Durham University

Department of Physics

Introduction to Programming in Python

A Self-Study Course

(Version 4.1.4 – November 2014)

Table of Contents

1 Chapter 1 - Introduction

The aim of this course is to introduce you to the development of computer programs and to provide you with a working knowledge of the Python language.

In addition to everyday usage, physicists use computers for:

- Controlling apparatus and taking readings automatically

- Processing data to extract conclusions

- Predicting the results of experiments

- Simulating systems that are too complex to describe analytically

All of these generally require using a computer programming language. The latter two can be described as 'Computational Physics'. There are computer 'applications' that allow you to address some of these problems – you will already have used Excel extensively. However, to obtain the flexibility to solve a particular problem the way YOU want to do it, you must write your own dedicated 'application' using a computer programming language. Thus skills in programming are extremely important to the modern physicist. They will be essential to you in later stages of your degree. This course provides you with an introduction to programming that should enable you to develop the more extensive skills required in this and later years. Such skills are also much sought after by a wide variety of employers.

There are many programming languages and much argument over which is 'best'. The skills that you acquire in this course should be easily transferable from the language that you will use, Python, to other languages. When preparing a course such as this, there is always a decision to be made over which programming language to use: in this case we have chosen Python, a relatively new (dating from the 1990s) but widely used language. Although not originally designed for scientific use, it is a powerful modern language, easy to learn and ideal for the rapid development of programs.

Python is widely used. For example, the installation scripts for many Linux distributions are written in Python, as is the BitTorrent software. Google uses Python extensively, as do Industrial Light and Magic, the company behind the special effects for Star Wars and many subsequent movies.

Python is interpreted, which means that the computer runs another program to read your input and then execute your instructions. Programs written in languages like FORTRAN or C++ are compiled, meaning that the computer translates them into a more basic machine-readable form that can be executed much faster. The speed disadvantage of interpreting will not be an issue in this course. In fact, not having to compile and link your code before running it will save you considerable time. Nevertheless, you should be aware of this fundamental difference between interpreted and compiled languages. Python's high-level language features also mean that programs will typically be much shorter than equivalent ones written in C or C++.

Python can be linked to compiled code from other languages, which can largely get over the speed penalty. This is done, for example, in the NumPy module that we will encounter during this course. NumPy provides fast vector and matrix handling, as well as Python interfaces to fast routines for tasks like integration, solving differential equations and linear algebra. If you become a Python expert, you'll be able to make links to your own compiled code.

Whilst one can use Python, rather like a powerful programmable calculator from the python prompt, you will soon want to write many lines of program and save them to a file. You then

tell Python to execute the file. For this you need an editor. You can use any ordinary text editor (NOT a word processor!!). It is however best to use an editor that knows the Python syntax. We will use notepad++ but there is also a very good editor bundled with the Enthought distribution of Python that you can download from the web.

 One advantage of Python is that, although it is a complex and rich language, it is easy to learn the basics. Once you have a grounding in Python, you can extend your knowledge later. Some topics in this guide are marked thus: *Advanced Topic*. They are not essential for this course. You can skip over them and come back to them later if you have time.

In this course you will be learning 'procedural programming'. That involves writing a set of functions to do what you want and arranging them in files called 'modules'. There is an additional technique called 'Object Oriented Programming' that you will hear mentioned and can read about in the recommended course textbook, Hetland (see section 2.6). That is an *Advanced Topic* and is not required in this course. However, Python is an object oriented language so such techniques will be mentioned both here, in the lectures and in the course textbook.

Procedural programming allows you to split your program into many small manageable functions. When you are writing programs to solve Physics problems, try to use functions to separate different bits of functionality within the program. In particular, keep clear the distinction between your algorithm or method of solution and the overall program. You can then easily change the algorithm function without affecting the way your main code works. An 'algorithm' is a procedure in the abstract sense, in that it's content is not linked to any particular way of expressing or implementing it.

All programming languages implement the same basic constructs, loops, decision-making and so on. Once you have seen them in one language (Python) it will be much quicker to pick up a different language later if you need to.

If you are keen to learn a compiled language, either now or later in your degree, the department maintains similar course notes for both Fortran and C++. These can be made available to you on request.

Note 1: Throughout this document, we have used the Courier font, which looks like this: `This is in the Courier font,` to represent Python code or input and output from the computer. Where there is a substantial block of code, it is enclosed in a box (or boxes if it runs over a page boundary).

Note 2: These course notes, and the associated examples and exercises, attempt to be both a tutorial and a reference document. DO skip the *advanced topics* on your first run through. DO look ahead to the later chapters when you need more detail on a subject.

2 Chapter 2 - Resources Required for the Course

2.1 Programming Language

Python!!

Python is 'open source' and free. You don't need a license to run it. It will run on most computer operating systems. It is available on the ITS networked PCs.

2.2 Computer Operating Systems

Python and all of the libraries that you need have been installed on the CIS networked PCs under both the Windows7 and the Linux operating systems. We will use Windows7 during the workshops. We will work in a terminal window; just type `python` at the Windows7 prompt.

Python is available for most operating systems including Windows, Linux and MacOS and can be easily installed on your own computer.

2.3 Additional Libraries (Modules)

You are learning to program in order to use computers to do, amongst other things, Computational Physics. Python does not provide sufficient facilities for this purpose built into the language. However, adding 'modules' to Python is easy, as you will find out when you add your own! There is a large number of modules (or libraries) available for Python. We will use only one of these extensively:

- NumPy – This provides fast vector and matrix handling. It extends the Python language to provide the data types that a Physicist needs – numpy arrays. Python is not a useful language for Physics problems without it!

We will briefly mention two further libraries that are very useful for computational physics:

- Matplotlib or Pylab – This is a standard plotting package for drawing graphs. (Import pylab – it will import the lower level matplotlib for you).

- SciPy: Provides large numbers of fast routines for scientific programming such as integration, solving differential equations and linear algebra

These modules have been installed for you on the CIS Windows7 service and are also available within the Enthought Python Distribution (EPD - see section 2.5).

The course textbook, Hetland, does NOT cover NumPy. Yes, this is annoying. I have not yet found a good Python textbook that does. Python is used by many types of people and only a few (Physicists etc) need NumPy.

There is good documentation for NumPy at: www.numpy.org and for SciPy at: http://www.scipy.org/.

Importing these or any other module into Python is easy:

EG   `import numpy  # Import the standard module numpy`

Or   `import mycleverlib  # Import my own module`

Not really too hard! We will return to the use of modules later in the course.

Note 1: All the text after a `#` on each line is treated as a comment and is not executed

Note 2: I describe modules with capitalised names like NumPy for clarity. When importing them, you must use their correct names (usually all lower case) EG: import numpy.

2.4 Editors

You can do a lot with Python just from the command prompt within a terminal window. However, at some point you will want to save loads of Python lines as a 'program', sometimes called a 'script'. To do this, you will need an editor.

You can use any ordinary text editor (NOT a word processor!!). It is however best to use an editor that knows the Python syntax. We will use **notepad++** on the CIS networked PCs but there is also a very good editor bundled with the Enthought distribution of Python that you can download from the web. There is also an editor called IDLE that is often bundled with Python.

To access the notepad++ editor on the CIS networked PCs, just type `notepad++` at the Windows7 command prompt. When you do this, it is best to first move to the folder in which you want to work. You will learn how to do this at the first workshop. By default, notepad++ will load and save your files from the directory you were in when you invoked it (directories are also known as folders). Your files are **stored on the J: drive** that can be accessed from any CIS PC.

2.5 Where to do the Work

Python will run happily on most operating systems. You could do all of the course work on a *Linux* machine or a Mac if you wish. However, you are expected, as a part of the course, to learn to run Python programs under Windows7. You should therefore start the course work using Windows7 on the networked PCs in the classroom where the workshops are held.

You are encouraged to transfer work to your own PC or laptop in order to continue to work anywhere. Whether your PC runs MS Windows, MacOS or Linux, you can download an excellent version of Python, including all of the additional modules you need for the course, from '**Enthought**'. The latest full version of Enthought Canopy can be downloaded only if you do so from the dur.ac.uk domain (IE From somewhere inside the university). You will have to sign up for an account to get an academic license. See:

https://www.enthought.com/products/canopy/academic/

You can download the 'light' version, Canopy Express, for free from anywhere and this will be adequate for this course. However, we recommend getting the full version to cover more advanced work in future years. Canopy Express is here:

https://www.enthought.com/products/epd/free/

Note that most MacOS and Linux distributions have Python pre-installed because it is used by the operating system. However, you are unlikely to find NumPy pre-installed.

2.6 Books

You don't need a book to learn Python! Just sit at a computer and write programs. Having said that, most programmers keep a good Python book nearby for reference. The first port of call however is the web. The Python documentation is very extensive. See: http://www.python.org/doc/. **This is the font of all knowledge for Python**. Be sure to use the documentation for the version of Python that you are using. **Do not use version 3 of Python**. The course is based on Python version 2.7.

You can download much of the documentation to your computer should you wish.

We recommend two books:

- Beginning Python by Magnus Lie Hetland. This is an excellent introduction to the language with loads of examples. (But does not cover NumPy).

- Python in a Nutshell by Alex Martelli. This is more of a handy reference book.

If you just want one good book to use, get 'Beginning Python'. We will refer to it in these notes as 'Hetland'. The course partly follows the early chapters of Hetland. In fact, chapters 1 to 11 (but not 7 and 9 which cover object oriented techniques) match the course well. However, Hetland does **not** cover scientific computing using NumPy and SciPy which is included in this course. You will need to use on-line documentation for these.

Note: There is a useful reference for much of the more commonly used aspects of Python in Hetland in appendix B.

3 Chapter 3 - Getting Started

As well as doing this course, you may also like to put Python on your computer and try the Python tutorial online at:

http://docs.python.org/tut/tut.html

(Try section 3 first).

Log on to the CIS Windows7 service and bring up a terminal. If there is no shortcut to this application on the desktop or start menu, it can be found at:

Start -> All Programmes -> Accessories -> Command Prompt

At the terminal prompt (also known as 'command line'), type `python.` You will get a new prompt in the form: `>>>.` The Python interpreter is now running. **To exit from Python, type <ctrl> d (hold down the ctrl key and press d)** or call the exit function: `exit().`

Try doing some calculations at this prompt using Python like a fancy calculator. We will do this in the first workshop.

As you learn the language, always have a Python prompt available. If you learn something new, try it at once!

Note: Much of this chapter is covered in more detail in chapter 1 of Hetland.

3.1 Numbers

The normal way to program in Python is to write a program in an editor and save it to a file. We then tell Python to run that file.

However, before we write any programs, we will start by using Python from the command line. It is rather like using a fancy calculator. Even after you have written a load of code, it is still very useful to quickly try things out at the command line.

Try this example:

```
>>> (2+4+6)*3-12/3
32
>>> 8**2 # ** is used for power 2.
64
```

It works! Note the use of a # to introduce a comment. The interpreter ignores all of the rest of the line after the #. Comments are not needed when you use the interpreter at the command line but are essential when you write a long program in a file.

These were integers. Let's try it with real numbers:

```
>>> (2.0+4.0+6.0)*3.0-12.0/3.0
32.0
>>> 8.0**2
64.0
```

This looks fine. But now try:

```
>>> 7.0/4.0
1.75
```

```
>>> 7/4
```

```
1
```

Not so fine! The lesson is that computers distinguish between exact integer arithmetic and real number (or floating point) arithmetic. In particular, whenever dividing two integers does not result in an integer, the result is truncated (rounded down) to the nearest integer. If you mix integer and floating point numbers, they are all converted to floating point by Python:

```
>>> 34+1.0
```

```
35.0
```

```
>>>
```

You can also convert from integer to float and back again, but note that conversion to integers is done by *truncating*

```
>>> float(1)
```

```
1.0
```

```
>>> int(3.678)
```

```
3
```

Python will do lots of auto converting for you that is very convenient. But beware – sometimes it will not do exactly what you wanted. It is good practice, if you want a floating point number to always include the decimal point. Thus `12.3+1.0` is better style than `12.3+1`.

Advanced Topic:

If you would like Python to NOT use integer division but to convert for you and give the answer you might expect, you can tell it to do so. How to do this is explained in Hetland chapter 1. You are using features from future versions of Python!

```
>>> from __future__ import division # Double underscores!
```

```
>>> 7/4
```

```
1.75
```

```
>>> 7//4  # Force python to do proper integer division
```

```
1
```

3.2 Assignments, Strings and Types

You can assign values to **variables.** The computer will save the value in some memory pointed to by the variable. If you do this, you won't see the value printed on the screen unless you ask for it. If you just type the variable name, Python will return the value (or you can use the print command:

```
>>> x=3
>>> x
3
>>> y=2
>>> z=7
>>> x*y*z
42
>>> a=x*y*z
```

```
>>> print a
42
```

Remember that the = sign is used for *assignment* in Python. This is NOT the same as its use in maths. Thus x=x+1 makes perfect sense. It means increment the value of the x variable by 1. It is NOT a maths equation!

Variables can be of types other than integer and float. For example, they can be strings of characters; in other words text. A string is a special case of a list object as we will see later. String handling is easy in Python. It is well explained in chapter 3 of Hetland.

A string is always enclosed in either single or double quotes (But not a mixture of both!).

Strings can be added (concatenated or joined):

```
>>> "spam"
'spam'
>>> "spam," + " " + "eggs and spam"
'spam, eggs and spam'
```

Strings can be duplicated by multiplying them by an integer:

```
>>> "Spam, "*4 + 'beans and spam'
'Spam, Spam, Spam, Spam, beans and spam'
```

Python does not require you to be strict about what type a variable is (the language is 'weakly typed'). However, this is sometimes important to a Physicist. If you must, you can find the type of a variable thus:

```
>>> a=1; b=1.0; c='1.0'
>>> type(a)
<type 'int'>
>>> type(b)
<type 'float'>
>>> type(c)
<type 'str'>
>>> print a,b,c
1 1.0 1.0
```

You can convert a string to a number (within reason) and visa versa:

```
>>> str(12.3)
'12.3'
>>> float('12.3')
12.300000000000001
```

Note that we didn't quite get what we expected from the float function. The computer has a limited precision. It stores the nearest number to what you want within its ability. We will come back to this issue of *precision*.

Note that you can put several short statements on the same line separated by semi colons. It is very poor style to do this for more than very simple lines as above.

3.2.1 A First Mention of Functions

The use, and creation, of functions will be handled later in this course. However, we have already mentioned the `str()` and `float()` functions above.

Python will provide you with a huge number of functions that are already written. Some of these are 'built-in' to the language like those above. Others are provided by modules that you 'import'. We will discuss functions in detail later on. For now, just note the format of a function call:

result = function(arguments)

For example, we can use 2**3 to raise 2 to the power 3. This is really just a shorthand for the function `pow()`:

```
>>> 2**3
8
>>> pow(2,3)
8
```

Note 1: The use of brackets of type () for function calls. Other shapes of brackets mean different things in Python. *Don't muddle them up!* We will come across [] and {} later on.

