

Undergraduate Topics in Computer Science

John Hunt

A Beginners Guide to Python3 Programming

Second Edition




Undergraduate Topics in Computer Science


Series Editor

Ian Mackie, University of Sussex, Brighton, UK


Advisory Editors

Samson Abramsky , Department of Computer Science, University of Oxford, Oxford, UK

Chris Hankin , Department of Computing, Imperial College London, London, UK


Mike Hinchey , Lero – The Irish Software Research Centre, University of Limerick, Limerick, Ireland

Dexter C. Kozen, Department of Computer Science, Cornell University, Ithaca, NY, USA

Andrew Pitts , Department of Computer Science and Technology, University of Cambridge, Cambridge, UK

Hanne Riis Nielson , Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kongens Lyngby, Denmark

Steven S. Skiena, Department of Computer Science, Stony Brook University, Stony Brook, NY, USA

Iain Stewart , Department of Computer Science, Durham University, Durham, UK

Joseph Migga Kizza, College of Engineering and Computer Science, The University of Tennessee-Chattanooga, Chattanooga, TN, USA

'Undergraduate Topics in Computer Science' (UTiCS) delivers high-quality instructional content for undergraduates studying in all areas of computing and information science. From core foundational and theoretical material to final-year topics and applications, UTiCS books take a fresh, concise, and modern approach and are ideal for self-study or for a one- or two-semester course. The texts are all authored by established experts in their fields, reviewed by an international advisory board, and contain numerous examples and problems, many of which include fully worked solutions.

The UTiCS concept relies on high-quality, concise books in softback format, and generally a maximum of 275–300 pages. For undergraduate textbooks that are likely to be longer, more expository, Springer continues to offer the highly regarded Texts in Computer Science series, to which we refer potential authors.

John Hunt

A Beginners Guide to Python 3 Programming

Second Edition

 Springer

John Hunt
Midmarsh Technology Ltd.
Chippenham, Wiltshire, UK

ISSN 1863-7310 ISSN 2197-1781 (electronic)
Undergraduate Topics in Computer Science
ISBN 978-3-031-35121-1 ISBN 978-3-031-35122-8 (eBook)
<https://doi.org/10.1007/978-3-031-35122-8>

1st edition: © Springer Nature Switzerland AG 2019, corrected publication 2020
2nd edition: © Springer Nature Switzerland AG 2023

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors, and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

*This book is simply for Denise, the love of my
life.*

Preface the Second Edition

I am writing this preface to the second edition of this book in the middle of 2023, and I am reflecting on what has happened to Python over the last 4 to 5 years. Python itself has gone from strength to strength, its popularity has not diminished and if anything has increased. I increasingly come across organizations in all walks of life that are actively using the language. This is true whether they are schools, colleges or universities, charities, governmental organization, research institutions or commercial enterprises. The uses being made of Python keeps on growing and the ecosystem of supporting libraries available keeps on growing.

This edition represents a significant update from Python 3.7 (which was the current version when the first edition of the book was written) to Python 3.11 (and a look ahead to Python 3.12). It also represents a significant extension to the original book with an additional eight chapters. These chapters represent topics and questions that those learning Python have expressed an interest in such as frozensets, enumerated types and sorting functions to Python virtual environments. It also reflects new features in the language such as pattern matching introduced in Python 3.10 and Exception Groups handling and Exception Notes introduced in Python 3.11.

The book also now includes a chapter entitled 'Python Bites Back' which highlights for someone learning Python, 14 of the most common mistakes made by a Python beginner. This chapter can be dropped into at any time but covers the whole of the language discussed in this book so some sections may be more relevant once the whole book has been read.

Chippenham, UK

John Hunt

Preface to the First Edition

There is currently huge interest in the Python programming language. This is driven by several factors; its use in schools with the Raspberry Pi platform, its ability to be used for DevOps scripts, its use in Data Science and Machine Learning and of course the language itself. Plus of course it is free! So anyone anywhere can download and install Python and get started.

There are many books on Python; however, most assume previous programming experience or are focused on particular aspects of Python use such as Data Science or Machine Learning or have a scientific flavor.

The aim of this book is to introduce Python to those with little or very little programming knowledge, but to take them through to become an experienced Python developer.

As such the earlier parts of the book introduce fundamental concepts such as what a *variable* is and how a *for loop* works. In contrast the later chapters introduce advanced concepts such as Functional Programming, object orientation and exception handling.

In between a wide range of topics are introduced and discussed from a Python point of view including functions, recursion, operators, Python properties, modules and packages, protocols and monkey patching, etc.

After the core elements of Python are established each new subject area is introduced by way of an introductory chapter presenting the topic in general, providing background on that subject, why it is of importance, etc. These introductions cover structure analysis, functional programming and object orientation.

Some of the key aspects of this book are:

1. It assumes very little knowledge or experience of Python or programming.
2. It provides a basic introduction to Python as well as advanced topics such as generators and coroutines.
3. This book provides extensive coverage of object orientation and the features in Python 3 supporting classes, inheritance and protocols.
4. Python's support for functional programming is also presented.

5. Following on from introducing the basic ideas behind functional programming, the book presents how advanced functional concepts such as closures, currying and higher-order functions work in Python.
6. The book includes exercises at the end of most chapters with online solutions
7. There are several case studies spread through the book that broaden understanding of preceding topics.
8. All code examples (and exercise solutions) are provided online in a GitHub repository.

Chapter Organization

Each chapter has a brief introduction, the main body of the chapter, followed by a list of (typically) online references that can be used for further reading.

Following this there is typically an *Exercises* section that lists one or more exercises that build on the skills you will have learnt in that chapter.

Sample solutions to the exercises are available in a GitHub online repository that supports this book.

What You Need

You can of course just read this book; however following the examples in this book will ensure that you get as much as possible out of the content.

