the `zeros` function, a set of numbers surrounded by round parentheses, is a data type called a *tuple*. Tuples behave a lot like lists, except they are immutable (i.e. after you create a tuple, you can't change the item values). We won't use tuples much in scientific computational, but you should be aware of what they are since you periodically need to create them as arguments for functions. You can also ask a NumPy array what size it is using the array's `shape` property, like this

```
s = A.shape
```

Notice that the `shape` property returns a tuple.

To create an array full of ones, you use this syntax:

```
B = np.ones((100,200))
```

You can also create an identity matrix like this

```
I = np.identity(100)
```

Since identity matrices must be square, it only requires one number as an argument rather than a tuple.

Math with Matrices

Once a 2D array is defined, it behaves as a matrix and you can perform all the standard linear algebra operations with SciPy's linear algebra module. The matrix multiplication operator is the `@` symbol. Here are a few examples illustrating matrix operations:

```
import numpy as np
import scipy.linalg as la
```



```

a = np.array([[1,2],[3,4]]) # Create 2 x 2 matrix
b = np.array([[5,6],[8,9]]) # Create 2 x 2 matrix
col = np.array([[3],[4]]) # Create 2 x 1 column vector

c = a.T # Transpose the matrix
e = a.conj().T # Find conjugate transpose of matrix
f = a @ b # Matrix multiplication
g = b @ col # Multiply matrix b to column vector
d = la.inv(a) # Find inverse of matrix

```

The `linalg` name space within SciPy provides lots of matrix functions. Pretty much anything you learned about in your linear algebra class is available there.

☞ NumPy also has a linear algebra module. The SciPy version can do all that the NumPy version can, plus it has more advanced functions. In addition, the SciPy linear algebra package has some optimizations to make it run faster. We recommend you always use the SciPy linear algebra package.

2.7 Solving a Set of Linear Equations

To illustrate the use of matrix computations, let's solve a set of linear equations. Here is an example of a set of two linear equations, with two unknowns:

$$3x + y = 9 \quad x + 2y = 8 \quad (2.6)$$

This problem can be represented in matrix form like this:

$$\mathbf{Ax} = \mathbf{b} \quad (2.7)$$

where

$$\mathbf{A} = \begin{pmatrix} 3 & 1 \\ 1 & 2 \end{pmatrix}, \quad (2.8)$$

$$\mathbf{b} = \begin{pmatrix} 9 \\ 8 \end{pmatrix}, \quad (2.9)$$

and

$$\mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix}. \quad (2.10)$$

SciPy has a function called `solve` that will solve this problem as follows:

```

import numpy as np
import scipy.linalg as la

a = np.array([[3,1],[1,2]])
b = np.array([[9],[8]])
x = la.solve(a,b)

```

2.8 NumPy Matrix

NumPy provides another object called a matrix that has been widely used for linear algebra to represent matrices. The matrix object is a subclass of the NumPy

array and provides a little cleaner syntax for manipulating matrices than the array syntax. However, as of Python 3.5, the developers of NumPy have indicated that they plan to deprecate the usage of the matrix type, so we haven't taught it here. But you may still see it in existing code bases. It is essentially a wrapper around the NumPy array object to make some linear algebra syntax cleaner. All of the functionality of the matrix object is available with NumPy arrays with some minor syntax changes.

Chapter 3

Functions

3.1 User-defined functions

Python's ecosystem is extremely rich, and you can usually find a pre-made library to do common manipulations. But you will also need to write your own functions. User-defined functions are created like this

```
def myFunction(a,b):  
    c = a + b  
    d = 3.0 * c  
    f = 5.0 * d**4  
    return f
```

The function above performs several simple calculations and then uses the `return` statement to pass the final result back out of the function. User-defined functions must begin with the keyword `def` followed by the function name (you can choose it), then a list of arguments enclosed in parentheses, and finally ending with a colon. Python does not use an end statement to signal the end of a function. Rather, it looks for indentation to determine where the function ends, just like it did with loops and logic.

You can integrate this function into a larger program like this

```
# This code just defines the function.  
# The function is not executed until it is called below.  
def myFunction(a,b):  
    c = a + b  
    d = 3.0 * c  
    f = 5.0 * d**4  
    return f  
  
#The rest of this code is not part of the function.  
r = 10  
t = 15  
x = myFunction(r,t)
```

In this case, when the function is called, the input variable `a` in the function definition gets assigned the value of 10 and input `b` gets assigned the value of 15. The result of this calculation (the variable `f` in the definition) is passed out of the function and stored in the variable `x`.

A word on local and global variables is in order here. In the example above, the variables `a`, `b`, `c`, `d`, and `f` are *local variables*. This means that these variables are used internally by the function when it is called and then immediately forgotten.

To see what I mean, add the following `print` statement at the end and observe the results

```
x = myFunction(r,t)
print(c)
```

The `print` statement just causes an error since Python does not remember that inside the function we wrote `c=a+b`.