Functions are fundamental to writing a program. They are like a 'black box'. Once written and tested, you can just call them knowing what they do. Every function may (or may not) require some input and may (or may not) return some output to you.

3.2.2 A Brief Mention of Methods

We will make brief mention of 'methods' here and there. Methods are a part of Object Oriented Programming which we are not considering in this course. We will use functions throughout the course rather than methods. (A method is just a function that is bound to an object). However, you will need to use some methods that are provided by Python so you must learn the syntax.

They can be thought of as an alternative way of calling some already-written code and take the form:

Result = object.method(arguments). Notice the . (dot) in the syntax.

An example is appending to a list (We will talk more about lists in section 6.1):

```
>>> a=[0,1,2] # Make a list
>>> a.append(3) # Use the append() method
>>> print a
[0, 1, 2, 3]
```

3.3 Complex numbers (*Advanced Topic*)

This is a type not often needed by an everyday Python user but much beloved of Physicists!

Python supports complex numbers. They are written in the form 3+5j, corresponding to 3+5i in usual mathematical notation. Here are some simple examples using complex numbers:

```
>>> d=3+5j
>>> type(d)
```

```
<type 'complex'>
>>> d.real
3.0
>>> d.imag
5.0
>>> abs(d) # absolute value or magnitude or modulus
5.8309518948452999
>>> e=1+1j
>>> print d*e
(-2+8j)
```

Note1: We used `d.real` to get the real part of a complex number. This is a way of getting an 'attribute' of the variable. It is a first hint of what is called 'object oriented programming'. This is not a part of this course but will be mentioned now and again. For now, just remember the syntax.

Note 2: There is no separate type for imaginary numbers in Python. They are treated as complex numbers whose real component is zero.

3.4 Errors and Exceptions

By now, you will have made some mistakes and will have seen your first Python 'Traceback'. Be aware that errors are referred to as 'exceptions' in Python.

For example:

```
>>> float('gumby')
Traceback (most recent call last):
  File "<pyshell#90>", line 1, in <module>
    float('gumby')
ValueError: invalid literal for float(): gumby
```

When an error occurs in a program, it may occur in a function that is called from another function in another module etc. The traceback gives very valuable information about the error that occurred AND where it occurred (eg the line number in a file).

In this case, the error was at the Python prompt (the Python shell). We supplied a string to the function `float()` which it really can't convert.

You may also get syntax errors. This is where you have written something that Python simply doesn't understand. For example:

```
>>> str(1.23a)
SyntaxError: invalid syntax
```

Now let's look at a more detailed traceback:

```
Traceback (most recent call last):
  File "\\elite\home\nad\My
Documents\AllDocs\teaching\Computing\2009\TestPrograms\test01.py", line 7,
in <module>
    T.theTest(theData)
  File "\\elite\home\nad\My
```

```
Documents\AllDocs\teaching\Computing\2009\TestPrograms\testModule.py", line
4, in theTest
    y = x.fred() # Cause an error
AttributeError: 'numpy.ndarray' object has no attribute 'fred'
```

The traceback starts with what was happening at the top level. An error occurred in my main program (called <module>) at line 7. I am told the name of the file and the line number. However, there is more… The actual error was in another module file (name given) that contains the function called 'theTest()' that I called from my main program. The line that caused the error is identified and the cause of the error is given. The function 'fred()' does not exist (No such attribute). The error descriptions become clearer as you get to know the terminology.

These tracebacks provide you with the main information you need to 'debug' your program. IE Find out what went wrong. Debugging is discussed further in chapter 9.

3.5 Precision and Overflow

As physicists, we are very often using floating point numbers (real numbers) to do calculations. There are some pitfalls with using floats of which you should be aware.

3.5.1 Precision

Computers have a limited dynamic range for storing numbers and also, for real numbers, a limited precision. This can be a significant problem in computational physics where the 'rounding errors' in real numbers may add up to a large error in your result.

As a trivial example:

```
>>> 1.0/3.0
```

```
0.33333333333333331
```

What is that 1 doing on the end?

Python has done its best but cannot give the answer to an infinite number of decimal places. The precision of Python floating point numbers is about 1 part in 10^{16}.

3.5.2 Overflow – Large Numbers

Python can store really large numbers. At a certain point it will automatically start using `long integers.` These require more storage.

Try:

```
 >>> 1000000*1000000
```

```
1000000000000L
```

The L on the end shows that Python has used a long integer but you don't really need to know.

For real numbers, Python can store enormous numbers. It will eventually give an error when the exponent is just too large:

```
>>> 1e200**2
Traceback (most recent call last):
  File "<pyshell#119>", line 1, in <module>
    1e200**2
```

```
OverflowError: (34, 'Result too large')
```

Note 1: We use ** to mean 'raise to the power'

Note 2: We can use 1e200 to represent the number 10^{200}.

3.6 Getting Help

There is a LOT of online help for Python. The full documentation set for Python is online at:

http://www.python.org/doc/.

There is a built-in help system. Use the function `help()`. Enter a Python keyword or function name between the brackets. Try just typing `help()` at the Python prompt to get started.

This help system depends on programmers putting 'documentation strings' into their programmes. You will be expected to do the same with the functions and modules that you write.

4 Chapter 4 - Output to the screen and file input and output

This chapter describes how to write data to the screen from Python and how to get data in and out of Python programs from files.

There is a much more detailed description of file input/output (IO) in Chapter 11. However, there is excellent support for simple file IO of numbers within the NumPy module. This technique is described here in section 4.2 so that you can do file IO without reading the gory details of chapter 11.

4.1 Screen Output

4.1.1 Output to the Screen

We have already seen how to get output from the command line. Generally, within a program, we use the `print` command to output things to the screen. If you don't tell Python what format to use on output, it will decide for you:

```
>>> a=1.123456789
>>> print a
1.123456789
```

- but you can tell Python exactly how to write to the screen using a format conversion specifier (%):

```
>>> print 'a= %.2f' % a, 'cms'
a= 1.12 cms
```

To output more than one thing, we separate them with commas. Note that there are two `%` symbols here and that they are performing two different functions. The `%` inside the string (i.e., inside the quotes) introduces the format specifier (`%.2f` in this case). The second % in this example says that the variable(s) whose contents are to be substituted into the format specifier(s) in the preceding string are to follow:

i.e, '<string with format specifier(s)' % variable(s)

More that one variable should be grouped into a tuple (i.e., surrounded by `()`)

4.1.2 The Format Conversion Specifier

The % sign in the above example is used to tell Python what format to use to output to the screen. The details of how to use % are in chapter 3 of Hetland. Here we have specified a float format (f) with 2 decimal places (.2). The % in the string tells the `print` function that a format specifier is coming next. There must be one format specifier for each item you output:

```
>>> a=1.23
>>> b=123.45
>>> print 'a = %5.2f' % a, 'b = %.4e' % b
a =  1.23 b = 1.2345e+002
```

The format specifier %5.2f means: Output the variable as a floating point number with a total field width (all digits plus the decimal point) of 5 and with 2 decimal places.

We can get a similar result like this:

```
>>> a=1.23
>>> b=123.45
>>> print 'a = %5.2f\tb = %.4e' % (a,b)
a =  1.23   b = 1.2345e+002
```

Here the string contains two conversion specifiers and the two variables have been put at the end as the tuple: `(a,b)`. (A tuple is just a list of variables – see section 6.1.5). The \t special character used here is described in section 4.1.3.

There are quite a lot of conversion specifier types. The most useful are:

- f or F – Floating point (EG 1.234)
- d or i – Integer (EG 123)
- e or E – Exponential format (EG 1.234e+002)

Beware: Just to be annoying (or so it seems), the % symbol is also used to get the remainder from an integer division. The context tells you which one is being used.

4.1.3 Special Characters in Output

There are some special characters that it is useful to output to the screen. They are represented within a string by a backslash followed by a character. The most useful examples are:

- \n – This inserts a 'new line'
- \t – This inserts a tab

We used the \t special character in the above example to insert a tab character.

Note: If you actually want the backslash character in your output string, you must enter it twice.

4.2 File Input and Output

If you need to write a lot of complicated mixed text and numbers to a file, you need to understand how file IO is done in Python (Chapter 11). However, saving or loading a NumPy array can be done very easily.

This section makes use of quite a few concepts that you have not met yet. Just try this example so that you can see how easy it is to input and output data to file. Before we start, what is a NumPy array? There is an introduction to NumPy in section 6.3. Take a quick look now or just try the example and learn more about NumPy later.

If you want to, just skip this section for now but return here once are used to using NumPy arrays and want to load/save them from/to a file.

4.2.1 Saving an Array to File

The NumPy module provides the `savetxt()` function to do this. Just get your data into an array of suitable size and shape and this function will write it to file for you. You should only do this with NumPy arrays. If you try to read from a file that is a mixture of types, NumPy will do its best but may give you confusing results.

Here is an example. The first line loads (imports) the NumPy module so that you can use it:

```
>>> import numpy
```

```
>>> data = numpy.zeros((5,2), dtype='float') # Make an empty 2-d array
>>> data[:,0] = numpy.arange(5.0) # Fill one column with numbers
>>> data[:,1] = data[:,0] ** 2.0 # Fill second column with squares
>>> numpy.savetxt('myFile.txt', data, fmt='%.3f') # Save to file
```

Try all this from the command line. Then look at the text in your file (for example, with the editor). It should look like this:

```
0.000 0.000
1.000 1.000
2.000 4.000
3.000 9.000
4.000 16.000
```

Note that savetxt() has used a blank space between entries in each row. You can change this (if you must) using the delimiter parameter. For example:

```
>>> N.savetxt('myFile.txt', data, fmt='%.3f', delimiter=',')
```

will use a comma as the delimiter.

4.2.2 Loading an Array from File

Loading the data back from your file to a NumPy array is equally simple using `loadtxt()`. Be sure to tell `loadtxt()` what the delimiter is (or it will assume a blank space):

```
>>> import numpy
>>> data=numpy.loadtxt('myFile.txt')
>>> print data
[[  0.    0.]
 [  1.    1.]
 [  2.    4.]
 [  3.    9.]
 [  4.   16.]]
```

5 Chapter 5 - Programs (Scripts) and conditional blocks

If you already know how to create a program in a file using an editor and run the program, you can skip this section, but do the exercises anyway.

5.1 My First Program – 'Hello world'

Before we go any further, its time to write a few 'scripts' or programs.

Do this with the notepad++ editor under Windows7. Type your program into the editor, save it as a file then run the file. (This can be done just as easily on Linux or MacOS).

The first program that everyone writes in every computer language is the 'Hello world' program. This just writes 'Hello world' on the screen when it is run. This can take 10 or even 20 lines to do in a language such as C++ or Java. In Python, it is 1 line. So let's do it…

When you log on to Windows7, you will be in your 'home' folder. First you should create a new folder in which to do your work, then move to that folder and run the editor from there. Notepad++ will by default load and save files from the folder or directory from which it was invoked:

- Open a terminal (Click on the relevant icon)
- Use the command `mkdir` to create a new directory – EG `mkdir mywork`
- You can use the command `dir` to list the files in a directory
- Now use the `cd` command to move to that directory – EG `cd mywork`
- Now run the editor – type `notepad++`
- Type in: `print "Hello world"` followed by a newline (Enter).
- Save your file. Give it a name with .py extension – EG myProgram.py
- The file will be saved in the directory from which you ran notepad++
- Now type: `dir` and check that your program file is there
- Type: `python myprogram.py` to run your program (With your own program name)
- It should write: `Hello world` to the screen.
- Congratulations! You wrote a working program.

Notes:

All Python program files should have the .py extension. EG myprogram.py

Don't muddle up the Windows7 command prompt with the Python prompt. If you are already running Python, you will see the >>> python prompt. Here you can try out python commands. If you are in Windows7, the prompt will normally look something like J:\>. You need to be at this operating system prompt to tell Python to run your program file by typing:

`python myprog.py`

where `myprog.py` is your program file name.

In all sections from now on, we provide an example of the relevant technique (contained in a

box) and then exercises for you to try. Model solutions will be provided on DUO for all of the exercises after the relevant workshops. Do write your own program first then look at the model solution.

5.2 Conditional blocks

You will often want to execute a different bit of your program depending on some condition or other. This is where the Boolean or logical variables come in. You can decide whether to execute some code depending on whether a condition is True or False.

Conditionals allow a program to make decisions. They are often called 'if/then' or 'if/then/else' for obvious reasons. Conditionals are dealt with in chapter 5 of Hetland.

Blocks of code in Python are defined by a line starting with a colon (:) and **must be indented**. The indent must be the same for every line of the block. Such blocks are used in conditionals, in loops and in function definitions.

Now that we have started writing programs, and there are quite a few lines of code involved, the examples should be typed into the editor, saved as a file and run. You therefore won't see the Python prompt (>>>) in the examples any more. Type this example in and try it:

```
a = float(raw_input('Please enter a positive number: '))
if a < 0:
    print 'That is a negative number'
    print 'That is not what I asked for'
else:
    print 'Thank you'
    print 'That is a positive number'
```

Note 1: The < symbol means 'is less than'.

Note 2: You execute a block of code lines after the `if`. Python knows how much to do because there is a semi colon (:) after the `if` statement *and* the block to be executed is *indented*. This indenting may be done for you by your editor but can also be put in by hand. Any number of spaces (or a tab) is an indent. Lines within the block to be executed *must all have the same indentation*.

Beware: Getting the indenting wrong is a very common error in Python.

You can test for more than one condition at a time using `elif` (short for else if):

```
a = float(raw_input('Please enter a positive number: '))
if a < 0:
    print 'That is a negative number'
    print 'That is not what I asked for'
elif a> 0:
    print 'Thank you'
    print 'That is a positive number'
else:
    print 'That looks like zero to me'
```

You can have conditionals within conditionals (an `if` within an `if`) but **be careful** to get the indentation right.

The comparison operators are not all so obvious as < and >. Here is a table of them:

Expression	Description
x = = y	x equals y
x < y	x is less than y
x > y	x is greater than y
x >= y	x is greater than or equal to y
x <= y	x is less than or equal to y
x != y	x is not equal to y
x is y	x and y are the same object
x is not y	x and y are different objects
x in y	x is a member of the sequence y
x not in y	x is not a member of the sequence y

Notes:

It is a very common error to use x=y to test if x is equal to y (Rather than x = = y). This usually gives a syntax error. There are situations where it will not! Python will just do as it is told and set x equal to y - VERY confusing!

The last 4 conditions in the table are less used but quite interesting. Try using them!

There is also an `and` operator and an `or` operator. These are used to test more than one condition at the same time:

```
if a < -10 or a > 10:
     print 'Magnitude is too large'


if a < 0 or b < 0:
     print 'One of the values is negative'
```

Note: These examples are code fragments. They won't work on their own.

5.3 Mathematical Functions and NumPy

We will meet NumPy in more detail in later chapters. However, it is so fundamental to doing Physics with Python that we introduce it very briefly here.

