
Introduction to Python for Computational Science and Engineering

Hans Fangohr

Apr 14, 2024

CONTENTS

1	Introduction	3
1.1	Computational Modelling	3
1.2	Why Python for scientific computing?	5
1.3	Python version	7
1.4	These documents	8
1.5	Your feedback	9
2	A powerful calculator	11
2.1	Python prompt and Read-Eval-Print Loop (REPL)	11
2.2	Calculator	11
2.3	Integer division	13
2.4	Mathematical functions	14
2.5	Variables	16
2.6	Impossible equations	18
3	Data Types and Data Structures	21
3.1	What type is it?	21
3.2	Numbers	21
3.3	Sequences	24
3.4	Passing arguments to functions	37
3.5	Equality and Identity/Sameness	43
4	Introspection	47
4.1	dir	47
4.2	type	51
4.3	isinstance	51
4.4	help	52
4.5	Docstrings	59
5	Input and Output	61
5.1	Printing to standard output (normally the screen)	61
5.2	Reading and writing files	67
5.3	Further reading	69
6	Control Flow	71
6.1	Basics	71
6.2	If-then-else	74
6.3	For loop	75
6.4	While loop	75
6.5	Relational operators (comparisons) in <code>if</code> and <code>while</code> statements	76
6.6	Exceptions	77

7	Functions and modules	81
7.1	Introduction	81
7.2	Using functions	81
7.3	Defining functions	82
7.4	Default values and optional parameters	84
7.5	Modules	85
7.6	Further Reading	89
8	Functional tools	91
8.1	Anonymous functions	91
8.2	Map	92
8.3	Filter	93
8.4	List comprehension	94
8.5	Reduce	95
8.6	Why not just use for-loops?	98
8.7	Speed	98
8.8	The <code>%%timeit</code> magic	100
9	Common tasks	103
9.1	Many ways to compute a series	103
9.2	Sorting	106
10	From Matlab to Python	109
10.1	Important commands	109
11	Python shells	111
11.1	IDLE	111
11.2	Python (command line)	111
11.3	Interactive Python (IPython)	111
11.4	Spyder	113
11.5	Editors	113
12	Symbolic computation	115
12.1	SymPy	115
12.2	Related tools	133
13	Numerical Computation	135
13.1	Numbers and numbers	135
14	Numerical Python (numpy): arrays	143
14.1	Numpy introduction	143
15	Visualising Data	151
15.1	Matplotlib – plotting $y=f(x)$, (and a bit more)	151
15.2	The pylab interface	158
15.3	The matplotlib.pyplot interface	165
15.4	Visual Python	173
15.5	Visualising higher dimensional data (VTK)	177
15.6	Further tools and developments	179
15.7	Jupyter Notebooks	182
16	Numerical Methods using Python (scipy)	183
16.1	Overview	183
16.2	SciPy	183
16.3	Numerical integration	185

16.4	Solving Ordinary Differential Equations (ODEs)	187
16.5	Root finding	197
16.6	Interpolation	199
16.7	Curve fitting	200
16.8	Fourier transforms	203
16.9	Optimisation	204
16.10	Other numerical methods	206
16.11	scipy.io: Scipy-input output	207
17	Pandas - Data Science with Python	209
17.1	Motivational example (Series)	209
17.2	Pandas Series	211
17.3	Create Series from list	217
17.4	Plotting data	217
17.5	Missing values	220
17.6	Series data access: explicit and implicit (<code>loc</code> and <code>iloc</code>)	221
17.7	Data Frame	224
17.8	Example: European population 2017	228
17.9	Further reading	240
18	Python packages and environments	241
18.1	Introduction	241
18.2	Python virtual environments	242
18.3	Python Package Index (PyPI)	244
18.4	Anaconda	253
18.5	Managing many different environments - <code>pyenv</code>	255
19	Where to go from here?	257
19.1	Advanced programming	257
19.2	Compiled programming language	257
19.3	Testing	257
19.4	Simulation models	258
19.5	Software engineering for research codes	258
19.6	Data and visualisation	258
19.7	Version control	258
19.8	Parallel execution	258
19.9	Acknowledgements	258
20	Change history	261

The content of [this book](#) is distributed into chapters, using one Jupyter Notebook for each chapter.

You can read the book in different formats: [html](#), [pdf](#), or you can use the [myBinder](#) environment, in which you can read the text and execute the examples in a browser (without having to install Python locally) using one Jupyter notebook per chapter, as mentioned above.

If you have not used the Jupyter Notebook before, please read the section “First steps with Jupyter Notebook” below before you proceed.

First steps with Jupyter Notebook

1. Navigating the notebook

When you open a notebook, you will find that you can move a highlighted block (with a blue line at the left) with the cursor keys to move up and down. This block highlights a *cell*. (You can also use the mouse to select a cell.) This is called the [Command mode](#).

2. Executing code

If you want to execute a cell (for example one that contains some Python code), you can press `Shift+ENTER`. If the cell creates some output, it will be displayed below the cell. (You may not notice if it just updates output that was displayed before, in particular if the new output is the same as the old output.)

3. Editing code

If you want to *change* the code in the currently highlighted cell, you need to press `ENTER`. You have now entered the [Editing mode](#), and the content of the cell can be edited. If you have completed your changes, and you want to execute them, use the `Shift+ENTER` short cut.

Note that you can also edit blocks of text (or go into edit mode for a text cell unintentionally). Just press `Shift+ENTER` to render the text again, and go back into command mode.


Warning: Changes on myBinder are temporary

If you use this text book interactively on the myBinder service, then you have been given a temporary resource in the cloud to execute the code examples. The changes you have made to the notebook *will be lost* when your session ends (which is when you close the window, or the patience of the service has expired). Thus, the interactive exploration of the notebooks is good to help learn Python, computing and data science, but you should not attempt to write any code in these notebooks that you want to re-use the next day or later.

Comments? Questions?

For feedback, corrections, and questions please refer to the [home page \(https://github.com/fangohr/introduction-to-python-for-computational-science-and-engineering/blob/master/Readme.md\)](https://github.com/fangohr/introduction-to-python-for-computational-science-and-engineering/blob/master/Readme.md) of the book. You can also find the most recent versions there.

Enjoy!



INTRODUCTION

This text summarises a number of core ideas relevant to Computational Engineering and Scientific Computing using Python. The emphasis is on introducing some basic Python (programming) concepts that are relevant for numerical algorithms. The later chapters touch upon numerical libraries such as `numpy` and `scipy` each of which deserves much more space than provided here. We aim to enable the reader to learn independently how to use other functionality of these libraries using the available documentation (online and through the packages itself).

1.1 Computational Modelling

1.1.1 Introduction

Increasingly, processes and systems are researched or developed through computer simulations: new aircraft prototypes such as for the recent A380 are first designed and tested virtually through computer simulations. With the ever increasing computational power available through supercomputers, clusters of computers and even desktop and laptop machines, this trend is likely to continue.

Computer simulations are routinely used in fundamental research to help understand experimental measurements, and to replace – for example – growth and fabrication of expensive samples/experiments where possible. In an industrial context, product and device design can often be done much more cost effectively if carried out virtually through simulation rather than through building and testing prototypes. This is in particular so in areas where samples are expensive such as nanoscience (where it is expensive to create small things) and aerospace industry (where it is expensive to build large things). There are also situations where certain experiments can only be carried out virtually (ranging from astrophysics to study of effects of large scale nuclear or chemical accidents). Computational modelling, including use of computational tools to post-process, analyse and visualise data, has been used in engineering, physics and chemistry for many decades but is becoming more important due to the cheap availability of computational resources. Computational Modelling is also starting to play a more important role in studies of biological systems, the economy, archeology, medicine, health care, and many other domains.

1.1.2 Computational Modelling

To study a process with a computer simulation we distinguish two steps: the first one is to develop a *model* of the real system. When studying the motion of a small object, such as a penny, say, under the influence of gravity, we may be able to ignore friction of air: our model — which might only consider the gravitational force and the penny's inertia, i.e. $a(t) = F/m = -9.81\text{m/s}^2$ — is an approximation of the real system. The model will normally allow us to express the behaviour of the system (in some approximated form) through mathematical equations, which often involve ordinary differential equations (ODEs) or partial differential equations (PDEs).

In the natural sciences such as physics, chemistry and related engineering, it is often not so difficult to find a suitable model, although the resulting equations tend to be very difficult to solve, and can in most cases not be solved analytically at all.

On the other hand, in subjects that are not as well described through a mathematical framework and depend on behaviour of objects whose actions are impossible to predict deterministically (such as humans), it is much more difficult to find a good model to describe reality. As a rule of thumb, in these disciplines the resulting equations are easier to solve, but they are harder to find and the validity of a model needs to be questioned much more. Typical examples are attempts to simulate the economy, the use of global resources, the behaviour of a panicking crowd, etc.

So far, we have just discussed the development of *models* to describe reality, and using these models does not necessarily involve any computers or numerical work at all. In fact, if a model's equation can be solved analytically, then one should do this and write down the solution to the equation.

In practice, hardly any model equations of systems of interest can be solved analytically, and this is where the computer comes in: using numerical methods, we can at least study the model *for a particular set of boundary conditions*. For the example considered above, we may not be able to easily see from a numerical solution that the penny's velocity under the influence of gravity will change linearly with time (which we can read easily from the analytical solution that is available for this simple system: $v(t) = t \cdot 9.81\text{m/s}^2 + v_0$).

The numerical solution that can be computed using a computer would consist of data that shows how the velocity changes over time for a particular initial velocity v_0 (v_0 is a boundary condition here). The computer program would report a long lists of two numbers keeping the (i) value of time t_i for which a particular (ii) value of the velocity v_i has been computed. By plotting all v_i against t_i , or by fitting a curve through the data, we may be able to understand the trend from the data (which we can just see from the analytical solution of course).

It is clearly desirable to find an analytical solutions wherever possible but the number of problems where this is possible is small. Usually, the obtaining numerical result of a computer simulation is very useful (despite the shortcomings of the numerical results in comparison to an analytical expression) because it is the only possible way to study the system at all.

The name *computational modelling* derives from the two steps: (i) *modelling*, i.e. finding a model description of a real system, and (ii) solving the resulting model equations using *computational* methods because this is the only way the equations can be solved at all.

1.1.3 Programming to support computational modelling

A large number of packages exist that provide computational modelling capabilities. If these satisfy the research or design needs, and any data processing and visualisation is appropriately supported through existing tools, one can carry out computational modelling studies without any deeper programming knowledge.

In a research environment – both in academia and research on new products/ideas/... in industry – one often reaches a point where existing packages will not be able to perform a required simulation task, or where more can be learned from analysing existing data in new ways etc.

At that point, programming skills are required. It is also generally useful to have a broad understanding of the building blocks of software and basic ideas of software engineering as we use more and more devices that are software-controlled.

It is often forgotten that there is nothing the computer can do that we as humans cannot do. The computer can do it much faster, though, and also with making far fewer mistakes. There is thus no magic in computations a computer carries out: they could have been done by humans, and – in fact – were for many years (see for example Wikipedia entry on [Human Computer](#)).

Understanding how to build a computer simulation comes roughly down to: (i) finding the model (often this means finding the right equations), (ii) knowing how to solve these equations numerically, (iii) to implement the methods to compute these solutions (this is the programming bit).

1.2 Why Python for scientific computing?

The design focus on the Python language is on productivity and code readability, for example through:

- Interactive python console
- Very clear, readable syntax through whitespace indentation
- Strong introspection capabilities
- Full modularity, supporting hierarchical packages
- Exception-based error handling
- Dynamic data types & automatic memory management

As Python is an interpreted language, and it runs many times slower than compiled code, one might ask why anybody should consider such a 'slow' language for computer simulations?

There are two replies to this criticism:

1. *Implementation time versus execution time:* It is not the execution time alone that contributes to the cost of a computational project: one also needs to consider the cost of the development and maintenance work.

In the early days of scientific computing (say in the 1960/70/80), compute time was so expensive that it made perfect sense to invest many person months of a programmer's time to improve the performance of a calculation by a few percent.

Nowadays, however, the CPU cycles have become much cheaper than the programmer's time. For research codes which often run only a small number of times (before the researchers move on to the next problem), it may be economic to accept that the code runs only at 25% of the expected possible speed if this saves, say, a month of a researcher's (or programmers) time. For example: if the execution time of the piece of code is 10 hours, and one can predict that it will run about 100 times, then the total execution time is approximately 1000 hours. It would be great if this could be reduced to 25% and one could save 750 (CPU) hours. On the other hand, is an extra wait (about a month) and the cost of 750 CPU hours worth investing one month of a person's time [who could do something else while the calculation is running]? Often, the answer is not.

Code readability & maintenance - short code, fewer bugs: A related issue is that a research code is not only used for one project, but carries on to be used again and again, evolves, grows, bifurcates etc. In this case, it is often justified to invest more time to make the code fast. At the same time, a significant amount of programmer time will go into (i) introducing the required changes, (ii) testing them even before work on speed optimisation of the changed version can start. To be able to maintain, extend and modify a code in often unforeseen ways, it can only be helpful to use a language that is easy to read and of great expressive power.

2. *Well-written Python code can be very fast* if time critical parts in executed through compiled language.

Typically, less than 5% percent of the code base of a simulation project need more than 95% of the execution time. As long as these calculations are done very efficiently, one doesn't need to worry about all other parts of the code as the overall time their execution takes is insignificant.

The compute intense part of the program should be tuned to reach optimal performance. Python offers a number of options.

- For example, the `numpy` Python extension provides a Python interface to the compiled and efficient LAPACK libraries that are the quasi-standard in numerical linear algebra. If the problems under study can be formulated such that eventually large systems of algebraic equations have to be solved, or eigenvalues computed, etc, then the compiled code in the LAPACK library can be used (through the Python-numpy package). At this stage, the calculations are carried out with the same performance of Fortran/C as it is essentially Fortran/C code that is used. Matlab, by the way, exploits exactly this: the Matlab scripting language is very slow (about 10 time slower than Python), but Matlab gains its power from delegating the matrix operation to the compiled LAPACK libraries.

- Existing numerical C/Fortran libraries can be interfaced to be usable from within Python (using for example Swig, Boost.Python and Cython).
- Python can be extended through compiled languages if the computationally demanding part of the problem is algorithmically non-standard and no existing libraries can be used.

Commonly used are C, Fortran and C++ to implement fast extensions.

- We list some tools that are used to use compiled code from Python:
 - The `scipy.weave` extension is useful if just a short expression needs to be expressed in C.
 - The Cython interface is growing in popularity to (i) semi-automatically declare variable types in Python code, to translate that code to C (automatically) and to then use the compiled C code from Python. Cython is also used to quickly wrap an existing C library with an interface so the C library can be used from Python.
 - Boost.Python is specialised for wrapping C++ code in Python.

The conclusion is that Python is “fast enough” for most computational tasks, and that its user friendly high-level language often makes up for reduced speed in comparison to compiled lower-level languages. Combining Python with tailor-written compiled code for the performance critical parts of the code, results in virtually optimal speed in most cases.

1.2.1 Optimisation strategies

We generally understand reduction of execution time when discussing “code optimisation” in the context of computational modelling, and we essentially like to carry out the required calculations as fast as possible. (Sometimes we need to reduce the amount of RAM, the amount of data input output to disk or the network.) At the same time, we need to make sure that we do not invest inappropriate amounts of programming time to achieve this speed up: as always there needs to be a balance between the programmers’ time and the improvement we can gain from this.

1.2.2 Get it right first, then make it fast

To write fast code effectively, we note that the right order is to (i) first write a program that carries out the correct calculation. For this, choose a language/approach that allows you to *write the code quickly and make it work quickly* — regardless of execution speed. Then (ii) either change the program or re-write it from scratch in the same language to make the execution faster. During the process, keep comparing results with the slow version written first to make sure the optimisation does not introduce errors. (Once we are familiar with the concept of regression tests, they should be used here to compare the new and hopefully faster code with the original code.)

A common pattern in Python is to start writing pure Python code, then start using Python libraries that use compiled code internally (such as the fast arrays Numpy provides, and routines from `scipy` that go back to established numerical codes such as ODEPACK, LAPACK and others). If required, one can – after careful profiling – start to replace parts of the Python code with a compiled language such as C and Fortran to improve execution speed further (as discussed above).

1.2.3 Prototyping in Python

It turns out that – even if a particular code has to be written in, say, C++ – it is (often) more time efficient to prototype the code in Python, and once an appropriate design (and class structure) has been found, to translate the code to C++.

1.2.4 Literature

While this text starts with an introduction of (some aspects of) the basic Python programming language, you may find - depending on your prior experience - that you need to refer to secondary sources to fully understand some ideas.

We repeatedly refer to the following documents:

- Allen Downey, *Think Python*. Available online in html and pdf at <https://www.greenteapress.com/thinkpython/thinkpython.html>, or from Amazon.
- The Python documentation <https://www.python.org/doc/>, and:
- The Python tutorial (<https://docs.python.org/3/tutorial/>)

You may also find the following links useful:

- The `numpy` home page (<https://numpy.org/>)
- The `scipy` home page (<https://www.scipy.org/>)
- The `matplotlib` home page (<https://matplotlib.org/>).
- The Python style guide (<https://www.python.org/dev/peps/pep-0008/>)

1.2.5 Recorded video lectures on Python for beginners

Do you like to listen/follow lectures? There is a series of 24 lectures titled *Introduction to Computer Science and Programming* delivered by Eric Grimson and John Guttag from the MIT available at <https://ocw.mit.edu/courses/electrical-engineering-and-computer-science/6-00-introduction-to-computer-science-and-programming-fall-2008/>

This is aimed at students with little or no programming experience. It aims to provide students with an understanding of the role computation can play in solving problems. It also aims to help students, regardless of their major, to feel justifiably confident of their ability to write small programs that allow them to accomplish useful goals.

An more recent [collection of topic specific \(and shorter\) tutorial videos](#) is available from Socratica .

1.2.6 Python tutor mailing list

There is also a Python tutor mailing list (<https://mail.python.org/mailman/listinfo/tutor>) where beginners are welcome to ask questions regarding Python. Both using the archives and posting your own queries (or in fact helping others) may help with understanding the language. Use the normal mailing list etiquette (i.e. be polite, concise, etc). You may want to read <https://www.freebsd.org/doc/en/articles/mailling-list-faq/etiquette.html> for some guidance on how to ask questions on mailing lists.

1.3 Python version

There are two version of the Python language out there: Python 2.x and Python 3.x. They are (slightly) different — the changes in Python 3.x were introduced to address shortcomings in the design of the language that were identified since Python's inception. A decision was made that some incompatibility should be accepted to achieve the higher goal of a better language for the future.

For scientific computation, it is crucial to make use of numerical libraries such as `numpy`, `scipy` and the plotting package `matplotlib`.

All of these are now available for Python 3, and we will use Python 3.x in this book.

However, there is a lot of code still in use that was written for Python 2, and it's useful to be aware of the differences. The most prominent example is that in Python 2.x, the `print` command is special, whereas in Python 3 it is an ordinary function. For example, in Python 2.7, we can write:

```
print "Hello World"
```

where as in Python 3, this would cause a `SyntaxError`. The right way to use `print` in Python 3 would be as a function, i.e.

```
print("Hello World")
```

```
Hello World
```

See *Chapter 5: Input and Output* for further details.

Fortunately, the function notation (i.e. with the parantheses) is also allowed in Python 2.7, so our examples should execute in Python 3.x and Python 2.7. (There are other differences.)

1.4 These documents

This material has been converted from Latex to a set of [Jupyter Notebooks](#), making it possible to interact with the examples. You can run any code block with an `In []:` prompt by clicking on it and pressing shift-enter, or by clicking the button in the toolbar.

Code blocks can be recognised (in the html and pdf version of this book) by having coloured items (to emphasise the syntactic role). For example:

```
for i in range(3):  
    print("Hello")
```

```
Hello  
Hello  
Hello
```

The *output* (here `Hello` repeated on three lines) that is produced by the code block, is shown below the code block and not coloured.

1.4.1 The `%%file` magic

We use some features in the notebook that are worth being aware of at this point: a cell starting with the special command `%%file FILENAME` will create (or override) a file with name `FILENAME` that contains the content that is shown in the cell below.

For example

```
%%file hello.txt  
This is the content of the file hello.txt
```

```
Overwriting hello.txt
```

To confirm the file has been written and contains, we use some Python commands (which you are not expected to understand at this point):

```
with open("hello.txt") as f:  
    print(f.read())
```

```
This is the content of the file hello.txt
```

1.4.2 The ! to execute shell commands

If we want to run a shell command, we can type it and precede it by the ! character. Here is an example: first we create a file that contains a Python hello world program, then we execute it:

```
%%file hello.py  
print("Hello World")
```

```
Overwriting hello.py
```

```
!python hello.py
```

```
Hello World
```

1.4.3 The #NBVAL tags

In some cells, you will find tags like #NBVAL_SKIP, #NBVAL_IGNORE_OUTPUT and #NBVAL_RAISES_EXCEPTION. You can ignore these.

(We use them to be able to [automatically execute all notebooks](#) to check that the output produced is the same as what is stored in the notebook. This is an advanced topic of testing, and you can read more about NBVAL at <https://github.com/computationalmodelling/nbval>).

See [Chapter 11](#) for more information on Jupyter and other Python interfaces.

1.5 Your feedback

is desired. If you find anything wrong in this text, or have suggestions how to change or extend it, please feel free to contact Hans at hans.fangohr@xfel.eu.

If you find a URL that is not working (or pointing to the wrong material), please let Hans know as well. As the content of the Internet is changing rapidly, it is difficult to keep up with these changes without feedback.

A POWERFUL CALCULATOR

2.1 Python prompt and Read-Eval-Print Loop (REPL)

Python is an *interpreted* language. We can collect sequences of commands into text files and save this to file as a *Python program*. It is convention that these files have the file extension “.py”, for example `hello.py`.

We can also enter individual commands at the Python prompt which are immediately evaluated and carried out by the Python interpreter. This is very useful for the programmer/learner to understand how to use certain commands (often before one puts these commands together in a longer Python program). Python’s role can be described as Reading the command, Evaluating it, Printing the evaluated value and repeating (Loop) the cycle – this is the origin of the REPL abbreviation.

Python comes with a basic terminal prompt; you may see examples from this with `>>>` marking the input:

```
>>> 2 + 2
4
```

We are using a more powerful REPL interface, the Jupyter Notebook. Blocks of code appear with an `In` prompt next to them:

```
In [ ]: 4 + 5
```

```
9
```

To edit the code, click inside the code area. You should get a green border around it. To run it, press Shift-Enter.

2.2 Calculator

Basic operations such as addition (+), subtraction (-), multiplication (*), division (/) and exponentiation (**) work (mostly) as expected:

```
10 + 10000
```

```
10010
```

```
42 - 1.5
```

```
40.5
```

```
47 * 11
```

```
517
```

```
10 / 0.5
```

```
20.0
```

```
2**2 # Exponentiation ('to the power of') is **, NOT ^
```

```
4
```

```
2**3
```

```
8
```

```
2**4
```

```
16
```

```
2 + 2
```

```
4
```

```
# This is a comment  
2 + 2
```

```
4
```

```
2 + 2 # and a comment on the same line as code
```

```
4
```

and, using the fact that $\sqrt[n]{x} = x^{1/n}$, we can compute the $\sqrt{3} = 1.732050 \dots$ using `**`:

```
3**0.5
```

```
1.7320508075688772
```

Parenthesis can be used for grouping:

```
2 * 10 + 5
```

```
25
```

```
2 * (10 + 5)
```

```
30
```

2.3 Integer division

In Python 3, division works as you'd expect:

```
15/6
```

```
2.5
```

In Python 2, however, `15/6` will give you `2`.

This phenomenon is known (in many programming languages, including C) as *integer division*: because we provide two integer numbers (15 and 6) to the division operator (`/`), the assumption is that we seek a return value of type integer. The mathematically correct answer is (the floating point number) 2.5. (→ numerical data types in [Chapter 13](#).)

The convention for integer division is to truncate the fractional digits and to return the integer part only (i.e. `2` in this example). It is also called “floor division”.

2.3.1 How to avoid integer division

There are two ways to avoid the problem of integer division:

1. Use Python 3 style division: this is available even in Python 2 with a special import statement:

```
>>> from __future__ import division
>>> 15/6
2.5
```

If you want to use the `from __future__ import division` feature in a python program, it would normally be included at the beginning of the file.

2. Alternatively, if we ensure that at least one number (numerator or denominator) is of type float (or complex), the division operator will return a floating point number. This can be done by writing `15.` instead of `15`, or by forcing conversion of the number to a float, i.e. use `float(15)` instead of `15`:

```
>>> 15./6
2.5
>>> float(15)/6
2.5
>>> 15/6.
2.5
>>> 15/float(6)
```

(continues on next page)

(continued from previous page)

```
2.5
>>> 15./6.
2.5
```

If we really want integer division, we can use `//`: `1//2` returns 0, in both Python 2 and 3.

2.3.2 Why should I care about this division problem?

Integer division can result in surprising bugs: suppose you are writing code to compute the mean value $m = (x + y)/2$ of two numbers x and y . The first attempt of writing this may read:

```
m = (x + y) / 2
```

Suppose this is tested with $x = 0.5, y = 0.5$, then the line above computes the correct answers $m = 0.5$ (because $0.5 + 0.5 = 1.0$, i.e. a 1.0 is a floating point number, and thus $1.0/2$ evaluates to 0.5). Or we could use $x = 10, y = 30$, and because $10 + 30 = 40$ and $40/2$ evaluates to 20, we get the correct answer $m = 20$. However, if the integers $x = 0$ and $y = 1$ would come up, then the code returns $m = 0$ (because $0 + 1 = 1$ and $1/2$ evaluates to 0) whereas $m = 0.5$ would have been the right answer.

We have many possibilities to change the line of code above to work safely, including these three versions:

```
m = (x + y) / 2.0
m = float(x + y) / 2
m = (x + y) * 0.5
```

This integer division behaviour is common amongst most programming languages (including the important ones C, C++ and Fortran), and it is important to be aware of the issue.

2.4 Mathematical functions

Because Python is a general purpose programming language, commonly used mathematical functions such as `sin`, `cos`, `exp`, `log` and many others are located in the mathematics module with name `math`. We can make use of this as soon as we *import* the `math` module:

```
import math
math.exp(1.0)
```

```
2.718281828459045
```

Using the `dir` function, we can see the directory of objects available in the `math` module:

```
# NBVAL_IGNORE_OUTPUT
dir(math)
```

```
['__doc__',
 '__loader__',
 '__name__',
```

(continues on next page)

(continued from previous page)

```
'__package__',  
'__spec__',  
'acos',  
'acosh',  
'asin',  
'asinh',  
'atan',  
'atan2',  
'atanh',  
'cbrt',  
'ceil',  
'comb',  
'copysign',  
'cos',  
'cosh',  
'degrees',  
'dist',  
'e',  
'erf',  
'erfc',  
'exp',  
'exp2',  
'expm1',  
'fabs',  
'factorial',  
'floor',  
'fmod',  
'frexp',  
'fsum',  
'gamma',  
'gcd',  
'hypot',  
'inf',  
'isclose',  
'isfinite',  
'isinf',  
'isnan',  
'isqrt',  
'lcm',  
'ldexp',  
'lgamma',  
'log',  
'log10',  
'log1p',  
'log2',  
'modf',  
'nan',  
'nextafter',  
'perm',  
'pi',  
'pow',  
'prod',  
'radians',  
'remainder',  
'sin',  
'sinh',
```

(continues on next page)

(continued from previous page)

```
'sqrt',  
'tan',  
'tanh',  
'tau',  
'trunc',  
'ulp']
```

As usual, the `help` function can provide more information about the module (`help(math)`) on individual objects:

```
# NBVAL_IGNORE_OUTPUT  
help(math.exp)
```

```
Help on built-in function exp in module math:
```

```
exp(x, /)  
Return e raised to the power of x.
```

The mathematics module defines two constants π and e :

```
math.pi
```

```
3.141592653589793
```

```
math.e
```

```
2.718281828459045
```

```
math.cos(math.pi)
```

```
-1.0
```

```
math.log(math.e)
```

```
1.0
```

2.5 Variables

A *variable* can be used to store a certain value or object. In Python, all numbers (and everything else, including functions, modules and files) are objects. A variable is created through assignment:

```
x = 0.5
```

Once the variable `x` has been created through assignment of 0.5 in this example, we can make use of it:

```
x*3
```

1.5

```
x**2
```

0.25

```
y = 111  
y + 222
```

333

A variable is overridden if a new value is assigned:

```
y = 0.7  
math.sin(y) ** 2 + math.cos(y) ** 2
```

1.0

The equal sign ('=') is used to assign a value to a variable.

```
width = 20  
height = 5 * 9  
width * height
```

900

A value can be assigned to several variables simultaneously:

```
x = y = z = 0 # initialise x, y and z with 0  
x
```

0

y

0

z

0

Variables must be created (assigned a value) before they can be used, or an error will occur:

```
# NBVAL_RAISES_EXCEPTION  
# try to access an undefined variable:  
n
```

```
-----  
NameError                                Traceback (most recent call last)  
Cell In[32], line 3  
      1 # NBVAL_RAISES_EXCEPTION  
      2 # try to access an undefined variable:  
----> 3 n  
  
NameError: name 'n' is not defined
```

In interactive mode, the last printed expression is assigned to the variable `_`. This means that when you are using Python as a desk calculator, it is somewhat easier to continue calculations, for example:

```
tax = 12.5 / 100  
price = 100.50  
price * tax
```

```
12.5625
```

```
price + _
```

```
113.0625
```

This variable should be treated as read-only by the user. Don't explicitly assign a value to it — you would create an independent local variable with the same name masking the built-in variable with its magic behavior.

2.5.1 Terminology

Strictly speaking, the following happens when we write

```
x = 0.5
```

First, Python creates the object `0.5`. Everything in Python is an object, and so is the floating point number `0.5`. This object is stored somewhere in memory. Next, Python *binds a name to the object*. The name is `x`, and we often refer casually to `x` as a variable, an object, or even the value `0.5`. However, technically, `x` is a name that is bound to the object `0.5`. Another way to say this is that `x` is a reference to the object.

While it is often sufficient to think about assigning `0.5` to a variable `x`, there are situations where we need to remember what actually happens. In particular, when we pass references to objects to functions, we need to realise that the function may operate on the object (rather than a copy of the object). This is discussed in more detail in *the next chapter*.

2.6 Impossible equations

In computer programs we often find statements like

```
x = x + 1
```

If we read this as an equation as we are use to from mathematics, $x = x + 1$ we could subtract x on both sides, to find that $0 = 1$. We know this is not true, so something is wrong here.

The answer is that “equations“ in computer codes are not equations but *assignments*. They always have to be read in the following way two-step way:

1. Evaluate the value on the right hand side of the equal sign
2. Assign this value to the variable name shown on the left hand side. (In Python: bind the name on the left hand side to the object shown on the right hand side.)

Some computer science literature uses the following notation to express assignments and to avoid the confusion with mathematical equations:

$$x \leftarrow x + 1$$

Let's apply our two-step rule to the assignment $x = x + 1$ given above:

1. Evaluate the value on the right hand side of the equal sign: for this we need to know what the current value of x is. Let's assume x is currently 4. In that case, the right hand side $x+1$ evaluates to 5.
2. Assign this value (i.e. 5) to the variable name shown on the left hand side x .

Let's confirm with the Python prompt that this is the correct interpretation:

```
x = 4
x = x + 1
x
```

5

2.6.1 The += notation

Because it is a quite a common operation to increase a variable x by some fixed amount c , we can write

```
x += c
```

instead of

```
x = x + c
```

Our initial example above could thus have been written

```
x = 4
x += 1
x
```

5

The same operators are defined for multiplication with a constant ($*=$), subtraction of a constant ($-=$) and division by a constant ($/=$).

Note that the order of $+$ and $=$ matters:

```
x += 1
```

will increase the variable x by one where as

```
x =+ 1
```

will assign the value $+1$ to the variable x .

DATA TYPES AND DATA STRUCTURES

3.1 What type is it?

Python knows different data types. To find the type of a variable, use the `type()` function:

```
a = 45  
type(a)
```

```
int
```

```
b = 'This is a string'  
type(b)
```

```
str
```

```
c = 2 + 1j  
type(c)
```

```
complex
```

```
d = [1, 3, 56]  
type(d)
```

```
list
```

3.2 Numbers

Further information

- Informal introduction to numbers. [Python tutorial, section 3.1.1](#)
- Python Library Reference: formal overview of numeric types, <https://docs.python.org/3.8/library/stdtypes.html#numeric-types-int-float-complex>
- Think Python, Sec 2.1

The in-built numerical types are integers and floating point numbers (see *floating point numbers*) and complex floating point numbers (*complex numbers*).

3.2.1 Integers

We have seen the use of integer numbers already in *Chapter 2*. Be aware of integer division problems (02 A Powerful Calculator, Integer Division).

If we need to convert string containing an integer number to an integer we can use `int()` function:

```
a = '34'      # a is a string containing the characters 3 and 4
x = int(a)    # x is in integer number
```

The function `int()` will also convert floating point numbers to integers:

```
int(7.0)
```

```
7
```

```
int(7.9)
```

```
7
```

Note that `int` will truncate any non-integer part of a floating point number. To round an floating point number to an integer, use the `round()` command:

```
round(7.9)
```

```
8
```

3.2.2 Integer limits

Integers in Python 3 are unlimited; Python will automatically assign more memory as needed as the numbers get bigger. This means we can calculate very large numbers with no special steps.

```
35**42
```

```
70934557307860443711736098025989133248003781773149967193603515625
```

In many other programming languages, such as C and FORTRAN, integers are a fixed size—most frequently 4 bytes, which allows 2^{32} different values—but different types are available with different sizes. For numbers that fit into these limits, calculations can be faster, but you may need to check that the numbers don't go beyond the limits. Calculating a number beyond the limits is called *integer overflow*, and may produce bizarre results.

Even in Python, we need to be aware of this when we use `numpy` (see *Chapter 14*). `numpy` uses integers with a fixed size, because it stores many of them together and needs to calculate with them efficiently. `numpy data types` include a range of integer types named for their size, so e.g. `int16` is a 16-bit integer, with 2^{16} possible values.

Integer types can also be *signed* or *unsigned*. Signed integers allow positive or negative values, unsigned integers only allow positive ones. For instance:

- `uint16` (unsigned) ranges from 0 to $2^{16} - 1$
- `int16` (signed) ranges from -2^{15} to $2^{15} - 1$

3.2.3 Floating Point numbers

A string containing a floating point number can be converted into a floating point number using the `float()` command:

```
a = '35.342'  
b = float(a)  
b
```

```
35.342
```

```
type(b)
```

```
float
```

3.2.4 Complex numbers

Python (as Fortran and Matlab) has built-in complex numbers. Here are some examples how to use these:

```
x = 1 + 3j  
x
```

```
(1+3j)
```

```
abs(x) # computes the absolute value
```

```
3.1622776601683795
```

```
x.imag
```

```
3.0
```

```
x.real
```

```
1.0
```

```
x * x
```

```
(-8+6j)
```

```
x * x.conjugate()
```

```
(10+0j)
```

```
3 * x
```

```
(3+9j)
```

Note that if you want to perform more complicated operations (such as taking the square root, etc) you have to use the `cmath` module (Complex MATHeMATics):

```
import cmath  
cmath.sqrt(x)
```

```
(1.442615274452683+1.0397782600555705j)
```

3.2.5 Functions applicable to all types of numbers

The `abs()` function returns the absolute value of a number (also called modulus):

```
a = -45.463  
abs(a)
```

```
45.463
```

Note that `abs()` also works for complex numbers (see above).

3.3 Sequences

Strings, lists and tuples are *sequences*. They can be *indexed* and *sliced* in the same way.

Tuples and strings are “immutable” (which basically means we can’t change individual elements within the tuple, and we cannot change individual characters within a string) whereas lists are “mutable” (*i.e* we can change elements in a list.)

Sequences share the following operations

- `a[i]` returns *i*-th element of `a`
- `a[i:j]` returns elements *i* up to *j*-1
- `len(a)` returns number of elements in sequence
- `min(a)` returns smallest value in sequence
- `max(a)` returns largest value in sequence
- `x in a` returns `True` if `x` is element in `a`
- `a + b` concatenates `a` and `b`
- `n * a` creates `n` copies of sequence `a`

3.3.1 Sequence type 1: String

Further information

- Introduction to strings, [Python tutorial 3.1.2](#)

A string is a (immutable) sequence of characters. A string can be defined using single quotes:

```
a = 'Hello World'
```

double quotes:

```
a = "Hello World"
```

or triple quotes of either kind

```
a = """Hello World"""  
a = '''Hello World'''
```

The type of a string is `str` and the empty string is given by `""`:

```
a = "Hello World"  
type(a)
```

```
str
```

```
b = ""  
type(b)
```

```
str
```

```
type("Hello World")
```

```
str
```

```
type("")
```

```
str
```

The number of characters in a string (that is its *length*) can be obtained using the `len()`-function:

```
a = "Hello Moon"  
len(a)
```

```
10
```

```
a = 'test'  
len(a)
```

```
4
```

```
len('another test')
```

```
12
```

You can combine (“concatenate”) two strings using the + operator:

```
'Hello ' + 'World'
```

```
'Hello World'
```

Strings have a number of useful methods, including for example `upper()` which returns the string in upper case:

```
a = "This is a test sentence."  
a.upper()
```

```
'THIS IS A TEST SENTENCE.'
```

A list of available string methods can be found in the Python reference documentation. If a Python prompt is available, one should use the `dir` and `help` function to retrieve this information, *i.e.* `dir()` provides the list of methods, `help` can be used to learn about each method.

A particularly useful method is `split()` which converts a string into a list of strings:

```
a = "This is a test sentence."  
a.split()
```

```
['This', 'is', 'a', 'test', 'sentence.']
```

The `split()` method will separate the string where it finds *white space*. White space means any character that is printed as white space, such as one space or several spaces or a tab.

By passing a separator character to the `split()` method, a string can split into different parts. Suppose, for example, we would like to obtain a list of complete sentences:

```
a = "The dog is hungry. The cat is bored. The snake is awake."  
a.split(". ")
```

```
['The dog is hungry', ' The cat is bored', ' The snake is awake', '']
```

The opposite string method to `split` is `join` which can be used as follows:

```
a = "The dog is hungry. The cat is bored. The snake is awake."  
s = a.split('.')  
s
```

```
['The dog is hungry', ' The cat is bored', ' The snake is awake', '']
```



```
".".join(s)
```

```
'The dog is hungry. The cat is bored. The snake is awake.'
```

```
" STOP".join(s)
```

```
'The dog is hungry STOP The cat is bored STOP The snake is awake STOP'
```

3.3.2 Sequence type 2: List

Further information

- Introduction to Lists, [Python tutorial, section 3.1.4](#)

A list is a sequence of objects. The objects can be of any type, for example integers:

```
a = [34, 12, 54]
```

or strings:

```
a = ['dog', 'cat', 'mouse']
```

An empty list is presented by []:

```
a = []
```

The type is list:

```
type(a)
```

```
list
```

```
type([])
```

```
list
```

As with strings, the number of elements in a list can be obtained using the `len()` function:

```
a = ['dog', 'cat', 'mouse']  
len(a)
```

```
3
```

It is also possible to *mix* different types in the same list:

```
a = [123, 'duck', -42, 17, 0, 'elephant']
```

In Python a list is an object. It is therefore possible for a list to contain other lists (because a list keeps a sequence of objects):

```
a = [1, 4, 56, [5, 3, 1], 300, 400]
```

You can combine (“concatenate”) two lists using the + operator:

```
[3, 4, 5] + [34, 35, 100]
```

```
[3, 4, 5, 34, 35, 100]
```

Or you can add one object to the end of a list using the `append()` method:

```
a = [34, 56, 23]
a.append(42)
a
```

```
[34, 56, 23, 42]
```

You can delete an object from a list by calling the `remove()` method and passing the object to delete. For example:

```
a = [34, 56, 23, 42]
a.remove(56)
a
```

```
[34, 23, 42]
```

The `range()` command

A special type of list is frequently required (often together with `for`-loops) and therefore a command exists to generate that list: the `range(n)` command generates integers starting from 0 and going up to *but not including* n. Here are a few examples:

```
list(range(3))
```

```
[0, 1, 2]
```

```
list(range(10))
```

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

This command is often used with `for` loops. For example, to print the numbers 0, 1, 2, 3, ..., 10, the following program can be used:

```
for i in range(11):
    print(i ** 2)
```

```
0
1
4
9
```

(continues on next page)

(continued from previous page)

```
16
25
36
49
64
81
100
```

The range command takes an optional parameter for the beginning of the integer sequence (start) and another optional parameter for the step size. This is often written as `range([start], stop, [step])` where the arguments in square brackets (*i.e.* start and step) are optional. Here are some examples:

```
list(range(3, 10))          # start=3
```

```
[3, 4, 5, 6, 7, 8, 9]
```

```
list(range(3, 10, 2))      # start=3, step=2
```

```
[3, 5, 7, 9]
```

```
list(range(10, 0, -1))     # start=10, step=-1
```

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

Why are we calling `list(range())`?

In Python 3, `range()` generates the numbers on demand. When you use `range()` in a for loop, this is more efficient, because it doesn't take up memory with a list of numbers. Passing it to `list()` forces it to generate all of its numbers, so we can see what it does.

To get the same efficient behaviour in Python 2, use `xrange()` instead of `range()`.

3.3.3 Sequence type 3: Tuples

A *tuple* is a (immutable) sequence of objects. Tuples are very similar in behaviour to lists with the exception that they cannot be modified (*i.e.* are immutable).

For example, the objects in a sequence can be of any type:

```
a = (12, 13, 'dog')
a
```

```
(12, 13, 'dog')
```

```
a[0]
```

```
12
```

The parentheses are not necessary to define a tuple: just a sequence of objects separated by commas is sufficient to define a tuple:

```
a = 100, 200, 'duck'  
a
```

```
(100, 200, 'duck')
```

although it is good practice to include the paranthesis where it helps to show that tuple is defined.

Tuples can also be used to make two assignments at the same time:

```
x, y = 10, 20  
x
```

```
10
```

```
y
```

```
20
```

This can be used to *swap* to objects within one line. For example

```
x = 1  
y = 2  
x, y = y, x  
x
```

```
2
```

```
y
```

```
1
```

The empty tuple is given by ()

```
t = ()  
len(t)
```

```
0
```

```
type(t)
```

```
tuple
```

The notation for a tuple containing one value may seem a bit odd at first:

```
t = (42,)  
type(t)
```

```
tuple
```

```
len(t)
```

```
1
```

The extra comma is required to distinguish `(42,)` from `(42)` where in the latter case the parenthesis would be read as defining operator precedence: `(42)` simplifies to `42` which is just a number:

```
t = (42)
type(t)
```

```
int
```

This example shows the immutability of a tuple:

```
a = (12, 13, 'dog')
a[0]
```

```
12
```

```
# NBVAL_RAISES_EXCEPTION
a[0] = 1
```

```
-----
TypeError                                Traceback (most recent call last)
Cell In[68], line 2
      1 # NBVAL_RAISES_EXCEPTION
----> 2 a[0] = 1

TypeError: 'tuple' object does not support item assignment
```

The immutability is the main difference between a tuple and a list (the latter being mutable). We should use tuples when we don't want the content to change.