For this you will need a computer.

Python is a cross platform programming language and as such you can use Python on a Windows PC, a Linux box or a Apple Mac, etc. So you are not tied to a particular type of operating system; you can use whatever you have available.

However you will need to install some software on that computer. At a minimum you will need Python.

This book focuses on Python 3 so you will need that. Some guidance on this is provided in Chap. 2 on setting up your environment.

You will also need some form of editor in which to write your programs. There are numerous generic programming editors available for different operating systems with VIM on Linux, Notepad++ on Windows and Sublime Text on Windows and Macs being popular choices.

However, using a Integrated Development Environment (IDE) editor such as PyCharm or potential Visual Studio Code will make writing and running your programs much easier.

Using an IDE

The IDE I prefer for Python is PyCharm, it is not the only IDE for Python by any means, but it is a very widely used one.

Other IDEs available for Python include:

- Spyder which is an IDE often used by scientists and engineers see <https://www.spyder-ide.org/>.
- Jupiter Notebook which is a web-based IDE and is particularly good for data scientists <https://jupyter.org/>.
- Visual Studio Code. This is a very good free editor from Microsoft that has really useful features <https://code.visualstudio.com>.
- Sublime Text is more of a text editor that color codes Python; however for a simple project it may be all you need <https://www.sublimetext.com>.

Mostly this book is editor or IDE agnostic only Chap. 3 ‘Setting up the Python Environment’ and Chap. 42 ‘Python Virtual Environments’ really focus on PyCharm. You can therefore use whatever editor you wish to use.

Conventions

Throughout this book you will find a number of conventions used for text styles. These text styles distinguish different kinds of information.

Code words, variables and Python values, used within the main body of the text, are shown using a Courier font. For example:

This program creates a top level window (the `wx.Frame`) and gives it a title. It also creates a label (a `wx.StaticText` object) to be displayed within the frame.

In the above paragraph `wx.Frame` and `wx.StaticText` are classes available in a Python graphical user interface library.

A block of Python code is set out as shown here:

```
num = int(input( 'Enter another number: ' ))
if num > 0:
    print(num, 'is positive')
    print(num, 'squared is ', num * num)

print('Bye')
```

Note that keywords and points of interest are shown in bold font.

Any command line or user input is shown in italics and is also bold:

```
> python hello.py  
  
Or  
  
Hello, world  
Enter your name: John  
Hello John
```

Example Code and Sample Solutions

The examples used in this book (along with sample solutions for the exercises at the end of most chapters) are available in a GitHub repository. GitHub provides a web interface and a server environment hosting Git.

Git is a version control system typically used to manage source code files (such as those used to create systems in programming languages such as Python but also Java, C#, C++, Scala, etc.). Systems such as Git are very useful for collaborative development as they allow multiple people to work on an implementation and to merge their work together. They also provide a useful historical view of the code (which also allows developers to roll back changes if modifications prove to be unsuitable).

If you already have Git installed on your computer then you can clone (obtain a copy of) the repository locally using:

```
git clone https://github.com/johnehunt/beginnerspython3\_2nd
```

If you do not have git then you can obtain a zip file of the examples using https://github.com/johnehunt/beginnerspython3_2nd/archive/master.zip

You can of course install Git yourself if you wish. To do this see <https://git-scm.com/downloads>. Versions of the Git client for Mac OS, Windows and Linux/Unix are available here.

However, many IDEs such as PyCharm come with Git support and so offer another approach to obtaining a Git repository.

For more information on Git see <http://git-scm.com/doc>. This Git guide provides a very good primer and is highly recommended.

Bath, UK

John Hunt

Contents

1	Introduction	1
1.1	What is Python?	1
1.2	Python Versions	2
1.3	Python History	3
1.4	Python Programming	4
1.5	Python Libraries	4
1.6	Python Execution Model	5
1.7	Running Python Programs	6
1.7.1	Interactively Using the Python Interpreter	6
1.7.2	Running a Python File	7
1.7.3	Executing a Python Script	9
1.7.4	Using Python in an IDE	11
1.8	Jupyter Notebooks	12
1.9	Useful Resources	13
2	Where is Python Used	15
2.1	Introduction	15
2.1.1	Data Analytics	15
2.1.2	Machine Learning and AI	17
2.1.3	Database Work	18
2.1.4	Python for Animation	19
2.1.5	Python for Film Making	19
2.1.6	Cross Platform UIs	19
2.1.7	Games Programming	20
2.1.8	Integration Testing Framework	20
2.1.9	Academic Research	20
2.1.10	Web Services	21
2.2	Useful Resources	21

- 3 Setting up the Python Environment** 23
 - 3.1 Introduction 23
 - 3.2 Check to See if Python is Installed 23
 - 3.3 Installing Python on a Windows PC 25
 - 3.3.1 Step 1: Downloading Python 25
 - 3.3.2 Step 2: Running the Installer 26
 - 3.3.3 Step 3: Set up the PATH (Optional) 27
 - 3.3.4 Step 4: Verify the Installation 29
 - 3.3.5 Step 5: Run Some Python 30
 - 3.3.6 Step 6: Exit the Python Interpreter 30
 - 3.4 Setting up on a Mac 31
 - 3.5 Setting up an IDE 33
 - 3.5.1 Downloading the PyCharm IDE 33
 - 3.5.2 Setting up the IDE 35
 - 3.6 Writing Python Programs 37
 - 3.7 A Note on Filenames 39
 - 3.8 Online Resources 40

- 4 A First Python Program** 41
 - 4.1 Introduction 41
 - 4.2 Hello World 41
 - 4.3 Interactive Hello World 42
 - 4.4 Variables 44
 - 4.5 Naming Conventions 45
 - 4.6 Assignment Operator 46
 - 4.7 Python Statements 47
 - 4.8 Comments in Code 47
 - 4.9 Scripts Versus Programs 48
 - 4.10 Online Resources 48
 - 4.11 Exercises 49