There are no mathematical functions built into Python. If we want to take the sin of a floating point number, we would expect to type

```
>>> y = sin(0.123)
```

Try this. It will give an error.

We add all of the mathematical functions that we need by 'importing' the NumPy module.

Try this:

```
>>> import numpy
>>> x=0.123
>>> y=numpy.sin(x)
>>> print y
>>>0.122690090024
```

5.4 Examples

5.4.1 Example 5.1

We will write a program that evaluates the function:

$y = 3x + 4$

at a single point. The program sets a fixed value of x. An alternative is provided that requests a value of x from the screen. This has been commented out with a #. The program prints out the equivalent value of y and stops:

```
''' Program to find the value of the function y=3x+4 at a single point '''


# x = float(raw_input('Enter a number: '))
x = 2.1
y = x*3.0+4.0   #This is the algorithm
print 'The result is: %.3f' % y
```

Notes:

The triple quoted string at the top of the program is called a documentation string. It records what the program does.

In the input from the screen, we have used the output of the raw_input() function (a string) as the input to the float() function.

We have used the %.3f format specifier to format the output to exactly 3 decimal places.

5.5 Exercises

5.5.1 Exercise 5.1 (Covered in workshop 1)

Write a similar program to the example to evaluate the function:

$y = x^2-1$

Suggestions:

Put a comment at the top of the program.

Make the program work with floating point numbers – much more useful than just integers.

Set x to 2.0. Print out the value of y. The output should look like this: The result is: 3.0

Test it with various values of x.

Model solution is in file: Exercise5.1.py on DUO.

5.5.2 Exercise 5.2 (Covered in workshop 1)

Write a program to evaluate the polynomial: $y = ax^2 + bx + c$

Set various values for a,b,c and x and test your program by checking that you get the correct answers.

Print out both the input and the output like this:

The coefficients are: 3.0 4.0 5.0

For x = 2.0, the result is: 25.0

Model solution is in file: Exercise5.2.py on DUO

5.5.3 Exercise 5.3 (Covered in workshop 1)

Write a program to find the square root of a number.

Use a conditional to check that the number is greater than or equal to zero.

Print the result like this:

The square root of 1.23 is 1.10905365064

If the number is negative, print an error message.

5.5.4 Exercise 5.4 (Covered in workshop 1)

Write a function to find the sin and cos of an angle in degrees.

Set a variable to the angle in degrees but remember to convert it to radians before calling the sin and cos functions.

Print the result in this format:

Sin of 60.0 degrees is 0.866025403784

Cos of 60.0 degrees is 0.5

Check your code! EG Find sin of 0 degrees, sin of 90 degrees and cos of 90 degrees.

Why is cos of 90 degrees not zero?

6 Chapter 6 - Sequences, Lists, Strings and Arrays

Suppose we want to evaluate the function in exercise 5.1 at more than 1 point and save all of the results? This is just one example of where you might need a list or array. We will soon see that what physicists really want much of the time is arrays that represent vectors and matrices which are represented by NumPy arrays.

Strictly speaking, Python provides various sorts of 'sequences' but the most important sequences within Python are lists and strings. The most important sequences for physicists however are NumPy arrays. These are introduced here.

So…Lists and strings are useful but don't spend too much time on them. All of the physics will be done using NumPy arrays (see section 6.3 and chapter 10).

Note 1: You can have 2-d lists but, if you are handling numbers, NumPy arrays are much better.

Note 2: You will come across a sequence called a 'tuple' (section 6.1.5). It is just a list whose elements cannot be changed.

6.1 Lists

You may not need lists for some time so feel free to skip over this section and return to it later.

We will look briefly here at how lists are created and used. For more detail, see chapter 2 of Hetland. Lists are just a collection or sequence of other objects.

You can create a list of objects. The objects don't have to be all of the same type and could even include other lists. **We use square brackets [] for lists**:

```
>>> a=[3,54,26,90]
>>> print a
[3, 54, 26, 90]
>>> b=[3,54,'llama',a] # Include the list a in the list b
>>> print b
[3, 54, 'llama', [3, 54, 26, 90]] # Note the nested brackets
```

The elements of a list are numbered from zero and you get an error if you ask for an element that is out of range. It is a very common error to forget that the *elements of a list start at zero* (not 1). So the list above called a has 4 elements numbered 0,1,2 and 3. The address of an element in a list is known as its 'index'. Do not muddle up the value of a list element and its index. In the above example the value of b[1] is 54; IE the index is 1.

You access one or more elements of a list by giving the index(s) of the element(s) you want:

```
>>> a=[3,54,26,90]
>>> b=[3,54,'llama',a]
>>> print b
[3, 54, 'llama', [3, 54, 26, 90]]
>>> a[0]
3
```

```
>>> b[3]
[3, 54, 26, 90]
>>> b[4]
Traceback (most recent call last):
  File "<pyshell#40>", line 1, in <module>
    b[4]
IndexError: list index out of range
```

Note 1: You can access elements of a list from the right hand end using negative indexes.

Note 2: You can get the length of a list (or any sequence) using the `len()` function.

Note 3: You can change any element of a list.

In the list called `a` above:

```
>>> a[-1]
90
>>> a[-2]
26
>>> len(a)
4
>>> a[2] = 'gecko'
>>> print a
[3, 54, 'gecko', 90]
```

6.1.1 Slicing Lists

You can get a 'slice' of a list. This is a very important technique which is also used for NumPy arrays. You will often want to access just one part of a list or array. To do this, use two indices (start and end) separated by a colon. Note that the 'end' index actually refers to the element *after* the last one you want!!

```
>>> a=[12,23,34,45,56,67,78]
>>> print a[3:6]
[45, 56, 67]
```

If you miss out one of the indices, Python will assume you want all the rest, from the beginning or up to the end:

```
>>> a[3:]
[45, 56, 67, 78]
>>> a[:3]
[12, 23, 34]
```

6.1.2 2-d Lists – *Advanced Topic*

You can have 2-d lists – it's just a list of lists! - but, especially when manipulating numbers, it is *far better to use NumPy arrays* (see section 6.3).

Addressing a 2-d list is done thus:

```
>>> a=[[0,1,2],[3,4,5],[6,7,8]] # Create a 2-d list
>>> print a
[[0, 1, 2], [3, 4, 5], [6, 7, 8]]
>>> print a[1][2] # Addressing a 2-d list
5
```

6.1.3 Basic List Operations

There are lots of operations that can be performed on lists. Here we mention just a few of the most useful:

You can *add* one list to another:

```
>>> a=[1,2,3]
>>> b=[4,5,6]
>>> c=a+b
>>> c
[1, 2, 3, 4, 5, 6]
```

You can *append* a new value to the end. This is one of the most useful features of a list. You cannot easily append to a NumPy array.

```
>>> a=[1,2,3]
>>> a.append(4) # Note the dot!  This is a method.
>>> print a
[1, 2, 3, 4]
```

You can *insert* a new value in the middle of a list:

```
>>> a=[1,2,3,4,5,6,7,8,9]
>>> a.insert(3,99)
>>> print a
[1, 2, 3, 99, 4, 5, 6, 7, 8, 9]
```

You can look to see if a list contains a particular value. For this we use the membership operator known as *in*:

```
>>> a=[1,2,3,4,5,6]
>>> 3 in a # Think of it as a question: Is 3 in a?
True
>>> 7 in a
False
```

What you get back is a Boolean or logical value. This can take only the value `True` or `False`. The `in` operator is most useful in loops, which we will look at later.

6.1.4 Fancy List Handling – zip() and map() – *Advanced Topic*

You can do some very clever things with lists. In general however, the 'array' type that we introduce later in the course (in module NumPy) is more useful to physicists.

Have a look at appendix B of Hetland where he has a summary of built in functions and methods. I will give just two examples here:

Example 1. Try the `zip()` function:

```
>>> a=[1,2,3]
>>> b=[4,5,6]
>>> c = zip(a,b)
```

Now print c and see what zip() has done to your 2 lists.

Example 2. Try the `map()` function:

You may have a list of integers and want to make them all into floats. Python won't let you do this in case the list contains types that cannot be converted. However, the map() function allows you to 'map' any function onto all the elements of a list. Of course, if the list contains unsuitable elements, you will get an error:

```
>>> a=[1,2,3,4,5]
>>> b=map(float,a)# Apply the float() function to all elements
>>> print b
[1.0, 2.0, 3.0, 4.0, 5.0]
>>> a.append('fred')
>>> print a
[1, 2, 3, 4, 5, 'fred']
>>> b=map(float,a)
Traceback (most recent call last):
  File "<pyshell#16>", line 1, in <module>
    b=map(float,a)
ValueError: invalid literal for float(): fred
```

6.1.5 Tuples

Tuples are lists that cannot be changed. In general, don't worry about them, just use lists! However, they can be useful and you will come across them.

If you use a library function that happens to return a tuple, just treat it like you would a list but *don't try to change it*! **Tuples use round brackets** for assignment and display:

```
>>> a=(1,2,3)
>>> a
(1, 2, 3)
>>> type(a)
<type 'tuple'>
>>> print a[1]
2
>>> a[1] = 99
Traceback (most recent call last):
  File "<pyshell#88>", line 1, in <module>
    a[1] = 99
TypeError: 'tuple' object does not support item assignment
```

6.2 Strings

We have already come across strings. They are very like a list in which every element is a character. You can access their individual characters by indexing or slicing.

There are some differences to lists. You cannot, for example, change the elements of a string:

```
>>> a='This parrot is dead'
>>> print a[5:11]
parrot
>>> a[5] = 'f'
Traceback (most recent call last):
  File "<pyshell#100>", line 1, in <module>
    a[5] = 'f'
TypeError: 'str' object does not support item assignment
```

Python is very good at handling strings. There is a rich set of methods to do things to them.

See chapter 3 of Hetland for details.

A few useful examples are:

`find()`: Find a substring within a larger string

`split()`: Splits a string up into its elements in a list.

`join()`: Joins a list of strings (The opposite of split)

For example:

```
>>> a="Monty Python's Flying Circus"
>>> a.find('Monty')
0
>>> a.find('Python')
6
>>> a.find('Gumby')
-1
```

Notes:

If you want to use a quote in a string (EG for an apostrophe), you must use the OTHER type of quotes around the string.

`find()` returns the index where the substring starts or -1 if the substring is not found.

6.3 Numpy arrays – An Introduction

Whilst Python lists can be very powerful, what the physicist really needs is arrays that can be used to represent vectors and matrices. Python was not specifically designed for physics so it has no such arrays. Nor does it have much in the way of mathematical functions. All of this is added to the language by importing an essential module called NumPy.

We have not yet covered the writing of functions or the importing of modules (which are essentially just a collection of functions). However, the arrays and functions provided by NumPy are so useful for solving physics problems that we will take a first look at them now.

More detail on writing functions and using modules are given in chapter 8. More details of the NumPy module are given in chapter 10.

There is an excellent introduction to NumPy (and its partner module SciPy) on the web at:

http://www.scipy.org/Getting_Started

Have a quick look at it now. However, some of the detail and examples in that introduction will mean more when you have looked at the later chapters of this course.

Arrays are very like lists. The most important differences are:

- All of their elements must be of the same type

- Their 'shape' can be changed (EG from 1-d to 2-d)

- Multi-dimensional arrays are addressed differently to lists

- You can't append to an array. Once its length is set, it's fixed.

- Manipulating arrays is much faster than manipulating lists

For the technically minded, NumPy is implemented in the C language and is a compiled library. All Python list handling is interpreted and thus much slower. Thus using NumPy is the secret of doing complex physics problems in a slow language (Python) very quickly.

For solving physics problems, if you are in doubt which to use, *use arrays, not lists.*

6.3.1 Using NumPy

The process of adding a module (or library) to Python is very simple and is called 'importing'.

At the Python prompt, type `import numpy` (Note the lower case name).

This will, as if by magic, add a vast number of functions to Python. If you want to see how many, try typing `dir(numpy)`. To use them, just prefix their names with the name of the module.

EG One way to create an array is using the NumPy function called `arange().`

Try this:

```
>>> a=numpy.arange(9, dtype='float')
>>> print a
[ 0.,  1.,  2.,  3.,  4.,  5.,  6.,  7.,  8.]
```

We have specifically told NumPy that we want floats using the `dtype` parameter. If you don't specify what you want you will get the default *on your system*. This may be integers or floats. You can also tell arrange you want floats by using a float as an argument:

```
a=numpy.arange(9.0)
```

To see how numpy.arange() works, just import numpy then type:

```
help(numpy.arange)
```

You can change this 9 element 1-d vector into a 3x3 matrix thus:

```
>>> a.shape=(3,3) # Use the shape method
>>> print a
[[ 0.  1.  2.]
```

```
    [ 3.    4.    5.]
    [ 6.    7.    8.]]
```

You can manipulate an entire array just like a variable. You can't (easily) do this with a list:

```
>>> b=a**2
>>> print b
[[  0.    1.    4.]
 [  9.   16.   25.]
 [ 36.   49.   64.]]
```

Another useful way to create an array (of all zeros) is using numpy.zeros():

```
>>> import numpy
>>> a = numpy.zeros((3,3), dtype=float) # Create a 3x3 array
>>> print a
[[ 0.   0.   0.]
 [ 0.   0.   0.]
 [ 0.   0.   0.]]
```

Here we have defined the shape of the matrix by giving numpy.zeros() the required shape as the tuple (3,3). We have also specified that we want floating point numbers (dtype=float).

6.3.2 Addressing and Slicing Arrays

Addressing and slicing arrays is done in a very similar way to lists. We use the same square brackets: []. So, in 1-d:

```
>>> a=numpy.arange(9, dtype='float')
>>> print a
[ 0.  1.  2.  3.  4.  5.  6.  7.  8.]
>>> print a[1:5] # A slice from element 1 to element 4
[ 1.  2.  3.  4.]
```

Just as for a list, we have taken a 'slice' from the array, `a[1:5]`. Remember that the slice starts at index 1 but ends at index 4. IE One less than the last index given.

The difference comes with arrays of more than one dimension. So in 2-d:

```
>>> a=numpy.arange(9,dtype='float')
>>> a.shape=(3,3) # Make it 2-d
>>> print a
[[ 0.   1.   2.]
 [ 3.   4.   5.]
 [ 6.   7.   8.]]
>>> print a[2,1] #Print one element of the matrix
7.0
```

We have addressed one element of the matrix at row 2, column 1 as: `a[2,1]`.

Note that we address a 2-d NumPy array with the 'y' coordinate first, then the 'x' coordinate. So think of the element that you want from your matrix 'a' as having an address a(y,x).

This is my preference! You can do it the other way round but, when you print an array to the screen, the x direction will be up the screen and y will be across. That confuses me!

This form of addressing, with a comma, is *not* valid for lists.