Note that Python functions that return more than one value, return these in tuples (which makes sense because you don't want these values be changed).

3.3.4 Indexing sequences

Further information

- Introduction to strings and indexing in [Python tutorial, section 3.1.2](#), the relevant section is starting after strings have been introduced.

Individual objects in lists can be accessed by using the index of the object and square brackets (`[` and `]`):

```
a = ['dog', 'cat', 'mouse']
a[0]
```

```
'dog'
```

```
a[1]
```

```
'cat'
```

```
a[2]
```

```
'mouse'
```

Note that Python (like C but unlike Fortran and unlike Matlab) starts counting indices from zero!

Python provides a handy shortcut to retrieve the last element in a list: for this one uses the index “-1” where the minus indicates that it is one element *from the back* of the list. Similarly, the index “-2” will return the 2nd last element:

```
a = ['dog', 'cat', 'mouse']  
a[-1]
```

```
'mouse'
```

```
a[-2]
```

```
'cat'
```

If you prefer, you can think of the index `a[-1]` to be a shorthand notation for `a[len(a) - 1]`.

Remember that strings (like lists) are also a sequence type and can be indexed in the same way:

```
a = "Hello World!"  
a[0]
```

```
'H'
```

```
a[1]
```

```
'e'
```

```
a[10]
```

```
'd'
```

```
a[-1]
```

```
'!'
```

```
a[-2]
```

```
'd'
```

3.3.5 Slicing sequences

Further information

- Introduction to strings, indexing and slicing in [Python tutorial, section 3.1.2](#)

Slicing of sequences can be used to retrieve more than one element. For example:

```
a = "Hello World!"
a[0:3]
```

```
'Hel'
```

By writing `a[0:3]` we request the first 3 elements starting from element 0. Similarly:

```
a[1:4]
```

```
'ell'
```

```
a[0:2]
```

```
'He'
```

```
a[0:6]
```

```
'Hello '
```

We can use negative indices to refer to the end of the sequence:

```
a[0:-1]
```

```
'Hello World'
```

It is also possible to leave out the start or the end index and this will return all elements up to the beginning or the end of the sequence. Here are some examples to make this clearer:

```
a = "Hello World!"
a[:5]
```

```
'Hello'
```

```
a[5:]
```

```
' World!'
```

```
a[-2:]
```

```
'd!'
```

```
a[:]
```

```
'Hello World!'
```

Note that `a[:]` will generate a *copy* of `a`. The use of indices in slicing is by some people experienced as counter intuitive. If you feel uncomfortable with slicing, have a look at this quotation from the [Python tutorial \(section 3.1.2\)](#):

The best way to remember how slices work is to think of the indices as pointing between characters, with the left edge of the first character numbered 0. Then the right edge of the last character of a string of 5 characters has index 5, for example:

```
+---+---+---+---+---+
| H | e | l | l | o |
+---+---+---+---+---+
0  1  2  3  4  5  <-- use for SLICING
-5 -4 -3 -2 -1   <-- use for SLICING
                    from the end
```

The first row of numbers gives the position of the slicing indices 0..5 in the string; the second row gives the corresponding negative indices. The slice from `i` to `j` consists of all characters between the edges labelled `i` and `j`, respectively.

So the important statement is that for *slicing* we should think of indices pointing between characters.

For *indexing* it is better to think of the indices referring to characters. Here is a little graph summarising these rules:

```
0  1  2  3  4  <-- use for INDEXING
-5 -4 -3 -2 -1  <-- use for INDEXING
                    from the end
+---+---+---+---+---+
| H | e | l | l | o |
+---+---+---+---+---+
0  1  2  3  4  5  <-- use for SLICING
-5 -4 -3 -2 -1   <-- use for SLICING
                    from the end
```

If you are not sure what the right index is, it is always a good technique to play around with a small example at the Python prompt to test things before or while you write your program.

3.3.6 Dictionaries

Dictionaries are also called “associative arrays” and “hash tables”. Dictionaries are *unordered* sets of *key-value pairs*.

An empty dictionary can be created using curly braces:

```
d = {}
```

Keyword-value pairs can be added like this:

```
d['today'] = '22 deg C'    # 'today' is the keyword
```

```
d['yesterday'] = '19 deg C'
```

`d.keys()` returns a list of all keys:

```
d.keys()
```

```
dict_keys(['today', 'yesterday'])
```

We can retrieve values by using the keyword as the index:

```
d['today']
```

```
'22 deg C'
```

Other ways of populating a dictionary if the data is known at creation time are:

```
d2 = {2:4, 3:9, 4:16, 5:25}
d2
```

```
{2: 4, 3: 9, 4: 16, 5: 25}
```

```
d3 = dict(a=1, b=2, c=3)
d3
```

```
{'a': 1, 'b': 2, 'c': 3}
```

The function `dict()` creates an empty dictionary.

Other useful dictionary methods include `values()`, `items()` and `get()`. You can use `in` to check for the presence of values.

```
d.values()
```

```
dict_values(['22 deg C', '19 deg C'])
```

```
d.items()
```

```
dict_items([('today', '22 deg C'), ('yesterday', '19 deg C')])
```

```
d.get('today', 'unknown')
```

```
'22 deg C'
```

```
d.get('tomorrow', 'unknown')
```

```
'unknown'
```

```
'today' in d
```

```
True
```

```
'tomorrow' in d
```

```
False
```

The method `get(key, default)` will provide the value for a given key if that key exists, otherwise it will return the default object.

Here is a more complex example:

```
# NBVAL_IGNORE_OUTPUT
order = {}          # create empty dictionary

#add orders as they come in
order['Peter'] = 'Pint of bitter'
order['Paul'] = 'Half pint of Hoegarden'
order['Mary'] = 'Gin Tonic'

#deliver order at bar
for person in order.keys():
    print(person, "requests", order[person])
```

```
Peter requests Pint of bitter
Paul requests Half pint of Hoegarden
Mary requests Gin Tonic
```

Some more technicalities:

- The keyword can be any (immutable) Python object. This includes:
 - numbers
 - strings
 - tuples.
- dictionaries are very fast in retrieving values (when given the key)

An other example to demonstrate an advantage of using dictionaries over pairs of lists:

```
# NBVAL_IGNORE_OUTPUT
dic = {}                                #create empty dictionary

dic["Hans"] = "room 1033"               #fill dictionary
dic["Andy C"] = "room 1031"            #"Andy C" is key
dic["Ken"] = "room 1027"               #"room 1027" is value

for key in dic.keys():
    print(key, "works in", dic[key])
```

```
Hans works in room 1033
Andy C works in room 1031
Ken works in room 1027
```

Without dictionary:

```
people = ["Hans", "Andy C", "Ken"]
rooms = ["room 1033", "room 1031", "room 1027"]

#possible inconsistency here since we have two lists
if not len( people ) == len( rooms ):
    raise RuntimeError("people and rooms differ in length")

for i in range( len( rooms ) ):
    print(people[i], "works in", rooms[i])
```

```
Hans works in room 1033
Andy C works in room 1031
Ken works in room 1027
```

3.4 Passing arguments to functions

This section contains some more advanced ideas and makes use of concepts that are only later introduced in this text. The section may be more easily accessible at a later stage.

When objects are passed to a function, Python always passes (the value of) the reference to the object to the function. Effectively this is calling a function by reference, although one could refer to it as calling by value (of the reference).

We review argument passing by value and reference before discussing the situation in Python in more detail.

3.4.1 Call by value

One might expect that if we pass an object by value to a function, that modifications of that value inside the function will not affect the object (because we don't pass the object itself, but only its value, which is a copy). Here is an example of this behaviour (in C):

```
#include <stdio.h>

void pass_by_value(int m) {
    printf("in pass_by_value: received m=%d\n", m);
    m=42;
```

(continues on next page)

(continued from previous page)

```
    printf("in pass_by_value: changed to m=%d\n",m);
}

int main(void) {
    int global_m = 1;
    printf("global_m=%d\n",global_m);
    pass_by_value(global_m);
    printf("global_m=%d\n",global_m);
    return 0;
}
```

together with the corresponding output:

```
global_m=1
in pass_by_value: received m=1
in pass_by_value: changed to m=42
global_m=1
```

The value 1 of the global variable `global_m` is not modified when the function `pass_by_value` changes its input argument to 42.

3.4.2 Call by reference

Calling a function by reference, on the other hand, means that the object given to a function is a reference to the object. This means that the function will see the same object as in the calling code (because they are referencing the same object: we can think of the reference as a pointer to the place in memory where the object is located). Any changes acting on the object inside the function, will then be visible in the object at the calling level (because the function does actually operate on the same object, not a copy of it).

Here is one example showing this using pointers in C:

```
#include <stdio.h>

void pass_by_reference(int *m) {
    printf("in pass_by_reference: received m=%d\n", *m);
    *m=42;
    printf("in pass_by_reference: changed to m=%d\n", *m);
}

int main(void) {
    int global_m = 1;
    printf("global_m=%d\n",global_m);
    pass_by_reference(&global_m);
    printf("global_m=%d\n",global_m);
    return 0;
}
```

together with the corresponding output:

```
global_m=1
in pass_by_reference: received m=1
in pass_by_reference: changed to m=42
global_m=42
```

C++ provides the ability to pass arguments as references by adding an ampersand in front of the argument name in the function definition:

```
#include <stdio.h>

void pass_by_reference(int &m) {
    printf("in pass_by_reference: received m=%d\n",m);
    m=42;
    printf("in pass_by_reference: changed to m=%d\n",m);
}

int main(void) {
    int global_m = 1;
    printf("global_m=%d\n",global_m);
    pass_by_reference(global_m);
    printf("global_m=%d\n",global_m);
    return 0;
}
```

together with the corresponding output:

```
global_m=1
in pass_by_reference: received m=1
in pass_by_reference: changed to m=42
global_m=42
```

3.4.3 Argument passing in Python

In Python, objects are passed as the value of a reference (think pointer) to the object. Depending on the way the reference is used in the function and depending on the type of object it references, this can result in pass-by-reference behaviour (where any changes to the object received as a function argument, are immediately reflected in the calling level).

Here are three examples to discuss this. We start by passing a list to a function which iterates through all elements in the sequence and doubles the value of each element:

```
def double_the_values(l):
    print("in double_the_values: l = %s" % l)
    for i in range(len(l)):
        l[i] = l[i] * 2
    print("in double_the_values: changed l to l = %s" % l)

l_global = [0, 1, 2, 3, 10]
print("In main: s=%s" % l_global)
double_the_values(l_global)
print("In main: s=%s" % l_global)
```

```
In main: s=[0, 1, 2, 3, 10]
in double_the_values: l = [0, 1, 2, 3, 10]
in double_the_values: changed l to l = [0, 2, 4, 6, 20]
In main: s=[0, 2, 4, 6, 20]
```

The variable `l` is a reference to the list object. The line `l[i] = l[i] * 2` first evaluates the right-hand side and reads the element with index `i`, then multiplies this by two. A reference to this new object is then stored in the list object `l` at position with index `i`. We have thus modified the list object, that is referenced through `l`.

The reference to the list object does never change: the line `l[i] = l[i] * 2` changes the elements `l[i]` of the list `l` but never changes the reference `l` for the list. Thus both the function and calling level are operating on the same object through the references `l` and `global_l`, respectively.

In contrast, here is an example where do not modify the elements of the list within the function: which produces this output:

```
def double_the_list(l):
    print("in double_the_list: l = %s" % l)
    l = l + l
    print("in double_the_list: changed l to l = %s" % l)

l_global = "Hello"
print("In main: l=%s" % l_global)
double_the_list(l_global)
print("In main: l=%s" % l_global)
```

```
In main: l=Hello
in double_the_list: l = Hello
in double_the_list: changed l to l = HelloHello
In main: l=Hello
```

What happens here is that during the evaluation of `l = l + l` a new object is created that holds `l + l`, and that we then bind the name `l` to it. In the process, we lose the references to the list object `l` that was given to the function (and thus we do not change the list object given to the function).

Finally, let's look at which produces this output:

```
def double_the_value(l):
    print("in double_the_value: l = %s" % l)
    l = 2 * l
    print("in double_the_values: changed l to l = %s" % l)

l_global = 42
print("In main: s=%s" % l_global)
double_the_value(l_global)
print("In main: s=%s" % l_global)
```

```
In main: s=42
in double_the_value: l = 42
in double_the_values: changed l to l = 84
In main: s=42
```

In this example, we also double the value (from 42 to 84) within the function. However, when we bind the object 84 to the python name `l` (that is the line `l = l * 2`) we have created a new object (84), and we bind the new object to `l`. In the process, we lose the reference to the object 42 within the function. This does not affect the object 42 itself, nor the reference `l_global` to it.

In summary, Python's behaviour of passing arguments to a function may appear to vary (if we view it from the pass by value versus pass by reference point of view). However, it is always call by value, where the value is a reference to the object in question, and the behaviour can be explained through the same reasoning in every case.

3.4.4 Performance considerations

Call by value function calls require copying of the value before it is passed to the function. From a performance point of view (both execution time and memory requirements), this can be an expensive process if the value is large. (Imagine the value is a `numpy.array` object which could be several Megabytes or Gigabytes in size.)

One generally prefers call by reference for large data objects as in this case only a pointer to the data objects is passed, independent of the actual size of the object, and thus this is generally faster than call-by-value.

Python's approach of (effectively) calling by reference is thus efficient. However, we need to be careful that our function do not modify the data they have been given where this is undesired.

3.4.5 Inadvertent modification of data

Generally, a function should not modify the data given as input to it.

For example, the following code demonstrates the attempt to determine the maximum value of a list, and – inadvertently – modifies the list in the process:

```
def mymax(s): # demonstrating side effect
    if len(s) == 0:
        raise ValueError('mymax() arg is an empty sequence')
    elif len(s) == 1:
        return s[0]
    else:
        for i in range(1, len(s)):
            if s[i] < s[i - 1]:
                s[i] = s[i - 1]
        return s[len(s) - 1]

s = [-45, 3, 6, 2, -1]
print("in main before calling mymax(s): s=%s" % s)
print("mymax(s)=%s" % mymax(s))
print("in main after calling mymax(s): s=%s" % s)
```

```
in main before calling mymax(s): s=[-45, 3, 6, 2, -1]
mymax(s)=6
in main after calling mymax(s): s=[-45, 3, 6, 6, 6]
```

The user of the `mymax()` function would not expect that the input argument is modified when the function executes. We should generally avoid this. There are several ways to find better solutions to the given problem:

- In this particular case, we could use the Python in-built function `max()` to obtain the maximum value of a sequence.
- If we felt we need to stick to storing temporary values inside the list [this is actually not necessary], we could create a copy of the incoming list `s` first, and then proceed with the algorithm (see *below* on Copying objects).
- Use another algorithm which uses an extra temporary variable rather than abusing the list for this. For example:
- We could pass a tuple (instead of a list) to the function: a tuple is *immutable* and can thus never be modified (this would result in an exception being raised when the function tries to write to elements in the tuple).

3.4.6 Copying objects

Python provides the `id()` function which returns an integer number that is unique for each object. (In the current CPython implementation, this is the memory address.) We can use this to identify whether two objects are the same.

To copy a sequence object (including lists), we can slice it, *i.e.* if `a` is a list, then `a[:]` will return a copy of `a`. Here is a demonstration:

```
a = list(range(10))
a
```

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

```
b = a
b[0] = 42
a           # changing b changes a
```

```
[42, 1, 2, 3, 4, 5, 6, 7, 8, 9]
```

```
# NBVAL_IGNORE_OUTPUT
id(a)
```

```
140195263471296
```

```
# NBVAL_IGNORE_OUTPUT
id(b)
```

```
140195263471296
```

```
# NBVAL_IGNORE_OUTPUT
c = a[:]
id(c)           # c is a different object
```

```
140195523088768
```

```
c[0] = 100
a           # changing c does not affect a
```

```
[42, 1, 2, 3, 4, 5, 6, 7, 8, 9]
```

Python's standard library provides the `copy` module, which provides copy functions that can be used to create copies of objects. We could have used `import copy; c = copy.deepcopy(a)` instead of `c = a[:]`.

3.5 Equality and Identity/Sameness

A related question concerns the equality of objects.

3.5.1 Equality

The operators `<`, `>`, `==`, `>=`, `<=`, and `!=` compare the *values* of two objects. The objects need not have the same type. For example:

```
a = 1.0; b = 1
type(a)
```

```
float
```

```
type(b)
```

```
int
```

```
a == b
```

```
True
```

So the `==` operator checks whether the values of two objects are equal.

3.5.2 Identity / Sameness

To see check whether two objects `a` and `b` are the same (i.e. `a` and `b` are references to the same place in memory), we can use the `is` operator (continued from example above):

```
a is b
```

```
False
```

Of course they are different here, as they are not of the same type.

We can also ask the `id` function which, according to the documentation string in Python 2.7 “Returns the identity of an object. This is guaranteed to be unique among simultaneously existing objects. (Hint: it’s the object’s memory address.)”

```
# NBVAL_IGNORE_OUTPUT
id(a)
```

```
140195523039280
```

```
# NBVAL_IGNORE_OUTPUT
id(b)
```

```
10861192
```

which shows that `a` and `b` are stored in different places in memory.

3.5.3 Example: Equality and identity

We close with an example involving lists:

```
x = [0, 1, 2]
y = x
x == y
```

```
True
```

```
x is y
```

```
True
```

```
# NBVAL_IGNORE_OUTPUT
id(x)
```

```
140195265280064
```

```
# NBVAL_IGNORE_OUTPUT
id(y)
```

```
140195265280064
```

Here, `x` and `y` are references to the same piece of memory, they are thus identical and the `is` operator confirms this. The important point to remember is that line 2 (`y=x`) creates a new reference `y` to the same list object that `x` is a reference for.

Accordingly, we can change elements of `x`, and `y` will change simultaneously as both `x` and `y` refer to the same object:

```
x
```

```
[0, 1, 2]
```

```
y
```

```
[0, 1, 2]
```

```
x is y
```

```
True
```

```
x[0] = 100
y
```

```
[100, 1, 2]
```

```
x
```

```
[100, 1, 2]
```

In contrast, if we use `z=x[:]` (instead of `z=x`) to create a new name `z`, then the slicing operation `x[:]` will actually create a copy of the list `x`, and the new reference `z` will point to the copy. The *value* of `x` and `z` is equal, but `x` and `z` are not the same object (they are not identical):

```
x
```

```
[100, 1, 2]
```

```
z = x[:]           # create copy of x before assigning to z
z == x            # same value
```

```
True
```

```
z is x           # are not the same object
```

```
False
```

```
# NBVAL_IGNORE_OUTPUT
id(z)           # confirm by looking at ids
```

```
140195263608512
```

```
# NBVAL_IGNORE_OUTPUT
id(x)
```

```
140195265280064
```

```
x
```

```
[100, 1, 2]
```

```
z
```

```
[100, 1, 2]
```

Consequently, we can change `x` without changing `z`, for example (continued)

```
x[0] = 42  
x
```

```
[42, 1, 2]
```

```
z
```

```
[100, 1, 2]
```

INTROSPECTION

A Python code can ask and answer questions about itself and the objects it is manipulating.

4.1 dir

`dir()` is a built-in function which returns a list of all the names belonging to some namespace.

- If no arguments are passed to `dir` (i.e. `dir()`), it inspects the namespace in which it was called.
- If `dir` is given an argument (i.e. `dir(<object>)`), then it inspects the namespace of the object which it was passed.

For example:

```
# NBVAL_IGNORE_OUTPUT
apples = ['Cox', 'Braeburn', 'Jazz']
dir(apples)
```

```
['__add__',
 '__class__',
 '__contains__',
 '__delattr__',
 '__delitem__',
 '__dir__',
 '__doc__',
 '__eq__',
 '__format__',
 '__ge__',
 '__getattr__',
 '__getitem__',
 '__gt__',
 '__hash__',
 '__iadd__',
 '__imul__',
 '__init__',
 '__init_subclass__',
 '__iter__',
 '__le__',
 '__len__',
 '__lt__',
 '__mul__',
 '__ne__',
```

(continues on next page)

(continued from previous page)

```
'__new__',
'__reduce__',
'__reduce_ex__',
'__repr__',
'__reversed__',
'__rmul__',
'__setattr__',
'__setitem__',
'__sizeof__',
'__str__',
'__subclasshook__',
'append',
'clear',
'copy',
'count',
'extend',
'index',
'insert',
'pop',
'remove',
'reverse',
'sort']
```

```
# NBVAL_IGNORE_OUTPUT
dir()
```

```
['In',
'Out',
'_',
'_1',
'_2',
'_3',
'__builtin__',
'__builtins__',
'__doc__',
'__loader__',
'__name__',
'__package__',
'__spec__',
'_dh',
'_i',
'_i1',
'_i2',
'_ih',
'_ii',
'_iii',
'_oh',
'apples',
'exit',
'get_ipython',
'quit']
```

```
# NBVAL_IGNORE_OUTPUT
name = "Peter"
```

(continues on next page)

(continued from previous page)

`dir(name)`

```
['__add__',
 '__class__',
 '__contains__',
 '__delattr__',
 '__dir__',
 '__doc__',
 '__eq__',
 '__format__',
 '__ge__',
 '__getattr__',
 '__getitem__',
 '__getnewargs__',
 '__gt__',
 '__hash__',
 '__init__',
 '__init_subclass__',
 '__iter__',
 '__le__',
 '__len__',
 '__lt__',
 '__mod__',
 '__mul__',
 '__ne__',
 '__new__',
 '__reduce__',
 '__reduce_ex__',
 '__repr__',
 '__rmod__',
 '__rmul__',
 '__setattr__',
 '__sizeof__',
 '__str__',
 '__subclasshook__',
 'capitalize',
 'casefold',
 'center',
 'count',
 'encode',
 'endswith',
 'expandtabs',
 'find',
 'format',
 'format_map',
 'index',
 'isalnum',
 'isalpha',
 'isdecimal',
 'isdigit',
 'isidentifier',
 'islower',
 'isnumeric',
 'isprintable',
 'isspace',
 'istitle',
```

(continues on next page)

(continued from previous page)

```
'isupper',
'join',
'ljust',
'lower',
'lstrip',
'maketrans',
'partition',
'replace',
'rfind',
'rindex',
'rjust',
'partition',
'rsplit',
'rstrip',
'split',
'splitlines',
'startswith',
'strip',
'swapcase',
'title',
'translate',
'upper',
'zfill']
```

4.1.1 Magic names

You will find many names which both start and end with a double underscore (e.g. `__name__`). These are called magic names. Functions with magic names provide the implementation of particular python functionality.

For example, the application of the `str` to an object `a`, i.e. `str(a)`, will – internally – result in the method `a.__str__()` being called. This method `__str__` generally needs to return a string. The idea is that the `__str__()` method should be defined for all objects (including those that derive from new classes that a programmer may create) so that all objects (independent of their type or class) can be printed using the `str()` function. The actual conversion of some object `x` to the string is then done via the object specific method `x.__str__()`.

We can demonstrate this by creating a class `my_int` which inherits from the Python's integer base class, and overrides the `__str__` method. (It requires more Python knowledge than provided up to this point in the text to be able to understand this example.)

```
class my_int(int):
    """Inherited from int"""
    def __str__(self):
        """Tailored str representation of my int"""
        return "my_int: %s" % (int.__str__(self))

a = my_int(3)
b = int(4)          # equivalent to b = 4
print("a * b = ", a * b)
print("Type a = ", type(a), "str(a) = ", str(a))
print("Type b = ", type(b), "str(b) = ", str(b))
```

```
a * b = 12
Type a = <class '__main__.my_int'> str(a) = my_int: 3
Type b = <class 'int'> str(b) = 4
```


Further Reading

See → [Python documentation](#), [Data Model](#)

4.2 type

The `type(<object>)` command returns the type of an object:

```
type(1)
```

```
int
```

```
type(1.0)
```

```
float
```

```
type("Python")
```

```
str
```

```
import math  
type(math)
```

```
module
```

```
type(math.sin)
```

```
builtin_function_or_method
```

4.3 isinstance

`isinstance(<object>, <typespec>)` returns True if the given object is an instance of the given type, or any of its superclasses. Use `help(isinstance)` for the full syntax.

```
isinstance(2, int)
```

```
True
```

```
isinstance(2., int)
```

```
False
```

```
isinstance(a,int)    # a is an instance of my_int
```

```
True
```

```
type(a)
```

```
__main__.my_int
```

4.4 help

- The `help(<object>)` function will report the docstring (magic attribute with name `__doc__`) of the object that it is given, sometimes complemented with additional information. In the case of functions, `help` will also show the list of arguments that the function accepts (but it cannot provide the return value).
- `help()` starts an interactive help environment.
- It is common to use the `help` command a lot to remind oneself of the syntax and semantic of commands.

```
help(isinstance)
```

```
Help on built-in function isinstance in module builtins:
```

```
isinstance(obj, class_or_tuple, /)
    Return whether an object is an instance of a class or of a subclass thereof.

    A tuple, as in ``isinstance(x, (A, B, ...))``, may be given as the target to
    check against. This is equivalent to ``isinstance(x, A) or isinstance(x, B)
    or ...`` etc.
```

```
# NBVAL_IGNORE_OUTPUT
import math
help(math.sin)
```

```
Help on built-in function sin in module math:
```

```
sin(...)
    sin(x)

    Return the sine of x (measured in radians).
```

```
# NBVAL_IGNORE_OUTPUT
help(math)
```

```
Help on module math:
```

```
NAME
    math
```

(continues on next page)

(continued from previous page)

MODULE REFERENCE

<https://docs.python.org/3.6/library/math>

The following documentation is automatically generated from the Python source files. It may be incomplete, incorrect or include features that are considered implementation detail and may vary between Python implementations. When in doubt, consult the module reference at the location listed above.

DESCRIPTION

This module is always available. It provides access to the mathematical functions defined by the C standard.

FUNCTIONS

`acos(...)`

`acos(x)`

Return the arc cosine (measured in radians) of x .

`acosh(...)`

`acosh(x)`

Return the inverse hyperbolic cosine of x .

`asin(...)`

`asin(x)`

Return the arc sine (measured in radians) of x .

`asinh(...)`

`asinh(x)`

Return the inverse hyperbolic sine of x .

`atan(...)`

`atan(x)`

Return the arc tangent (measured in radians) of x .

`atan2(...)`

`atan2(y, x)`

Return the arc tangent (measured in radians) of y/x . Unlike `atan(y/x)`, the signs of both x and y are considered.

`atanh(...)`

`atanh(x)`

Return the inverse hyperbolic tangent of x .

`ceil(...)`

`ceil(x)`

Return the ceiling of x as an Integral. This is the smallest integer $\geq x$.

(continues on next page)

(continued from previous page)

```
copysign(...)  
copysign(x, y)
```

Return a float with the magnitude (absolute value) of x but the sign of y . On platforms that support signed zeros, `copysign(1.0, -0.0)` returns `-1.0`.

```
cos(...)  
cos(x)
```

Return the cosine of x (measured in radians).

```
cosh(...)  
cosh(x)
```

Return the hyperbolic cosine of x .

```
degrees(...)  
degrees(x)
```

Convert angle x from radians to degrees.

```
erf(...)  
erf(x)
```

Error function at x .

```
erfc(...)  
erfc(x)
```

Complementary error function at x .

```
exp(...)  
exp(x)
```

Return e raised to the power of x .

```
expm1(...)  
expm1(x)
```

Return $\exp(x)-1$.

This function avoids the loss of precision involved in the direct evaluation of $\exp(x)-1$ for small x .

```
fabs(...)  
fabs(x)
```

Return the absolute value of the float x .

```
factorial(...)  
factorial(x) -> Integral
```

Find $x!$. Raise a `ValueError` if x is negative or non-integral.

```
floor(...)  
floor(x)
```

(continues on next page)

(continued from previous page)

```

Return the floor of x as an Integral.
This is the largest integer <= x.

fmod(...)
    fmod(x, y)

Return fmod(x, y), according to platform C.  x % y may differ.

frexp(...)
    frexp(x)

Return the mantissa and exponent of x, as pair (m, e).
m is a float and e is an int, such that x = m * 2.**e.
If x is 0, m and e are both 0.  Else 0.5 <= abs(m) < 1.0.

fsum(...)
    fsum(iterable)

Return an accurate floating point sum of values in the iterable.
Assumes IEEE-754 floating point arithmetic.

gamma(...)
    gamma(x)

Gamma function at x.

gcd(...)
    gcd(x, y) -> int
    greatest common divisor of x and y

hypot(...)
    hypot(x, y)

Return the Euclidean distance, sqrt(x*x + y*y).

isclose(...)
    isclose(a, b, *, rel_tol=1e-09, abs_tol=0.0) -> bool

Determine whether two floating point numbers are close in value.

    rel_tol
        maximum difference for being considered "close", relative to the
        magnitude of the input values
    abs_tol
        maximum difference for being considered "close", regardless of the
        magnitude of the input values

Return True if a is close in value to b, and False otherwise.

For the values to be considered close, the difference between them
must be smaller than at least one of the tolerances.

-inf, inf and NaN behave similarly to the IEEE 754 Standard.  That
is, NaN is not close to anything, even itself.  inf and -inf are
only close to themselves.

```

(continues on next page)

(continued from previous page)

```
isfinite(...)
    isfinite(x) -> bool

    Return True if x is neither an infinity nor a NaN, and False otherwise.

isinf(...)
    isinf(x) -> bool

    Return True if x is a positive or negative infinity, and False otherwise.

isnan(...)
    isnan(x) -> bool

    Return True if x is a NaN (not a number), and False otherwise.

ldexp(...)
    ldexp(x, i)

    Return x * (2**i).

lgamma(...)
    lgamma(x)

    Natural logarithm of absolute value of Gamma function at x.

log(...)
    log(x[, base])

    Return the logarithm of x to the given base.
    If the base not specified, returns the natural logarithm (base e) of x.

log10(...)
    log10(x)

    Return the base 10 logarithm of x.

log1p(...)
    log1p(x)

    Return the natural logarithm of 1+x (base e).
    The result is computed in a way which is accurate for x near zero.

log2(...)
    log2(x)

    Return the base 2 logarithm of x.

modf(...)
    modf(x)

    Return the fractional and integer parts of x. Both results carry the sign
    of x and are floats.

pow(...)
    pow(x, y)
```

(continues on next page)

(continued from previous page)

```

    Return x**y (x to the power of y).

radians(...)
    radians(x)

    Convert angle x from degrees to radians.

sin(...)
    sin(x)

    Return the sine of x (measured in radians).

sinh(...)
    sinh(x)

    Return the hyperbolic sine of x.

sqrt(...)
    sqrt(x)

    Return the square root of x.

tan(...)
    tan(x)

    Return the tangent of x (measured in radians).

tanh(...)
    tanh(x)

    Return the hyperbolic tangent of x.

trunc(...)
    trunc(x:Real) -> Integral

    Truncates x to the nearest Integral toward 0. Uses the __trunc__ magic
    ↪method.

DATA
    e = 2.718281828459045
    inf = inf
    nan = nan
    pi = 3.141592653589793
    tau = 6.283185307179586

FILE
    /Users/fangohr/anaconda3/lib/python3.6/lib-dynload/math.cpython-36m-darwin.so

```

The help function needs to be given the name of an object (which must exist in the current name space). For example `help(math.sqrt)` will not work if the `math` module has not been imported before.

```

# NBVAL_IGNORE_OUTPUT
help(math.sqrt)

```

```
Help on built-in function sqrt in module math:
```

```
sqrt(...)
    sqrt(x)

    Return the square root of x.
```

```
# NBVAL_IGNORE_OUTPUT
import math
help(math.sqrt)
```

```
Help on built-in function sqrt in module math:
```

```
sqrt(...)
    sqrt(x)

    Return the square root of x.
```

Instead of importing the module, we could also have given the *string* of `math.sqrt` to the help function, i.e.:

```
# NBVAL_IGNORE_OUTPUT
help('math.sqrt')
```

```
Help on built-in function sqrt in math:
```

```
math.sqrt = sqrt(...)
    sqrt(x)

    Return the square root of x.
```

`help` is a function which gives information about the object which is passed as its argument. Most things in Python (classes, functions, modules, etc.) are objects, and therefore can be passed to `help`. There are, however, some things on which you might like to ask for help, which are not existing Python objects. In such cases it is often possible to pass a string containing the name of the thing or concept to help, for example

- `help('modules')` will generate a list of all modules which can be imported into the current interpreter. Note that `help(modules)` (note absence of quotes) will result in a `NameError` (unless you are unlucky enough to have a variable called `modules` floating around, in which case you will get help on whatever that variable happens to refer to.)
- `help('some_module')`, where `some_module` is a module which has not been imported yet (and therefore isn't an object yet), will give you that module's help information.
- `help('some_keyword')`: For example `and`, `if` or `print` (i.e. `help('and')`, `help('if')` and `help('print')`). These are special words recognized by Python: they are not objects and thus cannot be passed as arguments to `help`. Passing the name of the keyword as a string to `help` works, but only if you have Python's HTML documentation installed, and the interpreter has been made aware of its location by setting the environment variable `PYTHONDOCS`.

4.5 Docstrings

The command `help(<object>)` accesses the documentation strings of objects.

Any literal string appearing as the first item in the definition of a class, function, method or module, is taken to be its *docstring*.

`help` includes the docstring in the information it displays about the object.

In addition to the docstring it may display some other information, for example, in the case of functions, it displays the function's signature.

The docstring is stored in the object's `__doc__` attribute.

```
# NBVAL_IGNORE_OUTPUT
help(math.sin)
```

```
Help on built-in function sin in module math:

sin(...)
    sin(x)

    Return the sine of x (measured in radians).
```

```
# NBVAL_IGNORE_OUTPUT
print(math.sin.__doc__)
```

```
sin(x)

Return the sine of x (measured in radians).
```

For user-defined functions, classes, types, modules, ...), one should always provide a docstring.

Documenting a user-provided function:

```
def power2and3(x):
    """Returns the tuple (x**2, x**3)"""
    return x**2, x**3
```

```
power2and3(2)
```

```
(4, 8)
```

```
power2and3(4.5)
```

```
(20.25, 91.125)
```

```
power2and3(0+1j)
```

```
((-1+0j), (-0-1j))
```

```
help(power2and3)
```

```
Help on function power2and3 in module __main__:
```

```
power2and3(x)  
    Returns the tuple (x**2, x**3)
```

```
print(power2and3.__doc__)
```

```
Returns the tuple (x**2, x**3)
```

INPUT AND OUTPUT

In this section, we describe printing, which includes the use of the `print` function, the old-style `%` format specifiers and the new style `{ }` format specifiers.

5.1 Printing to standard output (normally the screen)

The `print` function is the most commonly used command to print information to the “standard output device” which is normally the screen.

There are two modes to use `print`.

5.1.1 Simple print

The easiest way to use the `print` command is to list the variables to be printed, separated by comma. Here are a few examples:

```
a = 10
b = 'test text'
print(a)
```

```
10
```

```
print(b)
```

```
test text
```

```
print(a, b)
```

```
10 test text
```

```
print("The answer is", a)
```

```
The answer is 10
```

```
print("The answer is", a, "and the string contains", b)
```

```
The answer is 10 and the string contains test text
```

```
print("The answer is", a, "and the string reads", b)
```

```
The answer is 10 and the string reads test text
```

Python adds a space between every object that is being printed.

Python prints a new line after every print call. To suppress that, use the `end=` parameter:

```
print("Printing in line one", end='')  
print("...still printing in line one.")
```

```
Printing in line one...still printing in line one.
```

5.1.2 Formatted printing

The more sophisticated way of formatting output uses a syntax very similar to Matlab's `fprintf` (and therefore also similar to C's `printf`).

The overall structure is that there is a string containing format specifiers, followed by a percentage sign and a tuple that contains the variables to be printed in place of the format specifiers.

```
print("a = %d b = %d" % (10,20))
```

```
a = 10 b = 20
```

A string can contain format identifiers (such as `%f` to format as a float, `%d` to format as an integer, and `%s` to format as a string):

```
from math import pi  
print("Pi = %5.2f" % pi)
```

```
Pi = 3.14
```

```
print("Pi = %10.3f" % pi)
```

```
Pi = 3.142
```

```
print("Pi = %10.8f" % pi)
```

```
Pi = 3.14159265
```

```
print("Pi = %d" % pi)
```

```
Pi = 3
```

The format specifier of type `%W.Df` means that a Float should be printed with a total Width of `W` characters and `D` digits behind the Decimal point. (This is identical to Matlab and C, for example.)

To print more than one object, provide multiple format specifiers and list several objects in the tuple:

```
print("Pi = %f, 142*pi = %f and pi^2 = %f." % (pi,142*pi,pi**2))
```

```
Pi = 3.141593, 142*pi = 446.106157 and pi^2 = 9.869604.
```

Note that the conversion of a format specifier and a tuple of variables into string does not rely on the `print` command:

```
from math import pi
"pi = %f" % pi
```

```
'pi = 3.141593'
```

This means that we can convert objects into strings wherever we need, and we can decide to print the strings later – there is no need to couple the formatting closely to the code that does the printing.

Overview of commonly used format specifiers using the astronomical unit as an example:

```
AU = 149597870700 # astronomical unit [m]
"%f" % AU # line 1 in table
```

```
'149597870700.000000'
```

specifier	style	Example output for AU
<code>%f</code>	floating point	149597870700.000000
<code>%e</code>	exponential notation	1.495979e+11
<code>%g</code>	shorter of <code>%e</code> or <code>%f</code>	1.49598e+11
<code>%d</code>	integer	149597870700
<code>%s</code>	<code>str()</code>	149597870700
<code>%r</code>	<code>repr()</code>	149597870700L

5.1.3 “str” and “__str__”

All objects in Python should provide a method `__str__` which returns a nice string representation of the object. This method `a.__str__()` is called when we apply the `str` function to object `a`:

```
a = 3.14
a.__str__()
```

```
'3.14'
```

```
str(a)
```

```
'3.14'
```

The `str` function is extremely convenient as it allows us to print more complicated objects, such as

```
b = [3, 4.2, ['apple', 'banana'], (0, 1)]
str(b)
```

```
"[3, 4.2, ['apple', 'banana'], (0, 1)]"
```

The way Python prints this is that it uses the `__str__` method of the list object. This will print the opening square bracket `[` and then call the `__str__` method of the first object, i.e. the integer 3. This will produce 3. Then the list object's `__str__` method prints the comma `,` and moves on to call the `__str__` method of the next element in the list (i.e. 4.2) to print itself. This way any composite object can be represented as a string by asking the objects it holds to convert themselves to strings.

The string method of object `x` is called implicitly, when we

- use the “%s” format specifier to print `x`
- pass the object `x` directly to the print command:

```
print(b)
```

```
[3, 4.2, ['apple', 'banana'], (0, 1)]
```

```
print("%s" % b)
```

```
[3, 4.2, ['apple', 'banana'], (0, 1)]
```

5.1.4 “repr” and “__repr__”

A second function, `repr`, should convert a given object into a string presentation *so that this string can be used to re-created the object using the `eval` function*. The `repr` function will generally provide a more detailed string than `str`. Applying `repr` to the object `x` will attempt to call `x.__repr__()`.

```
from math import pi as a1
str(a1)
```

```
'3.141592653589793'
```

```
repr(a1)
```

```
'3.141592653589793'
```

```
number_as_string = repr(a1)
a2 = eval(number_as_string) # evaluate string
a2
```

```
3.141592653589793
```

```
a2-a1 # -> repr is exact representation
```

```
0.0
```

```
a1-eval(repr(a1))
```

```
0.0
```

```
a1-eval(str(a1))           # -> str has lost a few digits
```

```
0.0
```

We can convert an object to its `str()` or `repr` presentation using the format specifiers `%s` and `%r`, respectively.

```
import math
"%s" % math.pi
```

```
'3.141592653589793'
```

```
"%r" % math.pi
```

```
'3.141592653589793'
```

5.1.5 New-style string formatting

A new system of built-in formatting allows more flexibility for complex cases, at the cost of being a bit longer.

Basic ideas in examples:

```
"{} needs {} pints".format('Peter', 4)   # insert values in order
```

```
'Peter needs 4 pints'
```

```
"{0} needs {1} pints".format('Peter', 4) # index which element
```

```
'Peter needs 4 pints'
```

```
"{1} needs {0} pints".format('Peter', 4)
```

```
'4 needs Peter pints'
```

```
"{name} needs {number} pints".format(    # reference element to
    name='Peter',                        # print by name
    number=4
)
```

```
'Peter needs 4 pints'
```

```
"Pi is approximately {:.f}.".format(math.pi)    # can use old-style format options_
↪for float
```

```
'Pi is approximately 3.141593.'
```

```
"Pi is approximately {:.2f}.".format(math.pi)    # and precision
```

```
'Pi is approximately 3.14.'
```

```
"Pi is approximately {:.6.2f}.".format(math.pi)  # and width
```

```
'Pi is approximately   3.14.'
```

This is a powerful and elegant way of string formatting, which is gradually being used more.

Further information

- Examples <https://docs.python.org/3/library/string.html#format-examples>
- Python Enhancement Proposal 3101
- Python library String Formatting Operations
- Introduction to Fancier Output Formatting, Python tutorial, section 7.1

5.1.6 Changes from Python 2 to Python 3: `print`

One (maybe the most obvious) change going from Python 2 to Python 3 is that the `print` command loses its special status. In Python 2, we could print “Hello World” using:

```
print "Hello world"    # valid in Python 2.x
```

Effectively, we call the function `print` with the argument `Hello World`. All other functions in Python are called such that the argument is enclosed in parentheses, i.e.

```
print("Hello World")    # valid in Python 3.x
```

```
Hello World
```

This is the new convention *required* in Python 3 (and *allowed* for recent version of Python 2.x.)

Everything we have learned about formatting strings using the percentage operator still works the same way:

```
import math
a = math.pi
"my pi = %f" % a    # string formatting
```

```
'my pi = 3.141593'
```



```
print("my pi = %f" % a)    # valid print in 2.7 and 3.x
```

```
my pi = 3.141593
```

```
"Short pi = %.2f, longer pi = %.12f." % (a, a)
```

```
'Short pi = 3.14, longer pi = 3.141592653590.'
```

```
print("Short pi = %.2f, longer pi = %.12f." % (a, a))
```

```
Short pi = 3.14, longer pi = 3.141592653590.
```

```
print("Short pi = %.2f, longer pi = %.12f." % (a, a))
```

```
Short pi = 3.14, longer pi = 3.141592653590.
```

```
# 1. Write a file
out_file = open("test.txt", "w")           #'w' stands for Writing
out_file.write("Writing text to file. This is the first line.\n"+\
               "And the second line.")
out_file.close()                           #close the file

# 2. Read a file
in_file = open("test.txt", "r")            #'r' stands for Reading
text = in_file.read()                      #read complete file into
                                           #string variable text
in_file.close()                             #close the file

# 3. Display data
print(text)
```

```
Writing text to file. This is the first line.
And the second line.
```

5.2 Reading and writing files

Here is a program that

1. writes some text to a file with name `test.txt`,
2. and then reads the text again and
3. prints it to the screen.