- 5 Python Strings** 51
 - 5.1 Introduction 51
 - 5.2 What Are Strings? 51
 - 5.3 Representing Strings 52
 - 5.4 What Type is String? 53
 - 5.5 What Can You Do with Strings? 53
 - 5.5.1 String Concatenation 54
 - 5.5.2 Length of a String 54
 - 5.5.3 Accessing a Character 54
 - 5.5.4 Accessing a Subset of Characters 55
 - 5.5.5 Repeating Strings 55
 - 5.5.6 Splitting Strings 55
 - 5.5.7 Counting Strings 56
 - 5.5.8 Replacing Strings 57
 - 5.5.9 Finding Sub Strings 57

5.5.10	Converting Other Types into Strings	57
5.5.11	Remove Prefix and Suffix	58
5.5.12	Comparing Strings	58
5.5.13	Other String Operations	59
5.6	Hints on Strings	60
5.6.1	Python Strings are Case Sensitive	60
5.6.2	Function/Method Names	60
5.6.3	Function/Method Invocations	60
5.7	String Formatting	61
5.8	String Templates	63
5.9	Formatting Using f-strings	66
5.10	Online Resources	67
5.11	Exercises	67
6	Numbers, Booleans and None	69
6.1	Introduction	69
6.2	Types of Numbers	69
6.3	Integers	70
6.3.1	Converting to Ints	71
6.4	Floating Point Numbers	71
6.4.1	Converting to Floats	72
6.4.2	Converting an Input String into a Floating Point Number	72
6.5	Complex Numbers	73
6.6	Boolean Values	73
6.7	Arithmetic Operators	74
6.7.1	Integer Operations	75
6.7.2	Negative Number Integer Division	77
6.7.3	Floating Point Number Operators	77
6.7.4	Integers and Floating Point Operations	78
6.7.5	Complex Number Operators	78
6.8	Assignment Operators	79
6.9	None Value	80
6.10	Online Resources	81
6.11	Exercises	81
6.11.1	General Exercise	81
6.11.2	Convert Kilometers to Miles	81
7	Flow of Control Using if Statements	83
7.1	Introduction	83
7.2	Comparison Operators	83
7.3	Logical Operators	84
7.4	The If Statement	84
7.4.1	Working with an If Statement	85
7.4.2	Else in an If Statement	86
7.4.3	The Use of elif	86

- 7.5 Nesting if Statements 87
- 7.6 If Expressions 88
- 7.7 A Note on True and False 89
- 7.8 Hints 89
- 7.9 Online Resources 89
- 7.10 Exercises 90
 - 7.10.1 Check Input is Positive or Negative 90
 - 7.10.2 Test if a Number is Odd or Even 90
 - 7.10.3 Kilometers to Miles Converter 90
- 8 Iteration/Looping 93**
 - 8.1 Introduction 93
 - 8.2 While Loop 93
 - 8.3 For Loop 95
 - 8.3.1 Python for Loops 96
 - 8.3.2 Looping by Non-default Increments 97
 - 8.3.3 Anonymous Loop Variables 98
 - 8.4 Negative Loops 98
 - 8.5 Break Loop Statement 99
 - 8.6 Continue Loop Statement 100
 - 8.7 For Loop with Else 102
 - 8.8 A Note on Loop Variable Naming 103
 - 8.9 Dice Roll Game 103
 - 8.10 Online Resources 104
 - 8.11 Exercises 104
 - 8.11.1 Calculate the Factorial of a Number 104
 - 8.11.2 Print All the Prime Numbers in a Range 105
- 9 Number Guessing Game 107**
 - 9.1 Introduction 107
 - 9.2 Setting up the Program 107
 - 9.2.1 Create a New Python File 107
 - 9.2.2 Add a Welcome Message 108
 - 9.2.3 Running the Program 109
 - 9.3 What Will the Program Do? 110
 - 9.4 Creating the Game 111
 - 9.4.1 Generate the Random Number 111
 - 9.4.2 Obtain an Input From the User 111
 - 9.4.3 Check to See if the Player Has Guessed
the Number 112
 - 9.4.4 Check They Haven't Exceeded Their Maximum
Number of Guess 112
 - 9.4.5 Notify the Player Whether Higher or Lower 113
 - 9.4.6 End of Game Status 114
 - 9.5 The Complete Listing 114
 - 9.6 Hints 116

- 9.6.1 Initializing Variables 116
- 9.6.2 Blank Lines Within a Block of Code 116
- 9.7 Exercises 116
- 10 Recursion 117**
 - 10.1 Introduction 117
 - 10.2 Recursive Behaviour 117
 - 10.3 Benefits of Recursion 118
 - 10.4 Recursively Searching a Tree 118
 - 10.5 Recursion in Python 119
 - 10.6 Calculating Factorial Recursively 120
 - 10.7 Disadvantages of Recursion 122
 - 10.8 Online Resources 122
 - 10.9 Exercises 123
- 11 Introduction to Structured Analysis 125**
 - 11.1 Introduction 125
 - 11.2 Structured Analysis and Function Identification 125
 - 11.3 Functional Decomposition 126
 - 11.3.1 Functional Decomposition Terminology 127
 - 11.3.2 Functional Decomposition Process 128
 - 11.3.3 Calculator Functional Decomposition Example 128
 - 11.4 Functional Flow 130
 - 11.5 Data Flow Diagrams 130
 - 11.6 Flowcharts 131
 - 11.7 Data Dictionary 133
 - 11.8 Online Resources 134
- 12 Functions in Python 135**
 - 12.1 Introduction 135
 - 12.2 What are Functions? 135
 - 12.3 How Functions Work 136
 - 12.4 Types of Functions 137
 - 12.5 Defining Functions 137
 - 12.5.1 An Example Function 138
 - 12.6 Returning Values from Functions 139
 - 12.7 Docstring 140
 - 12.8 ReStructured Text 141
 - 12.9 Function Parameters 146
 - 12.9.1 Multiple Parameter Functions 146
 - 12.9.2 Default Parameter Values 147
 - 12.9.3 Named Arguments 148
 - 12.9.4 Arbitrary Arguments 149
 - 12.9.5 Positional and Keyword Arguments 150
 - 12.10 Anonymous Functions 151