Slicing a 2-d array in summary:

- a[2:4, 7:9] – Slice out a small piece of an array (y from 2 to 3 and x from 7 to 8)
- a[:,0] – Slice out the whole of the first column
- a[0,:] – Slice out the whole of the first row

This section was just a taster. We will return to arrays in chapter 10. Remember, if you are manipulating sequences of numbers, especially in more than one dimension, *use arrays not lists*.

6.4 Dictionaries – *Advanced Topic*

Dictionaries are a very powerful type in Python. However, they are not essential for this course and are not covered in any detail here. They are covered in chapter 4 of Hetland.

They are another form of sequence, like a list, but the order of the elements is not fixed. Thus there is a *key* for every element. They use curly brackets like {}. They can be used to form a sort of mini database. For example, you could store your friend's 'phone numbers in a dictionary:

```
>>> phonebook = {}
>>> phonebook['Fred'] = '12345'
>>> phonebook['Claire'] = '0234 7432'
>>> print phonebook
{'Claire': '0234 7432', 'Fred': '12345'}
>>> print phonebook['Claire']
0234 7432
>>> phonebook.has_key('Fred')
True
>>> len(phonebook)
2
```

The entries in a dictionary are *key/value* pairs separated by commas. Both the key and the value can be of most types so the 'phone numbers could have been integers. The key and the value are separated by a colon.

You can add an entry that is not already there. There is no fixed index for the elements. There is no need for an 'append' method.

Note: the method `has_key()` is used to check if the dictionary has such a key in it. If you try to access a non-existent key, you will get an error. Check it exists first.

```
>>> print phonebook['Clive']
Traceback (most recent call last):
  File "<pyshell#118>", line 1, in <module>
    print phonebook['Clive']
```

```
KeyError: 'Clive'
>>>
```

6.5 Indexing Sequences - Summary

Remember that **all** types of sequences are indexed using square brackets.

Remember that all indexes start at zero

The data in any type of sequence are like items in pigeonholes. The index is the label on the pigeonhole. Here is an example using a string:

```
>>> theName = 'Monty Python'
>>> print theName[2] # Index using square brackets
>>> n
```

These are the data:

M	o	n	t	y		P	y	t	h	o	n

These are the indices:

0	1	2	3	4	5	6	7	8	9	10	11

6.6 Examples

These examples are too long to try at the python prompt (>>>). If you want to try them out, type them into an editor and run them. The same is true for the exercises that follow.

6.6.1 Example 6.1

We will write a program to hold a small database of breakfast items and their costs using two lists. We will allow the user to query the database by item and print out the cost:

```
''' Program to hold a small database and look up entries in it '''

# The list of items
items = ['Spam', 'Ham', 'Egg', 'Beans', 'Bacon',\
        'Sausage'] # Use a \ to extend to the next line
# The list of prices
prices = [0.7, 1.2, 0.2, 0.52, 0.75, 0.6]


item = raw_input('Enter an item: ')
theIndex = items.index(item) # Find the index of item in the list items
# Let print format it as it likes:
print items[theIndex], 'costs ', prices[theIndex], ' pounds'
# or format it exactly as we want it
print '%s costs %.2f pounds' % (items[theIndex],prices[theIndex])
```

Notes:

If a line is too long, use a backslash (\) to continue to the next line

Two different ways to format the output are shown. The first is simpler but the print

statement will format the output as it likes. The second uses format specifiers (%) and prints the line exactly as you tell it.

If you use the format specifier, the items to be printed must be held in a tuple (in () type brackets).

6.6.2 Example 6.2

We will create a 2-d array of zero integers, fill it with some numbers, print out the full array then print out some slices from the array.

```
'''Program to create an array and print a slice from it'''
import numpy
a = numpy.zeros((3,3),dtype='int') # Create a 3x3 array of integers
a[0,0]=9; a[0,1]=8; a[0,2]=7 # First row
a[1,0]=6; a[1,1]=5; a[1,2]=4 # Second row
a[2,0]=3; a[2,1]=2; a[2,2]=1 # Third row
print "The full matrix:"
print a # We could also use: print a[:,:]
print "Slice out the second column and print it:"
print a[:,1] # Slice out a column
print "Slice out part of the second row and print it:"
print a[1,1:3] # Slice out part of a row
print "Slice out a 2x2 part of the array and print it:"
print a[:2,:2] # Slice out a sub-matrix
```

Notes:

Usually an array will be filled with data in some kind of loop rather than 1 element at a time. We will deal with loops in chapter 7.

The general format for a slice is for example: `a[1:3, 2:5]`. This means elements 1 to 2 in the first dimension and 2 to 4 in the second dimension. If one number is left out, the default is to use all of the elements from the first or up to the last in that dimension.

6.7 Exercises

6.7.1 Exercise 6.1 - Lists

Write a small database and some code to query it using lists.

Note: This example requires input from the screen, which is covered in chapter 14.

Create a list with the names of 10 'friends'. Create a second matching list with their years of birth. Ask the user to input a name. Tell him the place (index) of that person in the list, how many friends he has in total and the year of birth of the person. The input and output should look like this:

```
Enter a name: Will
You have 10 friends
Will is number 8 in your list
Will was born in 1991
```

Try entering a name that is not in the list. The program will crash. We will look at how to handle such conditions later.

Model solution is in file: Exercise6.1.py on DUO.

6.7.2 Exercise 6.2 – Dictionaries - *Advanced Topic*

Note: This example requires input from the screen, which is covered in chapter 14.

The program for exercise 6.1 would have been better done with a dictionary rather than two lists. Repeat the exercise but using a single dictionary to hold the data. Which of the outputs from example 6.1 cannot be generated from data stored in a dictionary and why? If you enter a friend who is not in the list, the program will crash with a different error to the version that used lists. Both errors are worth looking at.

Model solution is in file: Exercise6.2.py on DUO.

6.7.3 Exercise 6.3 – (Covered in workshop 2)

Create a 1-d array of length 10 array.

Use numpy.zeros() to do this. Use help(numpy.zeros) to learn how.

Arrays have a 'shape' which is a tuple. In 1-d this is just their length

Now add 1 to every element of the array.

This can be done in 1 line.

Print the array of 1s. It should look like this:

[1. 1. 1. 1. 1. 1. 1. 1. 1. 1.]

Model solution is in file: Exercise6.3.py on DUO.

6.7.4 Exercise 6.4 – (Covered in workshop 2)

Find the mean of an array of real numbers.

Create an array containing 4 specific real numbers (you choose).

Make a list and use numpy.array() to convert the list to an array.

Now find the mean of these numbers. Not surprisingly, you can use numpy.mean().

However, there is a better syntax for this. *The mean() method of the array object can be used directly*. See if you can do it this way.

Print out the result. Is it what you expect? Test your code using numbers of which you know the mean.

Model solution is in file: Exercise6.4.py on DUO.

6.7.5 Exercise 6.5 – (Covered in Workshop 3)

Create a 10x10 matrix of integers starting at 1 and ending at 100 as a NumPy array. The first row is the integers 1 to 10, the second 11 to 20 etc. (Use numpy.arange() and the `shape` attribute). Print the matrix to the screen.

Now slice out the column from the matrix where all of the numbers end in the digit 5 (the 5th column). Print this to the screen. It should look like this:

```
[ 5 15 25 35 45 55 65 75 85 95]
```

Now slice out a sub-matrix from the original matrix to look like this:

```
[[35 36 37 38]
 [45 46 47 48]
 [55 56 57 58]]
```

Print it to the screen.

Now convert the sub matrix to a vector to look like:

```
[35 36 37 38 45 46 47 48 55 56 57 58]
```

This last stage is best done using the `flatten()` method. If you try to convert the sub-matrix to a vector using the `shape()` method you will run into problems. You have to make a *copy* of the sub-matrix first and `flatten()` does this for you.

It is important when using arrays to be aware of when you need a copy. Python will try to not make a copy if it can avoid it because it wastes memory and is slow. In our example of making a sub-matrix, it will just point to parts of the original matrix.

Model solution is in file: Exercise6.5.py on DUO.

7 Chapter 7 - Loops

Loops allow the computer to loop over a set of data. You might want to evaluate a function at 100 different values for example.

Loops are dealt with in chapter 5 of Hetland.

7.1 Loops

Let's go back to the example of evaluating a function. We may want to evaluate it for 100 values of x and save all of the values somewhere. (Later, we will also draw a graph of it).

There are two sorts of loops in Python that allow you to execute a block of code many times: `for` loops and `while` loops.

In general, when doing Physics, you should use a `for` loop if you can. You need to know how many times to run around the loop. In general you will know this. If you really don't know when the loop will stop (EG when waiting for the user to provide some valid input) then use a `while` loop.

7.1.1 Loops - The 'while' loop

These do exactly as they say. They will execute a block of code while some condition is true:

```
prompt = 'Please enter a positive number less than 10: '
x=-1 # Initialise x to an invalid number
while x<=0 or x>9:
    x = float(raw_input(prompt)) # This is indented
print 'Thank you. That number is fine.'
```

Try this. The while loop will repeat the indented line until the user inputs a valid number.

Note: We put the prompt string into a variable so that we can more easily change it.

7.1.2 Loops – The 'for' loop

The `while` loop is very versatile and useful. However, physicists are normally looping over a set of numbers as in the example of evaluating a function for many values. For these applications, we use a `for` loop to loop over all of the elements of a list or array.

There are examples of looping over arrays in section 7.2 below.

First, let's loop over a list of strings. Type this in and try it:

```
cheeseList = ['Wensleydale', 'Stilton', 'Danish Blue', 'Red Leicester',
'Brie']
for cheese in range(0, len(cheeseList)):
    print "Do you have some", cheeseList[cheese], "?"
    print "Not as such"
```

Note: This could have been written more concisely as:

```
cheeseList = ['Wensleydale', 'Stilton', 'Danish Blue', 'Red Leicester',
'Brie']
for cheese in cheeseList:
```

```
    print "Do you have some", cheese, "?"
    print "Not as such"
```

Here you can more clearly see the power of the python keyword `in`. The `for` loop sets the variable `cheese` to each of the elements of the list `cheeseList.` A variable used in this way is known as an iterator. If you would like to read about the magic of iterators, see Hetland chapter 9. (That is an *advanced topic*).

In the first version of the above example, we have used a very useful Python function `range()`. Because we so often want to iterate (or loop) over a list of integers, Python provides this function to create a suitable list of integers.

The function can take up to 3 parameters: range(start, stop, step). As always, use `help(range)` for details.

Try this function out from the python prompt:

```
>>> range(10) # Only the stop parameter is given
[0, 1, 2, 3, 4, 5, 6, 7, 8, 9]
>>> range(1,10) # Using start and stop
[1, 2, 3, 4, 5, 6, 7, 8, 9]
>>> range(0,10,2) # Using start, stop and step
[0, 2, 4, 6, 8]
```

Note: You can give it a start and end value and an increment, or just an end value. It returns a list up to *one less than* the stop parameter.

So now let's use it to loop over a set of numbers:

```
for number in range(10):
    print "The number:", number, "It's square:", number**2
```

Note 1: The variable 'number' (the loop variable) is just used to store the element from the list generated by the range function that we use on each go round the loop.

Note 2: As you can see, Python is happy to square the value of 'number' before it prints it.

7.1.3 Getting out of Infinite Loops - break

It is often useful to use an infinite loop, typically using `while True:.` You can break out of such a loop on a condition using the break command. You should not use this more than is necessary as it can make code somewhat less clear to read. Here is an example in 'pseudo code'. This is a useful way of showing some code structure without sticking to the language syntax:

```
while True: # This is pseudo code
    Do something
    if <some test>:
        break
    else
        do some other stuff
carry on # The break causes a jump to here
```

You can also escape from an infinite loop by entering <CTRL C> from the keyboard. (Hold

down the CTRL key and press C).

7.2 Examples

You will usually be looping over arrays so examples are given here using arrays. Don't muddle up the index of an array element with the value of that element. The loop iterator can be used as the index of an array

7.2.1 Example 7.1

Create an array of floats from a list. Use a `for` loop to square each element of the array and put the results in a new array:

```
'''Program to put some numbers into an array then create a new array of
 their squares'''
import numpy
x = numpy.array([1.2,6.,15.,21.,3.1]) # Create an array from a list
y = numpy.zeros(len(x), dtype=float)  # Create an array to hold the squares
for i in range(len(x)): # Loop over a range from 0 up to the length of x
    y[i] = x[i]**2 # Square each element of x and put result in y
print 'Original array: ', x
print 'Squared array:  ', y
```

We used the variable `i` as the iterator in the `for` loop. We then used it to index each element of the arrays x and y. Read the indented line as: ' The ith element of y becomes equal to the ith element of x squared'.

This is a useful example of a `for` loop but can be done better! The power of arrays means that you can square an entire array like this:

```
x = numpy.array([1.2,6.,15.,21.,3.1]) # Create an array from a list
y = x**2
print 'Original array: ', x
print 'Squared array:  ', y
```

This code will also run about 100 times faster than the version using a `for` loop!

7.2.2 Example 7.2

Find the sum of an array using a `for` loop:

```
'''Program to put some numbers into an array then create a new array of
 their squares.'''
import numpy
x = numpy.array([1.2,6.,15.,21.,3.1]) # Create an array from a list
total = 0 # Initialise the total variable
for i in range(len(x)):  # Loop over a range from 0 up to the length of x
    total = total + x[i] # Add the next value in the array into the total
print 'The array: ', x
print 'The total: ', total
```

The indented line can be written using a very common shorthand as:

```
total += x[i] # The same as  total = total + x[i]
```

As before, this very common process of finding the sum of an array can be done in one go using NumPy:

```
x = numpy.array([1.2,6.,15.,21.,3.1]) # Create an array from a list
total = numpy.sum(x) # Sum the entire array using the sum() method in numpy
print 'The array: ', x
print 'The total: ', total
```

The writers have implemented many such functions to help you so if in doubt, take a look in the documentation to see if NumPy will do it for you before re-inventing the wheel as I have here!

7.2.3 Example 7.3

Restrict the data in an array to less than 2π:

```
'''Program to restrict the values of an array to less than 2*pi'''
import numpy
x = numpy.arange(0.,10.,0.2) # Create an array using numpy.arange()
print 'The original array:\n', x
for i in range(len(x)):  # Loop over a range from 0 up to the length of x
    if (x[i] > 2*numpy.pi):
        x[i] = x[i] - 2*numpy.pi # Restrict the data to < 2*pi
print 'The restricted array:\n', x # would this work for all input values?
```

Note that NumPy has provided the value of π for you as an attribute (`numpy.pi`).