The data stored in the file `test.txt` is:

```
Writing text to file. This is the first line.
And the second line.
```

In more detail, you have opened a file with the `open` command, and assigned this open file object to the variable `out_file`. We have then written data to the file using the `out_file.write` method. Note that in the example above, we have given a string to the `write` method. We can, of course, use all the formatting that we have discussed before—see *formatted printing* and *new style formatting*. For example, to write this file with name `table.txt` we can use this Python program. It is good practice to `close()` files when we have finished reading and writing. If a Python program is left in a controlled way (i.e. not through a power cut or an unlikely bug deep in the Python language or the operating system) then it will close all open files as soon as the file objects are destroyed. However, closing them actively as soon as possible is better style.

5.2.1 File reading examples

We use a file named `myfile.txt` containing the following 3 lines of text for the examples below:

```
This is the first line.
This is the second line.
This is a third and last line.
```

```
f = open('myfile.txt', 'w')
f.write('This is the first line.\n'
       'This is the second line.\n'
       'This is a third and last line.')
f.close()
```

`fileobject.read()`

The `fileobject.read()` method reads the whole file, and returns it as one string (including new line characters).

```
f = open('myfile.txt', 'r')
f.read()
```

```
'This is the first line.\nThis is the second line.\nThis is a third and last line.'
```

```
f.close()
```

`fileobject.readlines()`

The `fileobject.readlines()` method returns a list of strings, where each element of the list corresponds to one line in the string:

```
f = open('myfile.txt', 'r')
f.readlines()
```

```
['This is the first line.\n',
 'This is the second line.\n',
 'This is a third and last line.']
```

```
f.close()
```

This is often used to iterate over the lines, and to do something with each line. For example:

```
f = open('myfile.txt', 'r')
for line in f.readlines():
    print("%d characters" % len(line))
f.close()
```

```
24 characters
25 characters
30 characters
```

Note that this will read the complete file into a list of strings when the `readlines()` method is called. This is no problem if we know that the file is small and will fit into the machine's memory.

If so, we can also close the file before we process the data, i.e.:

```
f = open('myfile.txt', 'r')
lines = f.readlines()
f.close()
for line in lines:
    print("%d characters" % len(line))
```

```
24 characters
25 characters
30 characters
```

Iterating over lines (file object)

There is a neater possibility to read a file line by line which (i) will only read one line at a time (and is thus suitable for large files as well) and (ii) results in more compact code:

```
f = open('myfile.txt', 'r')
for line in f:
    print("%d characters" % len(line))
f.close()
```

```
24 characters
25 characters
30 characters
```

Here, the file handler `f` acts as an iterator and will return the next line in every subsequent iteration of the for-loop until the end of the file is reached (and then the for-loop is terminated).

5.3 Further reading

Methods of File objects, Tutorial, Section 7.2.1

CONTROL FLOW

6.1 Basics

For a given file with a python program, the python interpreter will start at the top and then process the file. We demonstrate this with a simple program, for example:

```
def f(x):  
    """function that computes and returns x*x"""  
    return x * x  
  
print("Main program starts here")  
print("4 * 4 = %s" % f(4))  
print("In last line of program -- bye")
```

```
Main program starts here  
4 * 4 = 16  
In last line of program -- bye
```

The basic rule is that commands in a file (or function or any sequence of commands) is processed from top to bottom. If several commands are given in the same line (separated by `;`), then these are processed from left to right (although it is discouraged to have multiple statements per line to maintain good readability of the code.)

In this example, the interpreter starts at the top (line 1). It finds the `def` keyword and remembers for the future that the function `f` is defined here. (It will not yet execute the function body, i.e. line 3 – this only happens when we call the function.) The interpreter can see from the indentation where the body of the function stops: the indentation in line 5 is different from that of the first line in the function body (line 2), and thus the function body has ended, and execution should carry on with that line. (Empty lines do not matter for this analysis.)

In line 5 the interpreter will print the output `Main program starts here`. Then line 6 is executed. This contains the expression `f(4)` which will call the function `f(x)` which is defined in line 1 where `x` will take the value 4. [Actually `x` is a reference to the object 4.] The function `f` is then executed and computes and returns `4*4` in line 3. This value 16 is used in line 6 to replace `f(4)` and then the string representation `%s` of the object 16 is printed as part of the print command in line 6.

The interpreter then moves on to line 7 before the program ends.

We will now learn about different possibilities to direct this control flow further.

6.1.1 Conditionals

The python values `True` and `False` are special inbuilt objects:

```
a = True
print(a)
```

True

```
type(a)
```

bool

```
b = False
print(b)
```

False

```
type(b)
```

bool

We can operate with these two logical values using boolean logic, for example the logical and operation (`and`):

```
True and True          #logical and operation
```

True

```
True and False
```

False

```
False and True
```

False

```
True and True
```

True

```
c = a and b
print(c)
```

False

There is also logical or (`or`) and the negation (`not`):

True or False

True

not True

False

not False

True

True and not False

True

In computer code, we often need to evaluate some expression that is either true or false (sometimes called a “predicate”). For example:

```
x = 30          # assign 30 to x
x > 15         # is x greater than 15
```

True

```
x > 42
```

False

```
x == 30        # is x the same as 30?
```

True

```
x == 42
```

False

```
not x == 42    # is x not the same as 42?
```

True

```
x != 42        # is x not the same as 42?
```

```
True
```

```
x > 30      # is x greater than 30?
```

```
False
```

```
x >= 30    # is x greater than or equal to 30?
```

```
True
```

6.2 If-then-else

Further information

- Introduction to If-then in Python tutorial, section 4.1

The `if` statement allows conditional execution of code, for example:

```
a = 34
if a > 0:
    print("a is positive")
```

```
a is positive
```

The `if`-statement can also have an `else` branch which is executed if the condition is wrong:

```
a = 34
if a > 0:
    print("a is positive")
else:
    print("a is non-positive (i.e. negative or zero)")
```

```
a is positive
```

Finally, there is the `elif` (read as “else if”) keyword that allows checking for several (exclusive) possibilities:

```
a = 17
if a == 0:
    print("a is zero")
elif a < 0:
    print("a is negative")
else:
    print("a is positive")
```

```
a is positive
```


6.3 For loop

Further information

- Introduction to for-loops in [Python tutorial, section 4.2](#)

The `for`-loop allows to iterate over a sequence (this could be a string or a list, for example). Here is an example:

```
for animal in ['dog', 'cat', 'mouse']:
    print(animal, animal.upper())
```

```
dog DOG
cat CAT
mouse MOUSE
```

Together with the `range()` command (03 Data Types Structures, The Range Command), one can iterate over increasing integers:

```
for i in range(5,10):
    print(i)
```

```
5
6
7
8
9
```

6.4 While loop

The `while` keyword allows to repeat an operation while a condition is true. Suppose we'd like to know for how many years we have to keep 100 pounds on a savings account to reach 200 pounds simply due to annual payment of interest at a rate of 5%. Here is a program to compute that this will take 15 years:

```
mymoney = 100          # in GBP
rate = 1.05            # 5% interest
years = 0
while mymoney < 200:  # repeat until 20 pounds reached
    mymoney = mymoney * rate
    years = years + 1
print('We need', years, 'years to reach', mymoney, 'pounds.')
```

We need 15 years to reach 207.89281794113688 pounds.

6.5 Relational operators (comparisons) in `if` and `while` statements

The general form of `if` statements and `while` loops is the same: following the keyword `if` or `while`, there is a *condition* followed by a colon. In the next line, a new (and thus indented!) block of commands starts that is executed if the condition is True).

For example, the condition could be equality of two variables `a1` and `a2` which is expressed as `a1==a2`:

```
a1 = 42
a2 = 42
if a1 == a2:
    print("a1 and a2 are the same")
```

a1 and a2 are the same

Another example is to test whether `a1` and `a2` are not the same. For this, we have two possibilities. Option number 1 uses the *inequality operator* `!=`:

```
if a1 != a2:
    print("a1 and a2 are different")
```

Option two uses the keyword `not` in front of the condition:

```
if not a1 == a2:
    print("a1 and a2 are different")
```

Comparisons for “greater” (`>`), “smaller” (`<`) and “greater equal” (`>=`) and “smaller equal” (`<=`) are straightforward.

Finally, we can use the logical operators “and” and “or” to combine conditions:

```
if a > 10 and b > 20:
    print("A is greater than 10 and b is greater than 20")
if a > 10 or b < -5:
    print("Either a is greater than 10, or "
          "b is smaller than -5, or both.")
```

Either a is greater than 10, or b is smaller than -5, or both.

Use the Python prompt to experiment with these comparisons and logical expressions. For example:

```
T = -12.5
if T < -20:
    print("very cold")

if T < -10:
    print("quite cold")
```

quite cold

```
T < -20
```

```
False
```

```
T < -10
```

```
True
```

6.6 Exceptions

Even if a statement or expression is syntactically correct, it may cause an error when an attempt is made to execute it. Errors detected during execution are called *exceptions* and are not necessarily fatal: exceptions can be *caught* and dealt with within the program. Most exceptions are not handled by programs, however, and result in error messages as shown here

```
# NBVAL_RAISES_EXCEPTION
10 * (1/0)
```

```
-----
ZeroDivisionError                                Traceback (most recent call last)
Cell In[36], line 2
      1 # NBVAL_RAISES_EXCEPTION
----> 2 10 * (1/0)

ZeroDivisionError: division by zero
```

```
# NBVAL_RAISES_EXCEPTION
4 + spam*3
```

```
-----
NameError                                         Traceback (most recent call last)
Cell In[37], line 2
      1 # NBVAL_RAISES_EXCEPTION
----> 2 4 + spam*3

NameError: name 'spam' is not defined
```

```
# NBVAL_SKIP
'2' + 2
```

```
-----
TypeError                                         Traceback (most recent call last)
Cell In[38], line 2
      1 # NBVAL_SKIP
----> 2 '2' + 2

TypeError: can only concatenate str (not "int") to str
```

Schematic exception catching with all options

```

try:
    # code body
    pass
except ArithmeticError:
    # what to do if arithmetic error
    pass
except IndexError as the_exception:
    # the_exception refers to the exception in this block
    pass
except:
    # what to do for ANY other exception
    pass
else: # optional
    # what to do if no exception raised
    pass

try:
    # code body
    pass
finally:
    # what to do ALWAYS
    pass

```

Starting with Python 2.5, you can use the `with` statement to simplify the writing of code for some predefined functions, in particular the `open` function to open files: see <https://docs.python.org/3/tutorial/errors.html#predefined-clean-up-actions>.

Example: We try to open a file that does not exist, and Python will raise an exception of type `FileNotFoundError` because the file cannot be found:

```

# NBVAL_RAISES_EXCEPTION
f = open("filenamethatdoesnotexist", "r")

```

```

-----
FileNotFoundError                                Traceback (most recent call last)
Cell In[40], line 2
      1 # NBVAL_RAISES_EXCEPTION
----> 2 f = open("filenamethatdoesnotexist", "r")

File /opt/venv/lib/python3.11/site-packages/IPython/core/interactiveshell.py:324, in
↳ _modified_open(file, *args, **kwargs)
    317 if file in {0, 1, 2}:
    318     raise ValueError(
    319         f"IPython won't let you open fd={file} by default "
    320         "as it is likely to crash IPython. If you know what you are doing,
↳ ""
    321         "you can use builtins' open."
    322     )
--> 324 return io_open(file, *args, **kwargs)

FileNotFoundError: [Errno 2] No such file or directory: 'filenamethatdoesnotexist'

```

If we were writing an application with a user interface where the user has to type or select a filename, we would not want to application to stop if the file does not exist. Instead, we need to catch this exception and act accordingly (for example by informing the user that a file with this filename does not exist and ask whether they want to try another file name).

Here is the skeleton for catching this exception:

```
try:
    f = open("filenamethatdoesnotexist", "r")
except FileNotFoundError:
    print("Could not open that file")
```

```
Could not open that file
```

There is a lot more to be said about exceptions and their use in larger programs. Start reading [Python Tutorial Chapter 8: Errors and Exceptions](#) if you are interested.

6.6.1 Raising Exceptions

Raising exception is also referred to as 'throwing an exception'.

Possibilities of raising an Exception

- `raise OverflowError`
- `raise OverflowError, "Bath is full"` (Old style, now discouraged)
- `raise OverflowError("Bath is full")`
- `e = OverflowError("Bath is full"); raise e`

Exception hierarchy

The standard exceptions are organized in an inheritance hierarchy e.g. `OverflowError` is a subclass of `ArithmeticError` (not `BathroomError`); this can be seen when looking at `help('exceptions')` for example.

You can derive your own exceptions from any of the standard ones. It is good style to have each module define its own base exception.

6.6.2 Creating our own exceptions

- You can and should derive your own exceptions from the built-in Exception.
- To see what built-in exceptions exist, look in the module `exceptions` (try `help('exceptions')`), or go to <https://docs.python.org/3/library/exceptions.html#builtin-exceptions>.

6.6.3 LBYL vs EAFP

- LBYL (Look Before You Leap) vs
- EAFP (Easier to ask forgiveness than permission)

```
numerator = 7
denominator = 0
```

Example for LBYL:

```
if denominator == 0:
    print("Oops")
else:
    print(numerator/denominator)
```

Oops

Easier to Ask for Forgiveness than Permission:

```
try:
    print(numerator/denominator)
except ZeroDivisionError:
    print("Oops")
```

Oops

The Python documentation says about EAFP:

Easier to ask for forgiveness than permission. This common Python coding style assumes the existence of valid keys or attributes and catches exceptions if the assumption proves false. This clean and fast style is characterized by the presence of many try and except statements. The technique contrasts with the LBYL style common to many other languages such as C.

Source: <https://docs.python.org/3/glossary.html#term-eafp>

The Python documentation says about LBYL:

Look before you leap. This coding style explicitly tests for pre-conditions before making calls or lookups. This style contrasts with the EAFP approach and is characterized by the presence of many if statements.

In a multi-threaded environment, the LBYL approach can risk introducing a race condition between “the looking” and “the leaping”. For example, the code, `if key in mapping: return mapping[key]` can fail if another thread removes key from mapping after the test, but before the lookup. This issue can be solved with locks or by using the EAFP approach.

Source: <https://docs.python.org/3/glossary.html#term-lbyl>

EAFP is the Pythonic way.

FUNCTIONS AND MODULES

7.1 Introduction

Functions allow us to group a number of statements into a logical block. We communicate with a function through a clearly defined interface, providing certain parameters to the function, and receiving some information back. Apart from this interface, we generally do not know exactly how a function does the work to obtain the value it returns.

For example the function `math.sqrt`: we do not know how exactly it computes the square root, but we know about the interface: if we pass x into the function, it will return (an approximation of) \sqrt{x} .

This abstraction is a useful thing: it is a common technique in engineering to break down a system into smaller (black-box) components that all work together through well defined interfaces, but which do not need to know about the internal realisations of each other's functionality. In fact, not having to care about these implementation details can help to have a clearer view of the system composed of many of these components.

Functions provide the basic building blocks of functionality in larger programs (and computer simulations), and help to control the inherent complexity of the process.

We can group functions together into a Python module (see *modules*), and in this way create our own libraries of functionality.

7.2 Using functions

The word “function” has different meanings in mathematics and programming. In programming it refers to a named sequence of operations that perform a computation. For example, the function `sqrt()` which is defined in the `math` module computes the square root of a given value:

```
from math import sqrt
sqrt(4)
```

```
2.0
```

The value we pass to the function `sqrt` is 4 in this example. This value is called the *argument* of the function. A function may have more than one argument.

The function returns the value 2.0 (the result of its computation) to the “calling context”. This value is called the *return value* of the function.

It is common to say that a function *takes* an argument and *returns* a result or return value.

Common confusion about printing and returning values

It is a common beginner's mistake to confuse the *printing* of values with *returning* values. In the following example it is hard to see whether the function `math.sin` returns a value or whether it prints the value:

```
import math
math.sin(2)
```

```
0.9092974268256817
```

We import the `math` module, and call the `math.sin` function with an argument of 2. The `math.sin(2)` call will actually *return* the value `0.909...` not print it. However, because we have not assigned the return value to a variable, the Python prompt will print the returned object.

The following alternative sequence works only if the value is returned:

```
x = math.sin(2)
print(x)
```

```
0.9092974268256817
```

The return value of the function call `math.sin(2)` is assigned to the variable `x`, and `x` is printed in the next line.

Generally, functions should execute “silently” (i.e. not print anything) and report the result of their computation through the return value.

Part of the confusion about printed versus return values at the Python prompt comes from the Python prompt printing (a representation) of returned objects *if* the returned objects are not assigned. Generally, seeing the returned objects is exactly what we want (as we normally care about the returned object), just when learning Python this may cause mild confusion about functions returning values or printing values.

Further information

- Think Python has a gentle introduction to functions (on which the previous paragraph is based) in [chapter 3 \(Functions\)](#) and [chapter 6 \(Fruitful functions\)](#).

7.3 Defining functions

The generic format of a function definitions:

```
def my_function(arg1, arg2, ..., argn):
    """Optional docstring."""

    # Implementation of the function

    return result # optional

# this is not part of the function
some_command
```

Allen Downey's terminology (in his book [Think Python](#)) of fruitful and fruitless functions distinguishes between functions that return a value, and those that do not return a value. The distinction refers to whether a function provides a return value (=fruitful) or whether the function does not explicitly return a value (=fruitless). If a function does not make use of the `return` statement, we tend to say that the function returns nothing (whereas in reality it will always return the `None` object when it terminates – even if the `return` statement is missing).

For example, the function `greeting` will print “Hello World” when called (and is fruitless as it does not return a value).


```
def greeting():  
    print("Hello World!")
```

If we call that function:

```
greeting()
```

```
Hello World!
```

it prints “Hello World” to stdout, as we would expect. If we assign the return value of the function to a variable `x`, we can inspect it subsequently:

```
x = greeting()
```

```
Hello World!
```

```
print(x)
```

```
None
```

and find that the `greeting` function has indeed returned the `None` object.

Another example for a function that does not return any value (that means there is no `return` keyword in the function) would be:

```
def printpluses(n):  
    print(n * "+")
```

Generally, functions that return values are more useful as these can be used to assemble code (maybe as another function) by combining them cleverly. Let’s look at some examples of functions that do return a value.

Suppose we need to define a function that computes the square of a given variable. The function source could be:

```
def square(x):  
    return x * x
```

The keyword `def` tells Python that we are *defining* a function at that point. The function takes one argument (`x`). The function returns `x*x` which is of course x^2 . Here is the listing of a file that shows how the function can be defined and used: (note that the numbers on the left are line numbers and are not part of the program)

```
def square(x):  
    return x * x  
  
for i in range(5):  
    i_squared = square(i)  
    print(i, '*', i, '=', i_squared)
```

```
0 * 0 = 0  
1 * 1 = 1  
2 * 2 = 4  
3 * 3 = 9  
4 * 4 = 16
```

It is worth mentioning that lines 1 and 2 define the square function whereas lines 4 to 6 are the main program.

We can define functions that take more than one argument:

```
import math

def hypot(x, y):
    return math.sqrt(x * x + y * y)
```

It is also possible to return more than one argument. Here is an example of a function that converts a given string into all characters uppercase and all characters lowercase and returns the two versions. We have included the main program to show how this function can be called:

```
def upperAndLower(string):
    return string.upper(), string.lower()

testword = 'Banana'

uppercase, lowercase = upperAndLower(testword)

print(testword, 'in lowercase:', lowercase,
      'and in uppercase', uppercase)
```

```
Banana in lowercase: banana and in uppercase BANANA
```

We can define multiple Python functions in one file. Here is an example with two functions:

```
def returnstars( n ):
    return n * '*'

def print_centred_in_stars( string ):
    linelength = 46
    starstring = returnstars((linelength - len(string)) // 2)

    print(starstring + string + starstring)

print_centred_in_stars('Hello world!')
```

```
*****Hello world!*****
```

Further reading

- [Python Tutorial: Section 4.6 Defining Functions](#)

7.4 Default values and optional parameters

Python allows to define *default* values for function parameters. Here is an example: This program will print the following output when executed: So how does it work? The function `print_mult_table` takes two arguments: `n` and `upto`. The first argument `n` is a “normal” variable. The second argument `upto` has a default value of 10. In other words: should the user of this function only provide one argument, then this provides the value for `n` and `upto` will default to 10. If two arguments are provided, the first one will be for `n` and the second for `upto` (as shown in the code example above).

7.5 Modules

Modules

- Group together functionality
- Provide namespaces
- Python’s standard library contains a vast collection of modules - “Batteries Included”
- Try `help('modules')`
- Means of extending Python

7.5.1 Importing modules

```
import math
```

This will introduce the name `math` into the namespace in which the `import` command was issued. The names within the `math` module will not appear in the enclosing namespace: they must be accessed through the name `math`. For example: `math.sin`.

```
import math, cmath
```

More than one module can be imported in the same statement, although the [Python Style Guide](#) recommends not to do this. Instead, we should write

```
import math
import cmath

import math as mathematics
```

The name by which the module is known locally can be different from its “official” name. Typical uses of this are

- To avoid name clashes with existing names
- To change the name to something more manageable. For example `import SimpleHTTPServer as shs`. This is discouraged for production code (as longer meaningful names make programs far more understandable than short cryptic ones), but for interactively testing out ideas, being able to use a short synonym can make your life much easier. Given that (imported) modules are first class objects, you can, of course, simply do `shs = SimpleHTTPServer` in order to obtain the more easily typable handle on the module.

```
from math import sin
```

This will import the `sin` function from the `math` module, but it will not introduce the name `math` into the current namespace. It will only introduce the name `sin` into the current namespace. It is possible to pull in more than one name from the module in one go:

```
from math import sin, cos
```

Finally, let’s look at this notation:

```
from math import *
```

Once again, this does not introduce the name `math` into the current namespace. It does however introduce *all public names* of the `math` module into the current namespace. Broadly speaking, it is a bad idea to do this:

- Lots of new names will be dumped into the current namespace.
- Are you sure they will not clobber any names already present?
- It will be very difficult to trace where these names came from
- Having said that, some modules (including ones in the standard library, recommend that they be imported in this way). Use with caution!
- This is fine for interactive quick and dirty testing or small calculations.

7.5.2 Creating modules

A module is in principle nothing else than a python file. We create an example of a module file which is saved in `module1.py`:

```
%%file module1.py
def someusefulfunction():
    pass

print("My name is", __name__)
```

```
Writing module1.py
```

We can execute this (module) file as a normal python program (for example `python module1.py`):

```
!python3 module1.py
```

```
My name is __main__
```

We note that the Python magic variable `__name__` takes the value `__main__` if the program file `module1.py` is executed.

On the other hand, we can *import* `module1.py` in another file (which could have the name `prog.py`), for example like this:

```
import module1          # in file prog.py
```

```
My name is module1
```

When Python comes across the `import module1` statement in `prog.py`, it looks for the file `module1.py` in the current working directory (and if it can't find it there in all the directories in `sys.path`) and opens the file `module1.py`. While parsing the file `module1.py` from top to bottom, it will add any function definitions in this file into the `module1` name space in the calling context (that is the main program in `prog.py`). In this example, there is only the function `someusefulfunction`. Once the import process is completed, we can make use of `module1.someusefulfunction` in `prog.py`. If Python comes across statements other than function (and class) definitions while importing `module1.py`, it carries those out immediately. In this case, it will thus come across the statement `print(My name is, __name__)`.

Note the difference to the output if we *import* `module1.py` rather than executing it on its own: `__name__` inside a module takes the value of the module name if the file is imported.

7.5.3 Use of `__name__`

In summary,

- `__name__` is `__main__` if the module file is run on its own
- `__name__` is the name of the module (i.e. the module filename without the `.py` suffix) if the module file is imported.

We can therefore use the following `if` statement in `module1.py` to write code that is *only run* when the module is executed on its own: This is useful to keep test programs or demonstrations of the abilities of a module in this “conditional” main program. It is common practice for any module files to have such a conditional main program which demonstrates its capabilities.

7.5.4 Example 1

The next example shows a main program for the another file `vectools.py` that is used to demonstrate the capabilities of the functions defined in that file:

```
%%file vectools.py
from __future__ import division
import math

import numpy as N

def norm(x):
    """returns the magnitude of a vector x"""
    return math.sqrt(sum(x ** 2))

def unitvector(x):
    """returns a unit vector x/|x|. x needs to be a numpy array."""
    xnorm = norm(x)
    if xnorm == 0:
        raise ValueError("Can't normalise vector with length 0")
    return x / norm(x)

if __name__ == "__main__":
    # a little demo of how the functions in this module can be used:
    x1 = N.array([0, 1, 2])
    print("The norm of " + str(x1) + " is " + str(norm(x1)) + ".")
    print("The unitvector in direction of " + str(x1) + " is " \
          + str(unitvector(x1)) + ".")
```

Writing `vectools.py`

If this file is executed using `python vectools.py`, then `__name__==__main__` is true, and the output reads

```
!python3 vectools.py
```

```
The norm of [0 1 2] is 2.23606797749979.
The unitvector in direction of [0 1 2] is [0.          0.4472136  0.89442719].
```

If this file is imported (i.e. used as a module) into another python file or the python prompt or in the Jupyter Notebook, then `__name__ == __main__` is false, and that statement block will not be executed.

This is quite a common way to conditionally execute code in files providing library-like functions. The code that is executed if the file is run on its own, often consists of a series of tests (to check that the file's functions carry out the right operations – *regression tests* or *unit tests*), or some examples of how the library functions in the file can be used.

7.5.5 Example 2

Even if a Python program is not intended to be used as a module file, it is good practice to always use a conditional main program:

- often, it turns out later that functions in the file can be reused (and saves work then)
- this is convenient for regression testing.

Suppose an exercise is given to write a function that returns the first 5 prime numbers, and in addition to print them. (There is of course a trivial solution to this as we know the prime numbers, and we should imagine that the required calculation is more complex). One might be tempted to write

```
def primes5():
    return (2, 3, 5, 7, 11)

for p in primes5():
    print("%d" % p, end=' ')
```

```
2 3 5 7 11
```

It is better style to use a conditional main function, i.e.:

```
def primes5():
    return (2, 3, 5, 7, 11)

if __name__ == "__main__":
    for p in primes5():
        print("%d" % p, end=' ')
```

```
2 3 5 7 11
```

A purist might argue that the following is even cleaner:

```
def primes5():
    return (2, 3, 5, 7, 11)

def main():
    for p in primes5():
        print("%d" % p, end=' ')

if __name__ == "__main__":
    main()
```

```
2 3 5 7 11
```

but either of the last two options is good.

The example in *Many ways to compute a series* demonstrates this technique. Including functions with names starting with `test_` is compatible with the very useful `py.test` regression testing framework (see <https://docs.pytest.org/en/stable/>).

7.6 Further Reading

- Python Tutorial Section 6

FUNCTIONAL TOOLS

Python provides a few in-built commands such as `map`, `filter`, `reduce` as well `lambda` (to create anonymous functions) and list comprehension. These are typical commands from functional languages of which LISP is probably best known.

Functional programming can be extremely powerful and one of the strengths of Python is that it allows to program using (i) imperative/procedural programming style, (ii) object oriented style and (iii) functional style. It is the programmers choice which tools to select from which style and how to mix them to best address a given problem.

In this chapter, we provide some examples for usage of the commands listed above.

8.1 Anonymous functions

All functions we have seen in Python so far have been defined through the `def` keyword, for example:

```
def f(x):  
    return x ** 2
```

This function has the name `f`. Once the function is defined (i.e. the Python interpreter has come across the `def` line), we can call the function using its name, for example

```
y = f(6)
```

Sometimes, we need to define a function that is only used once, or we want to create a function but don't need a name for it (as for creating closures). In this case, this is called *anonymous* function as it does not have a name. In Python, the `lambda` keyword can create an anonymous function.

We create a (named) function first, check it's type and behaviour:

```
def f(x):  
    return x ** 2
```

```
f
```

```
<function __main__.f(x)>
```

```
type(f)
```

```
function
```

```
f(10)
```

```
100
```

Now we do the same with an anonymous function:

```
lambda x: x ** 2
```

```
<function __main__.<lambda>(x)>
```

```
type(lambda x: x ** 2)
```

```
function
```

```
(lambda x: x ** 2)(10)
```

```
100
```

This works exactly in the same way but – as the anonymous function does not have a name – we need to define the function (through the `lambda` expression) – every time we need it.

Anonymous functions can take more than one argument:

```
(lambda x, y: x + y)(10, 20)
```

```
30
```

```
(lambda x, y, z: (x + y) * z)(10, 20, 2)
```

```
60
```

We will see some examples using `lambda` which will clarify typical use cases.

8.2 Map

The `map` function `lst2 = map(f, s)` applies a function `f` to all elements in a sequence `s`. The result of `map` can be turned into a list with the same length as `s`:

```
def f(x):  
    return x ** 2  
lst2 = list(map(f, range(10)))  
lst2
```

```
[0, 1, 4, 9, 16, 25, 36, 49, 64, 81]
```

```
list(map(str.capitalize, ['banana', 'apple', 'orange']))
```

```
['Banana', 'Apple', 'Orange']
```

Often, this is combined with the anonymous function `lambda`:

```
list(map(lambda x: x ** 2, range(10) ))
```

```
[0, 1, 4, 9, 16, 25, 36, 49, 64, 81]
```

```
list(map(lambda s: s.capitalize(), ['banana', 'apple', 'orange']))
```

```
['Banana', 'Apple', 'Orange']
```

8.3 Filter

The filter function `lst2 = filter(f, lst)` applies the function `f` to all elements in a sequence `s`. The function `f` should return `True` or `False`. This makes a list which will contain only those elements `si` of the sequence `s` for which `f(si)` has returned `True`.

```
def greater_than_5(x):
    if x > 5:
        return True
    else:
        return False
```

```
list(filter(greater_than_5, range(11)))
```

```
[6, 7, 8, 9, 10]
```

The usage of `lambda` can simplify this significantly:

```
list(filter(lambda x: x > 5, range(11)))
```

```
[6, 7, 8, 9, 10]
```

```
known_names = ['smith', 'miller', 'bob']
list(filter(lambda name : name in known_names, \
            ['ago', 'smith', 'bob', 'carl']))
```

```
['smith', 'bob']
```

8.4 List comprehension

List comprehensions provide a concise way to create and modify lists without resorting to use of `map()`, `filter()` and/or `lambda`. The resulting list definition tends often to be clearer than lists built using those constructs. Each list comprehension consists of an expression followed by a `for` clause, then zero or more `for` or `if` clauses. The result will be a list resulting from evaluating the expression in the context of the `for` and `if` clauses which follow it. If the expression would evaluate to a tuple, it must be parenthesized.

Some examples will make this clearer:

```
freshfruit = [' banana', ' loganberry ', 'passion fruit ']
[weapon.strip() for weapon in freshfruit]
```

```
['banana', 'loganberry', 'passion fruit']
```

```
vec = [2, 4, 6]
[3 * x for x in vec]
```

```
[6, 12, 18]
```

```
[3 * x for x in vec if x > 3]
```

```
[12, 18]
```

```
[3 * x for x in vec if x < 2]
```

```
[]
```

```
[[x, x ** 2] for x in vec]
```

```
[[2, 4], [4, 16], [6, 36]]
```

We can also use list comprehension to modify the list of integers returned by the `range` command so that our subsequent elements in the list increase by non-integer fractions:

```
[x*0.5 for x in range(10)]
```

```
[0.0, 0.5, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0, 4.5]
```

Let's now revisit the examples from the section on `filter`

```
[x for x in range(11) if x>5 ]
```

```
[6, 7, 8, 9, 10]
```

```
[name for name in ['ago','smith','bob','carl'] \
 if name in known_names]
```

```
['smith', 'bob']
```

and the examples from the map section

```
[x ** 2 for x in range(10) ]
```

```
[0, 1, 4, 9, 16, 25, 36, 49, 64, 81]
```

```
[fruit.capitalize() for fruit in ['banana', 'apple', 'orange'] ]
```

```
['Banana', 'Apple', 'Orange']
```

all of which can be expressed through list comprehensions.

More details

- [Python Tutorial 5.1.4 List comprehensions](#)

8.5 Reduce

The `reduce` function takes a binary function $f(x, y)$, a sequence s , and a start value a_0 . It then applies the function f to the start value a_0 and the first element in the sequence: $a_1 = f(a_0, s[0])$. The second element ($s[1]$) of the sequence is then processed as follows: the function f is called with arguments a_1 and $s[1]$, i.e. $a_2 = f(a_1, s[1])$. In this fashion, the whole sequence is processed. `Reduce` returns a single element.

This can be used, for example, to compute a sum of numbers in a sequence if the function $f(x, y)$ returns $x+y$:

```
from functools import reduce
```

```
def add(x, y):
    return x + y
```

```
reduce(add, [1, 2, 3, 4, 5, 6, 7, 8, 9, 10], 0)
```

```
55
```

```
reduce(add, [1, 2, 3, 4, 5, 6, 7, 8, 9, 10], 100)
```

```
155
```

We can modify the function `add` to provide some more detail about the process:

```
def add_verbose(x, y):
    print("add(x=%s, y=%s) -> %s" % (x, y, x+y))
    return x+y
```

```
reduce(add_verbose, [1, 2, 3, 4, 5, 6, 7, 8, 9, 10], 0)
```

```
add(x=0, y=1) -> 1
add(x=1, y=2) -> 3
add(x=3, y=3) -> 6
add(x=6, y=4) -> 10
add(x=10, y=5) -> 15
add(x=15, y=6) -> 21
add(x=21, y=7) -> 28
add(x=28, y=8) -> 36
add(x=36, y=9) -> 45
add(x=45, y=10) -> 55
```

55

It may be instructive to use an asymmetric function f , such as `add_len(n, s)` where s is a sequence and the function returns $n+\text{len}(s)$ (suggestion from Thomas Fischbacher):

```
def add_len(n, s):
    return n + len(s)

reduce(add_len, ["This", "is", "a", "test."], 0)
```

12

As before, we'll use a more verbose version of the binary function to see what is happening:

```
def add_len_verbose(n, s):
    print("add_len(n=%d, s=%s) -> %d" % (n, s, n+len(s)))
    return n+len(s)

reduce(add_len_verbose, ["This", "is", "a", "test."], 0)
```

```
add_len(n=0, s=This) -> 4
add_len(n=4, s=is) -> 6
add_len(n=6, s=a) -> 7
add_len(n=7, s=test.) -> 12
```

12

Another way to understand what the `reduce` function does is to look at the following function (kindly provided by Thomas Fischbacher) which behaves like `reduce` but explains what it does:

Here is an example using the `explain_reduce` function:

```
def explain_reduce(f, xs, start=None):
    """This function behaves like reduce, but explains what it does,
    step-by-step.
    (Author: Thomas Fischbacher, modifications Hans Fangohr)"""
    nr_xs = len(xs)
    if start == None:
        if nr_xs == 0:
            raise ValueError("No starting value given - cannot " + \
                             "process empty list!")
        if nr_xs == 1:
```

(continues on next page)

(continued from previous page)

```

    print("reducing over 1-element list without starting " + \
          "value: returning that element.")
    return xs[0]
else:
    print("reducing over list with >= 2 elements without " + \
          "starting value: using the first element as a " + \
          "start value.")
    return explain_reduce(f, xs[1:], xs[0])
else:
    s = start
    for n in range(len(xs)):
        x = xs[n]
        print("Step %d: value-so-far=%s next-list-element=%s" \
              % (n, str(s), str(x)))
        s = f(s, x)
    print("Done. Final result=%s" % str(s))
    return s

```

```

def f(a, b):
    return a + b

```

```
reduce(f, [1, 2, 3, 4, 5], 0)
```

15

```
explain_reduce(f, [1, 2, 3, 4, 5], 0)
```

```

Step 0: value-so-far=0 next-list-element=1
Step 1: value-so-far=1 next-list-element=2
Step 2: value-so-far=3 next-list-element=3
Step 3: value-so-far=6 next-list-element=4
Step 4: value-so-far=10 next-list-element=5
Done. Final result=15

```

15

Reduce is often combined with lambda:

```
reduce(lambda x, y: x + y, [1, 2, 3, 4, 5], 0)
```

15

There is also the `operator` module which provides standard Python operators as functions. For example, the function `operator.__add__(a,b)` is executed when Python evaluates code such as `a+b`. These are generally faster than lambda expressions. We could write the example above as

```

import operator
reduce(operator.__add__, [1, 2, 3, 4, 5], 0)

```

15

Use `help('operator')` to see the complete list of operator functions.

8.6 Why not just use for-loops?

Let's compare the example introduced at the beginning of the chapter written (i) using a for-loop and (ii) list comprehension. Again, we want to compute the numbers $0^2, 1^2, 2^2, 3^2, \dots$ up to $(n-1)^2$ for a given n .

Implementation (i) using a for-loop with $n=10$:

```
y = []
for i in range(10):
    y.append(i**2)
```

Implementation (ii) using list comprehension:

```
y = [x**2 for x in range(10)]
```

or using `map`:

```
y = map(lambda x: x**2, range(10))
```

The versions using list comprehension and `map` fit into one line of code whereas the for-loop needs 3. This example shows that functional code result in very *concise* expressions. Typically, the number of mistakes a programmer makes is per line of code written, so the fewer lines of code we have, the fewer bugs we need to find.

Often programmers find that initially the list-processing tools introduced in this chapter seem less intuitive than using for-loops to process every element in a list individually, but that – over time – they come to value a more functional programming style.

8.7 Speed

The functional tools described in this chapter can also be faster than using explicit (for or while) loops over list elements.

The program `list_comprehension_speed.py` below computes $\text{sum}_i = 0^{N-1}i^2$ for a large value of N using 4 different methods and records execution time:

- Method 1: for-loop (with pre-allocated list, storing of i^2 in list, then using in-built `sum` function)
- Method 2: for-loop without list (updating sum as the for-loop progresses)
- Method 3: using list comprehension
- Method 4: using `numpy`. (`Numpy` is covered in *chapter 14*)

Here is a possible program computing this:

```
# NBVAL_IGNORE_OUTPUT
"""Compare calculation of \sum_i x_i^2 with
i going from zero to N-1.

We use (i) for loops and list, (ii) for-loop, (iii) list comprehension
and (iv) numpy.

We use floating numbers to avoid using Python's long int (which would
```

(continues on next page)

(continued from previous page)

```

be likely to make the timings less representative).
"""

import time
import numpy
N = 10000000

def timeit(f, args):
    """Given a function f and a tuple args containing
    the arguments for f, this function calls f(*args),
    and measures and returns the execution time in
    seconds.

    Return value is tuple: entry 0 is the time,
    entry 1 is the return value of f."""

    starttime = time.time()
    y = f(*args)    # use tuple args as input arguments
    endtime = time.time()
    return endtime - starttime, y

def forloop1(N):
    s = 0
    for i in range(N):
        s += float(i) * float(i)
    return s

def forloop2(N):
    y = [0] * N
    for i in range(N):
        y[i] = float(i) ** 2
    return sum(y)

def listcomp(N):
    return sum([float(x) * x for x in range(N)])

def numpy_(N):
    return numpy.sum(numpy.arange(0, N, dtype='d') ** 2)

# main program starts
timings = []
print("N =", N)
forloop1_time, f1_res = timeit(forloop1, (N,))
timings.append(forloop1_time)
print("for-loop1: {:.3f}s".format(forloop1_time))
forloop2_time, f2_res = timeit(forloop2, (N,))
timings.append(forloop2_time)
print("for-loop2: {:.3f}s".format(forloop2_time))
listcomp_time, lc_res = timeit(listcomp, (N,))
timings.append(listcomp_time)
print("listcomp : {:.3f}s".format(listcomp_time))

```

(continues on next page)

(continued from previous page)

```

numpy_time, n_res = timeit(numpy_, (N,))
timings.append(numpy_time)
print("numpy      : {:.3f}s".format(numpy_time))

# ensure that different methods provide identical results
assert f1_res == f2_res
assert f1_res == lc_res

# Allow a bit of difference for the numpy calculation
numpy.testing.assert_approx_equal(f1_res, n_res)

print("Slowest method is {:.1f} times slower than the fastest method."
      .format(max(timings)/min(timings)))

```

```
N = 10000000
```

```
for-loop1: 0.647s
```

```
for-loop2: 0.978s
```

```
listcomp : 0.703s
numpy     : 0.020s
Slowest method is 48.7 times slower than the fastest method.
```

The actual execution performance depends on the computer. The relative performance may depend on versions of Python and its support libraries (such as numpy) we use.

With the current version (python 3.6, numpy 1.11, on a x84 machine running OS X), we see that methods 1 and 2 (for-loop without list and with pre-allocated list) are slowest, somewhat closely followed by the slightly faster list comprehension. The fastest method is number 4 (using numpy).

8.8 The %%timeit magic

If we are using IPython as our shell (or a cell in a Jupyter notebook running a python kernel), there is a much more sophisticated way to measure timings that what is done above: if a cell starts with %%timeit, then IPython will run the commands in that cell repeatedly and obtain (averaged) timings. This is particularly useful for timing of commands that execute relatively quickly.

Let's compare a list comprehension with an explicit loop using the timeit magic:

```
%%timeit
y = [x**2 for x in range(100)]
```

```
3.72 µs ± 56.3 ns per loop (mean ± std. dev. of 7 runs, 100,000 loops each)
```

```
%%timeit
y = []
for x in range(100):
    y.append(x**2)
```

3.82 μ s \pm 9.57 ns per loop (mean \pm std. dev. of 7 runs, 100,000 loops each)

COMMON TASKS

Here we provide a selection of small example programs addressing some common tasks and just providing some more Python code that can be read if seeking inspiration how to address a given problem.