In contrast, the variables `r`, `t`, and `x` are *global variables*, which means that Python remembers these assignments anywhere, including inside of functions. So, technically, you could do the following:

```
g = 9.8          # <--- g defined to be a global variable
def myFunction(a,b):
    c = a + g    # <--- Notice the reference to g here
    d = 3.0 * c
    f = 5.0 * d**4
    return f
#The rest of this code is not part of this function.
r = 10
t = 15
x = myFunction(r,t)
```

and there would be no error. Notice that `g` has been defined as a global variable, and the function `myFunction` knows its value and can use it in a calculation. Using global variables inside functions is usually considered bad practice and can confuse those reading the code. In general, every variable used in a function ought to be either passed in or defined inside of the function.

3.2 Importing User Defined Functions

If you write some functions that you find yourself using over and over again, or if you've written so many functions for a program that it makes your program hard to read, you can save your functions in a separate file and import them just like a Python library.

As an example, create a blank python file and type a couple of functions to calculate some of the parameters of projectile motion that you learned about in Newtonian physics:

```
import numpy as np

def maxRange(v0,theta):
    R = v0**2 * np.sin(2*theta) / 9.8
    return R

def maxHeight(v0,theta):
    vy = v0*np.sin(theta)
    h = vy**2 / (2*9.8)
```

```
    return h
```

Then you save the code above in a file called `projectile.py`. Now create another file in that same directory, and practice importing your projectile functions using the following four methods (just use one method at a time):

```
import numpy as np

#Method #1
import projectile
projectile.maxRange(10,np.pi/4)

#Method #2
from projectile import maxRange
maxRange(10,np.pi/4)

#Method #3
import projectile as pf
pf.maxRange(10,np.pi/4)

#Method #4
from projectile import *
maxRange(10,np.pi/4)
```

Using these methods, you can organize your functions into neatly divided units that you can reuse in multiple programs. By default Spyder doesn't give you great hints about user-defined functions. You can change this by going to Tools→Preferences→Help, and clicking the boxes under "Automatic connections."

It is possible to save your user-created libraries in a different folder, and then map to them during your import, or even make your functions available to any program using Python on your computer. When you are ready to work on larger projects where these techniques are important, you can learn about modules and default search paths online. However, to keep things simple in this course, just keep your files together in the same folder.

3.3 Writing Readable Code

You likely have had several programming courses before this one, so I'll spare you the lecture about commenting your code. For your own sake, and for the sake of those who may someday need to read your code, comment your code a little more than you think is necessary at the time you write it.

Comments

Any text written after the `#` symbol is ignored by Python as a comment:

```
# A full line of commenting
E=m*c**2 # comments after a short line of code
```

Docstrings

Python also has an environment like a comment that is called a docstring, delimited by a triple quote (either single or double):

```
"""
    This is a docstring. It can be automatically parsed to generate
    documentation. It is not the same as a comment.
    """
```

Docstrings are placed immediately after the beginning of a module or a function definition, and can be used to automatically generate documentation for your code. You'll see some people use them as a multi-line comment, but this is considered bad practice. If you just want to comment your code, use the `#` symbol in front of multiple lines. Docstrings are important when writing a large codebase to keep things well documented, but we won't use them much in this course.

Line Continuation

To keep your lines of code from becoming too long, Python provides two ways of continuing a logical line of code on the next physical line. The preferred method is an implicit continuation which happens when you put a line break in your code before closing all parentheses, like this:

```
x = (a + b
     + c)
```

Normally you wouldn't wrap lines of code this short, but this is just an example. You should indent the continued line of code so that it starts right after the unclosed opening parenthesis. You can almost always wrap lines of code this way simply by added parentheses. You can also use the backslash "`\`" to continue a line of code, like this

```
x = a + b \
     + c
```

if you don't like adding parentheses.

Code cells

Spyder will allow you to split your code into individual cells that can be run separately, without executing the whole program like this

```
%% This is the start of a cell
x=5
y=6
print(x*y)
%% This is the start of another cell
a=3
```

```
b=4  
print(a/b)
```

If you place your cursor in one of the cells and press Ctrl+Enter (or push the button to the right of the green arrow), Spyder will execute just the code in the current cell. We generally discourage writing code using cells. When you do, the current state of your program depends on which cells you've previously executed. This situation can be painful to debug and it is usually best to just avoid it for the type of coding we do in this class. However, there are times that cells can be useful.

Index

2D arrays, 18

Anaconda, 1

arrays, 7

assigning values, 3

boolean variables, 5

case sensitive, 3

data types, 3

floating point variables, 4

`for` loops, 14

help, 4

IDE, customizing, 11

integers, 3

linear algebra, 20

lists, 13

logic, 16

matrix operations, 19

NumPy, 6

plotting, y vs. x , 11

`range` function, 14

SciPy, 20

shorthand notation, 15

spyder console, 2

strings, 5

strings, formatting, 5

`while` loops, 17