Again, this can be done faster using NumPy. However, it requires the use of array conditionals. These can be confusing until you are used to them. It is fine to do it in a `for` loop. For those interested, the faster way is:

```
x = numpy.arange(0.,10.,0.2) # Create an array using numpy.arange()
print 'The original array:\n', x
x = numpy.where(x>2*numpy.pi, x-2*numpy.pi, x) # Use an array conditional
print 'The restricted array:\n', x # would this work for all input values?
```

7.2.4 Example 7.4

Implement a program to sort the values in an array into ascending order:

```
'''Program to sort some numbers into ascending order'''
import numpy
# Create some floating numbers as a list and convert to an array
x = numpy.array([23.2,12.1,1.2,34.6,0.1,15.2,19.3,45.6,3.4])
print 'The original array:\n', x
for cycle in range(len(x)):  # Cycle enough times so that sort is complete
    for i in range(len(x)-1): # Correct order of each pair of values
        if x[i] > x[i+1]:
```

```
            temp = x[i]     # Swap x[i] and x[i+1] by using another
            x[i] = x[i+1]  #     variable as a temporary store
            x[i+1] = temp
    print 'The array after cycle ', cycle, ':\n', x
print 'The ordered array:\n', x
```

Notes:

We have used two nested `for` loops to do this. The inner loop, with loop variable i, orders each pair of values. However, this is not enough, as the print lines will show. We must run through the process a number of times equal to the length of the array to ensure that it is fully ordered. The outer loop, with loop variable `cycle`, does this.

The inner loop variable i is used as the index into the array (EG `x[i+1]`). The outer loop variable `cycle` just causes the loop to happen. It is not used within the loop.

The code works but is quite inefficient. If you needed to order a huge array, you would have to find a more efficient ordering technique.

This process is usually called 'sorting'. Not surprisingly, NumPy provides a sort method for you. You can also sort multi-dimensional arrays but you must specify the axis. As usual, type help(numpy.sort) at a Python prompt to get the syntax. This is the code using NumPy:

```
print 'Use the numpy method'
x = numpy.array([23.2,12.1,1.2,34.6,0.1,15.2,19.3,45.6,3.4])
print 'The original array:\n', x
x = numpy.sort(x)
print 'The ordered array:\n', x
```

7.3 Exercises

7.3.1 Exercise 7.1 – (Covered in workshop 2)

Use a 'for' loop to print out the squares and cubes of the real whole numbers from 0.0 to 5.0.

Put the numbers, their squares and their cubes in three numpy arrays, then print the arrays.

The output should look like this:

```
[ 0.  1.  2.  3.  4.  5.]
[ 0.  1.  4.  9.  16.  25.]
[ 0.  1.  8.  27.  64.  125.]
```

If you use numpy.zeros() to create the arrays, the default type will be 'float'.

Try changing it to 'int' using dtype=int and re-running the program.

Model solution is in file: Exercise7.1.py on DUO.

7.3.2 Exercise 7.2 – (Covered in workshop 2)

Re-write exercise 7.1 without the `for` loop, by using array operators.

If we first populate an array with a range of real numbers, we can then **operate on the whole array** to find the square and the cube. This is much faster than using a `for` loop.

The built-in function `range()` returns a list of integers. However, there is a numpy equivalent of `range()` called `numpy.arange()` which returns an array of either integers or real numbers depending on the type of argument that you use. Use this to generate the arrays of numbers to be squared and cubed:

```
>>> print numpy.arange(5)

[0 1 2 3 4]

>>> print numpy.arange(5.0)

[ 0.  1.  2.  3.  4.]
```

Model solution is in file: Exercise7.2.py on DUO.

7.3.3 Exercise 7.3 – (Covered in workshop 2)

Now let's use a `for` loop to do something more useful. Write a program to generate the first 10 elements of a Fibonacci series of integers. The first 2 elements are 0 and 1 by definition. Each element is then the sum of the previous 2.

IE The series F_n where $F_0 = 0$, $F_1 = 1$ then $F_n = F_{n-1} + F_{n-2}$

Remember that you can set any element of an array using its index. If `a` is an array, `a[0]` is the first element.

Model solution is in file: Exercise7.3.py on DUO.

7.3.4 Exercise 7.4 – (Covered in workshop 2)

The ratio of consecutive elements of the Fibonacci series converges to a number known as the 'Golden Ratio'. This is useful in mathematics and architecture!

Modify your Fibonacci program to calculate this ratio for consecutive elements until the difference from the previous value of the ratio is less than 10^{-10}. (10^{-10} is written 10e-10 in Python).

Use a while loop to do this. The output should look like this:

Value of Golden Ration is 1.61803398896 after 23 iterations

Hint: You will need to initialise various variables before the loop

Note: You could fill an array with the Fibonacci numbers, but you don't need to. You only need three of the values at a time.

Use an array first then try to make it work without the array.

Model solution using an array is in file: Exercise7.4.py and using no array is in Exercise7.4a on DUO.

7.3.5 Exercise 7.5

Write a program similar to example 7.4 that will sort each row of a 2 dimensional array into ascending order. Print the initial array and the sorted array. Use numpy.zeros() to create an empty array of type float of shape (3,5) – 3 rows and 5 columns - then fill it with data by setting each row as a slice (EG `x[1,:] = [2.1,3.4,0.9,7.2,8.2]` sets the second row to these values.

You can find the size of any dimension of an array either using `numpy.size()` or `a.shape` where a is an array.

EG

`numpy.size(a,0)` gives the size of the first dimension of an array `a`

`a.shape[0]` will gives the same thing. Try them both!

Do this exercise first by adding using nested loops; you may need 3! You will have to address array elements in 2-d EG `x[0,2]`. Now repeat the exercise using numpy.sort(). You must specify the axis. You need only one line to do the whole job. Try using numpy.sort() on your 2-d array without specifying the axis. What does it do?

numpy.sort() also takes an optional parameter to change the sorting algorithm used. There are several available. help(numpy.sort) will tell you the characteristics (EG speed) of the available sort algorithms.

Model solution is in file: Exercise7.5.py on DUO.

8 Chapter 8 - Functions and Modules

The whole principal of writing software for Physics depends on good program structure. Much of the functionality that you require will be provided by modules (libraries) that have already been written and you import into your programs. In a similar way, nearly all of the functionality of your own code should be built as a set of modules (just a separate file), each of which contains a group of related functions.

Every function that you write is a separate entity and can, and should, be tested independently to be sure that it works as you expect it to.

Note 1: You can include one or two functions with your main program in these simple examples, but note that the function definition must occur before the main program. Once all of your functions are in modules, you import them at the top of the main program file.

Note 2: In his wisdom, Hetland calls his chapters on functions *Abstraction;* a very correct name but a touch confusing. His chapter 6 covers the writing and use of *functions*. His chapter 7 covers the alternative technique of using *methods* (a part of object-oriented programming). These are not included in this course but feel free to learn such techniques if you have time. *Beware:* Because the use of methods is becoming more and more popular, many programmers will refer to their functions as methods.

8.1 A First Function

So, let us now look at writing some code as a function and then calling that function from a main program. We want to evaluate the polynomial $x^3 - 7x^2 + 14x - 5$. The function is defined using the `def` keyword, an indented block and an optional `return` statement. Type this example into an editor, save it to file and try running it:

```
def poly(x):
    '''

    Evaluate the polynomial x^3 - 7x^2 + 14x - 5

    '''

    return x**3 - 7*x**2 + 14*x - 5.0


# Test program for the function poly()
while True:
    xValue = float(raw_input("Enter a value for x:"))
    yValue = poly(xValue)
    print "The value of the polynomial is:", yValue
```

There are various points to note here:

Syntax: Note that the line with the `def` keyword on it ends in a colon. You must then indent the block of code that is the function, just as you did for conditional blocks.

Warning: Be very careful with your indentation. Python relies heavily on indentation to know when a block of code starts and ends. Your indentation must be correct.

You can pass parameters to functions (separated by commas for more than one). The

corresponding variable within the function is a different variable. Its value will not affect that in the calling program. In general, use a different name to remind you that they are different. The variable x in the function above is said to be *local* to the function `poly()`. (See section 8.1.1)

If you want to return some value to the calling program, you must use the `return` statement.

We could have used another variable for the value of the polynomial in the function and returned that variable but Python is happy not to bother. It works out the value and returns it. This code would be just as valid:

```
y = x**3 - 7*x**2 + 14*x - 5.0

return y
```

This program will run forever! The value of `True` is always True. To escape from the program, just type <CTRL C>, meaning hold down the CTRL key and type C. This will interrupt your program and return you to the Python prompt. It is better to build in a way to exit cleanly as we saw in exercise 7.3.

8.1.1 The `return` Statement

The return statement is used in a function to return a value to the calling program:

```
def increment(x):
      ''' Increment function '''
      x += 1 # Equivalent to x = x+1
      return x # Return the incremented value
# Main program
value = 2
y = increment(value)
```

A function doesn't have to have a `return` statement. It's often handy to have a few lines of code in a function that do something useful but don't return a result:

```
def printMessage(message):
      ''' Print the message '''
      print 'The message is: ', message


# Main program
myMessage = 'Hello'
printMessage('Hello')
```

Not a very useful function but it works. Note that I have not allocated the return from this function to a variable. If I did so, it would return `None` (A special Python constant meaning 'nothing').

So, depending on their input and output, function calls can look like:

```
result = functionA() # Takes no arguments but returns a value
result = functionB(12, value) # Takes arguments and returns a value.
functionC() # Takes no arguments and returns no value
functionD(value1, value2) # Takes arguments but returning no value.
```

You may wish to return more than one value from a function. This is very easy. The return statement will take a list of values, separated by commas. This should match a set of values in the calling program. Alternatively, the returned values can be put in one variable. This will be a tuple. Just treat it like a list to access the results:

```
def cleverFunction(<Parameters>):
    return a,b
first, second = cleverFunction(<Parameters>)
```

or:

```
def cleverFunction(<Parameters>):
    return a,b
results = cleverFunction(<Parameters>)
a = results[0]
b = results[1]
```

Note: The above is what is known as 'pseudo code'. It is not working Python. It just gives an idea of how some real Python might be written. The non-Python bits are often put inside <>.

8.1.2 Name Spaces and Scope

Before we go on, it is important to introduce the concept of name spaces or the scope of variables. This sounds complex but is not. The two things refer to much the same concept.

All of the variables that you define within a function belong to the namespace (or scope) of that function. They are said to be 'local' to that function and can only be accessed inside the function. Of course you can pass copies of variables from outside the function into it by passing them as parameters. You can also return variables from a function using the return statement.

If you declare (set) a variable in a function, its scope will be only inside that function. If you declare a variable in your main program, that variable will be 'global'. It can be referenced anywhere. However, *beware* of using more than a very few global variables. It is poor programming style and causes confusion.

Here is an example to show the difference between the global scope and the local scope inside a function:

```
''' Function to demonstrate the scope of variables'''
def myFunction(x):
    ''' Variables declared inside the function are local '''
    x = x+1
    myValue = 99
    print 'myValue in the function: ', myValue
    print 'aGlobalVariable in the function: ', aGlobalVariable
    return x


# Main program. Variables declared here are global
aGlobalVariable = 123
```

```
myValue = 1

print 'myValue in the global scope: ', myValue

print 'aGlobalVariable in the global scope: ', aGlobalVariable

result = myFunction(myValue)

print 'the result is:', result
```

The output of this program is:

```
myValue in the global scope:  1

aGlobalVariable in the global scope:  123

myValue in the function:  99

aGlobalVariable in the function:  123

the result is: 2
```

I have deliberately used the same name (myValue) for two different variables – one is global, the other local (bad practice!). The local variable is different to the global one.

I have passed the global variable `myValue` to the function where it has been copied (or mapped onto) the local variable `x.` This has been modified then returned to the main program into the global variable `result.`

Finally, you might possibly want to create a global variable in a function. This can be done by declaring it as global inside the function: EG `global x.`

8.1.3 Default values for parameters in a function

In general, the number of arguments sent to a function should match the number that it is expecting. However, you can pass less than the number it is expecting, provided you give defaults for those that are not passed. This is done as follows:

Using the function definition:

```
def myCleverFunction(start, stop, inc=1):
```

you MUST provide `start` and `stop` but, if you miss out `inc` it will be set to 1.

In this call, the default value of `inc=1` is overridden by the value of 5 provided:

```
start = 1; stop = 10; increment = 5

a = myCleverFunction(start, stop, increment)
```

In this call, no value is provided for `inc`, so it is set to 1:

```
start = 1; stop = 10

a = myCleverFunction(start, stop)
```

8.1.4 Documentation Strings

As well as using a # for comments, you can just put a string on a line on its own as a comment. It is usual to put a *documentation* string at the start of every program, at the start of every function and at the top of every module. These documentation strings are important. Python itself will use them to provide help to the programmer. We normally use *triple* quotes (of either type) for this. This allows there to be newlines within the documentation string. The doc string must be the **first line** of the module, or the **first line** of a function definition.

You will find that you can retrieve the documentation strings that you have provided by using the Python help() function.

Note: In your course work, you will be *expected* to provide documentation strings as well as comments. You will lose marks if they are omitted.

8.2 Modules and import

So, what is a module?

It is just a library of useful functions or methods gathered together into a single file. In general, we gather together functions with related functionality into a module. There are very useful standard modules that you can use, and of course you can write your own. Once a module exists, any program can 'import' it and make use of its functions.

Whether it is a standard module or your own module, the ways of importing a module (making it available to your own program) are the same.

Note: A module can import other modules. If you write your own module of clever maths routines, it will probably have to import the `numpy` module in order to have some maths functionality.

Modules are so important they are mentioned in chapter 1 of Hetland. However, the main chapter on modules is his chapter 10 ('Batteries Included'). This gives a neat introduction to using standard modules in Python.

We give just a brief introduction to how to use modules here.

There are several different ways in which modules can be imported:

8.2.1 Import

This is the default way to import a module. You should always use this plain import. Every call to a function in the module must use the module name as a prefix to the function name, separated by a dot(.). The function call, with its module prefix, will be unique which is essential to avoid confusion. By using this syntax you are addressing function names that are in the scope of the module.

```
import myCleverLib # import the module

# Now call a function in that module

theAnswer = myCleverLib.miracle(<parameters>)
```

8.2.2 from <module> import <function>

As always, there are shortcuts and you will see them used. *These are poor practice so don't use them!* If you do, be careful. It's easy to get confused.

If you are sure the name won't clash with one in your program scope, then you can import one or all of the functions from a module so that they look just like functions in your own program or module:

```
# Import just the one function from a module

from myCleverLib import miracle

# Now call a function in the module

theAnswer = miracle(<parameters>)
```

```
#Import all of the functions from a module

from myCleverLib import *
```

```
# Now call two functions from the module
theAnswer = miracle(<parameters>)
theOtherAnswer = theOtherWay(<parameters>)
```

If you have a function with the same name in two different modules and one imports the other like this, you won't know which one you are calling! Even worse, Python will probably decide which one to use and assume that you know. IE It won't give you an error. So beware.

8.2.3 import <module> as <name>

This is just a shorthand. If the module has a long name, you can replace that name at import time with a shorter one. Again, although often done, *this is poor practice so avoid it.*

```
#import the module with an alternative name
import myCleverLib as lib
#Now call a function in the module
theAnswer = lib.miracle(<parameters>) # A bit shorter
```

Advanced topic:

Just for fun, there is another shortcut technique. You can use a variable to 'point to' a function:

```
import myCleverLib
clever = myCleverLib.miracle # A pointer to the function
theAnswer = clever(<parameters>)
```

The variable clever is a pointer to the function just as myCleverLib.miracle (without any brackets!) is.