9.1 Many ways to compute a series

As an example, we compute the sum of odd numbers in different ways.

```
def compute_sum1(n):
    """computes and returns the sum of 2,4,6, ..., m
    where m is the largest even number smaller than n.

    For example, with n = 7, we compute 0+2+4+6 = 12.

    This implementation uses a variable 'mysum' that is
    increased in every iteration of the for-loop."""

    mysum = 0
    for i in range(0, n, 2):
        mysum = mysum + i
    return mysum

def compute_sum2(n):
    """computes and returns ...

    This implementation uses a while-loop:
    """

    counter = 0
    mysum = 0
    while counter < n:
        mysum = mysum + counter
        counter = counter + 2

    return mysum

def compute_sum3(n, startfrom=0):
    """computes and returns ...

    This is a recursive implementation:"""
```

(continues on next page)

(continued from previous page)

```
    if n <= startfrom:
        return 0
    else:
        return startfrom + compute_sum3(n, startfrom + 2)

def compute_sum4a(n):
    """A functional approach ... this seems to be
    the shortest and most concise code.
    """
    return sum(range(0, n, 2))

from functools import reduce
def compute_sum4b(n):
    """A functional approach ... not making use of 'sum' which
    happens to exist and is of course convenient here.
    """
    return reduce(lambda a, b: a + b, range(0, n, 2))

def compute_sum4c(n):
    """A functional approach ... a bit faster than compute_sum4b
    as we avoid using lambda.
    """
    import operator
    return reduce(operator.__add__, range(0, n, 2))

def compute_sum4d(n):
    """Using list comprehension."""
    return sum([k for k in range(0, n, 2)])

def compute_sum4e(n):
    """Using another variation of list comprehension."""
    return sum([k for k in range(0, n) if k % 2 == 0])

def compute_sum5(n):
    """Using numerical python (numpy). This is very fast
    (but would only pay off if n >> 10)."""
    import numpy
    return numpy.sum(2 * numpy.arange(0, (n + 1) // 2))

def test_consistency():
    """Check that all compute_sum?? functions in this file produce
    the same answer for all n>=2 and <N.
    """
    def check_one_n(n):
        """Compare the output of compute_sum1 with all other functions
        for a given n>=2. Raise AssertionError if outputs disagree."""
        funcs = [compute_sum1, compute_sum2, compute_sum3,
                 compute_sum4a, compute_sum4b, compute_sum4c,
```

(continues on next page)

(continued from previous page)

```

        compute_sum4d, compute_sum4e, compute_sum5]
    ans1 = compute_sum1(n)
    for f in funcs[1:]:
        assert ans1 == f(n), "%s(n)=%d not the same as %s(n)=%d " \
            % (funcs[0], funcs[0](n), f, f(n))

    #main testing loop in test_consistency function
    for n in range(2, 1000):
        check_one_n(n)

if __name__ == "__main__":
    m = 7
    correct_result = 12
    thisresult = compute_sum1(m)
    print("this result is {}, expected to be {}".format(
        thisresult, correct_result))
    # compare with correct result
    assert thisresult == correct_result
    # also check all other methods
    assert compute_sum2(m) == correct_result
    assert compute_sum3(m) == correct_result
    assert compute_sum4a(m) == correct_result
    assert compute_sum4b(m) == correct_result
    assert compute_sum4c(m) == correct_result
    assert compute_sum4d(m) == correct_result
    assert compute_sum4e(m) == correct_result
    assert compute_sum5(m) == correct_result

    # a more systematic check for many values
    test_consistency()

```

```
this result is 12, expected to be 12
```

All the different implementations shown above compute the same result. There are a number of things to be learned from this:

- There are a large (probably an infinite) number of solutions for one given problem. (This means that writing programs is a task that requires creativity!)
- These may achieve the same 'result' (in this case computation of a number).
- Different solutions may have different characteristics. They might:
 - be faster or slower
 - use less or more memory
 - are easier or more difficult to understand (when reading the source code)
 - can be considered more or less elegant.

9.2 Sorting

Suppose we need to sort a list of 2-tuples of user-ids and names, i.e.

```
mylist = [("fangohr", "Hans Fangohr"),
          ("admin", "The Administrator"),
          ("guest", "The Guest")]
```

which we want to sort in increasing order of user-ids. If there are two or more identical user-ids, they should be ordered by the order of the names associated with these user-ids. This behaviour is just the default behaviour of `sort` (which goes back to how to sequences are compared).

```
stuff = mylist # collect your data
stuff.sort()  # sort the data in place
print(stuff)  # inspect the sorted data
```

```
[('admin', 'The Administrator'), ('fangohr', 'Hans Fangohr'), ('guest', 'The Guest
↩️')]
```

Sequences are compared by initially comparing the first elements only. If they differ, then a decision is reached on the basis of those elements only. If the elements are equal, only then are the next elements in the sequence compared ... and so on, until a difference is found, or we run out of elements. For example:

```
(2,0) > (1,0)
```

```
True
```

```
(2,1) > (1,3)
```

```
True
```

```
(2,1) > (2,1)
```

```
False
```

```
(2,2) > (2,1)
```

```
True
```

It is also possible to do:

```
stuff = sorted(stuff)
```

Where the list is not particularly large, it is generally advisable to use the `sorted` function (which *returns a sorted copy* of the list) over the `sort` method of a list (which changes the list into sorted order of elements, and returns `None`).

However, what if the data we have is stored such that in each tuple in the list, the name comes first, followed by the id, i.e.:


```
mylist2 = [("Hans Fangohr", "fangohr"),  
          ("The Administrator", "admin"),  
          ("The Guest", "guest")]
```

We want to sort with the id as the primary key. The first approach to do this is to change the order of `mylist2` to that of `mylist`, and use `sort` as shown above.

The second, neater approach relies on being able to decypher the cryptic help for the sorted function. `list.sort()` has the same options, but its help is less helpful.

```
# NBVAL_IGNORE_OUTPUT  
help(sorted)
```

```
Help on built-in function sorted in module builtins:
```

```
sorted(iterable, /, *, key=None, reverse=False)  
Return a new list containing all items from the iterable in ascending order.
```

```
A custom key function can be supplied to customize the sort order, and the  
reverse flag can be set to request the result in descending order.
```

You should notice that `sorted` and `list.sort` have two keyword parameters. The first of these is called `key`. You can use this to supply a *key function*, which will be used to transform the items for sort to compare.

Let's illustrate this in the context of our exercise, by assuming that we have stored a list of pairs like this

```
pair = name, id
```

(i.e. as in `mylist2`) and that we want to sort according to id and ignore name. We can achieve this by writing a function that retrieves only the second element of the pair it receives:

```
def my_key(pair):  
    return pair[1]
```

```
mylist2.sort(key=my_key)
```

This also works with an anonymous function:

```
mylist2.sort(key=lambda p: p[1])
```

9.2.1 Efficiency

The key function will be called exactly once for every element in the list. This is much more efficient than calling a function on every *comparison* (which was how you customised sorting in older versions of Python). But if you have a large list to sort, the overhead of calling a Python function (which is relatively large compared to the C function overhead) might be noticeable.

If efficiency is really important (and you have proven that a significant proportion of time is spent in these functions) then you have the option of re-coding them in C (or another low-level language).

FROM MATLAB TO PYTHON

10.1 Important commands

10.1.1 The for-loop

Matlab:

```
for i = 1:10
    disp(i)
end
```

Matlab requires the `end` key-word at the end of the block belonging to the for-loop.

Python:

```
for i in range(1,11):
    print(i)
```

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

Python requires a colon (“:”) at the of the `for`-line. (This is important and often forgotten when you have programmed in Matlab before.) Python requires the commands to be executed within the for-loop to be indented.

10.1.2 The if-then statement

Matlab:

```
if a==0
    disp('a is zero')
elseif a<0
    disp('a is negative')
elseif a==42
    disp('a is 42')
else
    disp('a is positive')
end
```

Matlab requires the `end` key-word at the very end of the block belonging to the for-loop.

Python:

```
a = -5

if a==0:
    print('a is zero')
elif a<0:
    print('a is negative')
elif a==42:
    print('a is 42')
else:
    print('a is positive')
```

```
a is negative
```

Python requires a colon (“:”) after every condition (i.e. at the of the lines starting with `if`, `elif`, `else`. Python requires the commands to be executed within each part of the if-then-else statement to be indented.

10.1.3 Indexing

Matlab’s indexing of matrices and vectors starts a 1 (similar to Fortran), whereas Python’s indexing starts at 0 (similar to C).

10.1.4 Matrices

In Matlab, every object is a matrix. In Python, there is a specialised extension library called `numpy` (see Sec. [cha:numer-pyth-numpy]) which provides the `array` object which in turns provides the corresponding functionality. Similar to Matlab, the `numpy` object is actually based on binary libraries and execution there very fast.

There is a dedicated introduction to `numpy` for Matlab users available at <https://numpy.org/doc/stable/user/numpy-for-matlab-users.html>.

PYTHON SHELLS

11.1 IDLE

IDLE comes with every Python distribution and is a useful tool for everyday programming. Its editor provides syntax highlighting.

There are two reasons why you might want to use another Python shell, for example:

- While working with the Python prompt, you like auto-completion of variable names, filenames and commands. In that case *IPython* is your tool of choice (see below). IPython does not provide an editor, but you can carry on using the IDLE editor to edit files, or any other editor you like.

IPython provides a number of nice features for the more experienced Python programmer, including convenient profiling of code (see <https://ipython.org/>).

Recently, some auto-completion has been added to Idle as well (press tab after having typed the first few letters of object names and keywords).

11.2 Python (command line)

This is the most basic face of the Python shell. It is very similar to the Python prompt in IDLE but there are no menus to click on and no facilities to edit files.

11.3 Interactive Python (IPython)

11.3.1 IPython console

IPython is an improved version of the Python command line. It is a valuable tool and worth exploring its capabilities (see <https://ipython.org/ipython-doc/stable/interactive/qtconsole.html>)

You will find the following features very useful:

- auto completion Suppose you want to type `a = range(10)`. Instead of typing all the letters, just type `a = r` and then press the “Tab” key. IPython will now show all the possible commands (and variable names) that start with `ra`. If you type a third letter, here `n` and press “Tab” again, IPython will auto complete and append `ge` automatically. This works also for variable names and modules.
- To obtain help on a command, we can use Python’s help command. For example: `help(range)`. IPython provides a shortcut. To achieve the same, it is sufficient to just type the command followed by a question mark: `range?`

- You can relatively easily navigate directories on your computer. For example,
 - `!dir` lists the content of the current directory (same as `ls`)
 - `pwd` shows the current working directory
 - `cd` allows to change directories
- In general, using an exclamation mark before the command will pass the command to the shell (not to the Python interpreter).
- You can execute Python programs from `ipython` using `%run`. Suppose you have a file `hello.py` in the current directory. You can then execute it by typing:
`%run hello`

Note that this differs from executing a python program in IDLE: IDLE restarts the Python interpreter session and thus deletes all existing objects before the execution starts. This is not the case with the `run` command in `ipython` (and neither when executing chunks of Python code from Emacs using the Emacs Python mode). In particular this can be very useful if one needs to setup a few objects which are needed to test the code one is working on. Using `ipython`'s `run` or Emacs instead of IDLE allows to keep these objects in the interpreter session and to only update the function/classes/... etc that are being developed.

- allows multi-line editing of command history
- provides on-the-fly syntax highlighting
- displays doc-strings on-the-fly
- can inline matplotlib figures (activate mode with if started with `%matplotlib inline`)
- `%load` loads file from disk or from URL for editing
- `%timeit` measures execution time for a given statement
- ...and a lot more.
- Read more at <https://ipython.org/ipython-doc/dev/interactive/qtconsole.html>

If you have access to this shell, you may want to consider it as your default Python prompt.

11.3.2 Jupyter Notebook

The Jupyter Notebook (formerly IPython Notebook) allows you to execute, store, load, re-execute a sequence of Python commands, and to include explanatory text, images and other media in between.

This is a recent and exciting development that has the potential to develop into a tool of great significance, for example for

- documenting calculations and data processing
- support learning and teaching of
 - Python itself
 - statistical methods
 - general data post-processing
 - ...
- documentation new code
- automatic regression testing by re-running `ipython` notebook and comparing stored output with computed output

Further reading

- Jupyter Notebook (<https://jupyter-notebook.readthedocs.io/en/latest/>).
- IPython (<https://ipython.org>).

11.4 Spyder

Spyder is the Scientific PYthon Development EnviRonment: a powerful interactive development environment for the Python language with advanced editing, interactive testing, debugging and introspection features and a numerical computing environment thanks to the support of IPython (enhanced interactive Python interpreter) and popular Python libraries such as NumPy (linear algebra), SciPy (signal and image processing) or matplotlib (interactive 2D/3D plotting). See <https://www.spyder-ide.org/> for more.

Some important features of Spyder:

- Within Spyder, the IPython console is the default Python interpreter, and
- code in the editor can be fully or partially be executed in this buffer.
- The editor supports automatic checking for Python errors using pyflakes, and
- the editor warns (if desired) if the code formatting deviates from the PEP8 style guide.
- The Ipython Debugger can be activated, and
- a profiler is provided.
- An object explorer shows documentation for functions, methods etc on the fly and a
- variable explorer displays names, size and values for numerical variables.

Spyder is currently (as of 2014) on the way to develop into a powerful and robust multi-platform integrated environment for Python development, with particular emphasis on Python for scientific computing and engineering.

11.5 Editors

All major editors that are used for programming, provide Python modes (such as Emacs, Vim, Sublime Text), some Integrated Development Enviroments (IDEs) come with their own editor (Spyder, Eclipse). Which of these is best, is partly a matter of choice.

For beginners, Spyder seems a sensible choice as it provides an IDE, allows execution of chunks of code in an interpreter session and is easy to pick up.

SYMBOLIC COMPUTATION

12.1 SymPy

In this section, we introduce some basic functionality of the SymPy (SYMbolic Python) library. In contrast to numerical computation (involving numbers), in symbolic calculation we are processing and transforming generic variables.

The SymPy home page is <https://www.sympy.org/>, and provides the full (and up-to-date) documentation for this library.

Symbolic calculation is very slow compared to floating point operation (see for example [symbolic calculation for decimals](13 Numeric Computation, Symbolic Calculation)), and thus generally not for direct simulation. However, it is a powerful tool to support the preparation of code and symbolic work. Occasionally, we use symbolic operations in simulations to work out the most efficient numerical code, before that is executed.

12.1.1 Output

Before we start using sympy, we'll call `init_printing`. This tells sympy to display expressions in a nicer format.

```
import sympy
sympy.init_printing(use_latex='mathjax')
```

12.1.2 Symbols

Before we can carry out any symbolic operations, we need to create symbolic variables using SymPy's `Symbol` function:

```
from sympy import Symbol
x = Symbol('x')
type(x)
```

```
sympy.core.symbol.Symbol
```

```
y = Symbol('y')
2 * x - x
```

x

```
x + y + x + 10*y
```

$$2x + 11y$$

```
y + x - y + 10
```

$$x + 10$$

We can abbreviate the creation of multiple symbolic variables using the `symbols` function. For example, to create the symbolic variables `x`, `y` and `z`, we can use

```
import sympy
x, y, z = sympy.symbols('x,y,z')
x + 2*y + 3*z - x
```

$$2y + 3z$$

Once we have completed our term manipulation, we sometimes like to insert numbers for variables. This can be done using the `subs` method.

```
from sympy import symbols
x, y = symbols('x,y')
x + 2*y
```

$$x + 2y$$

```
(x + 2*y).subs(x, 10)
```

$$x + 2y$$

```
(x + 2*y).subs(x, 10)
```

$$2y + 10$$

```
(x + 2*y).subs(x, 10).subs(y, 3)
```

```
(x + 2*y).subs({x:10, y:3})
```

16

We can also substitute a symbolic variable for another one such as in this example where y is replaced with x before we substitute x with the number 2.

```
myterm = 3*x + y**2
myterm
```

$$3x + y^2$$

```
myterm.subs(x, y)
```

$$y^2 + 3y$$

```
myterm.subs(x, y).subs(y, 2)
```

10

From this point onward, some of the code fragments and examples we present will assume that the required symbols have already been defined. If you try an example and SymPy gives a message like `NameError: name 'x' is not defined` it is probably because you need to define the symbol using one of the methods above.

12.1.3 isympy

The `isympy` executable is a wrapper around `ipython` which creates the symbolic (real) variables x , y and z , the symbolic integer variables k , m and n and the symbolic function variables f , g and h , and imports all objects from the SymPy `toplevel`.

This is convenient to figure out new features or experimenting interactively

```
$> isympy
Python 2.6.5 console for SymPy 0.6.7

These commands were executed:
>>> from __future__ import division
>>> from sympy import *
>>> x, y, z = symbols('xyz')
>>> k, m, n = symbols('kmn', integer=True)
>>> f, g, h = map(Function, 'fgh')

Documentation can be found at https://www.sympy.org/
```

```
In [1]:
```

12.1.4 Numeric types

SymPy has the numeric types `Rational` and `RealNumber`. The `Rational` class represents a rational number as a pair of two integers: the numerator and the denominator, so `Rational(1, 2)` represents $1/2$, `Rational(5, 2)` represents $5/2$ and so on.

```
from sympy import Rational
```

```
a = Rational(1, 10)
a
```

$$\frac{1}{10}$$

```
b = Rational(45, 67)
b
```

$$\frac{45}{67}$$

```
a * b
```

$$\frac{9}{134}$$

```
a - b
```

$$\frac{383}{670}$$

```
a + b
```

$$\frac{517}{670}$$

Note that the `Rational` class works with rational expressions *exactly*. This is in contrast to Python's standard `float` data type which uses floating point representation to *approximate* (rational) numbers.

We can convert the `sympy.Rational` type into a Python floating point variable using `float` or the `evalf` method of the `Rational` object. The `evalf` method can take an argument that specifies how many digits should be computed for the floating point approximation (not all of those may be used by Python's floating point type of course).

```
c = Rational(2, 3)
c
```

$$\frac{2}{3}$$


```
diff(10 + 3*x + 4*y + 10*x**2 + x**9, x).subs(x,1)
```

32

```
diff(10 + 3*x + 4*y + 10*x**2 + x**9, x).subs(x,1.5)
```

263.66015625

```
diff(exp(x), x)
```

e^x

```
diff(exp(-x ** 2 / 2), x)
```

$-xe^{-\frac{x^2}{2}}$

The SymPy `diff()` function takes a minimum of two arguments: the function to be differentiated and the variable with respect to which the differentiation is performed. Higher derivatives may be calculated by specifying additional variables, or by adding an optional integer argument:

```
diff(3*x**4, x)
```

$12x^3$

```
diff(3*x**4, x, x, x)
```

$72x$

```
diff(3*x**4, x, 3)
```

$72x$

```
diff(3*x**4*y**7, x, 2, y, 2)
```

$1512x^2y^5$

```
diff(diff(3*x**4*y**7, x, x), y, y)
```

$$1512x^2y^5$$

At times, SymPy may return a result in an unfamiliar form. If, for example, you wish to use SymPy to check that you differentiated something correctly, a technique that might be of use is to subtract the SymPy result from your result, and check that the answer is zero.

Taking the simple example of a multiquadric radial basis function, $\phi(r) = \frac{1}{\sqrt{r^2 + \sigma^2}}$ with $r = \sqrt{x^2 + y^2}$ and σ a constant, we can verify that the first derivative in x is $\partial\phi/\partial x = x/\sqrt{r^2 + \sigma^2}$.

In this example, we first ask SymPy to print the derivative. See that it is printed in a different form to our trial derivative, but the subtraction verifies that they are identical:

```
r = sqrt(x**2 + y**2)
sigma = Symbol('σ')
def phi(x, y, sigma):
    return sqrt(x**2 + y**2 + sigma**2)

mydfdx = x / sqrt(r**2 + sigma**2)
print(diff(phi(x, y, sigma), x))
```

```
x/sqrt(x**2 + y**2 + σ**2)
```

```
print(mydfdx - diff(phi(x, y, sigma), x))
```

```
0
```

Here it is trivial to tell that the expressions are identical without SymPy's help, but in more complicated examples there may be many more terms and it would become increasingly difficult, time consuming and error-prone to attempt to rearrange our trial derivative and SymPy's answer into the same form. It is in such cases that this subtraction technique is of most use.

Integration uses a similar syntax. For the indefinite case, specify the function and the variable with respect to which the integration is performed:

```
from sympy import integrate
integrate(x**2, x)
```

$$\frac{x^3}{3}$$

```
integrate(x**2, y)
```

$$x^2y$$

```
integrate(sin(x), y)
```

$$y \sin(x)$$

```
integrate(sin(x), x)
```

$$-\cos(x)$$

```
integrate(-x*exp(-x**2/2), x)
```

$$e^{-\frac{x^2}{2}}$$

We can calculate definite integrals by providing `integrate()` with a tuple containing the variable of interest, the lower and the upper bounds. If several variables are specified, multiple integration is performed. When SymPy returns a result in the `Rational` class, it is possible to evaluate it to a floating-point representation at any desired precision (see *numeric types*).

```
integrate(x**2, (x, 0, 1))
```

$$1$$

```
integrate(x**2, x)
```

$$\frac{x^3}{3}$$

```
integrate(x**2, x, x)
```

$$\frac{x^4}{12}$$

```
integrate(x**2, x, x, y)
```

$$\frac{x^4 y}{12}$$

```
integrate(x**2, (x, 0, 2))
```

$$\frac{8}{3}$$


```
dsolve(y_ + 5*y(x), y(x))
```

$$y(x) = C_1 e^{-5x}$$

```
dsolve(Eq(y_ + 5*y(x), 0), y(x))
```

$$y(x) = C_1 e^{-5x}$$

```
dsolve(Eq(y_ + 5*y(x), 12), y(x))
```

$$y(x) = C_1 e^{-5x} + \frac{12}{5}$$

The results from `dsolve` are an instance of the `Equality` class. This has consequences when we wish to numerically evaluate the function and use the result elsewhere (e.g. if we wanted to plot $y(x)$ against x), because even after using `subs()` and `evalf()`, we still have an `Equality`, not any sort of scalar. The way to evaluate the function to a number is via the `rhs` attribute of the `Equality`.

Note that, here, we use `z` to store the `Equality` returned by `dsolve`, even though it is an expression for a function called $y(x)$, to emphasise the distinction between the `Equality` itself and the data that it contains.

```
z = dsolve(y_ + 5*y(x), y(x))
z
```

$$y(x) = C_1 e^{-5x}$$

```
type(z)
```

```
sympy.core.relational.Equality
```

```
z.rhs
```

$$C_1 e^{-5x}$$

```
C1=Symbol('C1')
y3 = z.subs({C1:2, x:3})
y3
```

$$y(3) = \frac{2}{e^{15}}$$

```
y3.evalf(10)
```

$$y(3) = 6.11804641 \cdot 10^{-7}$$

```
y3.rhs
```

$$\frac{2}{e^{15}}$$

```
y3.rhs.evalf(10)
```

$$6.11804641 \cdot 10^{-7}$$

```
z.rhs.subs({C1:2, x:4}).evalf(10)
```

$$4.122307245 \cdot 10^{-9}$$

```
z.rhs.subs({C1:2, x:5}).evalf(10)
```

$$2.777588773 \cdot 10^{-11}$$

```
type(z.rhs.subs({C1:2, x:5}).evalf(10))
```

```
sympy.core.numbers.Float
```

At times, `dsolve` may return too general a solution. One example is when there is a possibility that some coefficients may be complex. If we know that, for example, they are always real and positive, we can provide `dsolve` this information to avoid the solution becoming unnecessarily complicated:

```
from sympy import *
a, x = symbols('a,x')
f = Function('f')
dsolve(Derivative(f(x), x, 2) + a**4*f(x), f(x))
```

$$f(x) = C_1 e^{-ia^2 x} + C_2 e^{ia^2 x}$$

```
a = Symbol('a', real=True, positive=True)
dsolve(Derivative(f(x), x, 2) + a**4*f(x), f(x))
```

$$f(x) = C_1 \sin(a^2 x) + C_2 \cos(a^2 x)$$

12.1.7 Series expansions and plotting

It is possible to expand many SymPy expressions as Taylor series. The `series` method makes this straightforward. At minimum, we must specify the expression and the variable in which to expand it. Optionally, we can also specify the point around which to expand, the maximum term number, and the direction of the expansion (try `help(Basic.series)` for more information).

```
from sympy import *
x = Symbol('x')
sin(x).series(x, 0)
```

$$x - \frac{x^3}{6} + \frac{x^5}{120} + O(x^6)$$

```
series(sin(x), x, 0)
```

$$x - \frac{x^3}{6} + \frac{x^5}{120} + O(x^6)$$

```
# NBVAL_IGNORE_OUTPUT
cos(x).series(x, 0.5, 10)
```

$$1.11729533119247 - 0.438791280945186 (x - 0.5)^2 + 0.0799042564340338 (x - 0.5)^3 + 0.0365659400787655 (x - 0.5)^4 - 0.0011729533119247 (x - 0.5)^5 + 0.000011729533119247 (x - 0.5)^6 - 0.00000011729533119247 (x - 0.5)^7 + 0.000000011729533119247 (x - 0.5)^8 - 0.0000000011729533119247 (x - 0.5)^9 + 0.00000000011729533119247 (x - 0.5)^{10}$$

In some cases, especially for numerical evaluation and plotting the results, it is necessary to remove the trailing $O(n)$ term:

```
# NBVAL_IGNORE_OUTPUT
cos(x).series(x, 0.5, 10).removeO()
```

$$-0.479425538604203x - 1.32116826114474 \cdot 10^{-6} (x - 0.5)^9 + 2.17654405230747 \cdot 10^{-5} (x - 0.5)^8 + 9.51241148024212 \cdot 10^{-5} (x - 0.5)^7 - 0.00011729533119247 (x - 0.5)^6 + 0.0000011729533119247 (x - 0.5)^5 - 0.00000011729533119247 (x - 0.5)^4 + 0.000000011729533119247 (x - 0.5)^3 - 0.0000000011729533119247 (x - 0.5)^2 + 0.00000000011729533119247 (x - 0.5)$$

SymPy provides two inbuilt plotting functions, `Plot()` from the `sympy.plotting` module, and `plot` from `sympy.mpmath.visualization`. At the time of writing, these functions lack the ability to add a key to the plot, which means they are unsuitable for most of our needs. Should you wish to use them nevertheless, their `help()` text is useful.

For most of our purposes, Matplotlib should be the plotting tool of choice. The details are in chapter [cha:visualisingdata]. Here we furnish just one example of how to plot the results of a SymPy computation.

```
%matplotlib inline
```

```
from sympy import sin, series, Symbol
import pylab
x = Symbol('x')
s10 = sin(x).series(x, 0, 10).removeO()
```

(continues on next page)

(continued from previous page)

```
s20 = sin(x).series(x,0,20).removeO()
s = sin(x)
xx = []
y10 = []
y20 = []
y = []
for i in range(1000):
    xx.append(i / 100.0)
    y10.append(float(s10.subs({x:i/100.0})))
    y20.append(float(s20.subs({x:i/100.0})))
    y.append(float(s.subs({x:i/100.0})))

pylab.figure()
```

<Figure size 640x480 with 0 Axes>


<Figure size 640x480 with 0 Axes>

```
pylab.plot(xx, y10, label='O(10)')
pylab.plot(xx, y20, label='O(20)')
pylab.plot(xx, y, label='sin(x)')

pylab.axis([0, 10, -4, 4])
pylab.xlabel('x')
pylab.ylabel('f(x)')

pylab.legend()
```

<matplotlib.legend.Legend at 0x7f027c60cb10>



12.1.8 Linear equations and matrix inversion

SymPy has a `Matrix` class and associated functions that allow the symbolic solution of systems of linear equations (and, of course, we can obtain numerical answers with `subs()` and `evalf()`). We shall consider the example of the following simple pair of linear equations:

$$\begin{aligned} 3x + 7y &= 12z \\ 4x - 2y &= 5z \end{aligned}$$

We may write this system in the form $A\vec{x} = \vec{b}$ (multiply A by \vec{x} if you want to verify that we recover the original equations), where

$$A = \begin{pmatrix} 3 & 7 \\ 4 & -2 \end{pmatrix}, \quad \vec{x} = \begin{pmatrix} x \\ y \end{pmatrix}, \quad \vec{b} = \begin{pmatrix} 12z \\ 5z \end{pmatrix}.$$

Here we included a symbol, z , on the right-hand side to demonstrate that symbols will be propagated into the solution. In many cases we would have $z = 1$, but there may still be benefit to using SymPy over a numerical solver even when the solution contains no symbols because of its ability to return exact fractions rather than approximate floats.

One strategy to solve for \vec{x} is to invert the matrix A and pre-multiply, *i.e.* $A^{-1}A\vec{x} = \vec{x} = A^{-1}\vec{b}$. SymPy's `Matrix` class has an `inv()` method that allows us to find the inverse, and `*` performs matrix multiplication for us, when appropriate:

```
from sympy import symbols, Matrix
x, y, z = symbols('x, y, z')
A = Matrix([[3, 7], [4, -2]])
A
```

$$\begin{bmatrix} 3 & 7 \\ 4 & -2 \end{bmatrix}$$

```
A.inv()
```

$$\begin{bmatrix} \frac{1}{17} & \frac{7}{34} \\ \frac{2}{17} & -\frac{3}{34} \end{bmatrix}$$

```
b = Matrix(( 12*z, 5*z ))
b
```

$$\begin{bmatrix} 12z \\ 5z \end{bmatrix}$$

```
x = A.inv()*b
x
```

$$\begin{bmatrix} \frac{59z}{34} \\ \frac{33z}{34} \end{bmatrix}$$

```
x.subs({z:3.3}).evalf(4)
```

$$\begin{bmatrix} 5.726 \\ 3.203 \end{bmatrix}$$

```
type(x)
```

```
sympy.matrices.dense.MutableDenseMatrix
```

An alternative method of solving the same problem is to construct the system as a matrix in augmented form; that is the form obtained by appending the columns of (in our example) A and \vec{b} together. The augmented matrix is [1]:

$$(A|\vec{b}) = \left(\begin{array}{cc|c} 3 & 7 & 12z \\ 4 & -2 & 5z \end{array} \right),$$

and as before we construct this as a SymPy Matrix object, but in this case we pass it to the `solve_linear_system()` function:

```
from sympy import Matrix, symbols, solve_linear_system
x, y, z = symbols('x, y, z')
system = Matrix([[3, 7, 12*z], [4, -2, 5*z]])
system
```

$$\begin{bmatrix} 3 & 7 & 12z \\ 4 & -2 & 5z \end{bmatrix}$$

```
sol = solve_linear_system(system, x, y)
sol
```

$$\left\{ x : \frac{59z}{34}, y : \frac{33z}{34} \right\}$$

```
type(sol)
```

```
dict
```

```
# NBVAL_IGNORE_OUTPUT
for k in sol.keys():
    print(k, '=', sol[k].subs({z:3.3}).evalf(4))
```

```
x = 5.726
y = 3.203
```

A third option is the `solve()` method, whose arguments include the individual symbolic equations, rather than any matrices. Like `dsolve()` (see *ODEs*), `solve()` expects either expressions which it will assume equal to zero, or Equality objects, which we can conveniently create with `Eq()`:

```
from sympy import symbols, solve, Eq
x, y, z = symbols('x, y, z')
solve((Eq(3*x+7*y-12*z), Eq(4*x-2*y-5*z)), x, y)
```

$$\left\{ x : \frac{59z}{34}, y : \frac{33z}{34} \right\}$$

```
solve((3*x+7*y-12*z, 4*x-2*y-5*z), x, y)
```

$$\left\{ x : \frac{59z}{34}, y : \frac{33z}{34} \right\}$$

For more information, see `help(solve)` and `help(solve_linear_system)`.

12.1.9 Non linear equations

Let's solve a simple equation such as $x = x^2$. There are two obvious solutions: $x = 0$ and $x = 1$. How can we ask Sympy to compute these for us?

```
import sympy
x, y, z = sympy.symbols('x, y, z')          # create some symbols
eq = x - x ** 2                             # define the equation
```

```
sympy.solve(eq, x)                          # solve eq = 0
```


$[0, 1]$

The `solve()` function expects an expression that is meant to be solved so that it evaluates to zero. For our example, we rewrite

$x = x^2$ as $x - x^2 = 0$ and then pass this to the solve function.

Let's repeat the same for the equation: $x = x^3$ and solve

```
eq = x - x ** 3          # define the equation
sympy.solve(eq, x)      # solve eq = 0
```

 $[-1, 0, 1]$

12.1.10 Output: LaTeX interface and pretty-printing

As is the case with many computer algebra systems, SymPy has the ability to format its output as LaTeX code, for easy inclusion into documents.

At the start of this chapter, we called:

```
sympy.init_printing()
```

Sympy detected that it was in Jupyter, and enabled Latex output. The Jupyter Notebook supports (some) Latex, so this gives us the nicely formatted output above.

We can also see the plain text output from Sympy, and the raw Latex code it creates:

```
print(series(1/(x+y), y, 0, 3))
```

```
y**2/x**3 - y/x**2 + 1/x + O(y**3)
```

```
print(latex(series(1/(x+y), y, 0, 3)))
```

```
\frac{y^2}{x^3} - \frac{y}{x^2} + \frac{1}{x} + O\left(y^3\right)
```

```
print(latex(series(1/(x+y), y, 0, 3), mode='inline'))
```

```
\frac{y^2}{x^3} - \frac{y}{x^2} + 1 / x + O\left(y^3\right)
```

Be aware that in its default mode, `latex()` outputs code that requires the `amsmath` package to be loaded via a `\backslashusepackage{amsmath}` command in the document preamble.

Sympy also supports a “pretty print” (`pprint()`) output routine, which produces better-formatted text output than the default printing routine, as illustrated below. Note features such as the subscripts for array elements whose names are of the form `Tn`, the italicised constant *e*, vertically-centred dots for multiplication, and the nicely-formed matrix borders and fractions.

```

File Edit View Terminal Help
>>> pprint(integrate(x**2*exp(y),x,x))
  4 y
x · e
-----
 12
>>> pprint(integrate(x**2*exp(x),x,x))
  x  2 x      x
6·e  + x ·e  - 4·x·e
>>> pprint(M)
[12·T1  T2  T3]
[14·T4  5·T5  T6]
>>> pprint(sin(x).series(x,0.5,10))
  3 5 7 9
x  x  x  x
x - — + — - — + — + 0(x**10)
  6 120 5040 362880
>>> pprint(cos(x).series(x,0.5,15))
  2 4 6 8 10 12 14
x  x  x  x  x  x  x
1 - — + — - — + — - — + — + 0(x**15)
  2 24 720 40320 3628800 479001600 87178291200
>>>

```

Finally, SymPy offers `preview()`, which displays rendered output on screen (check `help(preview)` for details).

12.1.11 Automatic generation of C code

A strong point of many symbolic libraries is that they can convert the symbolic expressions to C-code (or other code) that can subsequently be compiled for high execution speed. Here is an example that demonstrates this:

```

from sympy import *
from sympy.utilities.codegen import codegen
x = Symbol('x')
sin(x).series(x, 0, 6)

```

$$x - \frac{x^3}{6} + \frac{x^5}{120} + O(x^6)$$

```

# NBVAL_IGNORE_OUTPUT
print(codegen(("taylor_sine", sin(x).series(x,0,6)), language='C')[0][1])

```

```

/*****
*
*           Code generated with SymPy 1.12
*
*           See http://www.sympy.org/ for more information.
*
*           This file is part of 'project'
*****/
#include "taylor_sine.h"
#include <math.h>

```

(continues on next page)

(continued from previous page)

```
double taylor_sine(double x) {  
  
    double taylor_sine_result;  
    taylor_sine_result = x - 1.0/6.0*pow(x, 3) + (1.0/120.0)*pow(x, 5) + O(x**6);  
    return taylor_sine_result;  
  
}
```

12.2 Related tools

It is worth noting that the SAGE initiative <https://www.sagemath.org/> is trying to “create a viable free open source alternative to Magma, Maple, Mathematica and Matlab.” and includes the SymPy library among many others. Its symbolic capabilities are more powerful than SymPy’s, and SAGE, but the SymPy features will already cover many of the needs arising in science and engineering. SAGE includes the computer algebra system Maxima, which is also available standalone from https://doc.sagemath.org/html/en/reference/interfaces/sage/interfaces/maxima_abstract.html.

NUMERICAL COMPUTATION

13.1 Numbers and numbers

We have already seen (03 Data Types Structures, Numbers) that Python knows different *types* of numbers:

- floating point numbers such as 3.14
- integers such as 42
- complex numbers such as $3.14 + 1j$

13.1.1 Limitations of number types

Limitations of `ints`

Mathematics provides the infinite set of natural numbers $\mathbb{N} = \{1, 2, 3, \dots\}$. Because the computer has *finite* size, it is impossible to represent all of these numbers in the computer. Instead, only a small subset of numbers is represented.

The `int`-type can (usually[3]) represent numbers between -2147483648 and +2147483647 and corresponds to 4 bytes (that's 4*8 bit, and $2^{32} = 4294967296$ which is the range from -2147483648 and +2147483647).

You can imagine that the hardware uses a table like this to encode integers using bits (suppose for simplicity we use only 8 bits for this):

natural number	bit-representation
0	00000000
1	00000001
2	00000010
3	00000011
4	00000100
5	00000101
⋮	⋮
254	11111110
255	11111111

Using 8 bit we can represent 256 natural numbers (for example from 0 to 255) because we have $2^8 = 256$ different ways of combining eight 0s and 1s.

We could also use a slightly different table to describe 256 integer numbers ranging, for example, from -127 to +128.

This is *in principle* how integers are represented in the computer. Depending on the number of bytes used, only integer numbers between a minimum and a maximum value can be represented. On today's hardware, it is common to use

4 or 8 bytes to represent one integer, which leads exactly to the minimum and maximum values of -2147483648 and $+2147483647$ as shown above for 4 bytes, and $+9223372036854775807$ as the maximum integer for 8 bytes (that's $\approx 9.2 \cdot 10^{18}$).

Limitations of floats

The floating point numbers in a computer are not the same as the mathematical floating point numbers. (This is exactly the same as the (mathematical) integer numbers not being the same as the integer numbers in a computer: only a *subset* of the infinite set of integer numbers can be represented by the `int` data type as shown in *Numbers and numbers*). So how are floating point numbers represented in the computer?

- Any real number x can be written as $x = a \cdot 10^b$ where a is the mantissa and b the exponent.
- Examples:

x	a	b
$123.45 = 1.23456 \cdot 10^2$	1.23456	2
$1000000 = 1.0 \cdot 10^6$	1.00000	6
$0.0000024 = 2.4 \cdot 10^{-6}$	2.40000	-6

- Therefore, we can use 2 integers to encode one floating point number!
 $x = a \cdot 10^b$
- Following (roughly) the IEEE-754 standard, one uses 8 bytes for one float x : these 64 bits are split as
 - 10 bit for the exponent b and
 - 54 bit for the mantissa a .

This results in

- largest possible float $\approx 10^{308}$ (quality measure for b)
- smallest possible (positive) float $\approx 10^{-308}$ (quality measure for b)
- distance between 1.0 and next larger number $\approx 10^{-16}$ (quality measure for a)

Note that this is *in principle* how floating point numbers are stored (it is actually a bit more complicated).

Limitations of complex numbers

The `complex` number type has essentially the same limitations as the `float` data type (see *limitations of floats*) because a complex number consists of two `floats`: one represents the real part, the other one the imaginary part.

...are these number types of practical value?

In practice, we do not usually find numbers in our daily life that exceed 10^{300} (this is a number with 300 zeros!), and therefore the floating point numbers cover the range of numbers we usually need.

However, be warned that in scientific computation small and large numbers are used which may (often in intermediate results) exceed the range of floating point numbers.

- Imagine for example, that we have to take the fourth power of the constant $\hbar = 1.0545716 \cdot 10^{-34} \text{kgm}^2/\text{s}$:
- $\hbar^4 = 1.2368136958909421 \cdot 10^{-136} \text{k}^4 \text{m}^8/\text{s}^4$ which is “halfway” to our representable smallest positive float of the order of 10^{-308} .

If there is any danger that we might exceed the range of the floating point numbers, we have to *rescale* our equations so that (ideally) all numbers are of order unity. Rescaling our equations so that all relevant numbers are approximately 1 is also useful in debugging our code: if numbers much greater or smaller than 1 appear, this may be an indication of an error.

13.1.2 Using floating point numbers (carelessly)

We know already that we need to take care that our floating point values do not exceed the range of floating point numbers that can be represented in the computer.

There is another complication due to the way floating point numbers have to be represented internally: not all floating point numbers can be represented exactly in the computer. The number 1.0 can be represented exactly but the numbers 0.1, 0.2 and 0.3 cannot:

```
'%.20f' % 1.0
```

```
'1.00000000000000000000'
```

```
'%.20f' % 0.1
```

```
'0.100000000000000000555'
```

```
'%.20f' % 0.2
```

```
'0.200000000000000001110'
```

```
'%.20f' % 0.3
```

```
'0.29999999999999998890'
```

Instead, the floating point number “nearest” to the real number is chosen.

This can cause problems. Suppose we need a loop where x takes values 0.1, 0.2, 0.3, ..., 0.9, 1.0. We might be tempted to write something like this:

```
x = 0.0
while not x == 1.0:
    x = x + 0.1
    print ( " x =%.17f" % ( x ) )
```

However, this loop will never terminate. Here are the first 11 lines of output of the program:

```
x=0.100000000000000001
x=0.200000000000000001
x=0.300000000000000004
x=0.400000000000000002
x=
    0.5
x=0.59999999999999998
x=0.69999999999999996
x=0.79999999999999993
```

(continues on next page)

(continued from previous page)

```
x=0.899999999999999991
x=0.99999999999999989
x=1.09999999999999987
```

Because the variable `x` never takes exactly the value 1.0, the while loop will continue forever.

Thus: *Never compare two floating point numbers for equality.*

13.1.3 Using floating point numbers carefully 1

There are a number of alternative ways to solve this problem. For example, we can compare the distance between two floating point numbers:

```
x = 0.0
while abs(x - 1.0) > 1e-8:
    x = x + 0.1
    print ( " x =%19.17f" % ( x ) )
```

```
x =0.100000000000000001
x =0.200000000000000001
x =0.300000000000000004
x =0.400000000000000002
x =0.500000000000000000
x =0.59999999999999998
x =0.69999999999999996
x =0.79999999999999993
x =0.89999999999999991
x =0.99999999999999989
```

13.1.4 Using floating point numbers carefully 2

Alternatively, we can (for this example) iterate over a sequence of integers and compute the floating point number from the integer:

```
for i in range (1 , 11):
    x = i * 0.1
    print(" x =%19.17f" % ( x ) )
```

```
x =0.100000000000000001
x =0.200000000000000001
x =0.300000000000000004
x =0.400000000000000002
x =0.500000000000000000
x =0.600000000000000009
x =0.700000000000000007
x =0.800000000000000004
x =0.900000000000000002
x =1.000000000000000000
```



```

x=0.100000000000000001
x=0.20000000000000001
x=0.30000000000000004
x=0.40000000000000002
x=          0.5
x=0.60000000000000009
x=0.70000000000000007
x=0.80000000000000004
x=0.90000000000000002
x=          1

```

If we compare this with the output from the program in *Using floating point numbers (carelessly)*, we can see that the floating point numbers differ. This means that – in a numerical calculation – it is not true that $0.1 + 0.1 + 0.1 + 0.1 + 0.1 + 0.1 + 0.1 + 0.1 + 0.1 + 0.1 + 0.1 + 0.1 = 1.0$.