- 12.11 Online Resources 152
- 12.12 Exercises 152
- 13 Scope and Lifetime of Variables 155**
 - 13.1 Introduction 155
 - 13.2 Local Variables 155
 - 13.3 The Global Keyword 157
 - 13.4 Nonlocal Variables 158
 - 13.5 Hints 160
 - 13.6 Online Resources 160
 - 13.7 Exercise 160
- 14 Implementing a Calculator Using Functions 161**
 - 14.1 Introduction 161
 - 14.2 What the Calculator Will Do 161
 - 14.3 Getting Started 162
 - 14.4 The Calculator Operations 162
 - 14.5 Behaviour of the Calculator 163
 - 14.6 Identifying Whether the User has Finished 164
 - 14.7 Selecting the Operation 166
 - 14.8 Obtaining the Input Numbers 168
 - 14.9 Determining the Operation to Execute 169
 - 14.10 Running the Calculator 170
 - 14.11 Exercises 170
- 15 Introduction to Functional Programming 171**
 - 15.1 Introduction 171
 - 15.2 What is Functional Programming? 171
 - 15.3 Advantages to Functional Programming 173
 - 15.4 Disadvantages of Functional Programming 175
 - 15.5 Referential Transparency 175
 - 15.6 Further Reading 177
- 16 Higher-Order Functions 179**
 - 16.1 Introduction 179
 - 16.2 Recap on Functions in Python 179
 - 16.3 Functions as Objects 180
 - 16.4 Higher-Order Function Concepts 182
 - 16.4.1 Higher-Order Function Example 183
 - 16.5 Python Higher-Order Functions 184
 - 16.5.1 Using Higher-Order Functions 185
 - 16.5.2 Functions Returning Functions 186
 - 16.6 Online Resources 187
 - 16.7 Exercises 188

- 17 Curried Functions** 189
 - 17.1 Introduction 189
 - 17.2 Currying Concepts 189
 - 17.3 Python and Curried Functions 190
 - 17.4 Closures 191
 - 17.5 Online Resources 194
 - 17.6 Exercises 194

- 18 Introduction to Object Orientation** 195
 - 18.1 Introduction 195
 - 18.2 Classes 195
 - 18.3 What Are Classes for? 196
 - 18.3.1 What Should a Class Do? 197
 - 18.3.2 Class Terminology 197
 - 18.4 How is an OO System Constructed? 198
 - 18.4.1 Where Do We Start? 199
 - 18.4.2 Identifying the Objects 200
 - 18.4.3 Identifying the Services or Methods 201
 - 18.4.4 Refining the Objects 202
 - 18.4.5 Bringing It All Together 203
 - 18.5 Where Is the Structure in an OO Program? 205
 - 18.6 Further Readings 207

- 19 Python Classes** 209
 - 19.1 Introduction 209
 - 19.2 Python and Classes 209
 - 19.3 Class Definitions 210
 - 19.4 Creating Examples of the Class Person 212
 - 19.5 Be Careful with Assignment 213
 - 19.6 Printing Out Objects 214
 - 19.6.1 Accessing Object Attributes 214
 - 19.6.2 Defining a Default String Representation 215
 - 19.6.3 Defining a Default Storage Representation 216
 - 19.7 Providing a Class Comment 217
 - 19.8 Adding a Birthday Method 218
 - 19.9 Defining Instance Methods 219
 - 19.10 Person Class Recap 220
 - 19.11 The Del Keyword 221
 - 19.12 Automatic Memory Management 222
 - 19.13 Intrinsic Attributes 223
 - 19.14 Online Resources 224
 - 19.15 Exercises 224

- 20 Class Side and Static Behaviour** 227
 - 20.1 Introduction 227
 - 20.2 Class-Side Data 227
 - 20.3 Class-Side Methods 228
 - 20.3.1 Why Class-Side Methods? 229
 - 20.4 Static Methods 229
 - 20.5 Hints 230
 - 20.6 Online Resources 230
 - 20.7 Exercises 231

- 21 Class Inheritance** 233
 - 21.1 Introduction 233
 - 21.2 What is Inheritance? 233
 - 21.3 Defining Inheritance Between Classes 234
 - 21.4 Terminology Around Inheritance 237
 - 21.4.1 Types of Hierarchy 237
 - 21.5 The Class Object and Inheritance 239
 - 21.6 The Built-in Object Class 240
 - 21.7 Purpose of Subclasses 240
 - 21.8 Overriding Methods 241
 - 21.9 Extending Superclass Methods 243
 - 21.10 Inheritance-Oriented Naming Conventions 243
 - 21.11 Python and Multiple Inheritance 244
 - 21.12 Multiple Inheritance Considered Harmful 246
 - 21.13 Summary 250
 - 21.14 Online Resources 251
 - 21.15 Exercises 251

- 22 Why Bother with Object Orientation?** 253
 - 22.1 Introduction 253
 - 22.2 The Procedural Approach 253
 - 22.2.1 Procedures for the Data Structure 254
 - 22.2.2 Packages 254
 - 22.3 Does Object Orientation Do Any Better? 255
 - 22.3.1 Packages Versus Classes 255
 - 22.3.2 Inheritance 257
 - 22.4 Summary 258