8.2.4 Allowing Python to find your Modules

Python must know where to find your modules. It keeps a list of where all of the standard modules live so you don't have to worry about them.

The best way to start working is to keep your modules in the same directory from which you are running your main program. This is one of the places where Python will look for modules.

If you decide to keep your own modules in a different directory, you will need to tell Python where to find them. You do this by setting an 'environment variable' called PYTHONPATH. How this is done depends on the operating system. In windows7, you do it in the control panel from:

Control panel -> System and Security -> System -> Advanced System settings -> Environment Variables.

Now Python will look in that directory as well as all of its usual places.

This process is described in more detail in Hetland chapter 10. He also describes the different process used if you are working under Linux.

Note: If you don't know if a module exists on your system, try to import it from the Python prompt. If it's there and Python can find it, it will import ok.

8.2.5 What's in a Module?

Once you find, or are told, about a useful module, how do you know what is in it? What

functions does it offer?

There is a useful built-in function in Python called dir().

This takes the name of a module and tells you what is in it. You will have to import it first of course!

We may come across the math module. Try importing it then use dir(math) to find out what it can do:

```
>>> import math
>>> dir(math)
['__doc__', '__name__', 'acos', 'asin', 'atan', 'atan2', 'ceil', 'cos',
'cosh', 'degrees', 'e', 'exp', 'fabs', 'floor', 'fmod', 'frexp', 'hypot',
'ldexp', 'log', 'log10', 'modf', 'pi', 'pow', 'radians', 'sin', 'sinh',
'sqrt', 'tan', 'tanh']
>>> print math.__doc__
This module is always available.  It provides access to the
mathematical functions defined by the C standard.
>>> print math.sin.__doc__
sin(x)
Return the sine of x (measured in radians).
```

Right away, we can see a list of the functions in the module. We have printed the 'attribute' called __doc__. This gives you the documentation string or comment that the programmer put (as you should put) at the head of his module.

Every function in the module should also have such a documentation attribute. We have printed the one for the `math.sin()` function.

Python built-in functions also have __doc__ attributes. Try typing this at the command line:

`print range.__doc__`

Note: Python uses some special names within modules and functions that begin and end with two underscore characters. Don't use such names for your own variables!

You can also get a clear description of what is in a module using help().

Try importing the math module then type `help(math)`. Once you know what function you need in the module, use `help()` to get a full description of how it works. Try typing `help(math.radians)`.

8.2.6 Testing Functions and Modules

Testing every function and every module that you write is essential. When you write a new function, test it with some test code then put it into a relevant module. A module file contains one or more functions and is designed to be imported, not run on its own. However, it is useful to the programmer who tests and maintains the module to be able to test the functions by just running the module. There is a standard way to do this. The module will have a (hidden) variable called __name__. If you just run a module, this name will be set to the default, which is "__main__". You can thus detect that the user is running the module, rather than importing it, and run a bit of test code. This bit of code should be at the end of the module. You can change it to cause a call to the function being tested, and to print some results as in this example:

```
''' Module to do some fancy stuff '''
def clever(x,y):
    ' Function that is pretty smart'
    < Do clever stuff > # The clever code goes here
    return result


if __name__ == '__main__':
    print 'Testing clever()'
    a = 1.2
    b = 3.4
    answer = clever(a,b)
    print 'a=%.3f, b=%.3f, answer = %.3f' % (a,b,answer)
```

This is a standard way of testing a function in a module. It is also fine of course to write a dedicated test program that imports the module and calls the function. If you leave test code in a module that is not subject to this test on __name__, it will be executed on import to your main program; probably NOT what you want.

8.2.7 pyc Files

When you import a module, in order to speed up execution, Python will 'compile' the module and put the compiled code in a file with the same name but a .pyc extension. 'But this is an interpreted language isn't it?' Well yes. This is not a real compilation but a conversion to something called 'byte code'. Next time you import the module, Python will use this compiled version. If you change the source code of the compiled module, Python will notice and recompile it next time you use it.

The point is that the byte code in the .pyc module will run much faster than if the module is interpreted every time.

You don't have to worry about any of this. Just ignore the .pyc files. If you delete them, Python will re-create them when it needs to

8.3 Examples

8.3.1 Example 8.1

Rather than give several small examples of functions, modules and how they should be organized and tested, here is an example module. It contains a few simple functions and some test code at the end. Here the test code has been put into a function called test() so that it can be called if required.

```
''' Example Module. This module contains some example functions
and some code for testing them '''
# Import system modules
import time
def swap(x, y):
    ''' Swaps values of x and y '''
    temp = x
```

```
    x = y
    y = temp
    return x, y # Returns 2 values


def increment(value, inc=1):
    ''' Increases value by the increment. Default increment is 1 '''
    value = value + inc
    return value # Returns one value


def printMessage(message):
    ''' Print messages to the screen with a time stamp '''
    now = time.asctime()
    print now, ': ', message
    # No return required


def test():
    ''' Perform standard tests on functions in this module '''
    # Test the sway() function.
    value1 = 123; value2 = 987
    printMessage('Testing swap()')
    print 'Initial values:', value1, value2
    value1, value2 = swap(value1, value2)
    print 'Swapped values:', value1, value2


    # Test the increment() function.
    start = 123; theIncrement = 2
    printMessage('Testing increment()')
    print 'Start= ', start, 'Increment= ', theIncrement
    new = increment(start, theIncrement)
    print 'Result= ', new


if __name__ == "__main__":
    # This code will run if the module is run directly.
    # It will not run if the module is imported.
    test() # Run the module standard test function.
    # Any other test code can go here.
```

If the module were saved as simple.py it could be imported into a test program as shown in the following short program that imports the module and runs its test function.

```
''' Program to import and test the module 'simple' '''
import simple
```

```
simple.test() # Call the module test() function
```

If this program were saved as testSimple.py, it could be run from the python prompt:

```
>>> python testSimple.py
```

This example is included just for reference to remind you of how to write functions, arrange them into a module and test them. The exercises will take you through this process step by step.

To get help on this module, one can just import it and use the Python help() function:

```
import simple
```

```
help(simple)
```

8.4 Exercises

8.4.1 Exercise 8.1 (Covered in workshop 3)

Write a function to evaluate the polynomial $x^3 - 7x^2 + 14x - 5$ as described in section 8.1. For this first example, just put the function in a file followed by some code to test it. The function should have x as a single argument. The output should look something like this:

For x= 1.23 y= 3.490567

Model solution is in file: exercise8.1.py on DUO.

8.4.2 Exercise 8.2 (Covered in workshop 3)

Improve the program you wrote for exercise 8.1 by moving the function into a separate file; a new module. Put a documentation string in the module and in the function. Add some standard test code as described in section 8.2.6. Then write a separate test program that imports your new module and uses the function. You can use much the same code for testing within the module and in the stand-alone test program.

Model solution for the module is in file: polyLib.py on DUO.

Model solution for the test code is in file: exercise8.2.py on DUO.

8.4.3 Exercise 8.3 (Covered in workshop 3)

Write a function that takes a list of real numbers and returns both the maximum and the minimum values. Put the function in a module file with some simple test code. Make it work for positive and negative floating point numbers.

Model solution for the module is in file: myLib.py on DUO.

Now write a test program for this function. Generate a set of test values in a list. Print the test data and the maximum and minimum to the screen in the test program. Check that the function works as expected. Test it by filling an array with a range of numbers using either numpy.arange() or numpy.linspace() (Use help!). You should also test it using random data. You can use numpy.random.random_sample()

A model test program is in the file exercise8.3.py on DUO.

Note: There are of course built-in functions called max() and min() in Python so don't call your own variables 'max' or 'min'. Python may not do what you expect! You should not use variable or function names that already exist in Python. If you don't know if Python already uses a word as a function or keyword, try it at the command line.

8.4.4 Exercise 8.4 (Covered in workshop 3)

Write a function that takes a 2-d array (which might represent the image from a CCD detector on a telescope) and returns a *copy* of a small square area of the image (which might contain a galaxy).

The call to the function should look like this:

galaxy = region(image, x, y, size)

where image is the 2-d array, x and y are the centre of the area to be sliced out and size is the size of the square region.

In this example, test your function with code in the same file.

Test it first with a small simple image and check that the area that you slice out of the image is correct by printing it out.

What will you do for a size that is an even number of pixels?

What happens if you hit the edge of the image?

A model program is in the file exercise8.4.py on DUO.

9 Chapter 9 - Debugging and Exceptions

At some point, your program will go wrong! This may be an error that crashed the program and gives you a traceback (the easy ones to find), or it may be an error that means that the program runs to completion but gives you the wrong answer (harder to find). There are various ways to find out where the bug is.

We tend to call all errors in programs 'bugs' (for historical reasons). In Python, errors are properly referred to as 'exceptions'. Dealing with exceptions is an important part of programming. Chapter 8 in Hetland covers this. We give a brief introduction here.

The handling of exceptions is an *advanced topic* so you can skip section 9.4 if you wish. Just note the main debugging methods mentioned in sections 9.1, 9.2 and 9.3.

9.1 Using `print` for debugging

The first and foremost method of debugging is to print out any variables that you are using that might have the wrong values. Just add a print statement to the program, re-run it and check that the variables all have the value that you are expecting:

```
print 'a=', a, ' b=', b # This is a debug line
```

9.2 Use the Command Line

If you get an error that suggests you have used the wrong syntax or called a function incorrectly, repeat what the program was doing when it crashed at the command line. All Python programmers, however skilled, do this. Get it to work at the command prompt then use it in your program.

9.3 Module Test Code

You can test and debug the functions in your module using test code included in the module.

This technique was described in section 8.2.6

9.4 Handling Exceptions: try / except – *Advanced Topic*

You should, of course, try to remove all of the bugs from your program before it is used to solve a problem. However, there are situations where an error could still occur and you don't want the program to crash when it does.

To prevent this happening, there is a construct called `try / except` that tells Python what to do when the error occurs.

You don't need to catch exceptions in this course, just fix every bug!! However, exception handling is a powerful advanced technique. Do try it out.

9.4.1 Catching ALL Exceptions

The handling of exceptions is covered in chapter 8 of Hetland. There is not time in this course to go into exception handling in any detail. However, let's just look at one simple example because the `try / except` construct can be very useful. Try this program:

```
a = input('Enter a number: ')
b = input('Enter another number: ')
print a/b
```

If you enter 0 for the second number, the program will stop and you should get this output:

```
Enter a number: 2

Enter another number: 0

Traceback (most recent call last):

  File "testException.py", line 3, in <module>

    print a/b

ZeroDivisionError: integer division or modulo by zero
```

You can stop the program from crashing by 'catching' the exception:

```
try:

    a = input('Enter a number: ')

    b = input('Enter another number: ')

    print a/b

except: # Catch ALL exceptions

    print 'The second number must not be zero'
```

Run this as before and you should get:

```
Enter a number: 2

Enter another number: 0

The second number must not be zero
```

Here, Python has seen the division by zero, but it has also seen your 'exception handler'. This is the indented code after the except which it has duly executed. The program did not crash! If there is no exception, the exception handling code is just ignored.

Note 1: Take note of the colon used after try and except and the fact that, like a conditional block, the code within the try or except block must be indented.

9.4.2 Catching Specific Exceptions

The above example will catch ALL exceptions that occur. This is generally bad practice. If there is a syntax error in your 'try' block, the exception handler will catch it and give you no traceback!! You should specify the exceptions that are likely to occur.

It is possible to catch only one (or more) specific exceptions. This is well described in Hetland, chapter 8.

Here is some pseudo code as an example:

```
try:

    <Some smart code>

except (<exception1>, <exception2>): # Catch these exceptions only

    print "An exception occurred"
```

Only the `try` block is normally executed. If either of the specified exceptions occurs, the `except` block is executed. If any non-specified exception occurs, the program will crash.

Note 1: If there is more than one specified exception, put them in a tuple, as shown.

Note 2: To find out the name of an exception, make it happen at the command line and look in the traceback:

```
>>> a=2/0
Traceback (most recent call last):
  File "<pyshell#1>", line 1, in <module>
    a=2/0
ZeroDivisionError: integer division or modulo by zero
```

The name of the exception here is 'ZeroDivisionError'.

The code from section 9.4.1 should therefore have been written as below. Note that we have put only the code that might be subject to the error inside the `try` block:

```
a = input('Enter a number: ')
b = input('Enter another number: ')
try:
    print a/b
except ZeroDivisionError: # Catch division by zero
    print 'The second number must not be zero'
```

9.5 Examples

9.5.1 Example 9.1

Here is an example of some data analysis. The user wants to analyse a number of samples from an experimental data set whilst not using values below some noise threshold. They may run out of data. We catch the appropriate error condition and just return to him what data there is along with the number of samples used:

```
''' Program to demonstrate try/except in action '''
import numpy


def analyse(data, num):
    # Analyse 'num' samples from the data.
    sample  = 0 # The index of the sample in the data array
    samplesUsed = 0 # Number of samples used
    try:
        while samplesUsed < num:
            if data[sample] > 0.2: # Apply a threshold.
                result = data[sample] * 1.23 + 2.34 # The algorithm
                samplesUsed += 1
```

```
            sample += 1 # Next sample
    except IndexError: # Catch the indexing error
        pass # Do nothing. Block cannot be empty
    return result, samplesUsed # Return result and samples used


def getTheData():
    # Get all the data to be analysed from file.
    # This is a dummy version that gets some random numbers
    return  numpy.random.rand(113) # Generate 113 random numbers (0 to 1)


if __name__ == "__main__":
    # Test code
    theData = getTheData() # Get the data set to be analysed
    sampleSize = 100
    answer, samples = analyse(theData, sampleSize)
    print 'Used ', samples, ' samples. The result is: ', answer
```

If there is enough data above the noise threshold, the while loop will complete. If not, an IndexError will occur. This is caught by the try/except and the loop is exited gracefully.

The `pass` statement must be used where there would otherwise be an empty block.

Note the use of the random number generator provided by NumPy.

9.6 Exercises

9.6.1 Exercise 9.1 – *Advanced Topic*

Note: This example requires input from the screen, which is covered in chapter 14.

Write a routine that requests input of several floating point values from the screen, catches errors on input and returns the values. Put the function in a module file with some simple test code. The function should take a prompt string as its single parameter. It should return all of the valid numbers that are input as a list of floats. If there is some invalid input, it should return an empty list.

Test the function with some valid input and some invalid input. The two tests should produce output like:

```
Enter the parameters separated by commas:1,2,3
The parameters are: [1.0, 2.0, 3.0]
>>>
Enter the parameters separated by commas:1,2,3,fred
The parameters are: []
```

Hint: Use the `try/except` construct to catch errors where they are likely to occur.

A model solution is in the file exercise9.1.py on DUO.