13.1.5 Symbolic calculation

Using the sympy package we have arbitrary precision. Using `sympy.Rational`, we can define the fraction 1/10 exactly symbolically. Adding this 10 times will lead exactly to the value 1, as demonstrated by this script

```

from sympy import Rational
dx = Rational (1 ,10)
x = 0
while x != 1.0:
    x = x + dx
    print("Current x=%4s = %3.1f " % (x , x . evalf ()))
    print(" Reached x=%s " % x)

```

```

Current x=1/10 = 0.1
Reached x=1/10
Current x= 1/5 = 0.2
Reached x=1/5
Current x=3/10 = 0.3
Reached x=3/10
Current x= 2/5 = 0.4
Reached x=2/5
Current x= 1/2 = 0.5
Reached x=1/2
Current x= 3/5 = 0.6
Reached x=3/5
Current x=7/10 = 0.7
Reached x=7/10
Current x= 4/5 = 0.8
Reached x=4/5
Current x=9/10 = 0.9
Reached x=9/10
Current x=  1 = 1.0
Reached x=1

```

However, this symbolic calculation is much slower as it is done through software rather than the CPU-based floating point operations. The next program approximates the relative performances:

```

# NBVAL_IGNORE_OUTPUT
from sympy import Rational

```

(continues on next page)

(continued from previous page)

```
dx_symbolic = Rational (1 ,10)
dx = 0.1

def loop_sympy (n):
    x = 0
    for i in range(n):
        x = x + dx_symbolic
    return x

def loop_float (n):
    x = 0
    for i in range(n):
        x = x + dx
    return x

def time_this (f, n):
    import time
    starttime = time.time()
    result = f(n)
    stoptime = time.time()
    print(" deviation is %16.15g" % ( n * dx_symbolic - result ))
    return stoptime - starttime

n = 100000
print("loop using float dx:")
time_float = time_this(loop_float, n)
print("float loop n=%d takes %6.5f seconds" % (n, time_float))
print("loop using sympy symbolic dx:")
time_sympy = time_this (loop_sympy, n)
print("sympy loop n =% d takes %6.5f seconds" % (n , time_sympy ))
print("Symbolic loop is a factor %.1f slower." % ( time_sympy / time_float ))
```

```
loop using float dx:
 deviation is -1.88483681995422e-08
float loop n=100000 takes 0.00235 seconds
loop using sympy symbolic dx:
 deviation is          0
sympy loop n = 100000 takes 0.19528 seconds
Symbolic loop is a factor 83.2 slower.
```

This is of course an artificial example: we have added the symbolic code to demonstrate that these round off errors originate from the approximative representation of floating point numbers in the hardware (and thus programming languages). We can, in principle, avoid these complications by computing using symbolic expressions, but this is in practice too slow.[4]

13.1.6 Summary

In summary, we have learned that

- floating point numbers and integers used in numeric computation are generally quite different from “mathematical numbers” (symbolic calculations are exact and use the “mathematical numbers”):
 - there is a maximum number and a minimum number that can be represented (for both integers and floating point numbers)
 - within this range, not every floating point number can be represented in the computer.
- We deal with this limitation by:
 - never comparing two floating point numbers for equality (instead we compute the absolute value of the difference)
 - use of algorithms that are *stable* (this means that small deviations from correct numbers can be corrected by the algorithm. We have not yet shown any such examples this document.)
- Note that there is a lot more to be said about numerical and algorithmic tricks and methods to make numeric computation as accurate as possible but this is outside the scope of this section.

13.1.7 Exercise: infinite or finite loop

1. What does the following piece of code compute? Will the loop ever finish? Why?

```
eps = 1.0
while 1.0 + eps > 1.0:
    eps = eps / 2.0
print(eps)
```


NUMERICAL PYTHON (NUMPY): ARRAYS

14.1 Numpy introduction

The NumPy package (read as NUMerical PYthon) provides access to

- a new data structure called `arrays` which allow
- efficient vector and matrix operations. It also provides
- a number of linear algebra operations (such as solving of systems of linear equations, computation of Eigenvectors and Eigenvalues).

14.1.1 History

Some background information: There are two other implementations that provide nearly the same functionality as NumPy. These are called “Numeric” and “numarray”:

- Numeric was the first provision of a set of numerical methods (similar to Matlab) for Python. It evolved from a PhD project.
- Numarray is a re-implementation of Numeric with certain improvements (but for our purposes both Numeric and Numarray behave virtually identical).
- Early in 2006 it was decided to merge the best aspects of Numeric and Numarray into the Scientific Python (`scipy`) package and to provide (a hopefully “final”) `array` data type under the module name “NumPy”.

We will use in the following materials the “NumPy” package as provided by (new) SciPy. If for some reason this doesn’t work for you, chances are that your SciPy is too old. In that case, you will find that either “Numeric” or “numarray” is installed and should provide nearly the same capabilities.[5]

14.1.2 Arrays

We introduce a new data type (provided by NumPy) which is called “`array`”. An array *appears* to be very similar to a list but an array can keep only elements of the same type (whereas a list can mix different kinds of objects). This means arrays are more efficient to store (because we don’t need to store the type for every element). It also makes arrays the data structure of choice for numerical calculations where we often deal with vectors and matrices.

Vectors and matrices (and matrices with more than two indices) are all called “arrays” in NumPy.

Vectors (1d-arrays)

The data structure we will need most often is a vector. Here are a few examples of how we can generate one:

- Conversion of a list (or tuple) into an array using `numpy.array`:

```
import numpy as np
x = np.array([0, 0.5, 1, 1.5])
print(x)
```

```
[0.  0.5 1.  1.5]
```

- Creation of a vector using “ArrayRANGE”:

```
x = np.arange(0, 2, 0.5)
print(x)
```

```
[0.  0.5 1.  1.5]
```

- Creation of vector with zeros

```
x = np.zeros(4)
print(x)
```

```
[0.  0.  0.  0.]
```

Once the array is established, we can set and retrieve individual values. For example:

```
x = np.zeros(4)
x[0] = 3.4
x[2] = 4
print(x)
print(x[0])
print(x[0:-1])
```

```
[3.4 0.  4.  0. ]
3.4
[3.4 0.  4. ]
```

Note that once we have a vector we can perform calculations on every element in the vector with a single statement:

```
x = np.arange(0, 2, 0.5)
print(x)
print(x + 10)
print(x ** 2)
print(np.sin(x))
```

```
[0.  0.5 1.  1.5]
[10.  10.5 11.  11.5]
[0.  0.25 1.  2.25]
[0.          0.47942554 0.84147098 0.99749499]
```

Matrices (2d-arrays)

Here are two ways to create a 2d-array:

- By converting a list of lists (or tuples) into an array:

```
x = np.array([[1, 2, 3], [4, 5, 6]])
x
```

```
array([[1, 2, 3],
       [4, 5, 6]])
```

- Using the `zeros` method to create a matrix with 5 rows and 4 columns

```
x = np.zeros((5, 4))
x
```

```
array([[0., 0., 0., 0.],
       [0., 0., 0., 0.],
       [0., 0., 0., 0.],
       [0., 0., 0., 0.],
       [0., 0., 0., 0.]])
```

The “shape” of a matrix can be queried like this (here we have 2 rows and 3 columns):

```
x = np.array([[1, 2, 3], [4, 5, 6]])
print(x)
x.shape
```

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

```
(2, 3)
```

Individual elements can be accessed and set using this syntax:

```
x = np.array([[1, 2, 3], [4, 5, 6]])
x[0, 0]
```

```
1
```

```
x[0, 1]
```

```
2
```

```
x[0, 2]
```

```
3
```

```
x[1, 0]
```

```
4
```

```
x[:, 0]
```

```
array([1, 4])
```

```
x[0, :]
```

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

14.1.3 Convert from array to list or tuple

To create an array back to a list or tuple, we can use the standard python functions `list(s)` and `tuple(s)` which take a sequence `s` as the input argument and return a list and tuple, respectively:

```
a = np.array([1, 4, 10])  
a
```

```
array([ 1,  4, 10])
```

```
list(a)
```

```
[1, 4, 10]
```

```
tuple(a)
```

```
(1, 4, 10)
```

14.1.4 Standard Linear Algebra operations

Maxtrix multiplication

Two arrays can be multiplied in the usual linear-algebra way using `numpy.matrixmultiply`. Here is an example:

```
import numpy as np  
import numpy.random  
A = numpy.random.rand(5, 5)      # generates a random 5 by 5 matrix  
x = numpy.random.rand(5)        # generates a 5-element vector  
b = np.dot(A, x)                # multiply matrix A with vector x
```


Solving systems of linear equations

To solve a system of equations $A\mathbf{x} = \mathbf{b}$ that is given in matrix form (*i.e.* A is a matrix and \mathbf{x} and \mathbf{b} are vectors where A and \mathbf{b} are known and we want to find the unknown vector \mathbf{x}), we can use the linear algebra package (`linalg`) of `numpy`:

```
import numpy.linalg as LA
x = LA.solve(A, b)
```

Computing Eigenvectors and Eigenvalues

Here is a small example that computes the [trivial] Eigenvectors and Eigenvalues (`eig`) of the unity matrix (`eye`):

```
import numpy
import numpy.linalg as LA
A = numpy.eye(3)      #'eye' -> I -> 1 (ones on the diagonal)
print(A)
```

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

```
evals, evectors = LA.eig(A)
print(evals)
```

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

```
print(evectors)
```

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

Note that each of these commands provides its own documentation. For example, `help(LA.eig)` will tell you all about the eigenvector and eigenvalue function (once you have imported `numpy.linalg` as `LA`).

Curve fitting of polynomials

Let's assume we have x-y data to which we like to fit a curve (to minimise the least square deviation of the fit from the data).

`numpy` provides the routine `polyfit(x, y, n)` (which is similar to Matlab's `polyfit` function which takes a list `x` of x-values for data points, a list `y` of y-values of the same data points and a desired order of the polynomial that will be determined to fit the data in the least-square sense as well as possible).

```
%matplotlib inline
import numpy as np

# demo curve fitting : xdata and ydata are input data
xdata = np.array([0.0 , 1.0 , 2.0 , 3.0 , 4.0 , 5.0])
ydata = np.array([0.0 , 0.8 , 0.9 , 0.1 , -0.8 , -1.0])
```

(continues on next page)

(continued from previous page)


```
# now do fit for cubic (order = 3) polynomial
z = np.polyfit(xdata, ydata, 3)

# z is an array of coefficients , highest first , i . e .
#           X^3           X^2           X           0
# z = array ([ 0.08703704 , -0.81349206 , 1.69312169 , -0.03968254])
# It is convenient to use 'poly1d' objects for dealing with polynomials:
p = np.poly1d(z) # creates a polynomial function p from coefficients
                 # and p can be evaluated for all x then .

# create plot
xs = [0.1 * i for i in range (50)]
ys = [p(x) for x in xs] # evaluate p(x) for all x in list xs

import pylab
pylab.plot(xdata, ydata, 'o', label='data')
pylab.plot(xs, ys, label='fitted curve')
pylab.ylabel('y')
pylab.xlabel('x')
```

```
Text(0.5, 0, 'x')
```



This shows the fitted curve (solid line) together with the precise computed data points.

14.1.5 More numpy examples...

...can be found here: <https://numpy.org/doc/stable/reference/routines.html>

14.1.6 Numpy for Matlab users

There is a dedicated webpage that explains Numpy from the perspective of a (experienced) Matlab user at <https://numpy.org/doc/stable/user/numpy-for-matlab-users.html>.

VISUALISING DATA

The purpose of scientific computation is insight not numbers: To understand the meaning of the (many) numbers we compute, we often need postprocessing, statistical analysis and graphical visualisation of our data. The following sections describe

- Matplotlib/PyLab — a tool to generate high quality graphs of the type $y = f(x)$ (and a bit more)
 - the `pylab` interface
 - the `pyplot` interface

We also touch on:

- Visual Python — a tool to quickly generate animations of time dependent processes taking place in 3d space.
- Tools to store and visualise vtk files

We close with a short outlook on

- Further tools and developments discussing other tools and emerging approaches for data visualisation and analysis.

15.1 Matplotlib – plotting $y=f(x)$, (and a bit more)

The Python library *Matplotlib* is a python 2D plotting library which produces publication quality figures in a variety of hardcopy formats and interactive environments. Matplotlib tries to make easy things easy and hard things possible. You can generate plots, histograms, power spectra, bar charts, errorcharts, scatterplots, etc, with just a few lines of code.

For more detailed information, check these links

- A very nice introduction in the object oriented Matplotlib interface, and summary of all important ways of changing style, figure size, linewidth, etc. This is a useful reference: <https://github.com/jrjohansson/scientific-python-lectures/blob/master/Lecture-4-Matplotlib.ipynb>
- [Matplotlib tutorial](#)
- [Matplotlib home page](#)
- Extended thumbnail gallery of examples <https://matplotlib.org/stable/gallery/index.html>

15.1.1 Matplotlib and Pylab

The Matplotlib package provides *an object oriented plotting library* under the name space `matplotlib.pyplot`.

The `pylab` interface is provided through the Matplotlib package. Internally it uses `matplotlib.pyplot` functionality but imitates the (state-driven) Matlab plotting interface.

The `pylab` interface is slightly more convenient to use for simple plots, and `matplotlib.pyplot` gives far more detailed control over how plots are created. If you routinely need to produce figures, we suggest to learn about the object oriented `matplotlib.pyplot` interface (instead of the `pylab` interface).

This chapter focusses on the Pylab interface, but also provides examples for the object-oriented `matplotlib.pyplot` interface.

An excellent introduction and overview of the `matplotlib.pyplot` plotting interface is available in <https://github.com/jrjohansson/scientific-python-lectures/blob/master/Lecture-4-Matplotlib.ipynb>.

For the purpose of this book and the [Jupyterbook package](#), we use some settings to create a `svg` file for the html version of the book, and a high-resolution `png` file for the pdf version:

```
%matplotlib inline
# settings for jupyter book: svg for html version, high-resolution png for pdf
import matplotlib_inline
matplotlib_inline.backend_inline.set_matplotlib_formats('svg', 'png')
import matplotlib as mpl
mpl.rcParams['figure.dpi'] = 400
```

15.1.2 First example


The pyplot interface

The recommended way of using Matplotlib in a simple example is shown here:

```
# example 1 a
import numpy as np           # get access to fast arrays
import matplotlib.pyplot as plt # the plotting functions

x = np.arange(-3.14, 3.14, 0.01) # create x-data
y = np.sin(x)                   # compute y-data
plt.plot(x, y)                  # create plot
```

```
[<matplotlib.lines.Line2D at 0x7fc840326890>]
```



15.1.3 How to import matplotlib, pylab, pyplot, numpy and all that

The submodule `matplotlib.pyplot` provides an object oriented interface to the plotting library. Many of the examples in the `matplotlib` documentation follow the import convention to import `matplotlib.pyplot` as `plt` and `numpy` as `np`. It is of the user's decision whether to import the `numpy` library under the name `np` (as often done in `matplotlib` examples) or `N` as occasionally done in this text (and in the early days when the predecessor of `numpy` was called "Numeric") or any other name you like. Similarly, it is a matter of taste whether the plotting submodule (`matplotlib.pyplot`) is imported as `plt` as is done in the `matplotlib` documentation or `plot` (which could be argued is slightly clearer) etc.

As always a balance has to be struck between personal preferences and consistency with common practice in choosing these name. Consistency with common use is more important if the code is likely to be used by others or published.

The Pylab interface


We introduce the `pylab` interface by translating the example 1a above to the following example 1b (which is identical in functionality to the example 1a and will create the same plot):

```
# example 1b
import pylab
import numpy as np

x = np.arange (-3.14, 3.14, 0.01)
y = np.sin(x)

pylab.plot(x, y)
```

```
[<matplotlib.lines.Line2D at 0x7fc84029b550>]
```



Plotting nearly always needs arrays of numerical data and it is for this reason that the `numpy` module is used a lot: it provides fast and memory efficient array handling for Python (see [chapter 14](#)). The `pylab` interface has taken this a step further and automatically imports all objects from `numpy` into the `pylab` name space:


Because the `numpy.arange` and `numpy.sin` objects have already been imported into the `pylab` namespace, we could also write it as `example 1c`:

```
# example 1c
import pylab as p

x = p.arange(-3.14, 3.14, 0.01)
y = p.sin(x)

p.plot(x, y)
```


```
[<matplotlib.lines.Line2D at 0x7fc84011b950>]
```

If we really want to cut down on characters to type, we could also import all the objects (*) from the `pylab` convenience module into our current namespace, and rewrite the code as example 1d:

```
# example 1 d
from pylab import * # not generally recommended
                    # okay for interactive testing

x = arange(-3.14, 3.14, 0.01)
y = sin(x)
plot(x, y)
show()
```



This can be extremely convenient, but comes with a big health warning:

- While using `from pylab import *` is acceptable at the command prompt to interactively create plots and analyse data, this should never be used in any plotting scripts.
- The `pylab` toplevel provides over 800 different objects which are all imported into the global name space when running `from pylab import *`. This is not good practice, and could conflict with other objects that exist already or are created later.
- As a rule of thumb: do never use `from somewhere import *` in programs we save. This may be okay at the command prompt for interactive data exploration.

15.1.4 IPython's inline mode

Within the Jupyter Notebook or Qtconsole (see the *Python shells notebook*) we can use the `%matplotlib inline` magic command to make further plots appear within our console or notebook. To force pop up windows instead, use `%matplotlib qt`.

If you enjoy the `pylab` interface, then you maybe interested in the `%pylab` magic, which will not only switch to inline plotting but also automatically execute `from pylab import *`.

15.1.5 Saving the figure to a file

Once you have created the figure (using the `plot` command) and added any labels, legends etc, you have two options to save the plot.

1. You can display the figure (using `show`) and *interactively* save it by clicking on the disk icon. (This does not work with inline plots as the icons are not available.)
2. You can (without displaying the figure) save it directly from your Python code. The command to use is `savefig`. The format is determined by the extension of the file name you provide. Here is an example (`pylabsavefig.py`) which saves a figure into different files.

```
# saving figure files with the pylab interface
import pylab
import numpy as np

x = np.arange(-3.14, 3.14, 0.01)
y = np.sin(x)

pylab.plot(x, y, label='sin(x)')
pylab.savefig('myplot.png') # saves png file
pylab.savefig('myplot.svg') # saves svg file
pylab.savefig('myplot.eps') # saves eps file
pylab.savefig('myplot.pdf') # saves pdf file
pylab.close()
```

A note on file formats:

The `pdf`, `eps` and `svg` file formats are vector file formats which means that one can zoom into the image without losing quality (lines will still be sharp). File formats such as `png` (and `jpg`, `gif`, `tif`, `bmp`) save the image in form of a bitmap (i.e. a matrix of colour values) and will appear blurry or pixelated when zooming in (or when printed in high resolution).

Accordingly, choose a vector file format where you can, and use the bitmap (`png` for example) if there are no other options. Choose the `eps` or `pdf` file format if you plan to include the figure in a Latex document – depending on whether you want to compile it using `latex` (needs `eps`) or `pdflatex` (can use `pdf` [better] or `png`). If the version of MS Word (or other text processing software you use) can handle `pdf` files, it is better to use `pdf` than `png` for that.

```
# saving figure files with the pyplot interface
from matplotlib import pyplot as plt
import numpy as np

x = np.arange(-3.14, 3.14, 0.01)
y = np.sin(x)

fig, ax = plt.subplots()
ax.plot(x, y, label='sin(x)')
fig.savefig('myplot.png') # saves png file
fig.savefig('myplot.svg') # saves svg file
fig.savefig('myplot.eps') # saves eps file
fig.savefig('myplot.pdf') # saves pdf file
plt.close(fig)
```

15.2 The pylab interface

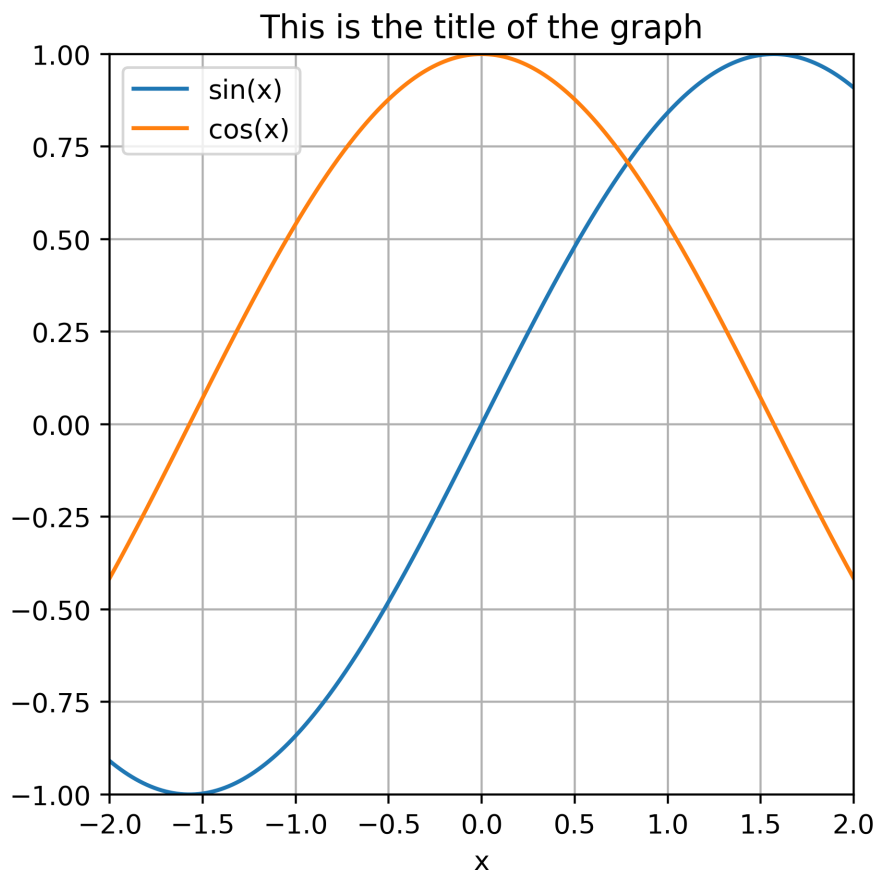
15.2.1 Fine tuning your plot

Matplotlib allows us to fine tune our plots in great detail. Here is an example:

```
import pylab
import numpy as N

x = N.arange(-3.14, 3.14, 0.01)
y1 = N.sin(x)
y2 = N.cos(x)
pylab.figure(figsize=(5, 5))
pylab.plot(x, y1, label='sin(x)')
pylab.plot(x, y2, label='cos(x)')
pylab.legend()
pylab.axis([-2, 2, -1, 1])
pylab.grid()
pylab.xlabel('x')
pylab.title('This is the title of the graph')
```

```
Text(0.5, 1.0, 'This is the title of the graph')
```



showing some other useful commands:

- `figure(figsize=(5, 5))` sets the figure size to 5inch by 5inch
- `plot(x, y1, label='sin(x)')` The “label” keyword defines the name of this line. The line label will be shown in the legend if the `legend()` command is used later.
- Note that calling the `plot` command repeatedly, allows you to overlay a number of curves.
- `axis([-2, 2, -1, 1])` This fixes the displayed area to go from `xmin=-2` to `xmax=2` in x-direction, and from `ymin=-1` to `ymax=1` in y-direction
- `legend()` This command will display a legend with the labels as defined in the plot command. Try `help(pylab.legend)` to learn more about the placement of the legend.
- `grid()` This command will display a grid on the backdrop.
- `xlabel('...')` and `ylabel('...')` allow labelling the axes.

Note further than you can chose different line styles, line thicknesses, symbols and colours for the data to be plotted. (The syntax is very similar to MATLAB.) For example:

- `plot(x, y, 'og')` will plot circles (o) in green (g)
- `plot(x, y, '-r')` will plot a line (-) in red (r)
- `plot(x, y, '-b', linewidth=2)` will plot a blue line (b) with two two pixel thickness `linewidth=2` which is twice as wide as the default.
- `plot(x, y, '--', alpha=0.5)` will plot a semi-transparent line (b).

The full list of options can be found when typing `help(pylab.plot)` at the Python prompt. Because this documentation is so useful, we repeat parts of it here:

```
plot(*args, **kwargs)
    Plot lines and/or markers to the
    :class:`matplotlib.axes.Axes`. *args* is a variable length
    argument, allowing for multiple *x*, *y* pairs with an
    optional format string. For example, each of the following is
    legal::

        plot(x, y)           # plot x and y using default line style and color
        plot(x, y, 'bo')     # plot x and y using blue circle markers
        plot(y)              # plot y using x as index array 0..N-1
        plot(y, 'r+')        # ditto, but with red plusses

    If *x* and/or *y* is 2-dimensional, then the corresponding columns
    will be plotted.

    An arbitrary number of *x*, *y*, *fmt* groups can be
    specified, as in::

        a.plot(x1, y1, 'g^', x2, y2, 'g-')

    Return value is a list of lines that were added.

    The following format string characters are accepted to control
    the line style or marker:

    =====
    character      description
    =====
    '-'            solid line style
```

(continues on next page)

(continued from previous page)

```

'--'          dashed line style
'-. '        dash-dot line style
':'          dotted line style
'.'          point marker
','          pixel marker
'o'          circle marker
'v'          triangle_down marker
'^'          triangle_up marker
'<'         triangle_left marker
'>'         triangle_right marker
'1'          tri_down marker
'2'          tri_up marker
'3'          tri_left marker
'4'          tri_right marker
's'          square marker
'p'          pentagon marker
'*'          star marker
'h'          hexagon1 marker
'H'          hexagon2 marker
'+'          plus marker
'x'          x marker
'D'          diamond marker
'd'          thin_diamond marker
'|'          vline marker
'_'          hline marker
=====

```

The following color abbreviations are supported:

```

=====
character  color
=====
'b'        blue
'g'        green
'r'        red
'c'        cyan
'm'        magenta
'y'        yellow
'k'        black
'w'        white
=====

```

In addition, you can specify colors in many weird and wonderful ways, including full names (`'green'`), hex strings (`'#008000'`), RGB or RGBA tuples (`(0,1,0,1)`) or grayscale intensities as a string (`'0.8'`). Of these, the string specifications can be used in place of a `fmt` group, but the tuple forms can be used only as `kwargs`.

Line styles and colors are combined in a single format string, as in `'bo'` for blue circles.

The `*kwargs` can be used to set line properties (any property that has a `set_*` method). You can use this to set a line label (for auto legends), linewidth, antialiasing, marker face color, etc. Here is an example::

(continues on next page)

(continued from previous page)

```

plot([1,2,3], [1,2,3], 'go-', label='line 1', linewidth=2)
plot([1,2,3], [1,4,9], 'rs', label='line 2')
axis([0, 4, 0, 10])
legend()

```

If you make multiple lines with one plot command, the kwargs apply to all those lines, e.g.::

```
plot(x1, y1, x2, y2, antialiased=False)
```

Neither line will be antialiased.

You do not need to use format strings, which are just abbreviations. All of the line properties can be controlled by keyword arguments. For example, you can set the color, marker, linestyle, and markercolor with::

```

plot(x, y, color='green', linestyle='dashed', marker='o',
      markerfacecolor='blue', markersize=12). See
      :class:`~matplotlib.lines.Line2D` for details.

```

The use of different line styles and thicknesses is particularly useful when colour cannot be used to distinguish lines (for example when graph will be used in document that is to be printed in black and white only).

15.2.2 Plotting more than one curve

There are three different methods to display more than one curve.

Two (or more) curves in one graph

By calling the `plot` command repeatedly, more than one curve can be drawn in the same graph. Example:


```

import numpy as np
t = np.arange(0, 2*np.pi, 0.01)

import pylab
pylab.plot(t, np.sin(t), label='sin(t)')
pylab.plot(t, np.cos(t), label='cos(t)')
pylab.legend()

```

```
<matplotlib.legend.Legend at 0x7fc82c917890>
```



Two (or more graphs) in one figure window

The `pylab.subplot` command allows to arrange several graphs within one figure window. The general syntax is

```
subplot(numRows, numCols, plotNum)
```

For example, to arrange 4 graphs in a 2-by-2 matrix, and to select the first graph for the next plot command, one can use:

```
subplot(2, 2, 1)
```


Here is a complete example plotting the sine and cosine curves in two graphs that are aligned underneath each other within the same window:

```
import numpy as np
t = np.arange(0, 2*np.pi, 0.01)

import pylab

pylab.subplot(2, 1, 1)
pylab.plot(t, np.sin(t))
pylab.xlabel('t')
pylab.ylabel('sin(t)')

pylab.subplot(2, 1, 2)
pylab.plot(t, np.cos(t))
pylab.xlabel('t')
pylab.ylabel('cos(t)');
```






Two (or more) figure windows

```
import pylab
pylab.figure(1)
pylab.plot(range(10), 'o')

pylab.figure(2)
pylab.plot(range(100), 'x')
```

```
[<matplotlib.lines.Line2D at 0x7fc82c819810>]
```



Note that you can use `pylab.close()` to close one, some or all figure windows (use `help(pylab.close)` to learn more). The closing of figures is not relevant for inline plots, but for plots that appear in pop-up windows, those windows will be closed when the figure is closed.

15.2.3 Interactive mode

Pylab can be run in two modes:

- non-interactive (this is the default)
- interactive.

In non-interactive mode, no plots will be displayed until the `show()` command has been issued. In this mode, the `show()` command should be the last statement of your program.

In interactive mode, plots will be immediately displayed after the plot command has been issued.

One can switch the interactive mode on using `pylab.ion()` and off using `pylab.ioff()`. IPython's `%matplotlib` magic also enables interactive mode.

If you use Jupyter notebooks with inline plots, then this feature is not so relevant.

15.3 The matplotlib.pyplot interface

This is the recommended way to use matplotlib for producing publication quality plots, or anything that needs some fine tuning: the object oriented approach of the `pyplot` interface makes it generally easier to tailor the plots than the state-driven `pylab` interface.

The central two commands to create `pyplot` figures are:

1. Create a figure object, and one (or more) axes objects within the figure.
2. Create some drawing inside the axes object.


Here is an example:

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

xs = np.linspace(0, 10, 100)
ys = np.sin(xs)

fig, ax = plt.subplots()
ax.plot(xs, ys)
```

```
[<matplotlib.lines.Line2D at 0x7fc82c882290>]
```




Below is a more complete example. We can see that the object oriented nature, for example the `ax` object, makes it possible to target our formatting instructions to that `ax` object. This becomes particularly useful if we have more than one axes object in the same figure.

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

xs = np.linspace(0, 10, 100)
ys = np.sin(xs)

fig, ax = plt.subplots(figsize=(6, 4))
ax.plot(xs, ys, 'x-', linewidth=2, color='orange')

ax.grid('on')
ax.set_xlabel('x')
ax.set_ylabel('y=f(x)')
fig.savefig("pyplot-demo2.pdf")
```



15.3.1 Histograms

The program below demonstrates how to create histograms from statistical data with matplotlib.

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


# create the data
mu, sigma = 100, 15
x = mu + sigma*np.random.randn(10000)

# create the figure and axes objects
fig, ax = plt.subplots()

# the histogram of the data
n, bins, patches = ax.hist(x, 50, density=1, facecolor='green', alpha=0.75)

# add a 'best fit' line
y = scipy.stats.norm.pdf(bins, mu, sigma)
l = ax.plot(bins, y, 'r--', linewidth=1)

# annotate the plot
ax.set_xlabel('Smarts')
ax.set_ylabel('Probability')
ax.set_title(r'$\mathrm{Histogram}$ of $IQ:$ \ $\mu=100,$ \ $\sigma=15$')
ax.axis([40, 160, 0, 0.03])
ax.grid(True)
```



Do not try to understand every single command in this file: some are rather specialised and have not been covered in this text. The intention is to provide a few examples to show what can – in principle – be done with Matplotlib. If you need a plot like this, the expectation is that you will need to experiment and possibly learn a bit more about Matplotlib.

15.3.2 Visualising matrix data

The program below demonstrates how to create a bitmap-plot of the entries of a matrix.

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

# Helper function (from https://github.com/matplotlib/matplotlib/blob/master/lib/
# matplotlib/mlab.py
# as of August 2018)
def bivariate_normal(X, Y, sigmax=1.0, sigmay=1.0,
                    mux=0.0, muy=0.0, sigmaxy=0.0):
    """
    Bivariate Gaussian distribution for equal shape *X*, *Y*.
    See `bivariate normal
    <https://mathworld.wolfram.com/BivariateNormalDistribution.html>`_
    at mathworld.
    """
    Xmu = X - mux
    Ymu = Y - muy

    rho = sigmaxy / (sigmax*sigmay)
```

(continues on next page)

(continued from previous page)

```

z = Xmu**2 / sigmax**2 + Ymu**2 / sigmay**2 - 2*rho*Xmu*Ymu/(sigmax*sigmay)
denom = 2*np.pi*sigmax*sigmay*np.sqrt(1-rho**2)
return np.exp(-z/(2*(1-rho**2))) / denom

# create matrix Z that contains some interesting data
delta = 0.1
x = y = np.arange(-3.0, 3.0, delta)
X, Y = np.meshgrid(x, y)
Z = bivariate_normal(X, Y, 3.0, 1.0, 0.0, 0.0)

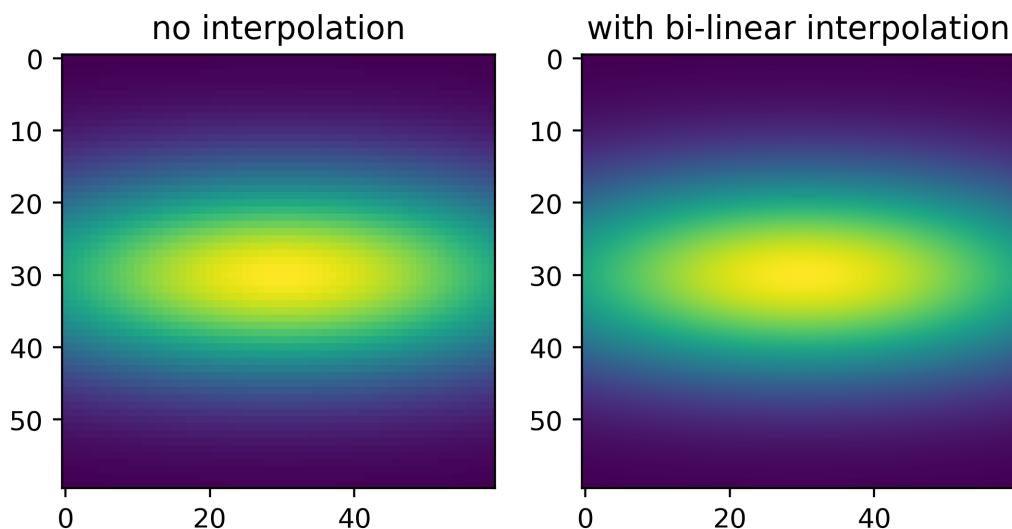
# display the 'raw' matrix data of Z in one set of axis
fig, axes = plt.subplots(ncols=2)
ax0, ax1 = axes
ax0.imshow(Z, interpolation='nearest')
ax0.set_title("no interpolation")

# display the data interpolated in other set of axis
im = ax1.imshow(Z, interpolation='bilinear')
ax1.set_title("with bi-linear interpolation")

fig.suptitle("imshow example")
fig.savefig("pylabimshow.pdf")

```

imshow example



To use different colourmaps, we make use of the `matplotlib.cm` module (where `cm` stands for Colour Map). The code below demonstrates how we can select colourmaps from the set of already provided maps, and how we can modify them (here by reducing the number of colours in the map). The last example mimics the behaviour of the more sophisticated `contour` command that also comes with `matplotlib`.

```
import numpy as np
```

(continues on next page)

(continued from previous page)

```

import matplotlib.pyplot as plt
import matplotlib.cm as cm      # Colour map submodule

# create matrix Z that contains some data interesting data
delta = 0.025
x = y = np.arange(-3.0, 3.0, delta)
X, Y = np.meshgrid(x, y)
Z = bivariate_normal(X, Y, 3.0, 1.0, 0.0, 0.0)

# Create a matrix of axes with 2 rows and 3 columns
fig, axes = plt.subplots(nrows=2, ncols=3)

ax = axes[0, 0]
ax.imshow(Z, cmap=cm.viridis) # viridis colourmap
ax.set_title("colourmap jet")

ax = axes[0, 1]
ax.imshow(Z, cmap=cm.viridis_r) # reverse viridis colourmap
ax.set_title("colourmap jet_r")

ax = axes[0, 2]
ax.imshow(Z, cmap=cm.gray)
ax.set_title("colourmap gray")

ax = axes[1, 0]
ax.imshow(Z, cmap=cm.hsv)
ax.set_title("colourmap hsv") # this one is periodic

ax = axes[1, 1]
ax.imshow(Z, cmap=cm.plasma)
ax.set_title("colourmap plasma")

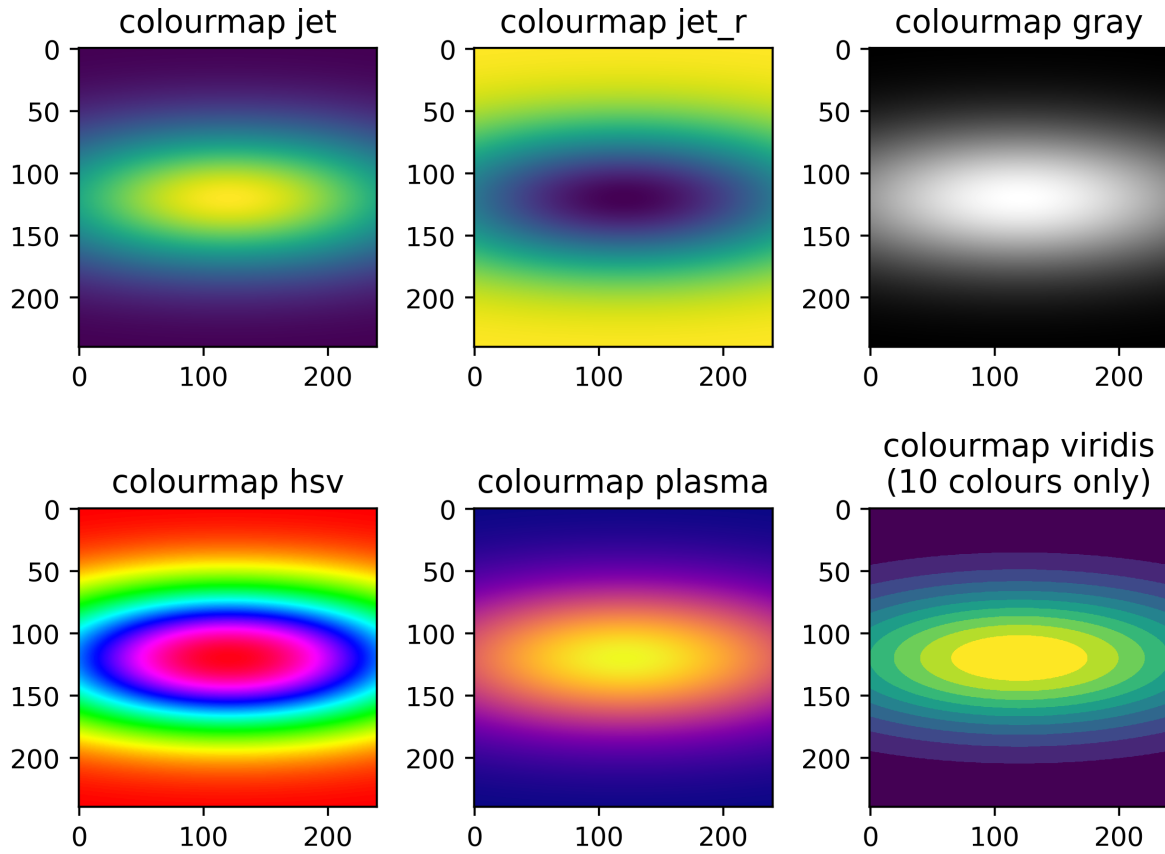
ax = axes[1, 2]
# make isolines by reducing number of colours to 10
mycmap = cm.get_cmap('viridis', 10) # 10 discrete colors
ax.imshow(Z, cmap=mycmap)
ax.set_title("colourmap viridis\n(10 colours only)")
fig.tight_layout() # avoid overlap of titles and axis labels
fig.savefig("pylabimshowcm.pdf")

```

```

/tmp/ipykernel_234/332518464.py:36: MatplotlibDeprecationWarning: The get_cmap_
↳function was deprecated in Matplotlib 3.7 and will be removed two minor releases_
↳later. Use ``matplotlib.colormaps[name]`` or ``matplotlib.colormaps.get_
↳cmap(obj)`` instead.
    mycmap = cm.get_cmap('viridis', 10) # 10 discrete colors

```

15.3.3 What colour map to choose?


It is a non-trivial question which colour map one should use. There is a useful discussion as part of the [matplotlib documentation](#).

By default a ‘perceptually uniform’ colourmap is a good choice: the perception of the colours follows the values we try to represent. Examples are “viridis”, “plasma”, “inferno”, “magma”, “cividis”.

This is a complex topic in its own right.

15.3.4 Plots of $z = f(x, y)$ and other features of Matplotlib

Matplotlib has a large number of features and can create all the standard (1d and 2d) plots such as histograms, pie charts, scatter plots, 2d-intensity plots (i.e. $z = f(x, y)$) and contour lines) and much more. The figure below shows such an example ([contour_demo.py] (https://matplotlib.org/stable/gallery/images_contours_and_fields/contour_demo.html)).



Some support for 3d plots is also available: <https://matplotlib.org/stable/gallery/index.html#d-plotting>

15.3.5 How to learn how to use Matplotlib?

A common strategy is to scan the examples at <https://matplotlib.org/stable/gallery> to find a plot similar to the desired one, and then to modify the given example code. In the process, it may be worth learning more about the commands used in the example through reading the documentation.

15.4 Visual Python

Visual Python is a Python module that makes it fairly easy to create and animate three-dimensional scenes.

Further information:

- The Visual Python home page
- The Visual Python documentation (explaining all objects with all their parameters)


Short videos introducing Visual Python:

- Shawn Weatherford, Jeff Polak (students of Ruth Chabay): <https://www.youtube.com/vpythonvideos>

15.4.1 Basics, rotating and zooming

Here is an example showing how to create a red and a blue sphere at two different positions together with a flat box (vpythondemo1.py):

```
import visual
sphere1 = visual.sphere(pos=[0, 0, 0], color=visual.color.blue)
sphere2 = visual.sphere(pos=[5, 0, 0], color=visual.color.red, radius=2)
base = visual.box(pos=(0, -2, 0), length=8, height=0.1, width=10)
```



Once you have created such a visual python scene, you can

- rotate the scene by pressing the right mouse button and moving the mouse
- zoom in and out by pressing the middle mouse button (this could be the wheel) and moving the mouse up and down. (On some (Windows?) installations, one has to press the left and the right mouse button simultaneously and then move the mouse up and down to zoom.)