- 23 Operator Overloading** 259
 - 23.1 Introduction 259
 - 23.2 Operator Overloading 259
 - 23.2.1 Why Have Operator Overloading? 259
 - 23.2.2 Why not Have Operator Overloading? 260
 - 23.2.3 Implementing Operator Overloading 260
 - 23.3 Numerical Operators 262
 - 23.4 Comparison Operators 264

- 23.5 Logical Operators 266
- 23.6 Summary 266
- 23.7 Online Resources 266
- 23.8 Exercises 267
- 24 Python Properties 269**
 - 24.1 Introduction 269
 - 24.2 Python Attributes 269
 - 24.3 Setter and Getter Style Methods 270
 - 24.4 Public Interface to Properties 272
 - 24.5 More Concise Property Definitions 273
 - 24.6 Online Resources 275
 - 24.7 Exercises 275
- 25 Abstract Base Classes 277**
 - 25.1 Introduction 277
 - 25.2 Abstract Classes as a Concept 277
 - 25.3 Abstract Base Classes in Python 278
 - 25.3.1 Subclassing an ABC 278
 - 25.3.2 Defining an Abstract Base Class 280
 - 25.4 Defining an Interface 282
 - 25.5 Virtual Subclasses 282
 - 25.6 Mixins 284
 - 25.7 Online Resources 286
- 26 Error and Exception Handling 287**
 - 26.1 Introduction 287
 - 26.2 Errors and Exceptions 287
 - 26.3 What is an Exception? 288
 - 26.4 What is Exception Handling? 289
 - 26.5 Handling an Exception 291
 - 26.5.1 Accessing the Exception Object 293
 - 26.5.2 Jumping to Exception Handlers 294
 - 26.5.3 Catch Any Exception 296
 - 26.5.4 The Else Clause 296
 - 26.5.5 The Finally Clause 297
 - 26.6 Raising an Exception 298
 - 26.7 Defining an Custom Exception 299
 - 26.8 Chaining Exceptions 301
 - 26.9 Exception Groups 302
 - 26.10 Exception Group Handling 303
 - 26.11 Exception Notes 305
 - 26.12 Generate a Traceback Programmatically 306
 - 26.13 Online Resources 307
 - 26.14 Exercises 308

- 27 Python Modules and Packages** 309
 - 27.1 Introduction 309
 - 27.2 Modules 309
 - 27.3 Python Modules 310
 - 27.4 Importing Python Modules 312
 - 27.4.1 Importing a Module 312
 - 27.4.2 Importing from a Module 313
 - 27.4.3 Importing Aliases 314
 - 27.4.4 Hiding Some Elements of a Module 315
 - 27.4.5 Importing Within a Function 315
 - 27.5 Module Properties 316
 - 27.6 Standard Modules 316
 - 27.7 Python Module Search Path 317
 - 27.8 Modules as Scripts 318
 - 27.9 Python Packages 321
 - 27.9.1 Package Organization 321
 - 27.9.2 Sub Packages 322
 - 27.10 Online Resources 323
 - 27.11 Exercise 323

- 28 Protocols, Polymorphism and Descriptors** 325
 - 28.1 Introduction 325
 - 28.2 Implicit Contracts 325
 - 28.3 Duck Typing 326
 - 28.4 Polymorphism 328
 - 28.5 Protocols 329
 - 28.6 The Length Protocol 330
 - 28.7 The Sequence Protocol 331
 - 28.8 The Context Manager Protocol 332
 - 28.9 The Descriptor Protocol 334
 - 28.10 Online Resources 336
 - 28.11 Exercises 337

- 29 Decorators** 339
 - 29.1 Introduction 339
 - 29.2 What Are Decorators? 339
 - 29.3 Defining a Decorator 340
 - 29.4 Using Decorators 341
 - 29.5 Functions with Parameters 341
 - 29.6 Stacked Decorators 342
 - 29.7 Parameterized Decorators 343
 - 29.8 Method Decorators 344
 - 29.8.1 Methods Without Parameters 344
 - 29.8.2 Methods with Parameters 345
 - 29.9 Class Decorators 346
 - 29.10 When is a Decorator Executed? 348

- 29.11 Built-In Decorators 348
- 29.12 FuncTools Wrap 349
- 29.13 Online Resources 350
- 29.14 Book Reference 350
- 29.15 Exercises 350
- 30 Iterables and Iterators 353**
 - 30.1 Introduction 353
 - 30.2 Iteration 353
 - 30.2.1 Iterables 353
 - 30.2.2 Iterators 354
 - 30.2.3 The Iteration Related Methods 354
 - 30.2.4 The Iterable Evens Class 354
 - 30.2.5 Using the Evens Class with a For Loop 355
 - 30.3 The Itertools Module 355
 - 30.4 Online Resources 356
- 31 Generators and Coroutines 357**
 - 31.1 Introduction 357
 - 31.2 Generators 357
 - 31.2.1 Defining a Generator Function 358
 - 31.2.2 Using a Generator Function in a for Loop 358
 - 31.2.3 When Do the Yield Statements Execute? 358
 - 31.2.4 An Even Number Generator 359
 - 31.2.5 Nesting Generator Functions 360
 - 31.2.6 Using Generators Outside a for Loop 360
 - 31.3 Coroutines 361
 - 31.4 Exercises 362
- 32 Collections, Tuples and Lists 365**
 - 32.1 Introduction 365
 - 32.2 Python Collection Types 365
 - 32.3 Characterising Collections 366
 - 32.4 Tuples 367
 - 32.4.1 Creating Tuples 367
 - 32.4.2 The Tuple() Constructor Function 367
 - 32.4.3 Accessing Elements of a Tuple 368
 - 32.4.4 Creating New Tuples from Existing Tuples 368
 - 32.4.5 Tuples Can Hold Different Types 369
 - 32.4.6 Iterating Over Tuples 370
 - 32.4.7 Tuple Related Functions 370
 - 32.4.8 Checking if an Element Exists 370
 - 32.4.9 Nested Tuples 371
 - 32.4.10 Things You Can't Do with Tuples 371
 - 32.5 Lists 372
 - 32.5.1 Creating Lists 372