10 Chapter 10 - Maths Modules: NumPy

As yet, we have not really done much mathematics! The Python language does not include built-in support for maths. However, there are of course modules that we can import to provide mathematics functions. The most basic is the math module. This is generally not enough for the Physicist. By all means load the math module and try it. However, we will use the far more powerful NumPy module.

NumPy and arrays were introduced briefly in section 6.3 but further details are given here.

10.1 The math Module

This provides basic maths functionality: sqrt(), sin(), cos() etc.

10.2 The NumPy Module

A huge amount of development has gone into the development of a mathematics module called NumPy.

Note 1: It is usually written as NumPy but when you import it use: `import numpy` (all lowercase).

The NumPy module has been installed for you on CIS Windows7. It is also bundled with the *Enthought* version of Python that we recommend for download to your own PC.

Hetland does not cover the NumPy module at all. However, as stated previously, there is an excellent web page that provides an introduction at: http://www.numpy.org/.

This web page also refers to another module, SciPy, which provides more complex mathematical functionality, but is heavily dependent on NumPy. Learn to use the facilities of NumPy first. You may like to try the tutorial that can be found at:

http://wiki.scipy.org/Tentative_NumPy_Tutorial

What NumPy gives you is new numerical array and matrix types and a set of basic operations on them. Unlike Python lists, this new *array* type is a set of elements, *all of the same type*. Thus these arrays can be used for vectors and matrices. They provide us with fast matrix and vector manipulation.

One thing to note is that these new types are implemented in pre-compiled code in the C language. The important consequence is that, once you define one, it has some memory allocated to it. You cannot therefore easily increase its overall size once it has been created.

So, as always, let's try some things at the command line:

10.2.1 Creating Arrays

There are various ways to create an array:

You can create an array from a list:

```
>>> import numpy
>>> a = numpy.array([1.0,2.0,3.0])
>>> type(a)
<type 'numpy.ndarray'>
>>> a
```

```
array([ 1.,  2.,  3.])
```

Note 1: The type of an array is actually *ndarray* for 'n-dimensional array'.

Note 2: We can have matrices of any number of dimensions.

- in 2 dimensions (with integers):

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

You can also create an array using the NumPy equivalent of `range()` which is called `arange():`

```
>>> y = numpy.arange(12)
>>> y
array([ 0,  1,  2,  3,  4,  5,  6,  7,  8,  9, 10, 11])
>>> y.shape = 3,4 # Make it a 3 by 4 matrix
>>> y
array([[ 0,  1,  2,  3],
       [ 4,  5,  6,  7],
       [ 8,  9, 10, 11]])
```

Note: You should specify the type of data that you want. This can be done in two ways:

- `y = numpy.arange(10, dtype=float) # Use the dtype parameter`
- `y = numpy.arange(10.0) # Use the type you want for the range`

So `arange()`, unlike `range()`, can take floats as parameters.

In the above example, we created `y` then changed its shape attribute. In maths terms, we started with a 1d vector and changed it into a 3x4 matrix. You can *only* do this if the total number of elements remains unchanged.

There is a function to create an array with all zeros. This can later be filled with the data we want:

```
>>> a=numpy.zeros((3,4), dtype=float)
>>> a
array([[ 0.,  0.,  0.,  0.],
       [ 0.,  0.,  0.,  0.],
       [ 0.,  0.,  0.,  0.]])
```

This function is very useful because we can create an array with the dimensions that we want. It takes the shape of the array as its first parameter as a tuple. We also specified the type. This is good practice. Some versions of NumPy will give you ints by default, others floats.

There is also a function to create a new array of ones:

```
>>> b=numpy.ones((2,2), dtype=float)
>>> b
```

```
array([[ 1.,   1.],
       [ 1.,   1.]])
```

Summary: Creating arrays:

- From a list: `a = numpy.array([1.1,2.2,3.3])`

- Using arange(): `a = numpy.arange(10, dtype=float)`

- Using zeros(): `a = numpy.zeros((10,20), dtype=float)`

10.2.2 Attributes and Methods of an array

Arrays are Python objects so they have attributes and methods. These are accessed using the name of the array followed by a dot followed by the attribute or method. Yes, I know we are not programming using objects, but you need to access the attributes and methods of existing objects. You were warned! Remember that attributes of an object have no brackets after them whereas methods do. Let's create an array and look at some of its attributes:

```
>>> a=numpy.zeros((3,4), dtype=float)
>>> a
array([[ 0.,   0.,   0.,   0.],
       [ 0.,   0.,   0.,   0.],
       [ 0.,   0.,   0.,   0.]])
>>> a.shape # The shape of a (No brackets!!)
(3, 4)
>>> a.ndim # The number of dimensions of a
2
>>> a.size # The overall size of a
12
```

What about methods. Well, there are a lot of them! Let's just look at a few:

First make an array:

```
>>> x=numpy.array([[1.1,2.2,3.3],[4.4,5.5,6.6]])
>>> x
array([[ 1.1,  2.2,  3.3],
       [ 4.4,  5.5,  6.6]])
```

Take the sum of an array (by default over the full array but can be just one axis):

```
>>> x.sum()
23.100000000000001
>>>
```

Hmm… A rounding error. Remember those?

Beware: If you use python's built in function `sum(x)` rather than `x.sum()`, it will handle an array but possibly not in the way you might expect! Try it.

Find the mean of an array:

```
>>> x.mean()
3.8500000000000001
```

Again the rounding error can be seen.

Flatten the array. In maths terms, convert a matrix to a vector:

```
>>> x.flatten()
array([ 1.1,  2.2,  3.3,  4.4,  5.5,  6.6])
```

The full list of array object attributes and methods is in the documentation. As always, use the help(). See section 10.2.6.

10.2.3 Using Arrays

The real power of the array objects provided by NumPy is that they can be manipulated in many ways just like simple variables. IE All of the elements can be manipulated at once and very fast.

We can do some maths on an existing matrix:

```
>>> y
array([[ 0,  1,  2,  3],
       [ 4,  5,  6,  7],
       [ 8,  9, 10, 11]])
>>> a = y*3 # Multiply the entire matrix by 3
>>> a
array([[ 0,  3,  6,  9],
       [12, 15, 18, 21],
       [24, 27, 30, 33]])
```

Without the matrix capabilities of NumPy, you would have to do this multiply one element at a time in two nested `for` loops. Try it! It's a useful exercise to try once.

So, in summary, there are basically three ways of manipulating arrays:

- Just like any other variable type: `y = x * 2`
- Using the methods of an array: `theMean = x.mean()`
- Using the methods of NumPy: `theSum = numpy.sum(x)`

As you may have noticed, there is some overlap between array methods and numpy methods:

`x.sum()` is the same as `numpy.sum(x)`.

10.2.4 Addressing and Slicing Arrays - Review

The elements of an array are addressed in a similar way to lists. However, once an array has more than one dimension, the addressing syntax is different. We saw this in section 6.3.2. Review that section now if you have forgotten the syntax. The general form for a slice of an array x is (in pseudo code):

theSlice = x[<start0>:<end0>:<step0>, <start1>:<end1>:<step1>]

where <start0> is the first element required in dimension 0 and <end0> is one more than the last element required in dimension 0; <step0> is the step or increment to use (default is 1). A comma is used to separate dimensions.

If you miss out <start0> you will get a slice from the first element of that dimension. If you miss out <end0> you will get elements up to the last element in that dimension. If you miss

out <step>, the step will be 1.

I will repeat here the summary of array slicing from section 6.3.2. The 'step' is omitted so is always 1:

Slicing a 2-d array in summary:

- `a[2:4, 7:9]` – Slice out a small piece of an array (y from 2 to 3 and x from 7 to 8)

- `a[:,0]` – Slice out the whole of the first column

- `a[0,:]` – Slice out the whole of the first row

Note that `a[:,:]` is the same as a[:] is the same as `a`!

You can also use negative indices:

`x[2:-1]` gives a slice from element 2 to the last element in that dimension.

Let's make an array and make a few slices:

```
>>> a=numpy.arange(11,2,-1) # Make an array
>>> a.shape=(3,3) # Make it 3x3
>>> a
array([[11, 10,  9],
       [ 8,  7,  6],
       [ 5,  4,  3]])
>>> a[:,2] # Slice out a column
array([9, 6, 3])
>>> a[1,:] # Slice out a row
array([8, 7, 6])
>>> a[1:3,2:3] # Slice a part of the array
array([[6],
       [3]])
```

Note that if you allocate a slice to a new variable, you will NOT get a new array. You will get a pointer to a part of the same array. If you want a copy, use `numpy.copy()` on the slice. This is explained in section 10.2.5.

10.2.5 Copying Arrays

The copying of arrays is slightly complex so, although it is a method like any other, I have given it a special section.

Let's make an array and then try to make a copy of it:

```
>>> x=numpy.array([1,2,3])
>>> y=x
>>> print x,y
[1 2 3] [1 2 3]
```

x and y are the same. We have a copy. So let's change x:

```
>>> x[0] = 99
>>> print x,y
```

```
[99  2  3] [99  2  3]
```

We changed x but y also changed. That is not what we wanted!

The problem is that we did NOT make a new array. In fact, x and y are the same object. Both x and y are actually *pointers* and they point at the same object. We can prove this by:

```
>>> x is y # Ask if x and y are the same object
True
```

Python does this to save on resources. It is much more efficient. Much of the time, different pointers to the same array are just what you need. If you really need an actual copy use numpy.copy():

```
>>> x=numpy.array([1,2,3])
>>> y=numpy.copy(x) # Make a copy of x
>>> print x,y
[1 2 3] [1 2 3]
>>> x[0]=99
>>> print x,y
[99  2  3] [1 2 3]
>>> x is y
False
```

As you can see, we changed x and y did not change. We have a real copy.

If you pass an array to a function as a parameter, Python will actually pass a pointer to the array. Thus if you change the array in the function, the original array in the calling program will also change! *That may or may not be what you want.*

So the rule is: If you need a copy of an array, use `numpy.copy()`. Otherwise, Python will use different pointers to the same array.

10.2.6 Getting Help on NumPy

If you need it, there is a tutorial on-line to get you started. It can be found at www.numpy.org under 'getting started'.

You can get online help on NumPy via the Python `help()` function. The documentation for the NumPy module is in the module along with the code. (This is how you must document your own code, using documentation strings). However, this help is very long and detailed. If you have an idea of the resource you want from NumPy, try `dir(numpy)`. The format of the output is a bit dense but it will list all of the methods.

Example:

I wonder if I can get the mean of all of the contents of an array?

I used dir(numpy) and spotted the function mean().

I then typed help(numpy.mean) and got to see how the function works.

There is also full documentation for NumPy online at www.numpy.org where you will find the NumPy reference Guide.

10.2.7 Linear Algebra

NumPy provides a LOT of useful matrix and vector manipulation routines; too many to list

here. Let's just look at one basic linear algebra routine that you might need:

If you use the normal * for multiplication, NumPy will multiply all of the elements of two matrices together:

```
>>> a = numpy.arange(1,5,dtype=numpy.float)
>>> a
array([ 1.,  2.,  3.,  4.])
>>> a.shape = 2,2
>>> a
array([[ 1.,  2.],
       [ 3.,  4.]])
>>> b = numpy.arange(1,5,dtype=numpy.float64)
>>> b
array([ 1.,  2.,  3.,  4.])
>>> b.shape = 2,2
>>> c=a*b
>>> c
array([[  1.,   4.],
       [  9.,  16.]])
```

This may not be what you want. If you want the dot product, use the function `numpy.dot()`:

```
>>> d = numpy.dot(a,b)
>>> d
array([[  7.,  10.],
       [ 15.,  22.]])
```

Check this is correct by hand.

Many of the more advanced linear algebra functions are contained in a sub-module within NumPy called `numpy.linalg`. Have a look at what is available using:

```
>>> import numpy
>>> dir(numpy.linalg)
['LinAlgError', '__builtins__', '__doc__', '__file__', '__name__',
'__path__', 'cholesky', 'det', 'eig', 'eigh', 'eigvals', 'eigvalsh',
'info', 'inv', 'lapack_lite', 'linalg', 'lstsq', 'norm', 'pinv', 'qr',
'solve', 'svd', 'tensorinv', 'tensorsolve', 'test']
```

or just type `help(numpy.linalg)`.

10.3 The SciPy Module – *Advanced Topic*

The NumPy module contains a large number of useful functions for numerical calculations. The SciPy module extends these facilities. Many of the functions in SciPy are higher level routines that build on the facilities provided by NumPy.

You don't need to use SciPy in this course but you will need to use SciPy routines in your

future courses on Computational Physics. There is documentation and examples on-line at: www.scipy.org. This page has considerable overlap with the numpy.org pages.

10.4 Examples

10.4.1 Example 10.1

I have taken an image with a digital camera. I want to smooth the data by binning it into bins of more than one pixel each. This binning of data is a very common requirement in data analysis. Let's write a function to bin a 2-d array (the image) into bins of any size. Of course, only certain bin sizes will be possible so we will return a zero size array if it can't be done:

```
''' Function to provide binning facilities for an image'''
import numpy


def bin(data, binSize):
    ''' Function to bin a 2-d array into square bins of size 'binSize'.
    If the array cannot be binned, return an empty array'''


    if data.shape[0] % binSize != 0 or data.shape[1] % binSize != 0:
        # This array cannot be binned with this bin size
        return numpy.array([])  # Return a zero length array
    xSize = data.shape[1]/binSize # Find size of binned array
    ySize = data.shape[0]/binSize
    binnedArray = numpy.zeros((ySize, xSize)) # Create the new array
    for y in range(ySize):
        for x in range(xSize):
            # Sum over each bin of the new array
            binnedArray[y, x] = numpy.sum(data[y*binSize:(y+1)*binSize,
                                             x*binSize:(x+1)*binSize])
    return binnedArray


if __name__ == "__main__":
    # Test code
    # Create the image as an array of floats
    image = numpy.array([[ 1., 2., 3., 4., 5., 6.],
                         [ 7., 8., 9.,10.,11.,12.],
                         [13.,14.,15.,16.,17.,18.],
                         [19.,20.,21.,22.,23.,24.]])
    print 'Original image:'
    print image
    result = bin(image, 2)
    if numpy.size(result) > 0:
```

```
        print 'Binned image:'

        print result

    else:

        print "Can't bin that array at that bin size"
```

We have summed a slice of the original array to get each bin value in the binned array. Note the use of an unexpected return value (an array of zero size) to indicate an error to the calling program. We used the `%` operator to get the remainder from an integer division. Beware, this symbol is also used in a different context as a format descriptor.

Python will send the image (an array) to the function as a pointer of course. However, the array returned was a new array created using numpy.zeros(). In some cases it is ok to perform functions on an array without making a copy or creating a new array. This is not one of them!

I have named the indexes into the two dimensions of the array in the opposite way to what we are used to. IE `data[y,x]` instead of `data[x,y]`. You can do either, but you must be consistent. I do it so that y is down the screen when an array is printed.