15.4.2 Setting the frame rate for animations

A particular strength of Visual Python is its ability to display time-dependent data:

- A very useful command is the `rate()` command which ensures that a loop is only executed at a certain frame rate. Here is an example printing exactly two “Hello World”s per second (`vpythondemo2.py`):

```
import visual

for i in range(10):
    visual.rate(2)
    print("Hello World (0.5 seconds per line)")
```

- All Visual Python objects (such as the spheres and the box in the example above) have a `.pos` attribute which contains the position (of the centre of the object [sphere,box] or one end of the object [cylinder, helix]). Here is an example showing a sphere moving up and down (`vpythondemo3.py`):

```
import visual, math


ball = visual.sphere()
box = visual.box( pos=[0,-1,0], width=4, length=4, height=0.5 )

#tell visual not to automatically scale the image
visual.scene.autoscale = False

for i in range(1000):
    t = i*0.1
    y = math.sin(t)

    #update the ball's position
    ball.pos = [0, y, 0]

    #ensure we have only 24 frames per second
    visual.rate(24)
```



15.4.3 Tracking trajectories

You can track the trajectory of an object using a “curve”. The basic idea is to append positions to that curve object as demonstrated in this example (`vpythondemo4.py`):

```
import visual, math

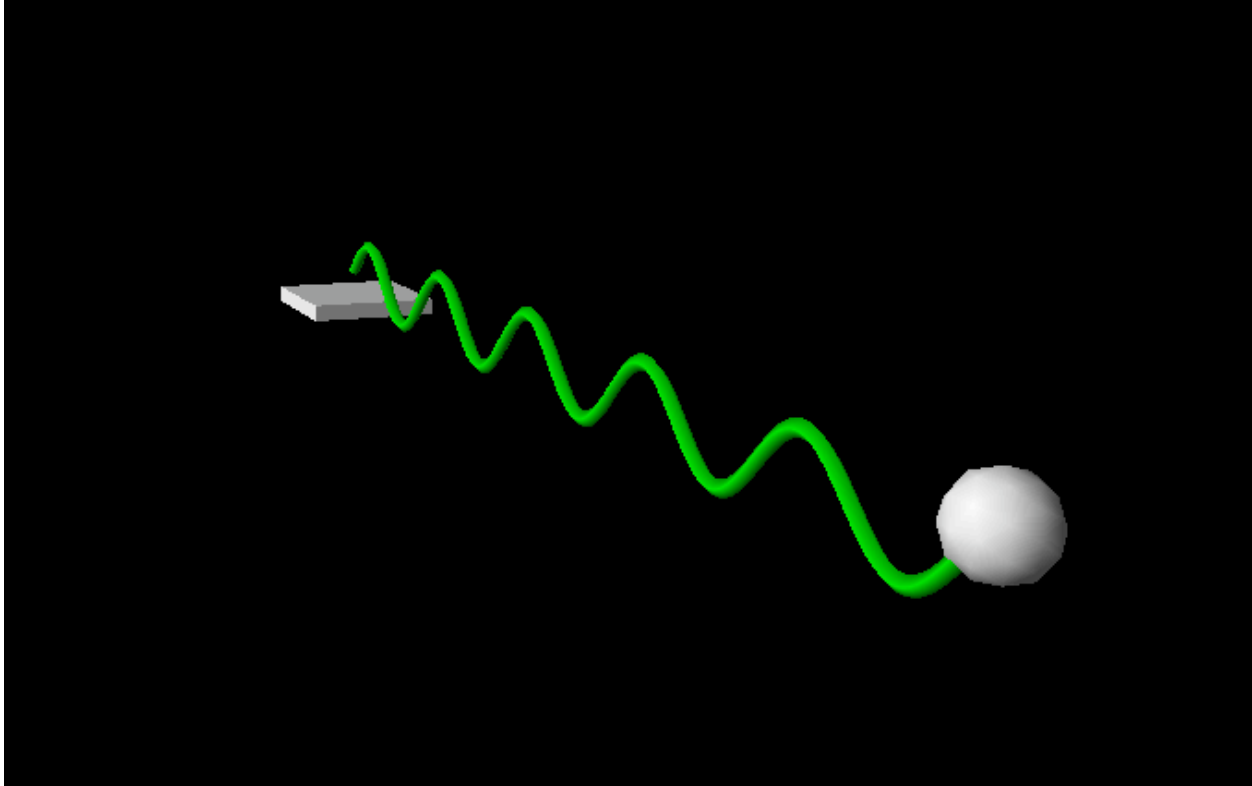
ball = visual.sphere()
box = visual.box( pos=[0,-1,0], width=4, length=4, height=0.5 )
trace=visual.curve( radius=0.2, color=visual.color.green)

for i in range(1000):
    t = i*0.1
    y = math.sin(t)

    #update the ball's position
    ball.pos = [t, y, 0]

    trace.append( ball.pos )

    #ensure we have only 24 frames per second
    visual.rate(24)
```



As with most visual Python objects, you can specify the colour of the curve (also per appended element!) and the radius.

15.4.4 Connecting objects (Cylinders, springs, ...)

Cylinders and helices can be used to “connect” two objects. In addition to the `pos` attribute (which stores the position of one end of the object), there is also an `axis` attribute which stores the vector pointing from `pos` to the other end of the object. Here is an example showing this for a cylinder: (`vpythondemo5.py`):

```
import visual, math

ball1 = visual.sphere( pos = (0,0,0), radius=2 )
ball2 = visual.sphere( pos = (5,0,0), radius=2 )
connection = visual.cylinder(pos = ball1.pos, \
                             axis = ball2.pos - ball1.pos)

for t in range(100):
    #move ball2
    ball2.pos = (-t,math.sin(t),math.cos(t))

    #keep cylinder connection between ball1 and ball2
    connection.axis = ball2.pos - ball1.pos

visual.rate(24)
```


15.4.5 3d vision

If you have access to “anaglyphic” (*i.e.* colored) glasses (best red-cyan but red-green or red-blue works as well), then you can switch visual python into this stereo mode by adding these two lines to the beginning of your program:

```
visual.scene.stereo='redcyan'
visual.scene.stereodepth=1
```

Note the effect of the `stereodepth` parameter:

- `stereodepth=0`: 3d scene “inside” the screen (default)
- `stereodepth=1`: 3d scene at screen surface (this often looks best)
- `stereodepth=2`: 3d scene sticking out of the screen



15.5 Visualising higher dimensional data (VTK)

Often, we need to understand data defined at 3d positions in space. The data itself is often a scalar field (such as temperature) or a 3d vector (such as velocity or magnetic field), or occasionally a tensor. For example for a 3d-vector field f defined in 3d-space ($\vec{f}(\vec{x})$ where $\vec{x} \in \mathbb{R}^3$ and $\vec{f}(\vec{x}) \in \mathbb{R}^3$) we could draw a 3d-arrow at every (grid) point in space. It is common for these data sets to be time dependent.

The probably most commonly used library in Science and Engineering to visualise such data sets is probably VTK, the Visualisation ToolKit (<https://vtk.org>). This is a substantial C++ library with interfaces to high level languages, including Python.

One can either call these routines directly from Python code, or write the data to disk in a format that the VTK library can read (so called vtk data files), and then use stand-alone programme such as Mayavi, ParaView and VisIt to read these


data files and manipulate them (often with a GUI). All three of these are using the VTK library internally, and can read vtk data files.


This package is very well suited to visualise static and time-dependent 2d and 3d-fields (scalar, vector and tensor fields). Two examples are shown below.

They can be used as stand-alone executables with a GUI to visualise VTK files. It can also be scripted from a Python program, or used interactively from a Python session.

15.5.1 Mayavi, Paraview, Visit

- Mayavi Home page <http://code.enthought.com/pages/mayavi-project.html>
- Paraview Home page <https://www.paraview.org>
- VisIt Home page <https://wci.llnl.gov/simulation/computer-codes/visit/>





Two examples from MayaVi visualisations.

15.5.2 Writing vtk files from Python (pyvtk)

A small but powerful Python library is `pyvtk` available at <https://github.com/pearu/pyvtk>. This allows to create vtk files from Python data structures very easily.

Given a finite element mesh or a finite difference data set in Python, one can use `pyvtk` to write such data into files, and then use one of the visualisation applications listed above to load the vtk files and to display and investigate them.

15.6 Further tools and developments

In addition to `matplotlib`, there are a number of other libraries with similar or related visualisation functionality.

`Plotly.py` and `Bokeh` – together with the veteran of python-based plotting `matplotlib` – form the basis for many tools that provide visualisation skills.

A beautiful summary and [categorisation](https://pyviz.org) of these and other libraries is available at <https://pyviz.org>.

15.6.1 Exploiting self-describing data for visualisation

Some libraries, such as [Pandas](#) (see also [Chapter on Pandas](#)), [Xarray](#), and [holoviews](#) make use of the idea of self-describing data to simplify the visualisation: while the data in a numpy array is ‘just’ a (multidimensional) matrix of data points, these libraries can store metadata – such as headings and coordinates – associated with these data points. We also talked about annotated or labelled data to describe the presence of such metadata.

What is the benefit of having this meta-data available? An xarray, for example, may store a 2d-array (like a numpy array) but have the metadata store that one dimension refers to the x-position and the other direction to time. The x-array object provides convenience methods to select and plot the data in the xarray.

15.6.2 The future of data visualisation

I would speculate that increasingly we will be using high-level plotting tools (such as [pandas](#), [xarray](#), [holoviews](#)) to explore data interactively.

We can see a trend in data analysis libraries that data objects can be converted to such high-level annotated data objects (such as European XFEL’s [extra-data](#) tools which can return a labelled xarray object). Other projects combine the metadata with the data in custom made objects to then provide convenience methods (such as [Ubermag](#)’s [discretisedfield](#) object).

Will we still need to learn the basics, such as the `matplotlib.pyplot` interface? Probably yes: the very least to fine tune the plots provided by these high level libraries:

15.6.3 Fine-tuning matplotlib plots that are generated by high level frame works

We show one example where [pandas](#) – as a representative for a high-level framework that can create plots – creates the plot, but we use `pyplot` commands to tailor the resulting plot.

Let’s define the [pandas](#) data series first (it is not important to understand the details of this now):


```
import pandas as pd
s = pd.Series(data=[10, 20, 1], index=['bananas', 'oranges', 'potatoes'])
s
```

```
bananas    10
oranges    20
potatoes     1
dtype: int64
```

We can use a convenience method from [pandas](#) to create a plot of the data series:

```
s.plot.bar()
```

```
<Axes: >
```



Note how the bar chart is labelled appropriately: the metadata (here the labels 'bananas', 'oranges' and 'potatoes') have been used to label the x-axis in the plot.

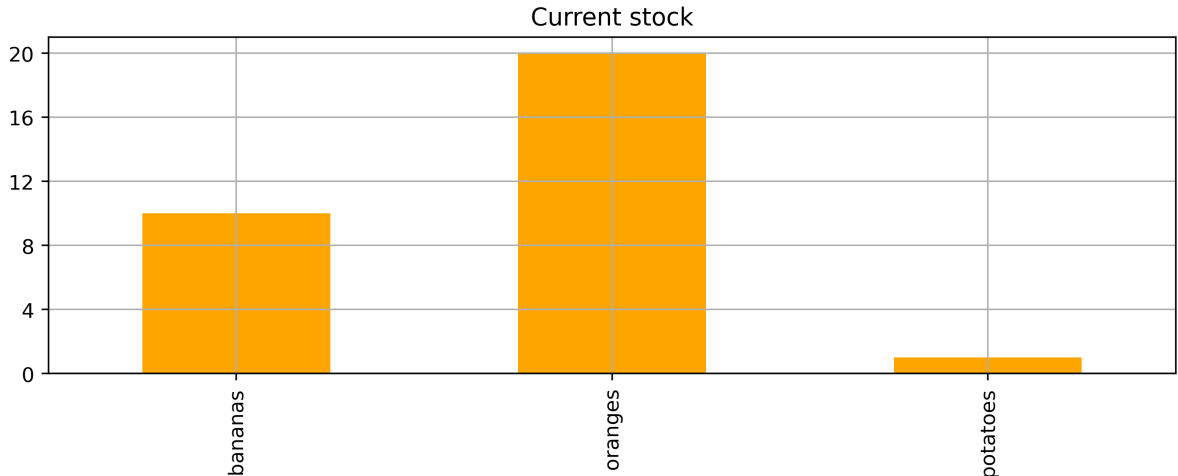
If we want to change this plot, the following strategy works, and is supported by other high-level frameworks as well:

- create an axes (and figure) object
- pass the axes object to the high level plotting framework
- use the axes object (and figure) to finetune the plot

The following example shows how to add a title, customise the labelling of the y-axis, and add a grid to the plot, and change the size of the figure to be 10 inches by 3 inches:

```
import matplotlib.pyplot as plt
fig, ax = plt.subplots(figsize=(10, 3))

s.plot.bar(ax=ax, color='orange')
ax.set_title("Current stock")
ax.set_yticks(range(0, 21, 4));
ax.grid('on')
```



15.7 Jupyter Notebooks

The [Jupyter Notebook](#) has become a central tool for interactive data exploration and data analysis. I would go so far to say that most data scientists will see the Jupyter notebook as the default place to start a data exploration, analysis and machine learning project.

Why is this so? The combination of annotation, code snippets, inlined results from computation or visualisation and the automatic logging of these steps in a notebook file can be of great use for research and development activities. A slightly longer summary is [available here](#).

Some recent publications on the topic:

- Brian Granger, Fernando Pérez. *Jupyter: Thinking and Storytelling With Code and Data*, Computing in Science & Engineering, vol. 23, no. 2, pp. 7-14, 1 March-April 2021, doi: 10.1109/MCSE.2021.3059263 [Authorea preprint](#) (2021)
- Hans Fangohr, Marijan Beg, et al, *Data exploration and analysis with Jupyter notebooks*, Proceedings of the 17th International Conference on Accelerator and Large Experimental Physics Control Systems ICALEPCS2019, TUCPR02, doi: 10.18429/JACoW-ICALEPCS2019-TUCPR02 ([pdf](#)) (2020)
- Marijan Beg; Juliette Belin; Thomas Kluyver; Alexander Kononov; Min Ragan-Kelley; Nicolas Thiery; Hans Fangohr. *Using Jupyter for Reproducible Scientific Workflows* in Computing in Science & Engineering, vol. 23, no. 2, pp. 36-46, 1 March-April 2021, doi: 10.1109/MCSE.2021.3052101 [arXiv preprint](#) (2021)

NUMERICAL METHODS USING PYTHON (SCIPY)

16.1 Overview

The core Python language (including the standard libraries) provide enough functionality to carry out computational research tasks. However, there are dedicated (third-party) Python libraries that provide extended functionality which

- provide numerical tools for frequently occurring tasks
- which are convenient to use
- and are more efficient in terms of CPU time and memory requirements than using the code Python functionality alone.

We list three such modules in particular:

- The `numpy` module provides a data type specialised for “number crunching” of vectors and matrices (this is the `array` type provided by “`numpy`” as introduced in [14-numpy.ipynb](#)), and linear algebra tools.
- The `matplotlib` package (also known as `pylab`) provides plotting and visualisation capabilities (see [15-visualising-data.ipynb](#)) and the
- `scipy` package (SCientific PYthon) which provides a multitude of numerical algorithms and which is introduced in this chapter.

Many of the numerical algorithms available through `scipy` and `numpy` are provided by established compiled libraries which are often written in Fortran or C. They will thus execute much faster than pure Python code (which is interpreted). As a rule of thumb, we expect compiled code to be two orders of magnitude faster than pure Python code.

You can use the `help` function for each numerical method to find out more about the source of the implementation.

16.2 SciPy

`Scipy` provides many scientific computing functions and is generally complementary to the the functionality of `numpy`.

First we need to import `scipy`:

```
import scipy
```

The `scipy` package provides information about its own structure when we use the `help` command:

```
help(scipy)
```

The output is very long, so we’re showing just a part of it here:

```

cluster          --- Vector Quantization / Kmeans
fft              --- Discrete Fourier transforms
fftpack         --- Legacy discrete Fourier transforms
integrate        --- Integration routines
interpolate      --- Interpolation Tools
io              --- Data input and output
linalg          --- Linear algebra routines
linalg.blas     --- Wrappers to BLAS library
linalg.lapack   --- Wrappers to LAPACK library
misc            --- Various utilities that don't have
                  another home.
ndimage         --- n-dimensional image package
odr             --- Orthogonal Distance Regression
optimize        --- Optimization Tools
signal         --- Signal Processing Tools
signal.windows  --- Window functions
sparse          --- Sparse Matrices
sparse.linalg   --- Sparse Linear Algebra
sparse.linalg.dsolve --- Linear Solvers
sparse.linalg.dsolve.umfpack --- :Interface to the UMFPACK library:
                                Conjugate Gradient Method (LOBPCG)
sparse.linalg.eigen --- Sparse Eigenvalue Solvers
sparse.linalg.eigen.lobpcg --- Locally Optimal Block Preconditioned
                                Conjugate Gradient Method (LOBPCG)
spatial         --- Spatial data structures and algorithms
special         --- Special functions
stats          --- Statistical Functions
    
```

If we are looking for an algorithm to integrate a function, we might explore the `integrate` package:

```
import scipy.integrate
```

```
scipy.integrate?
```

produces:

```

=====
Integration and ODEs (:mod:`scipy.integrate`)
=====

.. currentmodule:: scipy.integrate

Integrating functions, given function object
=====

.. autosummary::
   :toctree: generated/

   quad          -- General purpose integration
   quad_vec      -- General purpose integration of vector-valued functions
   dblquad       -- General purpose double integration
   tplquad       -- General purpose triple integration
   nquad         -- General purpose n-dimensional integration
   fixed_quad    -- Integrate func(x) using Gaussian quadrature of order n
   quadrature    -- Integrate with given tolerance using Gaussian quadrature
   romberg       -- Integrate func using Romberg integration
    
```

(continues on next page)

(continued from previous page)

```

quad_explain  -- Print information for use of quad
newton_cotes  -- Weights and error coefficient for Newton-Cotes integration
IntegrationWarning -- Warning on issues during integration

Integrating functions, given fixed samples
=====

.. autosummary::
   :toctree: generated/

   trapz      -- Use trapezoidal rule to compute integral.
   cumtrapz   -- Use trapezoidal rule to cumulatively compute integral.
   simps      -- Use Simpson's rule to compute integral from samples.
   romb       -- Use Romberg Integration to compute integral from
               -- (2**k + 1) evenly-spaced samples.

.. seealso::

   :mod:`scipy.special` for orthogonal polynomials (special) for Gaussian
   quadrature roots and weights for other weighting factors and regions.

Solving initial value problems for ODE systems
=====

The solvers are implemented as individual classes which can be used directly
(low-level usage) or through a convenience function.

.. autosummary::
   :toctree: generated/

   solve_ivp  -- Convenient function for ODE integration.
   RK23       -- Explicit Runge-Kutta solver of order 3(2).
   RK45       -- Explicit Runge-Kutta solver of order 5(4).
   DOP853     -- Explicit Runge-Kutta solver of order 8.
   Radau      -- Implicit Runge-Kutta solver of order 5.
   BDF        -- Implicit multi-step variable order (1 to 5) solver.
   LSODA      -- LSODA solver from ODEPACK Fortran package.
   OdeSolver  -- Base class for ODE solvers.
   DenseOutput -- Local interpolant for computing a dense output.
   OdeSolution -- Class which represents a continuous ODE solution.

```

The following sections show examples which demonstrate how to employ the algorithms provided by `scipy`.

16.3 Numerical integration

Scientific Python provides a number of integration routines. A general purpose tool to solve integrals I of the kind

$$I = \int_a^b f(x)dx$$

is provided by the `quad()` function of the `scipy.integrate` module.

It takes as input arguments the function $f(x)$ to be integrated (the “integrand”), and the lower and upper limits a and b . It returns two values (in a tuple): the first one is the computed results and the second one is an estimation of the numerical error of that result.

Here is an example: which produces this output:

```
# NBVAL_IGNORE_OUTPUT
from math import cos, exp, pi
from scipy.integrate import quad

# function we want to integrate
def f(x):
    return exp(cos(-2 * x * pi)) + 3.2

# call quad to integrate f from -2 to 2
res, err = quad(f, -2, 2)

print("The numerical result is {:f} (+-{:g})"
      .format(res, err))
```

```
The numerical result is 17.864264 (+-1.55117e-11)
```

Note that `quad()` takes optional parameters `epsabs` and `epsrel` to increase or decrease the accuracy of its computation. (Use `help(quad)` to learn more.) The default values are `epsabs=1.5e-8` and `epsrel=1.5e-8`. For the next exercise, the default values are sufficient.

16.3.1 Exercise: integrate a function

1. Using `scipy`'s `quad` function, write a program that solves the following integral numerically: $I = \int_0^1 \cos(2\pi x) dx$.
2. Find the analytical integral and compare it with the numerical solution.
3. Why is it important to have an estimate of the accuracy (or the error) of the numerical integral?

16.3.2 Exercise: plot before you integrate

It is good practice to plot the integrand function to check whether it is “well behaved” before you attempt to integrate. Singularities (i.e. x values where the $f(x)$ tends towards minus or plus infinity) or other irregular behaviour (such as $f(x) = \sin(\frac{1}{x})$ close to $x = 0$) are difficult to handle numerically.

1. Write a function with name `plotquad` which takes the same arguments as the `quad` command (i.e. f , a and b) and which
 - (i) creates a plot of the integrand $f(x)$ and
 - (ii) computes the integral numerically using the `quad` function. The return values should be as for the `quad` function.

```
%matplotlib inline
# settings for jupyter book: svg for html version, high-resolution png for pdf
import matplotlib_inline
matplotlib_inline.backend_inline.set_matplotlib_formats('svg', 'png')
# from IPython.display import set_matplotlib_formats
# set_matplotlib_formats('svg', 'png')
import matplotlib as mpl
mpl.rcParams['figure.dpi'] = 400
```


16.4 Solving Ordinary Differential Equations (ODEs)

To solve an ordinary differential equation of the type $\frac{dy}{dt}(t) = f(t, y)$

with a given $y(t_0) = y_0$, we can use `scipy`'s `solve_ivp` function. Here is a (self explaining) example program (`usesolve_ivp.py`) to find

$$y(t) \quad \text{for } t \in [0, 2]$$

given this differential equation: $\frac{dy}{dt}(t) = -2yt$ with $y(0) = 1$.

```
from scipy.integrate import solve_ivp
import numpy as np


def f(t, y):
    """this is the rhs of the ODE to integrate, i.e. dy/dt=f(y,t)"""
    return -2 * y * t

y0 = [1]           # initial value y0=y(t0)
t0 = 0            # integration limits for t: start at t0=0
tf = 2           # and finish at tf=2

sol = solve_ivp(fun=f, t_span=[t0, tf], y0=y0) # computation of SOLution

import pylab      # plotting of results
pylab.plot(sol.t, sol.y[0], 'o-')
pylab.xlabel('t'); pylab.ylabel('y(t)')
```

```
Text(0, 0.5, 'y(t)')
```



We have not given the `solve_ivp` command any guidance for which values of t we would like to know the solution $y(t)$: we have only specified that $t_0 = 0$ and that we would like to know the solution between $t_0 = 0$ and $t_y = 2$. The solver itself has determined the number of required function evaluations, and returns the corresponding values in `sol.t` and `sol.y[0]`.


We can obtain more data points in a number of ways:

1. Increase the default error tolerance. The relative tolerance (`rtol`) and absolute tolerance (`atol`) default to $1e-3$ each. If we increase them, we typically enforce the use of a larger number of intermediate points:

```
sol = solve_ivp(fun=f, t_span=[t0, tf], y0=y0, atol=1e-8, rtol=1e-8)

pylab.plot(sol.t, sol.y[0], '.')
pylab.xlabel('t'); pylab.ylabel('y(t)')
```

```
Text(0, 0.5, 'y(t)')
```




2. We can also prescribe the precise locations for which we like to know the solutions $y(t)$:

```
y0 = [1]           # initial value
t0 = 0             # integration limits for t
tf = 2
ts = np.linspace(t0, tf, 100) # 100 points between t0 and tf

sol = solve_ivp(fun=f, t_span=[t0, tf], y0=y0, t_eval=ts)

pylab.plot(sol.t, sol.y[0], '.')
pylab.xlabel('t'); pylab.ylabel('y(t)')
```

```
Text(0, 0.5, 'y(t)')
```



If we use `t_eval` - and thus request values of the solution at particular points - `solve_ivp` will not generally change the way it computes the solution, but rather use interpolation to map the way it has internally computed the solution to the values of `t` for which we would like to know the solution. There is thus no (significant) computational penalty if we use `t_eval` to get smoother looking plots.

The `solve_ivp` command returns a `OdeResult` object, which we have called `sol` in the example above.

```
type(sol)
```

```
scipy.integrate._ivp.ivp.OdeResult
```

We have already seen that the solution can be found in `sol.y` and `sol.t`:

```
type(sol.t)
```

```
numpy.ndarray
```

```
sol.t.shape
```

```
(100,)
```

```
type(sol.y)
```

```
numpy.ndarray
```

```
sol.y.shape
```

```
(1, 100)
```

Because `solve_ivp` is designed to integrate *systems* of ordinary differential equations, `sol.y` is a matrix, where each row contains the values for one degree of freedom. In our simple example above, we only have one degree of freedom (y). This is the reason, why we had to use `sol.y[0]` to access the solution values.

Other interesting attributes of the `OdeResult` object are the number of function evaluations that were necessary (where the function is the function `f` which computes the right hand side of the ODE).

```
sol.nfev
```

```
68
```

There is also a human-readable string, providing - for this example - a re-assuring message:

```
sol.message
```

```
'The solver successfully reached the end of the integration interval.'
```

A machine readable status is available in the `sol.status` attribute (0 is good):

```
sol.status
```

```
0
```

The `solve_ivp` command takes a number of optional parameters - we have already seen `atol` and `rtol` to change the default error tolerance of the integration. We can use the `help` command to explore these. The help string also explains the attributes of the solution object `OdeResult` in more detail:

```
help(scipy.integrate.solve_ivp)
```

will show:

```
Help on function solve_ivp in module scipy.integrate._ivp.ivp:
```

```
solve_ivp(fun, t_span, y0, method='RK45', t_eval=None, dense_output=False,
events=None, vectorized=False, args=None, **options)
```

```
    Solve an initial value problem for a system of ODEs.
```

```
    This function numerically integrates a system of ordinary differential
    equations given an initial value::
```

```
        dy / dt = f(t, y)
        y(t0) = y0
```

```
    Here t is a 1-D independent variable (time), y(t) is an
    N-D vector-valued function (state), and an N-D
```

(continues on next page)

(continued from previous page)

vector-valued function $f(t, y)$ determines the differential equations. The goal is to find $y(t)$ approximately satisfying the differential equations, given an initial value $y(t_0)=y_0$.

Some of the solvers support integration in the complex domain, but note that for stiff ODE solvers, the right-hand side must be complex-differentiable (satisfy Cauchy-Riemann equations [11]). To solve a problem in the complex domain, pass y_0 with a complex data type. Another option always available is to rewrite your problem for real and imaginary parts separately.

Parameters

fun : callable

Right-hand side of the system. The calling signature is ``fun(t, y)``. Here `t` is a scalar, and there are two options for the ndarray `y`: It can either have shape (n,); then `fun` must return array_like with shape (n,). Alternatively, it can have shape (n, k); then `fun` must return an array_like with shape (n, k), i.e., each column corresponds to a single column in `y`. The choice between the two options is determined by `vectorized` argument (see below). The vectorized implementation allows a faster approximation of the Jacobian by finite differences (required for stiff solvers).

t_span : 2-tuple of floats

Interval of integration (t0, tf). The solver starts with t=t0 and integrates until it reaches t=tf.

y0 : array_like, shape (n,)

Initial state. For problems in the complex domain, pass `y0` with a complex data type (even if the initial value is purely real).

method : string or `OdeSolver`, optional

Integration method to use:

- * 'RK45' (default): Explicit Runge-Kutta method of order 5(4) [1]_. The error is controlled assuming accuracy of the fourth-order method, but steps are taken using the fifth-order accurate formula (local extrapolation is done). A quartic interpolation polynomial is used for the dense output [2]_. Can be applied in the complex domain.
- * 'RK23': Explicit Runge-Kutta method of order 3(2) [3]_. The error is controlled assuming accuracy of the second-order method, but steps are taken using the third-order accurate formula (local extrapolation is done). A cubic Hermite polynomial is used for the dense output. Can be applied in the complex domain.
- * 'DOP853': Explicit Runge-Kutta method of order 8 [13]_. Python implementation of the "DOP853" algorithm originally written in Fortran [14]_. A 7-th order interpolation polynomial accurate to 7-th order is used for the dense output. Can be applied in the complex domain.
- * 'Radau': Implicit Runge-Kutta method of the Radau IIA family of order 5 [4]_. The error is controlled with a third-order accurate embedded formula. A cubic polynomial which satisfies the collocation conditions is used for the dense output.
- * 'BDF': Implicit multi-step variable-order (1 to 5) method based on a backward differentiation formula for the derivative approximation [5]_. The implementation follows the one described in [6]_. A quasi-constant step scheme is used and accuracy is

(continues on next page)

(continued from previous page)

enhanced using the NDF modification. Can be applied in the complex domain.

- * 'LSODA': Adams/BDF method with automatic stiffness detection and switching [7]_, [8]_. This is a wrapper of the Fortran solver from ODEPACK.

Explicit Runge-Kutta methods ('RK23', 'RK45', 'DOP853') should be used for non-stiff problems and implicit methods ('Radau', 'BDF') for stiff problems [9]_. Among Runge-Kutta methods, 'DOP853' is recommended for solving with high precision (low values of `rtol` and `atol`).

If not sure, first try to run 'RK45'. If it makes unusually many iterations, diverges, or fails, your problem is likely to be stiff and you should use 'Radau' or 'BDF'. 'LSODA' can also be a good universal choice, but it might be somewhat less convenient to work with as it wraps old Fortran code.

You can also pass an arbitrary class derived from `OdeSolver` which implements the solver.

`t_eval` : array_like or None, optional

Times at which to store the computed solution, must be sorted and lie within `t_span`. If None (default), use points selected by the solver.

`dense_output` : bool, optional

Whether to compute a continuous solution. Default is False.

`events` : callable, or list of callables, optional

Events to track. If None (default), no events will be tracked.

Each event occurs at the zeros of a continuous function of time and state. Each function must have the signature ``event(t, y)`` and return a float. The solver will find an accurate value of `t` at which ``event(t, y(t)) = 0`` using a root-finding algorithm. By default, all zeros will be found. The solver looks for a sign change over each step, so if multiple zero crossings occur within one step, events may be missed. Additionally each `event` function might have the following attributes:

`terminal`: bool, optional

Whether to terminate integration if this event occurs.

Implicitly False if not assigned.

`direction`: float, optional

Direction of a zero crossing. If `direction` is positive, `event` will only trigger when going from negative to positive, and vice versa if `direction` is negative. If 0, then either direction will trigger event. Implicitly 0 if not assigned.

You can assign attributes like ``event.terminal = True`` to any function in Python.

`vectorized` : bool, optional

Whether `fun` is implemented in a vectorized fashion. Default is False.

`args` : tuple, optional

Additional arguments to pass to the user-defined functions. If given, the additional arguments are passed to all user-defined functions.

So if, for example, `fun` has the signature ``fun(t, y, a, b, c)`` , then `jac` (if given) and any event functions must have the same signature, and `args` must be a tuple of length 3.

`options`

Options passed to a chosen solver. All options available for already

(continues on next page)

(continued from previous page)

implemented solvers are listed below.

`first_step` : float or None, optional
Initial step size. Default is `None` which means that the algorithm should choose.

`max_step` : float, optional
Maximum allowed step size. Default is `np.inf`, i.e., the step size is not bounded and determined solely by the solver.

`rtol`, `atol` : float or array_like, optional
Relative and absolute tolerances. The solver keeps the local error estimates less than `atol + rtol * abs(y)`. Here `rtol` controls a relative accuracy (number of correct digits). But if a component of `y` is approximately below `atol`, the error only needs to fall within the same `atol` threshold, and the number of correct digits is not guaranteed. If components of `y` have different scales, it might be beneficial to set different `atol` values for different components by passing array_like with shape `(n,)` for `atol`. Default values are `1e-3` for `rtol` and `1e-6` for `atol`.

`jac` : array_like, sparse_matrix, callable or None, optional
Jacobian matrix of the right-hand side of the system with respect to `y`, required by the 'Radau', 'BDF' and 'LSODA' method. The Jacobian matrix has shape `(n, n)` and its element `(i, j)` is equal to `d f_i / d y_j`. There are three ways to define the Jacobian:

- * If array_like or sparse_matrix, the Jacobian is assumed to be constant. Not supported by 'LSODA'.
- * If callable, the Jacobian is assumed to depend on both `t` and `y`; it will be called as `jac(t, y)`, as necessary. For 'Radau' and 'BDF' methods, the return value might be a sparse matrix.
- * If None (default), the Jacobian will be approximated by finite differences.

It is generally recommended to provide the Jacobian rather than relying on a finite-difference approximation.

`jac_sparsity` : array_like, sparse matrix or None, optional
Defines a sparsity structure of the Jacobian matrix for a finite-difference approximation. Its shape must be `(n, n)`. This argument is ignored if `jac` is not `None`. If the Jacobian has only few non-zero elements in *each* row, providing the sparsity structure will greatly speed up the computations [10]_. A zero entry means that a corresponding element in the Jacobian is always zero. If None (default), the Jacobian is assumed to be dense.
Not supported by 'LSODA', see `lband` and `uband` instead.

`lband`, `uband` : int or None, optional
Parameters defining the bandwidth of the Jacobian for the 'LSODA' method, i.e., `jac[i, j] != 0` only for `i - lband <= j <= i + uband`. Default is None. Setting these requires your `jac` routine to return the Jacobian in the packed format: the returned array must have `n` columns and `uband + lband + 1` rows in which Jacobian diagonals are written. Specifically `jac_packed[uband + i - j, j] = jac[i, j]`. The same format is used in `scipy.linalg.solve_banded` (check for an illustration). These parameters can be also used with `jac=None` to reduce the number of Jacobian elements estimated by finite differences.

`min_step` : float, optional
The minimum allowed step size for 'LSODA' method.
By default `min_step` is zero.

(continues on next page)

(continued from previous page)

```

Returns
-----
Bunch object with the following fields defined:
t : ndarray, shape (n_points,)
    Time points.
y : ndarray, shape (n, n_points)
    Values of the solution at `t`.
sol : `OdeSolution` or None
    Found solution as `OdeSolution` instance; None if `dense_output` was
    set to False.
t_events : list of ndarray or None
    Contains for each event type a list of arrays at which an event of
    that type event was detected. None if `events` was None.
y_events : list of ndarray or None
    For each value of `t_events`, the corresponding value of the solution.
    None if `events` was None.
nfev : int
    Number of evaluations of the right-hand side.
njev : int
    Number of evaluations of the Jacobian.
nlu : int
    Number of LU decompositions.
status : int
    Reason for algorithm termination:

        * -1: Integration step failed.
        *  0: The solver successfully reached the end of `tspan`.
        *  1: A termination event occurred.

message : string
    Human-readable description of the termination reason.
success : bool
    True if the solver reached the interval end or a termination event
    occurred (`status >= 0`).

```

16.4.1 Systems of coupled ODEs

We want to show one example of two first-order ODEs that are coupled. This helps to understand why the initial value y_0 in the above example had to be provided in a list ($[y_0]$) and why the solution is $\text{sol}.y[0]$ rather than just $\text{sol}.y$.

We use the [Predator and prey example](#). Let

- $p_1(t)$ be the number of rabbits and
- $p_2(t)$ be the number of foxes at a given time t

To compute the time dependence of p_1 and p_2 :

- Assume that rabbits proliferate at a rate a . Per unit time a number ap_1 of rabbits are born.
- Assume that the number of rabbits is reduced by collisions with foxes: per unit time cp_1p_2 rabbits are eaten.
- Assume that birth rate of foxes depends only on food intake in form of rabbits.
- Assume that foxes die a natural death at a rate b .

We put this together into the system of coupled ordinary differential equations:
$$\begin{cases} \frac{dp_1}{dt} = a p_1 - c p_1 p_2 \\ \frac{dp_2}{dt} = c p_1 p_2 - b p_2 \end{cases}$$

We use the following parameters:

- rabbit birth rate $a = 0.7$
- rabbit-fox-collision rate $c = 0.007$
- fox death rate $b = 1$

We want to solve this for $p_1(t_0) = 70$ and $p_2(t_0) = 50$ as initial values, starting at $t_0 = 0$ for 30 units of time.

```
import pylab
import numpy as np
from scipy.integrate import solve_ivp

def rhs(t, y):
    a = 0.7
    c = 0.007
    b = 1
    p1 = y[0]
    p2 = y[1]

    dp1dt = a * p1 - c * p1 * p2
    dp2dt = c * p1 * p2 - b * p2

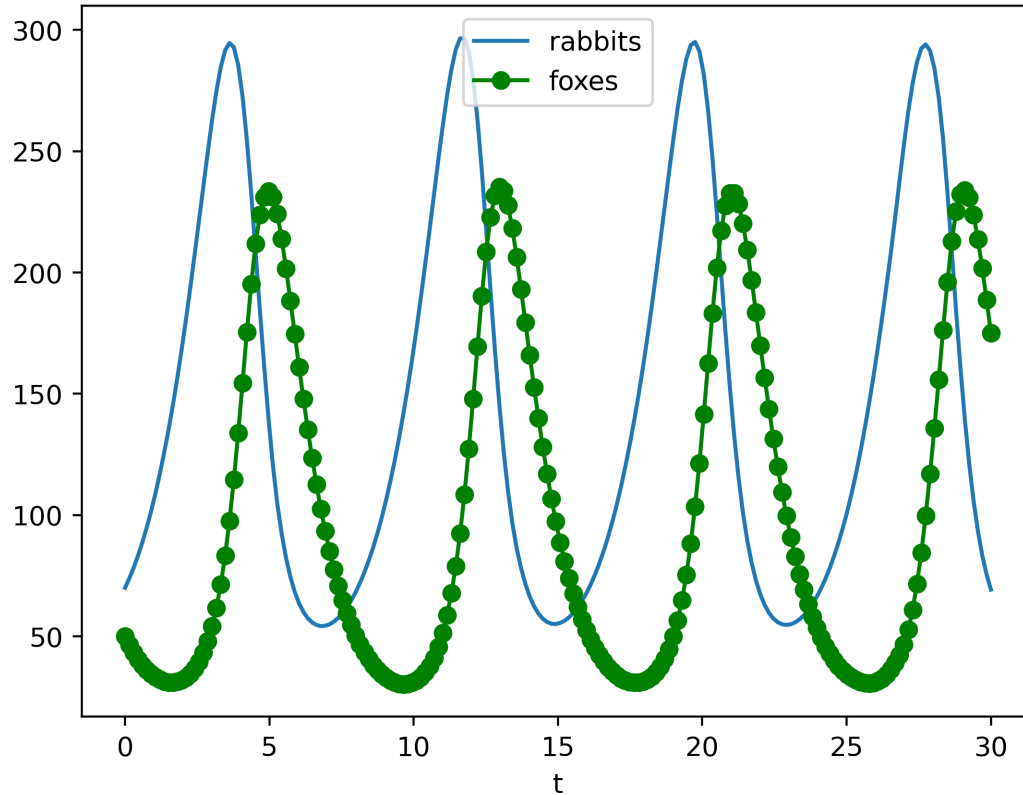
    return [dp1dt, dp2dt]

p0 = [70, 50]      # initial condition
t0 = 0
tfinal = 30
ts = np.linspace(t0, tfinal, 200)

sol = solve_ivp(rhs, [t0, tfinal], p0, t_eval=ts)

p1 = sol.y[0]
p2 = sol.y[1]

pylab.plot(sol.t, p1, label='rabbits')
pylab.plot(sol.t, p2, '-og', label='foxes')
pylab.legend()
pylab.xlabel('t')
pylab.savefig('predprey.pdf')
pylab.savefig('predprey.png')
```



16.5 Root finding

If you try to find a x such that $f(x) = 0$ then this is called *root finding*. Note that problems like $g(x) = h(x)$ fall in this category as you can rewrite them as $f(x) = g(x) - h(x) = 0$.

A number of root finding tools are available in `scipy`'s `optimize` module.

16.5.1 Root finding using the bisection method

First we introduce the `bisect` algorithm which is (i) robust and (ii) slow but conceptually very simple.

Suppose we need to compute the roots of $f(x) = x^3 - 2x^2$. This function has a (double) root at $x = 0$ (this is trivial to see) and another root which is located between $x = 1.5$ (where $f(1.5) = -1.125$) and $x = 3$ (where $f(3) = 9$). It is pretty straightforward to see that this other root is located at $x = 2$. Here is a program that determines this root numerically:

```
from scipy.optimize import bisect

def f(x):
    """returns f(x)=x^3-2x^2. Has roots at
    x=0 (double root) and x=2"""
    return x ** 3 - 2 * x ** 2

# main program starts here
x = bisect(f, 1.5, 3, xtol=1e-6)
```

(continues on next page)

(continued from previous page)

```
print("The root x is approximately x=%14.12g,\n"
      "the error is less than 1e-6." % (x))
print("The exact error is %g." % (2 - x))
```

```
The root x is approximately x= 2.00000023842,
the error is less than 1e-6.
The exact error is -2.38419e-07.
```

The `bisect()` method takes three compulsory arguments: (i) the function $f(x)$, (ii) a lower limit a (for which we have chosen 1.5 in our example) and (iii) an upper limit b (for which we have chosen 3). The optional parameter `xtol` determines the maximum error of the method.

One of the requirements of the bisection method is that the interval $[a, b]$ has to be chosen such that the function is either positive at a and negative at b , or that the function is negative at a and positive at b . In other words: a and b have to enclose a root.

16.5.2 Exercise: root finding using the bisect method

1. Write a program with name `sqrttwo.py` to determine an approximation of $\sqrt{2}$ by finding a root x of the function $f(x) = 2 - x^2$ using the bisection algorithm. Choose a tolerance for the approximation of the root of 10^{-8} .
2. Document your choice of the initial bracket $[a, b]$ for the root: which values have you chosen for a and for b and why?
3. Study the results:
 - Which value for the root x does the bisection algorithm return?
 - Compute the value of `sqrt2` using `math.sqrt(2)` and compare this with the approximation of the root. How big is the absolute error of x ? How does this compare with `xtol`?

16.5.3 Root finding using the `fsolve` function

A (often) better (in the sense of “more efficient”) algorithm than the bisection algorithm is implemented in the general purpose `fsolve()` function for root finding of (multidimensional) functions. This algorithm needs only one starting point close to the suspected location of the root (but is not guaranteed to converge).

Here is an example:

```
from scipy.optimize import fsolve

def f(x):
    return x ** 3 - 2 * x ** 2

x = fsolve(f, 3)          # one root is at x=2.0

print("The root x is approximately x=%21.19g" % x)
print("The exact error is %g." % (2 - x))
```

```
The root x is approximately x= 2.0000000000000006661
The exact error is -6.66134e-15.
```

```

/tmp/ipykernel_264/1678381742.py:8: DeprecationWarning: Conversion of an array
↳with ndim > 0 to a scalar is deprecated, and will error in future. Ensure you
↳extract a single element from your array before performing this operation.
↳(Deprecated NumPy 1.25.)
    print("The root x is approximately x=%21.19g" % x)
/tmp/ipykernel_264/1678381742.py:9: DeprecationWarning: Conversion of an array
↳with ndim > 0 to a scalar is deprecated, and will error in future. Ensure you
↳extract a single element from your array before performing this operation.
↳(Deprecated NumPy 1.25.)
    print("The exact error is %g." % (2 - x))

```

The return value[6] of `fsolve` is a numpy array of length n for a root finding problem with n variables. In the example above, we have $n = 1$.

16.6 Interpolation

Given a set of N points (x_i, y_i) with $i = 1, 2, \dots, N$, we sometimes need a function $\hat{f}(x)$ which returns $y_i = f(x_i)$ where $x = x_i$, and which in addition provides some interpolation of the data (x_i, y_i) for all x .

The function `y0 = scipy.interpolate.interp1d(x, y, kind='nearest')` does this interpolation based on splines of varying order. Note that the function `interp1d` returns a function `y0` which will then interpolate the x-y data for any given x when called as `y0(x)`.