- 32.5.2 List Constructor Function 374
- 32.5.3 Accessing Elements from a List 374
- 32.5.4 Adding to a List 375
- 32.5.5 Inserting into a List 376
- 32.5.6 List Concatenation 376
- 32.5.7 Removing from a List 377
- 32.5.8 The Pop() Method 377
- 32.5.9 Deleting from a List 378
- 32.5.10 List Methods 378
- 32.6 Online Resources 379
- 32.7 Exercises 379
- 33 Sets** 381
- 33.1 Introduction 381
- 33.2 Creating a Set 381
- 33.3 The Set() Constructor Function 382
- 33.4 Accessing Elements in a Set 382
- 33.5 Working with Sets 382
- 33.5.1 Checking for Presence of an Element 382
- 33.5.2 Adding Items to a Set 383
- 33.5.3 Changing Items in a Set 383
- 33.5.4 Obtaining the Length of a Set 383
- 33.5.5 Obtaining the Max and Min Values in a Set 384
- 33.5.6 Removing an Item 384
- 33.6 Nesting Lists Within Sets 384
- 33.7 Set Operations 385
- 33.8 Set Methods 387
- 33.9 Online Resources 387
- 33.10 Exercises 388
- 34 Dictionaries** 389
- 34.1 Introduction 389
- 34.2 What is a Dictionary 389
- 34.3 Creating a Dictionary 390
- 34.3.1 The dict() Constructor Function 390
- 34.4 Working with Dictionaries 391
- 34.4.1 Accessing Items via Keys 391
- 34.4.2 Adding a New Entry 391
- 34.4.3 Changing a Keys Value 391
- 34.4.4 Removing an Entry 392
- 34.4.5 Iterating Over Keys 393
- 34.4.6 Values, Keys and Items 393
- 34.4.7 Checking Key Membership 394
- 34.4.8 Obtaining the Length of a Dictionary 394
- 34.4.9 Nesting Dictionaries 395
- 34.4.10 Checking to See if a Key is Present 395

34.5	Dictionary Operators Added in 3.9	395
34.6	A Note on Dictionary Key Objects	396
34.7	Dictionary Methods	397
34.8	Online Resources	398
34.9	Exercises	398
35	Frozensets	401
35.1	Introduction	401
35.2	Frozenset Characteristics	401
35.3	Creating a Frozenset	401
35.4	Frozenset Operations	403
35.5	Using Frozensets with Sets	404
36	Collection Related Modules	407
36.1	Introduction	407
36.2	List Comprehension	407
36.3	The Collections Module	408
	36.3.1 The Counter Container	409
	36.3.2 The Defaultdict Container	410
	36.3.3 The Namedtuple Container	412
36.4	The Itertools Module	413
36.5	Online Resources	414
36.6	Exercises	414
37	ADTs, Queues and Stacks	417
37.1	Introduction	417
37.2	Abstract Data Types	417
37.3	Data Structures	418
37.4	Queues	418
	37.4.1 Python List as a Queue	419
	37.4.2 Defining a Queue Class	419
37.5	Stacks	421
	37.5.1 Python List as a Stack	422
37.6	Online Resources	423
37.7	Exercises	423
38	Map, Filter and Reduce	425
38.1	Introduction	425
38.2	Filter	425
38.3	Map	427
38.4	Reduce	429
38.5	Online Resources	430
38.6	Exercises	431

- 39 Sorting and Higher-Order Functions** 433
 - 39.1 Introduction 433
 - 39.2 The sorted() Function 433
 - 39.3 Sorting in Reverse Order 436
 - 39.4 The sort() Method 437
 - 39.5 Reversing the sort() Method 438
 - 39.6 Changing the Sort Key 439
 - 39.7 Sorting Containers of Containers 440
 - 39.8 Exercises 442

- 40 Python Enumerated Values** 445
 - 40.1 Introduction 445
 - 40.2 What is an Enumeration? 445
 - 40.3 Creating an Enumeration Using Class Syntax 446
 - 40.4 Enum Values 447
 - 40.5 Using the Name and Value Attributes 448
 - 40.6 Creating an Enumeration Using the Function Call Syntax 449
 - 40.7 Creating Enumerations from Automatic Values 450
 - 40.8 Working with Enum Aliases 451
 - 40.9 Enumerations Unique Values 452
 - 40.10 Iterating Through Enumeration Values 453
 - 40.11 Comparing Enumerations 453
 - 40.12 Integer Enumerations with IntEnum 454
 - 40.13 Exercises 455

- 41 Structural Pattern Matching** 457
 - 41.1 Introduction 457
 - 41.2 Basic Pattern Matching 457
 - 41.3 Default Case 459
 - 41.4 Capture Default Case 460
 - 41.5 Alternative Values in Case Condition 461
 - 41.6 Matching Sequences/Containers 461
 - 41.7 Structural Matching and Lists 463
 - 41.8 Matching Alternative Sequences/Containers 464
 - 41.9 Matching Multiple Options Within a Container 465
 - 41.10 Guards on a Case Condition 466
 - 41.11 Exercises 467

- 42 Python Virtual Environments** 469
 - 42.1 Introduction 469
 - 42.2 Third-Party Libraries 469
 - 42.3 The Pip Tool 471
 - 42.3.1 Finding the Pip Version 472
 - 42.3.2 Installing Pip 472
 - 42.3.3 Updating Pip 472
 - 42.4 Default Installation 473