10.5 Exercises

10.5.1 Exercise 10.1

Write your own routine to find the dot product of two matrices *without* using `numpy.dot()`. This can be done using three nested loops. Test it with two 2x2 matrices and check your result by hand. Test it with two 3x3 matrices and check your result using `numpy.dot()`.

A model solution is in the file exercise10.1.py on DUO.

10.5.2 Exercise 10.2

This is an example of some code that is practical and useful. It is a little harder than previous examples.

When we take an astronomical image, the pupil and image field of view of the telescope is circular due to the circular primary mirror. However, the CCD detector is square. When we come to analyse the CCD data, we don't want to bother with the pixels that are outside the pupil. We therefore need a mask that has a '1' for every pixel inside the field of view, and a '0' outside it.

Write a program that will make such a mask for any square array of data and any pupil radius. There are many ways of doing this. One way is to generate an array that, for every element or pixel, contains the distance of that pixel from the centre of the array. The mask elements are then just '1' when this value is less than the pupil radius and '0' if it is greater.

Hints: There is a useful NumPy method called `numpy.where()`. This sets elements of an array to one value or another depending on a condition.

The 2-d array of distances from the centre can be filled by using two nested loops. This is rather slow. There are various faster ways of doing this without loops using NumPy techniques. To map values from a 1-d array onto a 2-d array, the 1-d array must be extended to 2-d. There is a NumPy facility called 'newaxis' that will add a dummy dimension to an array.

A model solution is in the file exercise10.2.py on DUO. This module contains two functions. One is based on using loops and the other, faster function uses NumPy facilities. If there is no

requirement to go fast then by all means do it the slow way with loops.

11 Chapter 11 - File Input and Output – The Details

I have not marked this chapter as an advanced topic because you should know how to do input/output from/to files. However, you can simply use the Numpy techniques of section 4.2 for now. *Read this chapter only when you need it.*

So far, we have looked mostly at input and output to the screen. You will often want to do some calculation and output the results to a file. Similarly, you may need to read some data from a file in order to analyse it. You don't have to worry about how the computer stores its data on the hard disk. The operating system will look after that. All you have to do is tell Python which file you want to use, and how you want to use it.

We gave a short introduction to doing this in section 4.2 that allowed you to read and write NumPy arrays to/from file. However, you may want to output quite complex mixtures of text and numbers to a file. The details of file IO are given here and there is a self-contained short chapter (11) in Hetland on 'Files and Stuff'.

11.1 Line Terminators – The \n character

It is often useful to read individual *lines* of text from a file. The end of a line is indicated by a special character or characters. The characters used vary from language to language and for different operating systems. Python uses the character \n (backslash n). This represents a single End of Line (EOL) character.

Try this at the command line:

```
>>> a='This is the first line. This is the second'
>>> print a
This is the first line. This is the second
>>> a='This is the first line.\nThis is the second.'
>>> print a
This is the first line.
This is the second.
```

As you can see, the print function translates the \n into a new line on the screen.

Python is very good at dealing with new line characters. Unix also uses \n so no problem there. Windows7 does not. However, Python on Windows7 will do all the translating for you so you don't have to worry. Just use \n when you want a new line.

When writing text to a file, separate the lines with a \n. You can then use the readline() and readlines() functions when you read the file.

11.2 Writing to File

Let's first write some text to a file, then we can try reading it back in various ways. First of all, you must open a file. Not surprisingly, we open a file using the `open()` function. It will return a file object. We use this object and its methods to access the file. We are again slipping into using objects and methods rather than functions. That is how Python works and you know the syntax so just use the methods provided.

Note 1: Before you write to a file, remember that the file will be written in the directory in which you are working unless you give a full path name. In Windows7, the current directory

is indicated in the system prompt. EG J:\myname\mywork>

To write to a file, try this at the python prompt:

```
>>> f = open('temp.txt', 'w')
>>> f
<open file 'temp.txt', mode 'w' at 0x0138D260>
```

The variable f is a file pointer object.

The second parameter of the open function tells Python how to use the file:

'r' => You may only read from this file. This is usually the default.

'w' => You may only write to the file

'a' => You may append to the file. IE You will write *after* any data already in the file.

Note 2: It is also possible to read and write to a file using the 'r+' or 'w+' modes in combination with the `seek()` method to read or write to any point within the data in an existing file. You probably won't need this yet so it is not described here.

Now write some lines to the file:

```
>>> f.write('This is the first line.\n')
>>> f.write('This is the second line.\n')
>>> f.write('This is yet another boring line.\n')
>>> f.close()
```

Notice that we close the file when we have finished with it. You should not leave files open.

You will find that you now have a file called temp.txt in your working directory. (Check under Windows7 with the `dir` command). Try printing the file to the screen– type `more temp.txt`.

Note 3: After each write to the file, we have put in a \n to split the file into lines.

Note 4: By default, Python will always write from the beginning of a file that is opened in 'w' mode. So *beware* of overwriting any existing contents.

Note 5: All of your data should be written to file as strings. Thus you must convert any numbers to strings before writing them to file. You can write numbers to file in 'binary' mode but this is not needed for this course.

11.3 Reading from File

Reading from a file is equally easy. You must however tell Python how much of the file you want to read. In general you can:

Read a fixed number of bytes from a file using the `read(num)` method where `num` is the number of bytes to read. This is not very useful!

Read a single line from a file using the `readline()` method.

Read all the lines in a file using the `readlines()` method. They will be returned as a Python list.

Note: After a read, Python keeps a pointer to where you got to in the file. If you do another read, it will start from there.

Try reading from the file that you created from the command line:

```
>>> f = open('temp.txt')
>>> f.readline()
'This is the first line.\n'
>>> f.readlines()
['This is the second line.\n', 'This is yet another boring line.\n']
```

11.4 Exercises

The first two exercises expect you to output some mixed data to a file. We output the number of records as an integer 'header' to the file. These exercises are intended to give you some practice at detailed file IO. Do the first two exercises *without* using the NumPy savetxt() and loadtxt() functions. Then repeat the exercises using the NumPy functions. Refer back to section 4.2 to remind you how to do this.

11.4.1 Exercise 11.1

Write a program to calculate the sine of some angles from 0 to 100 degrees. (Remember to convert to radians before taking the sine). Use numpy arrays for this. Output the data to the screen so that you can see it as it is generated. Output the data to file. On the first line of the file, output the number of records that you are writing to file (100). This acts as a useful 'header' so that you can easily find out how much data is in the file. Output the data to file with one angle in degrees and its sin on each line. Separate the data with a comma. You can use any separator, but a comma is usual. You are creating a 'comma separated variable' file or CSV file. Look at the file you have produced and check it is correct.

Model solution is in file exercise11.1.py on DUO.

11.4.2 Exercise 11.2

Write a program to read your angle and sin data from the file into numpy arrays and write every 10[th] angle and its sine to the screen with one angle and its sine on each line in the format:

```
angle =  0.000 degrees      Sin = 0.000
angle = 10.000 degrees      Sin = 0.174
```

etc

Read the 'header' to find out how much data is in the file and use this to size your arrays.

Try to get the output data to line up as shown using the % format descriptor.

Model solution is in file exercise11.2.py on DUO.

11.4.3 Exercise 11.3

Repeat exercise 11.1 using the NumPy savetxt() function but do not write the number of records as a header since this is not needed.

Model solution is in file exercise11.3.py on DUO.

11.4.4 Exercise 11.4

Repeat exercise 11.2 using the NumPy loadtxt() function.

Model solution is in file exercise11.4.py on DUO.

12 Chapter 12 - Plotting Graphs

Python does not have any facilities built in for drawing graphs. However, everyone wants to draw graphs so there are many modules around that do so. The matplotlib package is one of the most common and has been provided for you on CIS Windows7. It goes by the name of PyLab so, if you want a graph, just `import pylab`. The module provides a nice Graphical User Interface (GUI) for you to use to display your graph. Pylab just provides some higher level facilities such as the GUI. It imports matplotlib so just import pylab and you will get both.

There is a good introduction to matplotlib on the web at:

http://matplotlib.org

Look at the 'Gallery' to see examples of what can be done using this package.

As with all things Python, it is easy to draw a simple graph very quickly. The module provides all the extra facilities that you might need such as labels for axes, different coloured symbols, different types of plot etc. You can slowly learn these and use them to produce good looking clear displays of your data. Just a few of the basics will be enough for this course. It is worth gaining some experience with this package as you will be expected to use it to draw graphs in later courses in level 2 and in level 3.

We do not explain any of the detail of PyLab here. You should get that from the documentation on the web. We include some general points of interest and an example for you to try. Plotting graphs is great fun but don't spend too much time on it!

12.1 PyLab: The absolute Basics

What do we always do with something new in Python? Try it at the Python prompt. The basic plot command is `plot(a,b)` where a and b are lists (or arrays) of data to plot in x and y respectively.

```
>>> import pylab
>>> a=[1,2,3]
>>> b=[2,4,6]
>>> pylab.plot(a,b)
[<matplotlib.lines.Line2D instance at 0x01A03170>]
>>> pylab.show()
```

You should now see a plot popup something like this:

This is an example of a Graphical User Interface or GUI. It has been written for you. You are just using it from your program. PyLab has plotted your two lists of data against each other. Kill the GUI window and you should get your Python prompt back.

Note: PyLab has been set up for you to work from a file rather than *interactively* from the command line so you should develop any real programs in a script file as usual.

12.2 The PyLab Graphical User Interfaces (GUI)

The last thing you do in your program, if you are going to plot a graph, is to issue the command `pylab.show()`. This causes pylab to draw the Graphical User Interface (GUI) that is used to display your graph. It then enters an infinite loop waiting for you to use the controls on the GUI. Thus you won't get your Python prompt back until you kill the GUI window. If you have problems, you can usually cause the program to end with a Control C (Hold down control and hit c).

Advanced topic:

Writing your own GUIs (like PyLab) is great fun but outside the remit of this course. Chapter 12 of Hetland is a good introduction. To write a GUI requires an extra module called a 'GUI toolkit' that we have not installed for you on CIS Windows7. *wxPython* is my current favourite. You can download it to your own computer if you want to try it. You really need to know about object oriented programming first. Warning: It can be addictive and time

consuming!

12.3 Exercises

12.3.1 Exercise 12.1

Modify your program that writes a file of angles and their sines to cover the range 0 to 360 degrees. Modify your program that reads the file to plot this full cycle of the sine function using PyLab. You should give the axes legends and add a title to the graph. Output the graph to a .png file. This could then be added as a figure in a report written in Word or another word processor.

Model solution is in file exercise12.1.py on DUO.

13 Chapter 13 - Random Numbers

There are many problems in Physics where the Physics is well understood but describing an entire system that we wish to model is not possible analytically. In this situation, the Physicist will often resort to a method known as 'Monte Carlo'. This involves using random numbers to repeatedly test what happens to, for example, a cosmic ray interacting with the Earth's atmosphere. (It's like throwing the dice in a casino, hence Monte Carlo).

Python (like all languages) incorporates a "random" number generator. Of course such numbers cannot be truly random. Python can however generate a "pseudo-random" sequence of numbers. We will not address the use of random numbers in any detail here, but you should be aware that there is a module, called 'random', that provides random numbers in a wide variety of different ways.

Import the random module and use dir(random) to see what calls are available. The module is described in detail in the Python documentation in the Module Index at:

http://docs.python.org/modindex.html

The most basic call is `random.random()`. This will return a random floating point number between 0.0 and 1.0.

However, the most useful is probably `random.uniform(<min>, <max>)` which returns a random floating point number between min and max.

The simplest distribution of random numbers you might want would be a 'flat' distribution. One could generate a large number of random numbers using say `uniform(0.0, 100.0),` bin them in some way and do some tests to see if they are truly random.

One of the other most useful distributions of random numbers is a Gaussian distribution. The module `random` provides the function `random.gauss( <mu>, <sigma>)` to give you just such a distribution where `mu` is the mean and `sigma` is the standard deviation.

Not to be left out of this technique, NumPy also provides random number facilities in a sub module called numpy.random.

13.1 Exercises

13.1.1 Exercise 13.1

Write a program to request a number of Gaussian distributed random numbers. Put them in an array and make a histogram of the array. Use a mean of 100.0 and a standard deviation of 15.0. You can normalise the output so that the integral is 1.0 and the Y axis will represent probability. This also makes the y axis scaling easier for different numbers of random numbers.

Hint: There is a pylab.hist() function to plot histograms.

Model solution is in file exercise13.1.py on DUO.

14   Chapter 14 – Input from the screen

Whilst input from the screen is not an advanced topic, *you do not need it for this course*. However, it is a very useful skill to acquire, so do return to this section when you find a need to enter data to your program from the screen.

14.1 Input from the Screen

Sometimes you will want the user of your program to type in some input from the terminal. You can 'prompt' them to do so by outputting a string and then wait for his input. To do this, there is an `input()` function. This can be used to input directly to a Python variable.

```
>>> a=input("Input a number: ")
Input a number: 123.45
>>> print a
123.45
```

This is fine for the input of numbers but not great for the input of strings. The problem is that `input()` assumes that what you type is valid Python.

```
>>> b=input("Input your name: ")
Input your name: Ron
Traceback (most recent call last):
  File "<pyshell#2>", line 1, in <module>
    b=input("Input your name: ")
  File "<string>", line 1, in <module>
NameError: name 'Ron' is not defined
```

To avoid this error, the user must enter the string with quotes around it thus:

```
>>> b=input('Input your name: ')
Input your name: "Ron"
>>> print b
Ron
```

To avoid this, you can use the function `raw_input()`. (See Hetland Page 26). Doing this however means that your input is **all** treated as a string and, if you want another type, you must convert it:

```
>>> a = float(raw_input('Type a number: '))
Type a number: 23
>>> print a
23.0
```

Which input function you use is up to you.

The `raw_input()` and `input()` functions will use the string that you give it as a parameter as a 'prompt'. It writes this on the screen and then waits for some input terminated by a 'new line' (The 'enter' key).

Note: You can apply a function to the result of another function directly as shown above. `raw_input()` returns a string that is the input to the function `float()`.

14.2 Input of Multiple values from the Screen

If you want to input say three different numbers from the screen, this can again be done using input() or raw_input. Using input() is easy but restrictive:

```
>>> b=input("Enter 3 numbers separated by commas: ")
Enter 3 numbers separated by commas: 2,3,4
>>> b
(2, 3, 4)
```

The numbers must be numbers and must be separated by commas.

Note that the 3 values are returned as a tuple.

If you want full control of the input, use raw_input() as before. The input will be a string. You can split up the string using the split() method (See section 6.2):

```
>>> a=raw_input("Enter 3 numbers: ")
Enter 3 numbers: 3 4 5
>>> a
'3 4 5'
>>> b=a.split(' ') # Split up the string using 'space' as a separator
>>> b # Print b. It is a list of strings
['3', '4', '5']
>>> first = float(b[0]) # Convert one element of the list into a float
>>> first
3.0
```

This may seem rather longwinded (and it is!) but it does give you more control over what the user inputs. You could use any character for the separator.