The code below demonstrates this, and shows the different interpolation kinds.

```

import numpy as np
import scipy.interpolate
import pylab

def create_data(n):
    """Given an integer n, returns n data points
    x and values y as a numpy.array."""
    xmax = 5.
    x = np.linspace(0, xmax, n)
    y = - x**2
    #make x-data somewhat irregular
    y += 1.5 * np.random.normal(size=len(x))
    return x, y

#main program
n = 10
x, y = create_data(n)

#use finer and regular mesh for plot
xfine = np.linspace(0.1, 4.9, n * 100)
#interpolate with piecewise constant function (p=0)
y0 = scipy.interpolate.interp1d(x, y, kind='nearest')
#interpolate with piecewise linear func (p=1)
y1 = scipy.interpolate.interp1d(x, y, kind='linear')
#interpolate with piecewise constant func (p=2)
y2 = scipy.interpolate.interp1d(x, y, kind='quadratic')

pylab.plot(x, y, 'o', label='data point')
pylab.plot(xfine, y0(xfine), label='nearest')

```

(continues on next page)


(continued from previous page)

```

pylab.plot(xfine, y1(xfine), label='linear')
pylab.plot(xfine, y2(xfine), label='cubic')
pylab.legend()
pylab.xlabel('x')

```

```
Text(0.5, 0, 'x')
```



16.7 Curve fitting

We have already seen in [the numpy chapter](#) that we can fit polynomial functions through a data set using the `numpy.polyfit` function. Here, we introduce a more generic curve fitting algorithm.

Scipy provides a somewhat generic function (based on the Levenburg-Marquardt algorithm) through `scipy.optimize.curve_fit` to fit a given (Python) function to a given data set. The assumption is that we have been given a set of data with points x_1, x_2, \dots, x_N and with corresponding function values y_i and a dependence of y_i on x_i such that $y_i = f(x_i, \vec{p})$. We want to determine the parameter vector $\vec{p} = (p_1, p_2, \dots, p_k)$ so that r , the sum of the residuals, is as small as possible:

$$r = \sum_{i=1}^N (y_i - f(x_i, \vec{p}))^2$$

Curve fitting is of particular use if the data is noisy: for a given x_i and $y_i = f(x_i, \vec{p})$ we have a (unknown) error term ϵ_i so that $y_i = f(x_i, \vec{p}) + \epsilon_i$.

We use the following example to clarify this: $f(x, \vec{p}) = a \exp(-bx) + c$, i.e. $\vec{p} = a, b, c$

```
# NBVAL_IGNORE_OUTPUT
import numpy as np
from scipy.optimize import curve_fit

def f(x, a, b, c):
    """Fit function  $y=f(x,p)$  with parameters  $p=(a,b,c)$ . """
    return a * np.exp(- b * x) + c


#create fake data
x = np.linspace(0, 4, 50)
y = f(x, a=2.5, b=1.3, c=0.5)
#add noise
yi = y + 0.2 * np.random.normal(size=len(x))

#call curve fit function
popt, pcov = curve_fit(f, x, yi)
a, b, c = popt
print("Optimal parameters are a=%g, b=%g, and c=%g" % (a, b, c))

#plotting
import pylab
yfitted = f(x, *popt) # equivalent to f(x, popt[0], popt[1], popt[2])
pylab.plot(x, yi, 'o', label='data $y_i$')
pylab.plot(x, yfitted, '-', label='fit $f(x_i)$')
pylab.xlabel('x')
pylab.legend()
```

```
Optimal parameters are a=2.60173, b=1.1264, and c=0.445222
```

```
<matplotlib.legend.Legend at 0x7f5f1848b250>
```



Note that in the source code above we define the fitting function $y = f(x)$ through Python code. We can thus fit (nearly) arbitrary functions using the `curve_fit` method.

The `curve_fit` function returns a tuple `popt, pcov`. The first entry `popt` contains a tuple of the OPTimal Parameters (in the sense that these minimise equation ([eq:1]). The second entry contains the covariance matrix for all parameters. The diagonals provide the variance of the parameter estimations.

For the curve fitting process to work, the Levenburg-Marquardt algorithm needs to start the fitting process with initial guesses for the final parameters. If these are not specified (as in the example above), the value “1.0” is used for the initial guess.

If the algorithm fails to fit a function to data (even though the function describes the data reasonably), we need to give the algorithm better estimates for the initial parameters. For the example shown above, we could give the estimates to the `curve_fit` function by changing the line

```
popt, pcov = curve_fit(f, x, yi)
```

to

```
popt, pcov = curve_fit(f, x, yi, p0=(2, 1, 0.6))
```

if our initial guesses would be $a = 2, b = 1$ and $c = 0.6$. Once we take the algorithm “roughly in the right area” in parameter space, the fitting usually works well.

16.8 Fourier transforms

In the next example, we create a signal as a superposition of a 50 Hz and 70 Hz sine wave (with a slight phase shift between them). We then Fourier transform the signal and plot the absolute value of the (complex) discrete Fourier transform coefficients against frequency, and expect to see peaks at 50Hz and 70Hz.

```
import scipy.fft
import numpy as np
import matplotlib.pyplot as plt
pi = np.pi

signal_length = 0.5      # [seconds]
sample_rate = 500        # sampling rate [Hz]
dt = 1. / sample_rate    # time between two samples [s]

df = 1 / signal_length   # frequency between points in
                        # in frequency domain [Hz]
t = np.arange(0, signal_length, dt) # the time vector
n_t = len(t)             # length of time vector


# create signal
y = np.sin(2*pi*50*t) + np.sin(2*pi*70*t+pi/4)

# compute Fourier transform
f = scipy.fft.fft(y)

# work out meaningful frequencies in Fourier transform
freqs = df * np.arange(0, (n_t-1)/2., dtype='d') # 'd'=double precision float
n_freq = len(freqs)

# plot input data y against time
plt.subplot(2, 1, 1)
plt.plot(t, y, label='input data')
plt.xlabel('time [s]')
plt.ylabel('signal')

#plot frequency spectrum
plt.subplot(2, 1, 2)
plt.plot(freqs, abs(f[0:n_freq]),
         label='abs(fourier transform)')
plt.xlabel('frequency [Hz]')
plt.ylabel('abs(DFT(signal))');
```



The lower plot shows the discrete Fourier transform computed from the data shown in the upper plot.

16.9 Optimisation

Often we need to find the maximum or minimum of a particular function $f(x)$ where f is a scalar function but x could be a vector. Typical applications are the minimisation of entities such as cost, risk and error, or the maximisation of productivity, efficiency and profit. Optimisation routines typically provide a method to minimise a given function: if we need to maximise $f(x)$ we create a new function $g(x)$ that reverses the sign of f , i.e. $g(x) = -f(x)$ and we minimise $g(x)$.

Below, we provide an example showing (i) the definition of the test function and (ii) the call of the `scipy.optimize.fmin` function which takes as argument a function f to minimise and an initial value x_0 from which to start the search for the minimum, and which returns the value of x for which $f(x)$ is (locally) minimised. Typically, the search for the minimum is a local search, i.e. the algorithm follows the local gradient. We repeat the search for the minimum for two values ($x_0 = 1.0$ and $x_0 = 2.0$, respectively) to demonstrate that depending on the starting value we may find different minima of the function f .

The majority of the commands (after the two calls to `fmin`) in the file `fmin1.py` creates the plot of the function, the start points for the searches and the minima obtained:

```
from numpy import arange, cos, exp
from scipy.optimize import fmin
import pylab

def f(x):
    return cos(x) - 3 * exp(-(x - 0.2) ** 2)
```

(continues on next page)

(continued from previous page)

```

# find minima of f(x),
# starting from 1.0 and 2.0 respectively
minimum1 = fmin(f, 1.0)
print("Start search at x=1., minimum is", minimum1)
minimum2 = fmin(f, 2.0)
print("Start search at x=2., minimum is", minimum2)

# plot function
x = arange(-10, 10, 0.1)
y = f(x)
pylab.plot(x, y, label='$\cos(x)-3e^{-(x-0.2)^2}$')
pylab.xlabel('x')
pylab.grid()
pylab.axis([-5, 5, -2.2, 0.5])

# add minimum1 to plot
pylab.plot(minimum1, f(minimum1), 'vr',
           label='minimum 1')
# add start1 to plot
pylab.plot(1.0, f(1.0), 'or', label='start 1')

# add minimum2 to plot
pylab.plot(minimum2, f(minimum2), 'vg', \
           label='minimum 2')
# add start2 to plot
pylab.plot(2.0, f(2.0), 'og', label='start 2')

pylab.legend(loc='lower left')


```

```

Optimization terminated successfully.
    Current function value: -2.023866
    Iterations: 16
    Function evaluations: 32
Start search at x=1., minimum is [0.23964844]
Optimization terminated successfully.
    Current function value: -1.000529
    Iterations: 16
    Function evaluations: 32
Start search at x=2., minimum is [3.13847656]

```

```
<matplotlib.legend.Legend at 0x7f5f18167810>
```



Calling the `fmin` function will produce some diagnostic output, which you can also see above.

Return value of `fmin`

Note that the return value from the `fmin` function is a `numpy array` which – for the example above – contains only one number as we have only one parameter (here x) to vary. In general, `fmin` can be used to find the minimum in a higher-dimensional parameter space if there are several parameters. In that case, the `numpy array` would contain those parameters that minimise the objective function. The objective function $f(x)$ has to return a scalar even if there are more parameters, i.e. even if x is a vector as in $f(\mathbf{x})$.

16.10 Other numerical methods

Scientific Python and Numpy provide access to a large number of other numerical algorithms including function interpolation, Fourier transforms, optimisation, special functions (such as Bessel functions), signal processing and filters, random number generation, and more. Start to explore `scipy`'s and `numpy`'s capabilities using the `help` function and the documentation provided on the web.

16.11 scipy.io: Scipy-input output

Scipy provides routines to read and write Matlab `mat` files. Here is an example where we create a Matlab compatible file storing a (1x11) matrix, and then read this data into a numpy array from Python using the scipy Input-Output library:

First we create a mat file in Octave (Octave is [mostly] compatible with Matlab):

```
octave:1> a=-1:0.5:4
a =
Columns 1 through 6:
  -1.0000  -0.5000   0.0000   0.5000   1.0000   1.5000
Columns 7 through 11:
   2.0000   2.5000   3.0000   3.5000   4.0000
octave:2> save -6 octave_a.mat a           %save as version 6
```

Then we load this array within python:

```
from scipy.io import loadmat
mat_contents = loadmat('static/data/octave_a.mat')
```

```
mat_contents
```

```
{'__header__': b'MATLAB 5.0 MAT-file Platform: posix, Created on: Mon Aug  8
↵12:21:36 2016',
  '__version__': '1.0',
  '__globals__': [],
  'a': array([[ -1. , -0.5,  0. ,  0.5,  1. ,  1.5,  2. ,  2.5,  3. ,  3.5,  4. ]])}
```

```
mat_contents['a']
```

```
array([[ -1. , -0.5,  0. ,  0.5,  1. ,  1.5,  2. ,  2.5,  3. ,  3.5,  4. ]])
```

The function `loadmat` returns a dictionary: the key for each item in the dictionary is a string which is the name of that array when it was saved in Matlab. The key is the actual array.

A Matlab matrix file can hold several arrays. Each of those is presented by one key-value pair in the dictionary.

Let's save two arrays from Python to demonstrate that:

```
import scipy.io
import numpy as np

# create two numpy arrays
a = np.linspace(0, 50, 11)
b = np.ones((4, 4))

# save as mat-file
# create dictionary for savemat
tmp_d = {'a': a,
         'b': b}
scipy.io.savemat('data.mat', tmp_d)
```

This program creates the file `data.mat`, which we can subsequently read using Matlab or here Octave:

```
HAL47:code fangohr$ octave
GNU Octave, version 3.2.4
Copyright (C) 2009 John W. Eaton and others.
<snip>

octave:1> whos
Variables in the current scope:

  Attr Name      Size          Bytes  Class
  ==== =====
      ans        1x11           92   cell

Total is 11 elements using 92 bytes

octave:2> load data.mat
octave:3> whos
Variables in the current scope:

  Attr Name      Size          Bytes  Class
  ==== =====
      a          11x1           88  double
      ans        1x11           92   cell
      b           4x4          128  double

Total is 38 elements using 308 bytes

octave:4> a
a =

    0
    5
   10
   15
   20
   25
   30
   35
   40
   45
   50

octave:5> b
b =

    1    1    1    1
    1    1    1    1
    1    1    1    1
    1    1    1    1
```

Note that there are other functions to read from and write to in formats as used by IDL, Netcdf and other formats in `scipy.io`.

More → see [Scipy tutorial](#).

PANDAS - DATA SCIENCE WITH PYTHON

Numpy and numpy arrays are our tool of choice for numeric data that resembles vectors, matrices (and higher dimensional tensors).

Where data is gathered from experiments, and in particular where we want to extract meaning from the combination of different data sources, and where data is often incomplete, the pandas library offers a number of useful tools (and has become a standard tool for data scientists).

In this section, we introduce the basics of Pandas.

In particular, we introduce the two key data types in Pandas: the `Series` and the `DataFrame` objects.

By convention, the pandas library is imported under the name `pd` (the same way that `numpy` is imported under the name `np`):

```
import pandas as pd
```

17.1 Motivational example (Series)

Imagine we are working on software for a greengrocer or supermarket, and need to track the number of apples (10), oranges(3) and bananas (22) that are available in the supermarket.

We could use a python list (or a numpy array) to track these numbers:

```
stock = [10, 3, 22]
```

However, we would need to remember separately that the entries are in the order of apples, oranges, and bananas. This could be achieved through a second list:

```
stocklabels = ['apple', 'orange', 'banana']
```

```
assert len(stocklabels) == len(stock) # check labels and
                                         # stock are consistent
for label, count in zip(stocklabels, stock):
    print(f'{label:10s} : {count:4d}')
```

```
apple      :   10
orange     :    3
banana     :   22
```

The above 2-list solution is a little awkward in two ways: firstly, we have use two lists to describe one set of data (and thus need to be carefully to update them simulatenously, for example), and secondly, the access to the data given a label is inconvenient: We need to find the index of the label with one list, then use this as the index to the other list, for example

```
index = stocklabels.index('banana')
bananas = stock[index]
print(f"There are {bananas} bananas [index={index}].")
```

```
There are 22 bananas [index=2].
```

We have come across similar examples in the section on dictionaries, and indeed a dictionary is a more convenient solution:

```
stock_dic = {'apple': 10,
            'orange': 3,
            'banana': 22}
```

The keys of the dictionary contain the stock labels and the values contain the actual values:

```
stock_dic.keys()
```

```
dict_keys(['apple', 'orange', 'banana'])
```

```
stock_dic.values()
```

```
dict_values([10, 3, 22])
```

To retrieve (or change) the value for `apple`, we use `apple` as the key and retrieve the value through the dictionary's indexing notation:

```
stock_dic['apple']
```

```
10
```

And we can summarise the stock as follows:

```
for label in stock_dic:
    print(f'{label:10s} : {stock_dic[label]:4d}')

```

```
apple      :   10
orange     :    3
banana     :   22
```

This is a vast improvement over the 2-lists solution: (i) we only maintain one structure, which contains a value for every key - so we don't need to check that the lists have the same length. (ii) we can access individual elements through the label (using it as a key for the dictionary).

The Pandas Series object address the requirements above. It is similar to a dictionary, but with improvements for the given problem:

- the order of the items is maintained
- the values have to have the same type (higher execution performance)

- a (large) number of convenience functionality, for example to deal with missing data, time series, sorting, plotting, and more

17.2 Pandas Series

17.2.1 Stock example - Series

We can create a `Series` object - for example - from a dictionary:

```
stock = pd.Series({'apple': 10,
                  'orange': 3,
                  'banana': 22})
```

The default presentation shows the entries one per row, with the label on the left, and the value on the right.

```
type(stock)
```

```
pandas.core.series.Series
```

```
stock
```

```
apple    10
orange    3
banana   22
dtype: int64
```

The items on the left are referred to as the `index` of the `Series`, and are available as the `index` attribute of the `series` object:

```
stock.index
```

```
Index(['apple', 'orange', 'banana'], dtype='object')
```

```
type(stock.index)
```

```
pandas.core.indexes.base.Index
```

We can also access the list of values for each item, using the `values` attribute:

```
stock.values
```

```
array([10,  3, 22])
```

Regarding data access, the `Series` object behaves like a dictionary:

```
stock['apple']
```

```
10
```

```
stock['potato'] = 101    # adding more values  
stock['cucumber'] = 1
```

```
print(stock)
```

```
apple      10  
orange     3  
banana    22  
potato    101  
cucumber   1  
dtype: int64
```

```
stock
```


```
apple      10  
orange     3  
banana    22  
potato    101  
cucumber   1  
dtype: int64
```

We can plot the data as a bar chart:

```
%matplotlib inline  
# settings for Jupyter book: svg for html version, high-resolution png for pdf  
import matplotlib  
import matplotlib_inline  
matplotlib_inline.backend_inline.set_matplotlib_formats('svg', 'png')  
matplotlib.rcParams['figure.dpi'] = 400
```

```
stock.plot(kind='bar')
```


```
<Axes: >
```



We can sort the data according to the values in the Series (and then plot to visualise):

```
stock.sort_values().plot(kind='bar')
```


<Axes: >



Or sort the index to get alphabetical order of our fruit and vegetables:

```
stock.sort_index().plot(kind='bar')
```

```
<Axes: >
```



The `Series` object has a number of numerical methods available, including `mean` and `sum`:

```
stock.sum()
```

```
137
```

```
stock.mean()
```

```
27.4
```

It also behaves like a sequence in that the `len` function returns the number of data points in the `Series` object:

```
len(stock)
```

```
5
```

17.2.2 memory usage

For larger data sets, it might be important to know how many bytes storing the Series costs. The bytes required to store the actual series data are available as

```
stock.nbytes
```

```
40
```

or from the underlying numpy array directly:

```
stock.values.nbytes
```

```
40
```

It is 40 bytes, because we have 5 elements stored as int64 (each needing 8 bytes):

```
stock.dtype
```

```
dtype('int64')
```

The Series object needs additional memory. This can be queried using:

```
stock.memory_usage()
```

```
252
```

17.2.3 Statistics

A number of statistical descriptors of the data in the `stock` Series object is available using `describe()`:

```
stock.describe()
```

```
count      5.000000
mean       27.400000
std        41.955929
min         1.000000
25%         3.000000
50%        10.000000
75%        22.000000
max        101.000000
dtype: float64
```

As usual, the documentation strings provide documentation (`help(stock.describe)`), and the pandas home page (<https://pandas.pydata.org>) provides links to the Pandas documentation.

17.3 Create Series from list

In the example above, we showed how to create a Series from a dictionary where the keys of the dictionary entries served as the index for the Series object.

We can also create a Series from a list, and provide an additional index:

```
stock = pd.Series([10, 3, 22], index=['apple', 'orange', 'banana'])
```

```
stock
```

```
apple    10
orange    3
banana   22
dtype: int64
```

If we omit the `index` argument, the Series will assume an integer index:

```
stock = pd.Series([10, 3, 22])
```

```
stock
```

```
0    10
1     3
2    22
dtype: int64
```

However, an index can be added subsequently:

```
stock.index = ['apple', 'orange', 'banana']
```

```
stock
```

```
apple    10
orange    3
banana   22
dtype: int64
```

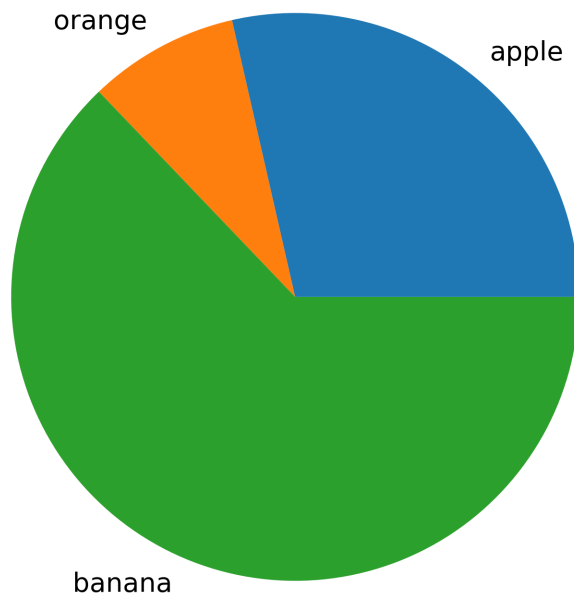
17.4 Plotting data

Commonly used plots are easily accessible via the `plot()` method of the Series object. We have seen a bar plot above already. The `Series.plot()` method accepts an argument `kind` such as `kind="bar"`, but there is an equivalent method `Series.plot.bar()` available.

Further examples:


```
stock.plot.pie()
```

<Axes: >




To tailor the plot, we can either get the axis object and modify it subsequently:

```
ax = stock.plot.pie()
ax.set_aspect(1)
ax.set_ylabel(None);
ax.set_title(None);
```

```
import matplotlib.pyplot as plt
fig, ax = plt.subplots(1, 1, figsize=(9, 3))
stock.plot.bar(ax=ax)
ax.set_title("Current stock");
```



We can also retrieve the data from the series and drive the plotting “manually” ourselves:

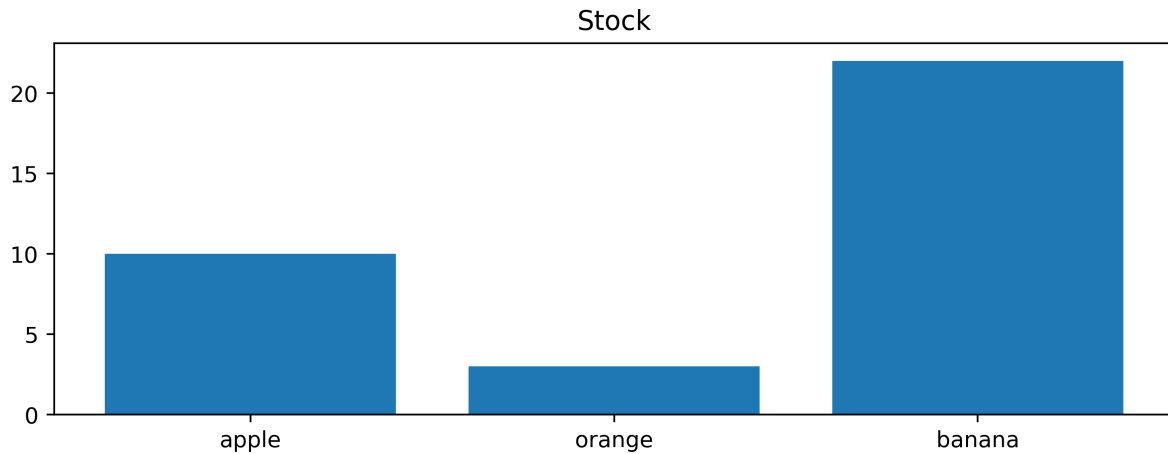
```
import matplotlib.pyplot as plt

names = list(stock.index) # conversion to list not necessary
values = list(stock.values) # conversion to list not necessary
```

(continues on next page)

(continued from previous page)

```
fig, ax = plt.subplots(1, 1, figsize=(9, 3))
ax.bar(names, values)
ax.set_title('Stock');
```



17.5 Missing values

“Real” data sets tend to be incomplete. Dealing with missing values is an important topic in data science. The agreement in Pandas is that the special floating point value “NaN” (standing for Not a Number) represents missing data points. For example, if we have a table for the stock, but we don’t know the value for `apple`, we would replace it with `NaN`.

The special `Nan` value in Python can be created using `float('nan')` or using `numpy.nan` if the `numpy` module is imported.

```
stock['apple'] = float('nan')
```

```
stock
```

```
apple      NaN
orange     3.0
banana     22.0
dtype: float64
```

Note that the `dtype` of the `stock` Series object has changed from `int64` to `float64` when we assigned `NaN` to `apple`: the whole series has been converted to `float`, because `NaN` is only defined for floating point numbers.

(There is a proposal to create a `NaN` object as part of `pandas` - this would overcome the above limitation.)

Assume we need to calculate how many items of stock we have in total using the `sum` function:

```
stock.values
```

```
array([nan,  3., 22.])
```

A common situation is that we have an incomplete Series or DataFrame (which are multiple Series with the same index) and we want to process with our analysis, but treat the missing values in a special way.

```
stock.sum()
```

```
25.0
```

The above example `sum` shows that NaN values are simply ignored, which can be convenient.

We can also ‘tidy up’ the Series object, by removing all entries that have a NaN value:

```
stock.dropna()
```

```
orange    3.0
banana   22.0
dtype: float64
```

17.6 Series data access: explicit and implicit (`loc` and `iloc`)

```
stock = pd.Series({'apple': 10,
                  'orange': 3,
                  'banana': 22,
                  'cucumber': 1,
                  'potato': 110})
```

```
stock
```

```
apple      10
orange      3
banana     22
cucumber    1
potato     110
dtype: int64
```

17.6.1 Indexing

We can access single values through their index as if the `stock` Series object would be a dictionary:

```
stock['banana']
```

```
22
```

There is an equivalent and recommended way of using this retrieval using the `loc` (for LOcation?) attribute:

```
stock.loc['banana']
```

```
22
```

For convenience, pandas also (!) allows us to use integer indexing into the Series object. This is called *implicit* indexing as the series Object doesn’t use integers as the index, but the name of the fruits.

For example, we can also retrieve the value for `banana` through its implicit index 2, because it is in row 3 of the Series object (which would need index 2 as we start counting from 0):

```
stock[2]
```

```
/tmp/ipykernel_293/4060357290.py:1: FutureWarning: Series.__getitem__ treating keys as positions is deprecated. In a future version, integer keys will always be treated as labels (consistent with DataFrame behavior). To access a value by position, use `ser.iloc[pos]`
stock[2]
```

```
22
```

In this example, this works fine and seems convenient, but can become very confusing if the actual index of the object consists of integers. For that reason, the explicit (and recommended way) of using the indirect indexing is through the `iloc` (ImplicitLOCation) attribute:

```
stock.iloc[2]
```

```
22
```

17.6.2 Slicing

```
stock
```

```
apple      10
orange      3
banana     22
cucumber    1
potato     110
dtype: int64
```

We can also slice the Series:

```
stock['orange':'potato']
```

```
orange      3
banana     22
cucumber    1
potato     110
dtype: int64
```

Or skip every second entry:

```
stock['orange':'potato':2]
```

```
orange      3
cucumber    1
dtype: int64
```

17.6.3 Data manipulation

Numerical operations on the series object can be carried for all data values at the same time in the same way that numpy arrays are processed:

```
stock - stock.mean()
```

```
apple      -19.2
orange     -26.2
banana     -7.2
cucumber   -28.2
potato      80.8
dtype: float64
```

```
import numpy as np
```

```
np.sqrt(stock)
```

```
apple      3.162278
orange     1.732051
banana     4.690416
cucumber   1.000000
potato     10.488088
dtype: float64
```

Where preferred, we can extract the numpy array and work with that:

```
data = stock.values
```

```
type(data)
```

```
numpy.ndarray
```

```
data - data.mean()
```

```
array([-19.2, -26.2, -7.2, -28.2, 80.8])
```

17.6.4 Import and Export

Pandas (and its objects `Series` and `DataFrame`) support export to and import from a number of useful formats.

For example, we can write a `Series` object into a comma separated value file:

```
stock.to_csv('stock.csv', header=False)
```

```
#NBVAL_IGNORE_OUTPUT
!cat stock.csv
```

```
apple, 10
orange, 3
banana, 22
cucumber, 1
potato, 110
```

We can also create a LaTeX representation of the table:

```
stock.to_latex()
```

```
'\\begin{tabular}{lr}\\n\\toprule\\n & 0 \\n\\midrule\\napple & 10 \\norange & 3 \\nbanana & 22 \\ncucumber & 1 \\npotato & 110 \\n\\bottomrule\\n\\end{tabular}\\n'
```

We'll come back to reading from files in the `DataFrame` section.

17.7 Data Frame

17.7.1 Stock Example - `DataFrame`

After having introduced the `Series` object above, we will focus on the second important type in pandas: the `DataFrame`.

As a first description, we could say that the `DataFrame` is similar to a (2d) spreadsheet: it contains rows and columns.

The series object we have studied above is a special case of the `DataFrame`, where the `DataFrame` has only one column.

We'll continue with our stock example:

```
stock
```

```
apple      10
orange      3
banana     22
cucumber    1
potato     110
dtype: int64
```

In addition to tracking how many objects of each type we have stocked, we have a second `Series` object that provides the price per item at which the item is sold:

```
price = pd.Series({'apple': 0.55, 'banana': 0.50, 'cucumber': 0.99, 'potato': 0.17,
                  ↪ 'orange': 1.76})
price
```

```
apple      0.55
banana     0.50
cucumber    0.99
potato     0.17
orange     1.76
dtype: float64
```

The `DataFrame` object allows us to treat the two series together. In fact, a convenient way to create the `DataFrame` object is to combine a number of series as follows:

```
shop = pd.DataFrame({'stock' : stock, 'price' : price})
shop
```

	stock	price
apple	10	0.55
banana	22	0.50
cucumber	1	0.99
orange	3	1.76
potato	110	0.17

Because both `Series` objects had the same index elements, our data is nicely aligned in the `DataFrame` with name `shop`, even though the data was stored in different order in the `price` and `stock`.

If one `Series` is missing a data point, pandas will insert a `NaN` entry into that field:

```
price2 = price.copy()
```

```
price2['grapefruit'] = 1.99
price2
```

apple	0.55
banana	0.50
cucumber	0.99
potato	0.17
orange	1.76
grapefruit	1.99

dtype: float64

```
pd.DataFrame({'stock' : stock, 'price' : price2})
```

	stock	price
apple	10.0	0.55
banana	22.0	0.50
cucumber	1.0	0.99
grapefruit	NaN	1.99
orange	3.0	1.76
potato	110.0	0.17

17.7.2 Accessing data in a DataFramea

```
shop
```

	stock	price
apple	10	0.55
banana	22	0.50
cucumber	1	0.99
orange	3	1.76
potato	110	0.17

The data frame has an *index* which is the same for all columns, and shown in bold in the left most column. We can also ask for it:

```
shop.index
```

```
Index(['apple', 'banana', 'cucumber', 'orange', 'potato'], dtype='object')
```

Each column has name (here `stock` and `price`):

```
shop.columns
```

```
Index(['stock', 'price'], dtype='object')
```

17.7.3 Extracting columns of data

Using the column names, we can extract one column into a Series object using the index operator (`[]`):

```
shop['stock']
```

```
apple      10
banana     22
cucumber    1
orange      3
potato     110
Name: stock, dtype: int64
```

```
shop['price']
```

```
apple      0.55
banana     0.50
cucumber   0.99
orange     1.76
potato     0.17
Name: price, dtype: float64
```

17.7.4 Extracting rows of data

We have two options of extracting a row of data.

First, explicit indexing using the label of the index in that row:

```
shop.loc['apple'] # single row is returned as series
```

```
stock      10.00
price       0.55
Name: apple, dtype: float64
```

```
shop.loc['banana':'cucumber'] # multiple rows are returned as DataFrame
```



```

      stock  price
banana     22   0.50
cucumber     1   0.99

```

Second, we can use the implicit indexing (as for Series objects):

```
shop.iloc[0]
```

```

stock     10.00
price      0.55
Name: apple, dtype: float64

```

```
shop.iloc[1:3]
```

```

      stock  price
banana     22   0.50
cucumber     1   0.99

```

Warning

Note that there are some inconsistencies here: the explicit slicing with index labels (such as `.loc['banana':'cucumber']`) is inclusive of `cucumber`, whereas in the implicit slicing (such as `.iloc[1:3]`) the row with index 3 is *not* included.

The behaviour of `.loc` is convenient and a good design choice if labels such as strings in our `stock` example are used. The behaviour of `.iloc` is reflecting the normal Python behaviour.

It is thus understandable how we have arrived at the situation.

17.7.5 Data manipulation with `shop`

The real strength of the DataFrames is that we can continue to process the data conveniently.

For example, we could work out the financial value of the items we have in stock, and add this as an extra column:

```
shop['value'] = shop['price'] * shop['stock']
shop
```

```

      stock  price  value
apple      10   0.55   5.50
banana     22   0.50  11.00
cucumber     1   0.99   0.99
orange       3   1.76   5.28
potato     110   0.17  18.70

```

Of course we can compute the sum, for example, to estimate the value of the total stock:

```
shop['value'].sum()
```

```
41.47
```

If, for whatever reason, we want to swap columns with rows, we can `transpose` the data frame like a numpy array:

```
shop.transpose()
```

	apple	banana	cucumber	orange	potato
stock	10.00	22.0	1.00	3.00	110.00
price	0.55	0.5	0.99	1.76	0.17
value	5.50	11.0	0.99	5.28	18.70

17.8 Example: European population 2017

Here is a second example to demonstrate some use cases of pandas DataFrames.

First, we get the data. It is originally from EUROSTAT (reference “demo_gind”)

```
#NBVAL_IGNORE_OUTPUT
!wget https://fangohr.github.io/data/eurostat/population2017/eu-pop-2017.csv
```

```
--2024-04-14 10:28:17-- https://fangohr.github.io/data/eurostat/population2017/eu-
↳pop-2017.csv
Resolving fangohr.github.io (fangohr.github.io)... 185.199.108.153, 185.199.109.
↳153, 185.199.111.153, ...
Connecting to fangohr.github.io (fangohr.github.io)|185.199.108.153|:443...
↳connected.
```

```
HTTP request sent, awaiting response...
```

```
200 OK
Length: 1087 (1.1K) [text/csv]
Saving to: 'eu-pop-2017.csv'
```

```
eu-pop-2017.csv      0%[          ]          0  --.-KB/s
eu-pop-2017.csv    100%[=====>]    1.06K  --.-KB/s   in 0s
```

```
2024-04-14 10:28:17 (66.2 MB/s) - 'eu-pop-2017.csv' saved [1087/1087]
```

The data source is a comma-separated-value file (CSV), which looks like this:

```
#NBVAL_IGNORE_OUTPUT
!head eu-pop-2017.csv
```

```
geo,pop17,pop18,births,deaths
Belgium ,11351727,11413058,119690,109666
Bulgaria,7101859,7050034,63955,109791
Czechia,10578820,10610055,114405,111443
Denmark,5748769,5781190,61397,53261
Germany,82521653,82850000,785000,933000
Estonia ,1315634,1319133,13784,15543
Ireland,4784383,4838259,62084,30324
Greece,10768193,10738868,88523,124530
Spain,46527039,46659302,390024,421269
```

Pandas has very strong support of reading files from different formats, including MS Excel, CSV, HDF5 and others. Each reading routine has a number of options to tailor the process.

Many data science projects leave the data in their original files, and use a few lines of Python code to import it.

```
df = pd.read_csv('eu-pop-2017.csv')
```

```
df
```

	geo	pop17	pop18	births	deaths
0	Belgium	11351727	11413058	119690	109666
1	Bulgaria	7101859	7050034	63955	109791
2	Czechia	10578820	10610055	114405	111443
3	Denmark	5748769	5781190	61397	53261
4	Germany	82521653	82850000	785000	933000
5	Estonia	1315634	1319133	13784	15543
6	Ireland	4784383	4838259	62084	30324
7	Greece	10768193	10738868	88523	124530
8	Spain	46527039	46659302	390024	421269
9	France	66989083	67221943	767691	603141
10	Croatia	4154212	4105493	36556	53477
11	Italy	60589445	60483973	458151	649061
12	Cyprus	854802	864236	9229	5997
13	Latvia	1950116	1934379	20828	28757
14	Lithuania	2847904	2808901	28696	40142
15	Luxembourg	590667	602005	6174	4263
16	Hungary	9797561	9778371	94646	131877
17	Malta	460297	475701	4319	3571
18	Netherlands	17081507	17181084	169200	150027
19	Austria	8772865	8822267	87633	83270
20	Poland	37972964	37976687	401982	402852
21	Portugal	10309573	10291027	86154	109586
22	Romania	19644350	19523621	189474	260599
23	Slovenia	2065895	2066880	20241	20509
24	Slovakia	5435343	5443120	57969	53914
25	Finland	5503297	5513130	50321	53722
26	Sweden	9995153	10120242	115416	91972
27	United Kingdom	65808573	66238007	755043	607172

We look at the dataframe as it is, and use the ‘head()’ command which will only show the first 5 lines of data:

```
df.head()
```

	geo	pop17	pop18	births	deaths
0	Belgium	11351727	11413058	119690	109666
1	Bulgaria	7101859	7050034	63955	109791
2	Czechia	10578820	10610055	114405	111443
3	Denmark	5748769	5781190	61397	53261
4	Germany	82521653	82850000	785000	933000

The meaning of the columns, we have to get from metadata information. In this case, we have the following description of the data:

- **geo**: the country in question
- **pop17**: the population count of that country as of 1 January 2017

- **pop18**: the population count of that country as of 1 January 2018
- **births**: the number of (live) births in the country during the year 2017
- **deaths**: the number of deaths in that country during the year 2017

The data is provided for all of the 28 European Union members (as of 2017).

We want to use the country as the country name as the index. We can achieve this either with

```
df2 = df.set_index('geo')
```

```
df2.head()
```

	pop17	pop18	births	deaths
geo				
Belgium	11351727	11413058	119690	109666
Bulgaria	7101859	7050034	63955	109791
Czechia	10578820	10610055	114405	111443
Denmark	5748769	5781190	61397	53261
Germany	82521653	82850000	785000	933000

Note that we cannot change the index in a given DataFrame, so the `set_index()` method returns a new DataFrame. (This happens for many operations.)

As an alternative, we can also modify the import statement to already indicate which column we want to use as the index:

```
df = pd.read_csv('eu-pop-2017.csv', index_col="geo")
```


```
df.head()
```

	pop17	pop18	births	deaths
geo				
Belgium	11351727	11413058	119690	109666
Bulgaria	7101859	7050034	63955	109791
Czechia	10578820	10610055	114405	111443
Denmark	5748769	5781190	61397	53261
Germany	82521653	82850000	785000	933000

We explore the data by plotting some of it:

```
df.plot(kind='bar', y='pop17')
```

```
<Axes: xlabel='geo'>
```



The above shows the population as of 1 Jan 2017.


We'll try to improve this in two ways:

- we want to count population in millions. We can do this by dividing all the data by 10^6 .
- it would be interesting to sort the countries in order of size for this plot.

```
df_millions = df / 1e6
```

```
df_millions['pop17'].sort_values(ascending=False).plot(kind='bar')
```

```
<Axes: xlabel='geo'>
```




The example above selects one column from the data frame (`['pop17']`) and that returns a `Series` object. Then we sort this `Series` object using `sort_values()` according to the values (that's the number of people in each country), then we plot this.

Alternatively, we could also create a plot for the whole data frame, but say that the `pop17` is the column for sorting, and that we want to plot only the column with `pop17`:


```
df_millions.sort_values(by='pop17').plot(kind='bar', y='pop17')
```

```
<Axes: xlabel='geo'>
```




We can also plot more than one column at the same time:

```
ax = df_millions.sort_values(by='pop17').plot(kind='bar', y=['pop17', 'pop18'])
```



We can also fine tune the plot with the usual `matplotlib` commands:

```
ax = df_millions.sort_values(by='pop17').plot(kind='bar', y='pop17', figsize=(10, 4))
ax.set_ylabel("population 2017 [in millions]")
ax.grid()
ax.set_xlabel(None); # get rid of default label for x-axis ('geo')
```

Based on the number of births and deaths, we can compute change in population for each country for 2017. This is sometimes called the “natural-change”:

```
df['natural-change'] = df['births'] - df['deaths']
```

```
df['natural-change'].sort_values()
```

```
geo
Italy          -190910
Germany        -148000
Romania         -71125
Bulgaria        -45836
Hungary         -37231
Greece          -36007
Spain           -31245
Portugal        -23432
Croatia         -16921
Lithuania       -11446
Latvia          -7929
Finland         -3401
Estonia         -1759
Poland           -870
Slovenia        -268
Malta            748
Luxembourg       1911
Czechia         2962
Cyprus           3232
Slovakia        4055
Austria         4363
Denmark         8136
Belgium        10024
Netherlands     19173
Sweden         23444
Ireland         31760
```

(continues on next page)

(continued from previous page)

```

United Kingdom    147871
France            164550
Name: natural-change, dtype: int64

```

From this, we can see that the population change due to births and deaths in Italy and Germany is decreasing most in absolute terms.