42.5	Pip Virtual Environments	474
42.6	Installing Modules	476
42.7	Listing and Freezing on Modules	478
42.7.1	The pip list Command	478
42.7.2	The pip inspect Command	479
42.7.3	The pip freeze Command	479
42.8	Uninstalling Modules	481
42.9	Module Path Search Order	482
42.10	Using PyCharm with Virtual Environments	483
42.11	Online Resources	486
43	Monkey Patching	487
43.1	Introduction	487
43.2	What is Monkey Patching?	487
43.2.1	How Does Monkey Patching Work?	488
43.2.2	Monkey Patching Example	488
43.2.3	The Self Parameter	489
43.2.4	Adding New Data to a Class	490
43.3	Online Resources	490
44	Attribute Lookup	491
44.1	Introduction	491
44.2	Attributes in Python	491
44.3	Handling Unknown Attribute Access	494
44.4	Handling Unknown Method Invocations	495
44.5	Intercepting Attribute Lookup	496
44.6	Intercepting Setting an Attribute	497
44.7	Exercises	498
45	The Python Bites Back	499
45.1	Fourteen Beginner Mistakes to Avoid in Python	499
45.2	Mistake #1: Indentation	499
45.3	Mistake #2: Case Sensitivity	501
45.4	Mistake #3: Default Initialization of Function/Method Parameters with Mutable Types	501
45.5	Mistake #4: Default Initialization of Function/Method Parameters Using a Function	503
45.6	Mistake #5: Out by 1 Ranges	504
45.7	Mistake #5: Assignment Operator Versus Equality Operator	505
45.8	Mistake #6: Falsey and Truthy	506
45.9	Mistake #7: Hiding Built-In Functions and Classes	507
45.10	Mistake #8: Hiding Built-In Modules	508
45.11	Mistake #9: Function Reference Versus Function Call	509
45.12	Mistake #10: Modifying a List While Iterating Over It	510
45.13	Mistake #11: Sets Can Only Contain Immutable Objects	511

- 45.14 Mistake #12: Local Versus Global Variables 512
- 45.15 Mistake #13: Referencing Nonlocal Variable 514
- 45.16 Mistake #14: Copying a Mutable Type 515
- 46 TicTacToe Game 517**
 - 46.1 Introduction 517
 - 46.2 Classes in the Game 517
 - 46.3 Counter Class 520
 - 46.4 Move Class 521
 - 46.5 The Player Class 521
 - 46.6 The HumanPlayer Class 522
 - 46.7 The ComputerPlayer Class 523
 - 46.8 The Board Class 524
 - 46.9 The Game Class 525
 - 46.10 Running the Game 527

Chapter 1

Introduction



1.1 What is Python?

Python is a general-purpose programming language in a similar vein to other programming languages that you might have heard of such as C++ , JavaScript or Microsoft’s C# and Java.

It has been around for some considerable time having been originally conceived back in the late 1980s by Guido van Rossum at Centrum Wiskunde and Informatica (CWI) in the Netherlands. The language is named after one of Guido’s favorite programs “Monty Pythons Flying Circus”, a classic and somewhat anarchic British comedy sketch show originally running from 1969 to 1974 (but which has been rerun on various stations ever since) and with several film spin offs. You will even find various references to this show in the documentation available with Python.

As a language it has gained in interest over recent years, particularly within the commercial world, with many people wanting to learn the language. This increased interest in Python is driven by several different factors:

1. Its flexibility and simplicity which makes it easy to learn.
2. Its use by the Data Science community where it provides a more standard programming language than some rivals such as R.
3. Its suitability as a scripting language for those working in the DevOps field where it provides a higher level of abstraction than alternative traditionally used languages.
4. Its ability to run on (almost) any operating system, but particularly the big three operating systems Windows, MacOS and Linux.
5. The availability of a wide range of libraries (modules) that can be used to extend the basic features of the language.
6. It is free!

Python itself is now managed by the *not-for-profit* Python Software Foundation (see https://en.wikipedia.org/wiki/Python_Software_Foundation) which was launched in March 2001. The mission of the foundation is to foster development of

the Python community; it is also responsible for various processes within the Python community, including developing the core Python distribution, managing intellectual rights and supporting developer conferences including PyCon.

1.2 Python Versions

There is now one main version of Python, that is Python 3. However, until the start of 2020 there were two main supported versions in circulation one called Python 2 as well as Python 3.

- Python 2 was launched in October 2000 and was very widely used. Support for this version of Python ended on the 1st of January 2020. However, you will still find examples of code, libraries and tutorials that focus on this version of the language. This is particularly true of the web so be careful when browser for examples or help—make sure they relate to Python 3. The more current the article or suggestion you are looking the more likely that it is Python 3.
- Python 3 was launched in December 2008 and is a major revision to the language that is not backward compatible. It is current the most widely used version of the language and represents the version we will use throughout this book.

The issue between the two versions can be highlighted by the simple print facility:

- In Python 2 this is written as `print 'Hello World'`
- In Python 3 this is written as `print('Hello World')`

It may not look like much of a difference but the inclusion of the `'()` marks a major change and means that any code written for one version of Python will probably not run on the other version. There are tools available, such as the 2–3 utility, that will (partially) automate translation from Python 2 to Python 3 but in general you are still left with significant work to do.

A survey of Python Developers carried out by JetBrains (the organization behind the PyCharm editor) in 2022 found that only 2% of developers now claimed to be using a Python 2.x version, with 2% unsure which version they used. The rest (96%) were using a version of the 3.x line mostly 3.8, 3.9 or 3.10. This survey was completed before 3.11 was made available, and this over 70% of developers were using one of the three most recent versions of Python. For more information see <https://www.jetbrains.com/lp/devecosystem-2022/python/>.