To relate this to the overall size of the population, one often uses rates per year and per 1000 people in the country, such as the birth rate per 1000 inhabitants [1] (and death rate accordingly):

[1] https://en.wikipedia.org/wiki/Birth_rate

```

df['birth-rate'] = df['births'] / df['pop17'] * 1000
df['death-rate'] = df['deaths'] / df['pop17'] * 1000
df['natural-change-rate'] = df['natural-change'] / df['pop17'] * 1000

```

```
df.head()
```

geo	pop17	pop18	births	deaths	natural-change	birth-rate	\
Belgium	11351727	11413058	119690	109666	10024	10.543770	
Bulgaria	7101859	7050034	63955	109791	-45836	9.005389	
Czechia	10578820	10610055	114405	111443	2962	10.814533	
Denmark	5748769	5781190	61397	53261	8136	10.680026	
Germany	82521653	82850000	785000	933000	-148000	9.512655	

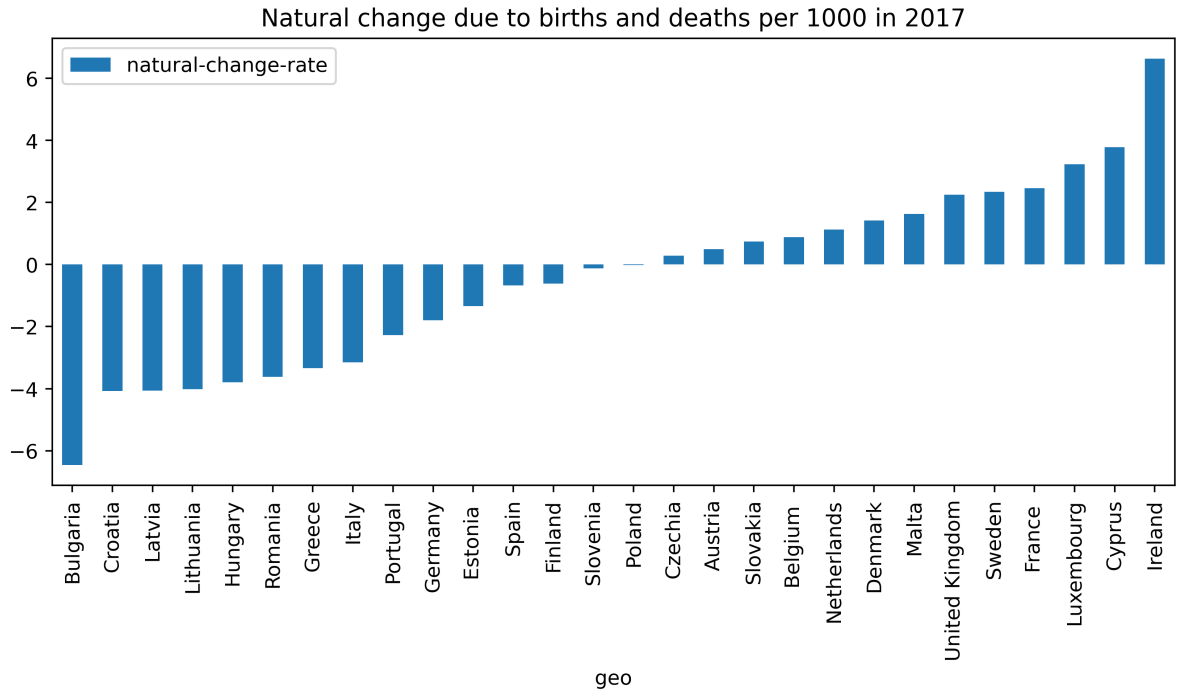
geo	death-rate	natural-change-rate
Belgium	9.660733	0.883037
Bulgaria	15.459473	-6.454085
Czechia	10.534540	0.279993
Denmark	9.264766	1.415260
Germany	11.306123	-1.793469

We can now look at the natural rate of change of population for each country, which is normalised by the population in that country.

```

ax = df.sort_values(by='natural-change-rate').plot(kind='bar', y='natural-change-rate
→', figsize=(10, 4))
ax.set_title("Natural change due to births and deaths per 1000 in 2017");

```



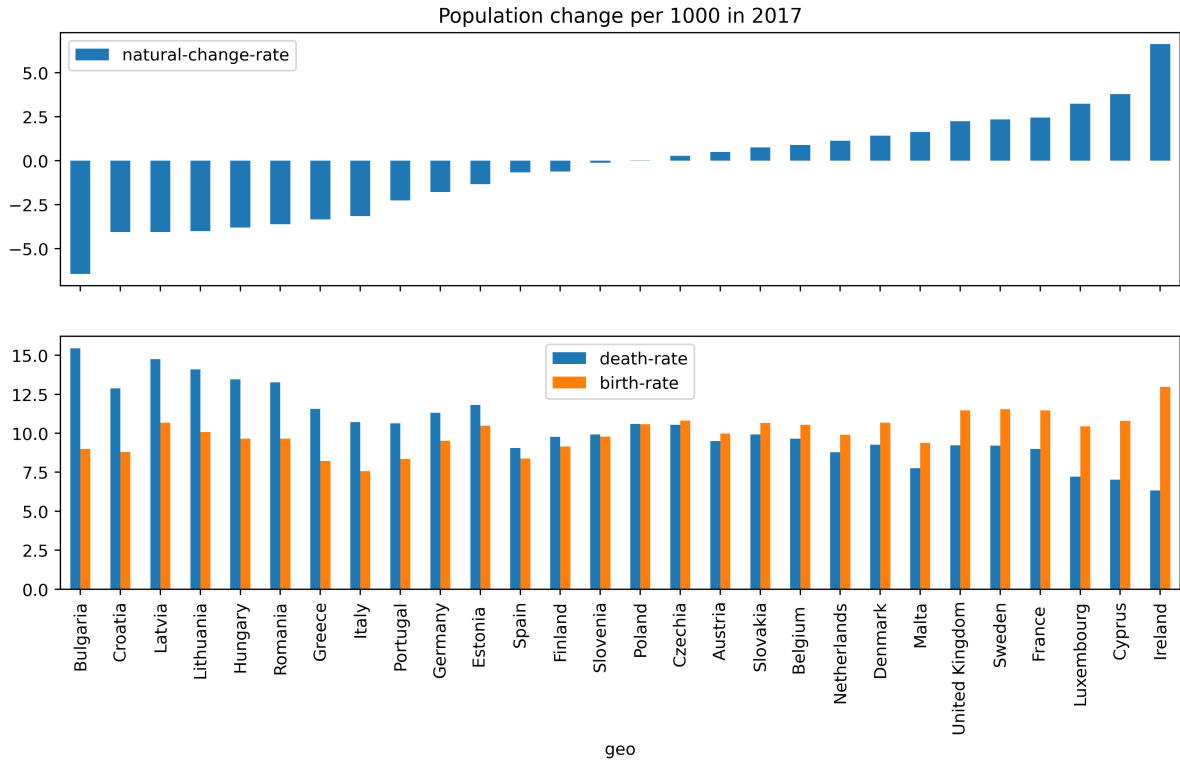
We can show the data together with the underlying birth and death rate data:

```
tmp = df.sort_values(by='natural-change-rate')

fig, axes = plt.subplots(2, 1, figsize=(12, 6))

tmp.plot(kind='bar', y=['natural-change-rate'], sharex=True, ax=axes[0])
axes[0].set_title("Population change per 1000 in 2017")
tmp.plot(kind='bar', y=['death-rate', 'birth-rate'], sharex=True, ax=axes[1])
```

```
<Axes: xlabel='geo'>
```

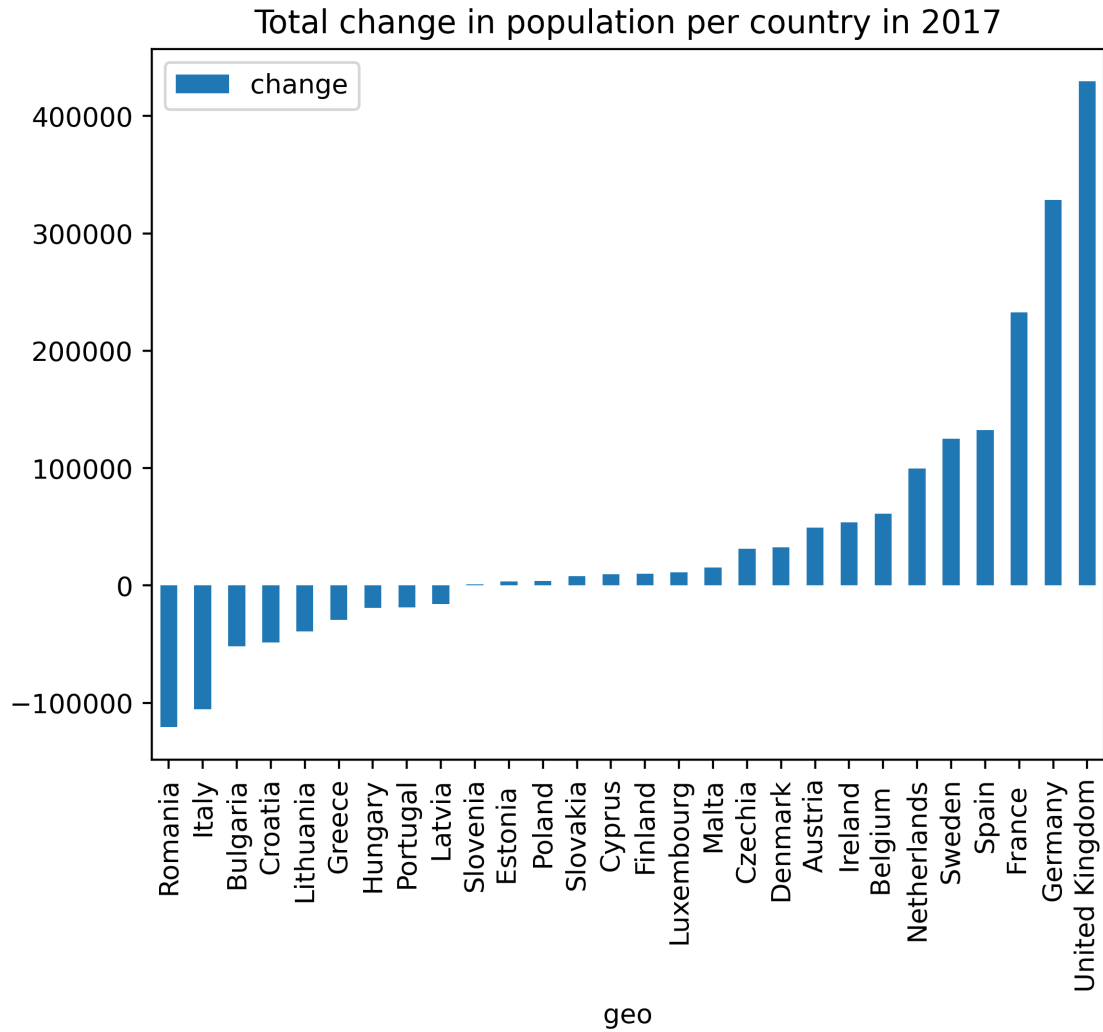


We haven't used the information we have about the population on 1 January 2018 yet.

Let's first look at the absolute changes in the population based on the (census?) data from 1 Jan 2017 and 1 Jan 2018:

```
df['change'] = df['pop18'] - df['pop17']
```

```
ax = df.sort_values(by='change').plot(y='change', kind='bar')
ax.set_title("Total change in population per country in 2017");
```



With that information, we can estimate migration. (It is important to note that this estimated number will also absorb all inaccuracies or changes of the data gathering method, in the original data described as “statistical adjustment”).

```
df['migration'] = df['change'] - df['natural-change']
```

```
df.head()
```

geo	pop17	pop18	births	deaths	natural-change	birth-rate	\
Belgium	11351727	11413058	119690	109666	10024	10.543770	
Bulgaria	7101859	7050034	63955	109791	-45836	9.005389	
Czechia	10578820	10610055	114405	111443	2962	10.814533	
Denmark	5748769	5781190	61397	53261	8136	10.680026	
Germany	82521653	82850000	785000	933000	-148000	9.512655	

geo	death-rate	natural-change-rate	change	migration
Belgium	9.660733	0.883037	61331	51307
Bulgaria	15.459473	-6.454085	-51825	-5989

(continues on next page)

(continued from previous page)

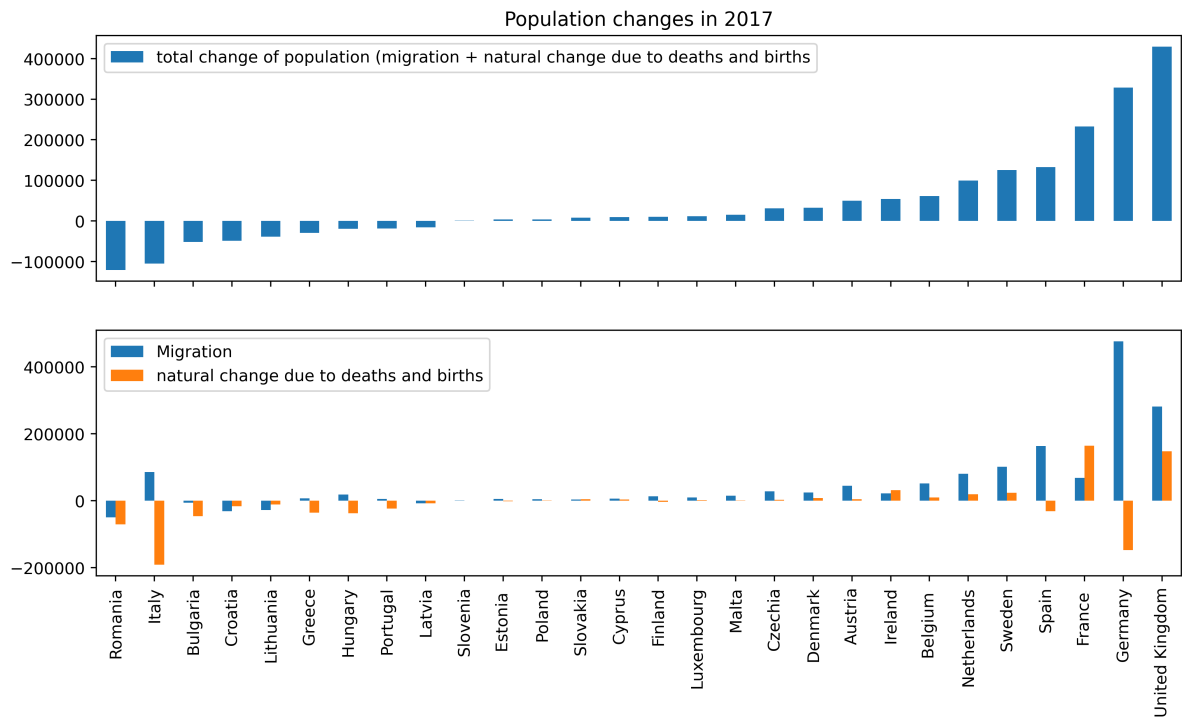
Czechia	10.534540	0.279993	31235	28273
Denmark	9.264766	1.415260	32421	24285
Germany	11.306123	-1.793469	328347	476347

Let's plot the total change of the population per country in the top subfigure, and the contribution from natural changes and migration in the lower subfigure:

```
tmp = df.sort_values(by='change')
fig, axes = plt.subplots(2, 1, figsize=(12, 6))

tmp.plot(kind='bar', y=['change'], sharex=True, ax=axes[0])
axes[0].set_title("Population changes in 2017")
axes[0].legend(['total change of population (migration + natural change due to deaths and births')])

tmp.plot(kind='bar', y=['migration', 'natural-change'], sharex=True, ax=axes[1])
axes[1].legend(['Migration', "natural change due to deaths and births"])
axes[1].set_xlabel(None);
```



17.9 Further reading

There is a lot more to say about Pandas. The following resources may be useful but there are countless others available:

- Further reading on `[]`, `.loc[]` and `.iloc[]` from Ted Petrou as a [Jupyter Notebook](#) and [blog entry](#).
- Jake VanderPlas: [Python Data Science Handbook online](#)

PYTHON PACKAGES AND ENVIRONMENTS

18.1 Introduction

This chapter provides an introduction to installing python packages and using python environments.

In the first part of the chapter you will learn how to

- search the database of Python packages
- install Python packages from the Python Packaging Index (PyPI)
- create virtual Python environments to separate different projects

In the second part, we provide additional information for users of the Anaconda distribution, in particular

- using environments within conda
- installing packages with conda
- interplay between conda and pip

We mention `pyenv` as an advanced tool at the very end.

This chapter does not comment on the creation of Python packages.

18.1.1 Shell commands in the Jupyter notebook

This chapter is written in a Jupyter notebook. This can be helpful for the reader as the notebook can be executed, and so the command can be replayed and varied easily.

In this particular chapter, we have lots of interaction with the shell of the operating system, and need to know two things:

1. We use the exclamation mark (!) to tell Jupyter to send the following command to the shell (rather than interpret in the Python environment of this notebook). Here is an example:

```
!date
```

```
Wed Jan 5 11:55:11 CET 2022
```

2. If we modify shell variables (such as the `PATH`), these are only set within the same cell. This will lead to repetition of some commands. This is mildly annoying (and will not be the case if the same commands are used outside the Jupyter Notebook).

Here is an example to illustrate the issue: First we set a variable value and then we display it:

```
!export NEW_VAR="test" && echo $NEW_VAR
```

```
test
```

The && operator instructs the shell to carry out the command to the right of && if the command on the left succeeded.

Repeating the “echo” command, we find the variable is not defined anymore:

```
!echo $NEW_VAR
```

So if we want to make use of those variable values, we need to repeat the setting of them:

```
!export NEW_VAR="test" && echo "The value of NEW_VAR is $NEW_VAR."
```

```
The value of NEW_VAR is test.
```

We need to make use of this when activating virtual environments (below).

18.1.2 Prerequisites

We assume that you have already python installed on your system. (And we assume that you are using Python3.) If you haven't got Python 3 yet, then either [install the Anaconda distribution](#), follow the instructions of the [Hitchhiker's Guide to Python](#) or take some other action.

Check that you have python installed:

```
!python --version
```

```
Python 3.9.7
```

We also assume you have a somewhat recent Python version (3.8 and above).

The commands below are tested for Linux and OSX operating system. If you use windows, please check the corresponding commands from here: <https://packaging.python.org/en/latest/tutorials/installing-packages/>

18.2 Python virtual environments

Before we install packages (be it our own or those from somebody else), we should create a new virtual environment. This is good practice because

- we can delete it when we don't need it any more
- we can not break other projects using python we may be working on
- we have no constraints about versions of particular libraries (it could be that one application needs version 2.x of a library and another application needs the version 1.8: if the two applications are installed in different environments, then this is no problem)

18.2.1 Creating virtual environments

We can create a virtual environment using this command

```
!python -m venv myvirtualenv
```

This command creates a subdirectory with name `myvirtualenv` in the current directory which contains the new virtual environment:

```
!ls -dl myvirtualenv/
```

```
drwxr-xr-x  6 fangohr  staff  192 Jan  5 11:55 myvirtualenv/
```

The virtual environment will use the same Python interpreter that we have used above when creating it. On Linux/OSX system, we can find out which interpreter this is by using the `which` command:

```
!which python
```

```
/Users/fangohr/anaconda3/bin/python
```

If the basics are sufficient for you, you can skip to the next section about activating a virtual environment.

We also be more specific, and choose a particular Python interpreter when using the `-m venv` command to force use of a particular python version. On a Mac (OSX), if `python3` was installed via `brew`, there is a `python` executable in `/usr/local/bin/python3`. To force creation of a virtual environment using this interpreter, we could use

```
/usr/local/bin/python3 -m venv myvirtualenv
```

There is no need to know what is happening inside this folder, but as we are curious, we'll have a very brief look anyway:

```
!ls myvirtualenv/
```

```
bin          include      lib          pyvenv.cfg
```

The `pyvenv.cfg` file contains information about the Python interpreter we are using

```
!cat myvirtualenv/pyvenv.cfg
```

```
home = /Users/fangohr/anaconda3/bin
include-system-site-packages = false
version = 3.9.7
```

Out of interest, we can check the total Disk Usage of that subdirectory:

```
!du -hs myvirtualenv/
```

```
15M          myvirtualenv/
```

To use this virtual environment, we need to activate it:

18.2.2 Activating a virtual environment

To activate the virtual environment, we need to know in which folder it is installed (in our case in `myvirtualenv`). On Linux and OSX, we run the following shell command:

```
!source myvirtualenv/bin/activate
```

This changes the `PATH` variable, which the operating system uses search for the `python` executable: it puts the directory that contains the `python` interpreter in our virtual environment to the beginning of the `PATH` variable. We can check that this works, by using the `which` command again:

```
!source myvirtualenv/bin/activate && which python
```

```
/Users/fangohr/git/introduction-to-python-for-computational-science-and-  
engineering/book/myvirtualenv/bin/python
```

(As outlined in the introduction, the repetition of the `activate` command is only necessary because we work from a Notebook here: if you go through these steps in a shell, you can ignore this, and write `which python` straight away.)

18.2.3 Using the virtual environment

Once we have activated the virtual environment, we can use it as we would use the default Python environment provided on the system.

For example to install some Python packages.

18.2.4 Name of the virtual environment

We have used `myvirtualenv` as the name of the virtual environment. In general, the name can be chosen freely. Commonly used names include `env` or `venv`. Occasionally, the environment is installed in a hidden subdirectory (such as `.env` or `.venv`).

We have not used the `venv` name for pedagogical to avoid confusion with the `venv` module.

18.3 Python Package Index (PyPI)

The Python Package Index provides a searchable web interface (<https://pypi.org>) that provides all Python packages registered with PyPI.

PyPI is the standard way of distributing (open source) python packages, and commonly used in science and engineering as well.

18.3.1 Installing packages with pip

The command to install one or more of these package is `pip`. We will activate our virtual environment and install some example packages:

```
!source myvirtualenv/bin/activate && pip install cowsay
```

```
Collecting cowsay
  Using cached cowsay-4.0-py2.py3-none-any.whl (24 kB)
Installing collected packages: cowsay
Successfully installed cowsay-4.0
WARNING: You are using pip version 21.2.3; however, version 21.3.1 is available.
You should consider upgrading via the '/Users/fangohr/git/introduction-to-python-
for-computational-science-and-engineering/book/myvirtualenv/bin/python -m pip
install --upgrade pip' command.
```

As we get a warning that suggests to upgrade the `pip` package itself, we shall follow the instructions and run the recommended command:

```
!source myvirtualenv/bin/activate && pip install --upgrade pip
```

```
Requirement already satisfied: pip in ./myvirtualenv/lib/python3.9/site-packages
(21.2.3)
Collecting pip
  Using cached pip-21.3.1-py3-none-any.whl (1.7 MB)
Installing collected packages: pip
  Attempting uninstall: pip
    Found existing installation: pip 21.2.3
    Uninstalling pip-21.2.3:
      Successfully uninstalled pip-21.2.3
  Successfully installed pip-21.3.1
```

```
!source myvirtualenv/bin/activate && cowsay Hellooo World
```

```

| Hellooo World |
|=====|
| \
|  \
|   ^__^
|  (oo)\_______
|  (__)\       )\/\
|    ||----w |
|    ||     ||
```

We can confirm the list of packages we have installed (together with their version number) using `pip list`:

```
!source myvirtualenv/bin/activate && pip list
```

```
Package    Version
-----
cowsay     4.0
```

(continues on next page)

(continued from previous page)

```
pip          21.3.1
setuptools  57.4.0
```

18.3.2 Learn more about an installed package using `pip show`

Once a package is installed, we can use `pip show` to learn more about it:

```
!source myvirtualenv/bin/activate && pip show cowsay
```

```
Name: cowsay
Version: 4.0
Summary: The famous cowsay for GNU/Linux is now available for python
Home-page: https://github.com/VaasuDevanS/cowsay-python
Author: Vaasudevan Srinivasan
Author-email: vaasuceg.96@gmail.com
License: GNU-GPL
Location: /Users/fangohr/git/introduction-to-python-for-computational-science-and-
engineering/book/myvirtualenv/lib/python3.9/site-packages
Requires:
Required-by:
```

For packages that are not installed yet, we need to search <https://pypi.org> to learn more about them. This includes the list of available packages (under “release history”).

(There are command line tools such as `pip-search` that can help to find package names, but they do not provide the same depth of information as the web page [at the time of writing]).

18.3.3 Uninstalling packages with `pip`

(The `-y` is short for `yes` and tells `pip uninstall` not to ask for confirmation if `cowsay` should be uninstalled.)

```
!source myvirtualenv/bin/activate && pip uninstall -y cowsay
```

```
Found existing installation: cowsay 4.0
Uninstalling cowsay-4.0:
  Successfully uninstalled cowsay-4.0
```

```
!source myvirtualenv/bin/activate && pip list
```

```
Package    Version
-----
pip        21.3.1
setuptools 57.4.0
```

18.3.4 Installing packages with additional dependencies

As a second example, we'll install the `wikipedia` package. We will see that it needs additional python packages as dependencies, which will be installed automatically:

```
!source myvirtualenv/bin/activate && pip install wikipedia
```

```
Collecting wikipedia
  Using cached wikipedia-1.4.0.tar.gz (27 kB)
  Preparing metadata (setup.py) ... ?25ldone
?25hCollecting beautifulsoup4
  Using cached beautifulsoup4-4.10.0-py3-none-any.whl (97 kB)
Collecting requests<3.0.0,>=2.0.0
  Using cached requests-2.27.0-py2.py3-none-any.whl (63 kB)
Collecting idna<4,>=2.5
  Using cached idna-3.3-py3-none-any.whl (61 kB)
Collecting charset-normalizer~=2.0.0
  Using cached charset_normalizer-2.0.10-py3-none-any.whl (39 kB)
Collecting urllib3<1.27,>=1.21.1
  Using cached urllib3-1.26.7-py2.py3-none-any.whl (138 kB)
Collecting certifi>=2017.4.17
  Using cached certifi-2021.10.8-py2.py3-none-any.whl (149 kB)
Collecting soupsieve>1.2
  Using cached soupsieve-2.3.1-py3-none-any.whl (37 kB)
Using legacy 'setup.py install' for wikipedia, since package 'wheel' is not
↳ installed.
Installing collected packages: urllib3, soupsieve, idna, charset-normalizer,
↳ certifi, requests, beautifulsoup4, wikipedia
  Running setup.py install for wikipedia ... ?25ldone
?25hSuccessfully installed beautifulsoup4-4.10.0 certifi-2021.10.8 charset-
↳ normalizer-2.0.10 idna-3.3 requests-2.27.0 soupsieve-2.3.1 urllib3-1.26.7
↳ wikipedia-1.4.0
```

```
!source myvirtualenv/bin/activate && python -c "import wikipedia; print(wikipedia.
↳ summary('cowsay')) "
```

```
cowsay is a program that generates ASCII art pictures of a cow with a message. It
↳ can also generate pictures using pre-made images of other animals, such as Tux
↳ the Penguin, the Linux mascot. It is written in Perl. There is also a related
↳ program called cowthink, with cows with thought bubbles rather than speech
↳ bubbles. .cow files for cowsay exist which are able to produce different
↳ variants of "cows", with different kinds of "eyes", and so forth. It is
↳ sometimes used on IRC, desktop screenshots, and in software documentation. It is
↳ more or less a joke within hacker culture, but has been around long enough that
↳ its use is rather widespread. In 2007, it was highlighted as a Debian package of
↳ the day.
```

It is worth noting that if we uninstall `wikipedia`, the dependencies that `wikipedia` needs (such as `beautifulsoup4`) are *not* uninstalled:

```
!source myvirtualenv/bin/activate && pip uninstall -y wikipedia
```

```
Found existing installation: wikipedia 1.4.0
Uninstalling wikipedia-1.4.0:
```

(continues on next page)

(continued from previous page)

```
Successfully uninstalled wikipedia-1.4.0
```

```
!source myvirtualenv/bin/activate && pip list
```

Package	Version
-----	-----
beautifulsoup4	4.10.0
certifi	2021.10.8
charset-normalizer	2.0.10
idna	3.3
pip	21.3.1
requests	2.27.0
setuptools	57.4.0
soupsieve	2.3.1
urllib3	1.26.7

This can lead to an accumulation of (partly unneeded) python packages. Also for this reason, it is good practice to create a virtual environment *from scratch* when starting a new project, and to discard it afterwards.

18.3.5 Installing particular versions with pip

Occasionally, we need to install a particular version of a package. For example, imagine we need version 2.0 of cowsay. In that case, we can use the == operator to specify this requirement:

```
!source myvirtualenv/bin/activate && pip install cowsay==3.0
```

```
Collecting cowsay==3.0
  Using cached cowsay-3.0-py2.py3-none-any.whl (19 kB)
Installing collected packages: cowsay
Successfully installed cowsay-3.0
```

```
!source myvirtualenv/bin/activate && cowsay --version
```

```
3.0
```

18.3.6 Upgrading a pip-installed package

```
!source myvirtualenv/bin/activate && pip install -U cowsay
```

```
Requirement already satisfied: cowsay in ./myvirtualenv/lib/python3.9/site-
packages (3.0)
Collecting cowsay
  Using cached cowsay-4.0-py2.py3-none-any.whl (24 kB)
Installing collected packages: cowsay
  Attempting uninstall: cowsay
    Found existing installation: cowsay 3.0
    Uninstalling cowsay-3.0:
```

(continues on next page)

(continued from previous page)

```
Successfully uninstalled cowsay-3.0  
Successfully installed cowsay-4.0
```

```
!source myvirtualenv/bin/activate && cowsay --version
```

```
4.0
```

Let's remove cowsay again:

```
!source myvirtualenv/bin/activate && pip uninstall -y cowsay
```

```
Found existing installation: cowsay 4.0  
Uninstalling cowsay-4.0:  
Successfully uninstalled cowsay-4.0
```

18.3.7 Installing a package from github

If we want to install the latest development version of the (Python) cowsay package, we have two options.

The first option is to pip install directly from the github. The github repository is at <https://github.com/VaasuDevanS/cowsay-python>

```
!source myvirtualenv/bin/activate && pip install git+https://github.com/VaasuDevanS/  
↪cowsay-python.git
```

```
Collecting git+https://github.com/VaasuDevanS/cowsay-python.git  
  Cloning https://github.com/VaasuDevanS/cowsay-python.git to /private/var/folders/  
↪wc/d1lyft3x2jn29b6yffrzh4vw0000gq/T/pip-req-build-7gxyjbt6  
  Running command git clone --filter=blob:none -q https://github.com/VaasuDevanS/  
↪cowsay-python.git /private/var/folders/wc/d1lyft3x2jn29b6yffrzh4vw0000gq/T/pip-  
↪req-build-7gxyjbt6  
  Resolved https://github.com/VaasuDevanS/cowsay-python.git to commit_  
↪767c09425d813b80d67cdebba02ce387ca2eb4e8  
  Preparing metadata (setup.py) ... ?25ldone  
?25hUsing legacy 'setup.py install' for cowsay, since package 'wheel' is not_  
↪installed.  
Installing collected packages: cowsay  
  Running setup.py install for cowsay ... ?25ldone  
?25hSuccessfully installed cowsay-4.0
```

```
!source myvirtualenv/bin/activate && cowsay --version
```

```
4.0
```

```
!source myvirtualenv/bin/activate && pip uninstall -y cowsay
```

```
Found existing installation: cowsay 4.0  
Uninstalling cowsay-4.0:  
Successfully uninstalled cowsay-4.0
```

The second option is to clone the git repository to our local machine, and then to install the package from that local directory:

```
!cd /tmp && git clone https://github.com/VaasuDevanS/cowsay-python.git
```

```
Cloning into 'cowsay-python'...
remote: Enumerating objects: 170, done.
remote: Counting objects: 100% (82/82), done.
remote: Compressing objects: 100% (40/40), done.
remote: Total 170 (delta 41), reused 77 (delta 40), pack-reused 88
Receiving objects: 100% (170/170), 79.19 KiB | 772.00 KiB/s, done.
Resolving deltas: 100% (72/72), done.
```

```
!source myvirtualenv/bin/activate && cd /tmp/cowsay-python && pip install .
```

```
Processing /private/tmp/cowsay-python
  Preparing metadata (setup.py) ... ?25ldone
?25hUsing legacy 'setup.py install' for cowsay, since package 'wheel' is not
  installed.
Installing collected packages: cowsay
  Running setup.py install for cowsay ... ?25ldone
?25hSuccessfully installed cowsay-4.0
```

```
!source myvirtualenv/bin/activate && cowsay --version
```

```
4.0
```

18.3.8 Pip install a user-editable package from a local directory

This example carries on from the `git clone` example above.

If we `pip-install` python packages, these are normally installed in the directory tree in the virtual environment. For example:

```
!ls myvirtualenv/lib/python3.*/site-packages/cowsay
```

```
__init__.py  __pycache__  main.py
__main__.py  characters.py test.py
```

If we intend to edit the python files in the package (for example because we want to develop it further, or explore it), and we want those edits to be visible in the ‘installed’ package, we can ask `pip` to carry out an `editable` install using the `-e` flag:

```
!source myvirtualenv/bin/activate && pip uninstall -y cowsay
```

```
Found existing installation: cowsay 4.0
Uninstalling cowsay-4.0:
  Successfully uninstalled cowsay-4.0
```

```
!source myvirtualenv/bin/activate && cd /tmp/cowsay-python && pip install -e .
```



```
Obtaining file:///private/tmp/cowsay-python
Preparing metadata (setup.py) ... ?25ldone
?25hInstalling collected packages: cowsay
Running setup.py develop for cowsay
Successfully installed cowsay-4.0
```

In this case, only a link to our local package is created:

```
!ls -l myvirtualenv/lib/python3.*/site-packages/cowsay*
```

```
-rw-r--r-- 1 fangohr staff 28 Jan 5 11:55 myvirtualenv/lib/python3.9/site-
  ↪packages/cowsay.egg-link
```

```
!cat myvirtualenv/lib/python3.*/site-packages/cowsay*
```

```
/private/tmp/cowsay-python
.
```

18.3.9 Advance pip use: freeze, -r requirements.txt and creating reproducible environments

If you want to record (and later re-use) a combination of python packages *with their specific version numbers*, you can use the pip freeze command to provide such a list.

```
!source myvirtualenv/bin/activate && pip freeze
```

```
beautifulsoup4==4.10.0
certifi==2021.10.8
charset-normalizer==2.0.10
-e git+https://github.com/VaasuDevanS/cowsay-python.
  ↪git@767c09425d813b80d67cdebba02ce387ca2eb4e8#egg=cowsay
idna==3.3
requests==2.27.0
soupsieve==2.3.1
urllib3==1.26.7
```

We can re-direct the output into a file (which by convention is called requirements.txt):

```
!source myvirtualenv/bin/activate && pip freeze > requirements.txt
```

```
!cat requirements.txt
```

```
beautifulsoup4==4.10.0
certifi==2021.10.8
charset-normalizer==2.0.10
-e git+https://github.com/VaasuDevanS/cowsay-python.
  ↪git@767c09425d813b80d67cdebba02ce387ca2eb4e8#egg=cowsay
idna==3.3
requests==2.27.0
```

(continues on next page)

(continued from previous page)

```
soupsieve==2.3.1
urllib3==1.26.7
```

We can now create a new virtual environment, and install all the packages listed in the `requirements.txt` file into this new virtual environment:

```
!python -m venv myvirtualenv-copy
```

```
!source myvirtualenv-copy/bin/activate && pip install -r requirements.txt
```

```
Obtaining cowsay from git+https://github.com/VaasuDevanS/cowsay-python.
↳git@767c09425d813b80d67cdebba02ce387ca2eb4e8#egg=cowsay (from -r requirements.
↳txt (line 4))
Cloning https://github.com/VaasuDevanS/cowsay-python.git (to revision
↳767c09425d813b80d67cdebba02ce387ca2eb4e8) to ./myvirtualenv-copy/src/cowsay
Running command git clone -q https://github.com/VaasuDevanS/cowsay-python.git /
↳Users/fangohr/git/introduction-to-python-for-computational-science-and-
↳engineering/book/myvirtualenv-copy/src/cowsay
Running command git rev-parse -q --verify 'sha^
↳767c09425d813b80d67cdebba02ce387ca2eb4e8'
Running command git fetch -q https://github.com/VaasuDevanS/cowsay-python.git
↳767c09425d813b80d67cdebba02ce387ca2eb4e8
Resolved https://github.com/VaasuDevanS/cowsay-python.git to commit
↳767c09425d813b80d67cdebba02ce387ca2eb4e8
Collecting beautifulsoup4==4.10.0
Using cached beautifulsoup4-4.10.0-py3-none-any.whl (97 kB)
Collecting certifi==2021.10.8
Using cached certifi-2021.10.8-py2.py3-none-any.whl (149 kB)
Collecting charset-normalizer==2.0.10
Using cached charset_normalizer-2.0.10-py3-none-any.whl (39 kB)
Collecting idna==3.3
Using cached idna-3.3-py3-none-any.whl (61 kB)
Collecting requests==2.27.0
Using cached requests-2.27.0-py2.py3-none-any.whl (63 kB)
Collecting soupsieve==2.3.1
Using cached soupsieve-2.3.1-py3-none-any.whl (37 kB)
Collecting urllib3==1.26.7
Using cached urllib3-1.26.7-py2.py3-none-any.whl (138 kB)
Installing collected packages: urllib3, soupsieve, idna, charset-normalizer,
↳certifi, requests, cowsay, beautifulsoup4
Running setup.py develop for cowsay
Successfully installed beautifulsoup4-4.10.0 certifi-2021.10.8 charset-normalizer-
↳2.0.10 cowsay-4.0 idna-3.3 requests-2.27.0 soupsieve-2.3.1 urllib3-1.26.7
WARNING: You are using pip version 21.2.3; however, version 21.3.1 is available.
You should consider upgrading via the '/Users/fangohr/git/introduction-to-python-
↳for-computational-science-and-engineering/book/myvirtualenv-copy/bin/python -m
↳pip install --upgrade pip' command.
```

It is good practice to use the `freeze` command to store the list of packages and versions required for important projects (including those such as scientific publications, reports, theses), and the `requirements.txt` file should be archived together with the data and software.

It is even better practice if the *creation* of the virtual environment is done in a scripted way (based on a `requirements.txt` file which should be part of the archived [and version-controlled] files of the analysis), and before all the required processing/simulation/analysis is in that environment.

In practice, achieving full and guaranteed reproducibility is difficult. There are a variety of problems that could occur, such as for example the disappearance of the `pypi.org` service. How to achieve full reproducibility is an active research area, and deserves a separate chapter or book.

In any case, recording the Python packages used is a very good first step.

18.3.10 Deactivate a virtual environment

To de-activate a virtual environment, use the `deactivate` command.

```
!source myvirtualenv/bin/activate && deactivate && which python
```

```
/Users/fangohr/anaconda3/bin/python
```

18.3.11 Deleting a virtual environment

To completely remove the virtual environment, we can delete the subfolder in which it was installed:

```
!rm -rf myvirtualenv
```

18.3.12 Further reading

- Installing Python packages: <https://packaging.python.org/en/latest/tutorials/installing-packages>
- `venv` module documentation: <https://docs.python.org/3/library/venv.html>

18.4 Anaconda

18.4.1 Introduction

The Anaconda software distribution brings its own packaging, which is controlled through the `conda` command.

Anaconda is well known as a Python distribution, but is in no way limited to Python packages: it is a generic package manager. Together with the community-run `conda-forge` project, there is a multitude of packages available. A particular bonus of the the `conda` packages is that they can be provided for the three major operating systems in use (Linux, OSX, Windows).

Conda provides **conda packages** for some Python packages that are available from the Python Package Index (PyPI). One thus needs to ask: should I install a package using the `conda` command (`conda install spyder`) or through `pip` (`pip install spyder`). See below for the answer.

`conda` provides its own (**conda**) **environments** (see <https://docs.conda.io/projects/conda/en/latest/user-guide/tasks/manage-environments.html>). They show many similarities with our (basic) discussion of (Python) virtual environments. We have no space to discuss the `conda` environments further here.

The following comments are meant to be helpful for those who have installed their Python interpreter through `anaconda`. If you are not using Anaconda, you can ignore this section.

18.4.2 Can I use Python virtual environments when using the anaconda distribution?

This can be done and is a good way to create virtual environments. (All the examples above in this chapter use the Python3 interpreter from an anaconda installation on an OSX system).

18.4.3 Should I install a python package through conda or pip?

The typical scenario is that one installs Anaconda, and most Python packages needed are available: somewhat standard tools such as `numpy`, `scipy`, `matplotlib`, `pandas`, `jupyter`, `ipython` and `spyder` already come with the anaconda distribution. Then some package is missing that needs to be installed additionally.

For example, the package `xarray`: this can be installed through conda or through pip.

Experience-based rough guidance is as follows if working within an anaconda installation of Python:

- avoid mixing `pip` installs with `conda` installs
- if `conda` can install the required packages, then use that
- if `conda` cannot install the required packages, we have to use `pip`. In that case:
 - install the requirements that need to come from `conda` (if any)
 - then install the desired packages through `pip`
 - after having used `pip`, do not use `conda` again to install more packages.

The reason for this is that `conda` and `pip` cannot interact perfectly, and so the changes that one package manager made, maybe overridden or accidentally repeated slightly differently by the other one.

A more detailed discussion is available on the [Anaconda blog](#).

18.4.4 Can I create a conda environment, and then create python virtual environments from this?

Yes.

This is also an option to install different Python versions.

For example: create a conda (!) environment providing Python 3.8:

```
conda create -y -n python38 python=3.8
```

Then

```
conda activate python38
```

and then create a virtual environment using

```
python -m venv myvirtualenv38
```

18.5 Managing many different environments - pyenv

If you use many different python environments, possibly with different interpreter versions, you may want to learn about pyenv (home page at <https://github.com/pyenv/pyenv>).

Pyenv can install a multitude of python interpreters, create virtual environments for each of those. It is further possible to define an a *per-directory* basis which environment should be used in that directory. This is convenient when using different enviroments for different projects, as one does not need to manually activate the virtual environments.

Tidy up: remove files created in this section

```
!rm -rf /tmp/cowsay-python
!rm -rf myvirtualenv-copy
!rm -f requirements.txt
```


WHERE TO GO FROM HERE?

Learning a programming language is the first step towards becoming a computationalist who advances science and engineering through computational modelling and simulation.

We list some additional skills that can be very beneficial for day-to-day computational science work, but is of course not exhaustive.

19.1 Advanced programming

This text has put emphasis on providing a robust foundation in terms of programming, covering control flow, data structures and elements from function and procedural programming. We have not touch Object Orientation in great detail, nor have we discussed some of Python's more advanced features such as iterators, and decorators, type hinting, nor many of the fantastic (standard) libraries available.

19.2 Compiled programming language

When performance starts to be the highest priority, we may need to use compiled code, and likely embed this in a Python code to carry out the computational that are the performance bottle neck.

Fortran, C and C++ are sensible choices here; maybe Rust in the near future.

We may also need to learn how to integrate the compiled code with Python using tools such as Cython, Boost, Ctypes and Swig.

With the rise of GPUs as cheap and powerful compute resources, it is likely we want to drive computation carried out on the GPU. This can be done through GPU-specific libraries and languages (CUDA and OpenCL, for example). For some use cases, it may be sufficient to use frameworks that translate computational work from a higher level language (ideally as high as Python) to the GPUs.

19.3 Testing

Good software development is supported by a range of unit and system tests that can be run routinely to check that the code works correctly. Tools such as pytest, doctest and others are invaluable, and we should at least learn at least how to use pytest for automated tests.

19.4 Simulation models

A number of standard simulation tools such as Monte Carlo, Molecular Dynamics, lattice based models, agents, finite difference and finite element models are commonly used to solve particular simulation challenges – it is useful to have at least a broad overview of these.

19.5 Software engineering for research codes

Research codes bring particular challenges: the requirements may change during the run time of the project, we need great flexibility yet reproducibility. A number of techniques are available to support effectively, including version control (see below), automatic tests and continuous integration.

19.6 Data and visualisation

Dealing with large amounts of data, processing and visualising it can be a challenge. Fundamental knowledge of database design, 3d visualisation and modern data processing tools such as the Pandas and xarray Python package help with this. For interactive 3d visualisation VTK remains an important tool, although WebGL starts to be an interesting alternative.

19.7 Version control

Using a version control tool, such as git, should be a standard approach and improves code writing effectiveness significantly, helps with working in teams, and - maybe most importantly - supports reproducibility of computational results.

19.8 Parallel execution

Parallel execution of code is a way to make it run orders of magnitude faster. This could be using MPI for inter-node communication or OpenMP for intra-node parallelisation or a hybrid mode bringing both together.

The recent rise of GPU computing provides yet another avenue of parallelisation.

19.9 Acknowledgements

Big thanks go to

- Marc Molinari for carefully proof reading this manuscript around 2007.
- Neil O'Brien for contributing to the SymPy section.
- Jacek Generowicz for introducing me to Python in the last millennium, and for kindly sharing countless ideas from his excellent Python course.
- EPSRC (GR/T09156/01 and EP/G03690X/1) and the European Union (OpenDreamKit Horizon 2020 European Research Infrastructures project, #676541) for support.
- Students and other readers who have provided feedback and pointed out typos and errors etc.
- Thomas Kluyver who helped to translate the Python 2 LaTeX based document into Python 3 Jupyter Notebooks and provided the machinery to create html and pdf versions. automatically (via his [bookbook package](#)).

- Robert Rosca who helped to create html and pdf files after using `jupyterbook` was released (2020).

[1] the vertical line is to show the division between the original components only; mathematically, the augmented matrix behaves like any other 2×3 matrix, and we code it in SymPy as we would any other.

[2] from the `help(preview)` documentation: “Currently this depends on `pexpect`, which is not available for windows.”

[3] The exact value for the upper limit is available in `sys.maxint`.

[4] We add for completeness, that a C-program (or C++ or Fortran) that executes the same loop will be about 100 times faster than the python float loop, and thus about $100 \times 200 = 20000$ faster than the symbolic loop.

[5] In this text, we usually import `numpy` under the name `N` like this: `import numpy as N`. If you don't have `numpy` on your machine, you can substitute this line by `import Numeric as N` or `import numarray as N`.

[6] Historical note: this has changed from `scipy` version 0.7 to 0.8. Before 0.8, the return value was a float if a one-dimensional problem was to solve.

CHANGE HISTORY

Since 2022

- 3 Jan 2022: review *visualisation chapter*
- 4 Jan 2022: change from `odeint` to `solve_ivp` in *scipy chapter*
- 5 Jan 2022: adding new section *virtual environments and pip*