Python 3 is the future of the Python language, and it is this version that has introduced many of the new and improved language and library features (that have admittedly been back ported to Python 2 in many cases). This book is solely focused on Python 3.

In the remainder of this book when we refer to Python we will always be referring to Python 3. Since then there have been numerous point and sub point releases, for example, as well as Python 3.1, 3.2, etc. there have also been intermediate releases such as 3.1.1, 3.1.2, etc.

1.3 Python History

As has been said, Python has been around along time, although for some people it can seem like Python was an overnight success. In fact Python originates from the very late 1980s when Guido Van Rossum started implementing the first version of Python. The first release of Python was on 20th February 1991. Guido actually remained the lead developer on Python until 12 July 2018 when he stepped down as the “benevolent dictator” of Python. Since January 2019 there has been a five-member Steering Council that leads the Python project.

Python 2.0 was initially released on October 16th 2000 with its end of life being the 1st of January 2020. Python 3.0 was initially released 3rd December 2008.

It is worth noting that in the Python World, the major version number represents an epoch (such that there has been a Python 1, a Python 2 and a Python 3 epoch). Within this each Python 3 release has a prefix of 3.1, thus 3.1 was the second Python 3 release, etc. This differs from many other programming language such as Java with Java 18 or Java 19 being the equivalent of 3.11 and 3.12 in Python, etc. The Python organization does not always release a new version of Python each year, although recent releases have tended to be yearly. However, although the language is evolving many of these releases do not represent major changes to the language, and thus there is relatively little difference between Python 3.10 and Python 3.11 for example. A quick history of the main Python 3 releases is give below:

- Python 3.0 released 3rd December 2008
- Python 3.1 released 27th June 2009
- Python 3.2 released 20th February 2011
- Python 3.3 released 29th September 2012
- Python 3.4 16th March 2014
- Python 3.5 13th September 2015
- Python 3.6 23rd December 2016
- Python 3.7 27th June 2018
- Python 3.8 14th October 2019
- Python 3.9 5th October 2020
- Python 3.10 4th October 2021
- Python 3.11 24th October 2022.

In general almost all of the features in the book are supported by Python 3.8 onwards. One or two of the features require a later version such as 3.10. Such requirements will be noted when each topic is introduced. Personally when teaching Python I try and use the latest stable version such as 3.11 or 3.12.

1.4 Python Programming

There are several different programming paradigms that a programming language may allow developers to code in, these are:

- **Procedural Programming** in which a program is represented as a sequence of instructions that tell the computer what it should do explicitly. Procedures and / or functions are used to provide structure to the program with control structures such as if statements and loop constructs to manage which steps are executed and how many times. Languages typifying this approach include C and Pascal.
- **Declarative Programming** languages, such as prolog, that allow developers to describe how a problem should be solved with the language / environment determining how the solution should be implemented. SQL (a database query language) is one of the most common declarative languages that you are likely to encounter and is introduced at the end of this book.
- **Object-Oriented Programming** approaches that represent a system in terms of the objects that form that system. Each object can hold its own data (also known as state) as well as define behaviour what the object can do. A computer program is formed from a set of these objects co-operating together. Languages such as Java and C# typify the object-oriented approach.
- **Functional Programming** languages decompose a problem into a set of functions. Each function is independent of any external state operating only on the inputs they received to generate their outputs. The programming language Haskell is an example of a functional programming language.

Some programming languages are considered to be *hybrid* languages; that is they allow developers to utilize a combination of difference approaches within the same program. Python is an example of a hybrid programming language as it allows you to write very procedural code, to use objects in an object-oriented manner and to write functional programs. Each of these approaches is covered in this book.

1.5 Python Libraries

As well as the core language there are very many libraries available for Python. These libraries extend the functionality of the language and make it much easier to develop applications. These libraries cover:

- web frameworks such as Django / Flask,
- email clients such as smtplib (a SMTP email client) and imaplib (an IMAP4 email client),
- content management operations such as the Zope library,
- lightweight concurrency (running multiple operations at the same time) using the Stackless library,
- the Generation of Microsoft Excel files using the openpyxl library,


- graphics libraries such as Matplotlib and PyOpenGL,
- Graphical User Interface (GUI) libraries such as Tkinter or PyQt5,
- Data Analytics using libraries such as Pandas and SciPy and NumPy,
- Machine learning using libraries such as SKLearn and TensorFlow.

A very useful resource to look at, which introduces many of these libraries (also known as modules), is the ‘Python 3 module of the Week’ web site which can be found at <https://pymotw.com/3>. This lists many of the libraries / modules available and provides a short introduction to what they do and how to use them. The chapter entitled ‘Python Virtual Environments’ explains how to install such libraries into your Python runtime environment.

1.6 Python Execution Model

Python is not a precompiled language in the way that some other languages you may have come across are (such as C++). Instead it is what is known as an interpreted language (although even this is not quite accurate). An interpreted language is one that does not require a separate compilation phase to convert the human readable format into something that can be executed by a computer. Instead the plain text version is fed into another program (generally referred to as the interpreter) which then executes the program for you.

Python actually uses an intermediate model in that it actually converts the plain text English style Python program into an intermediate ‘pseudo’ machine code format and it is this intermediate format that is executed. This is illustrated below:



The way in which the Python interpreter processes a Python program is broken down into several steps. The steps shown here are illustrative (and simplified) but the general idea is correct.

1. First the program is checked to make sure that it is valid Python, that is a check is made that the program follows all the rules of the language and that each of the commands and operations, etc. is understood by the Python environment.

2. It then translates the plain text, English like commands, into a more concise intermediate format that is easier to execute on a computer. Python can store this intermediate version in a file which is named after the original file but with a `.pyc` extension instead of a `.py` extension (the 'c' in the extension indicates it contains the compiled version of the code).
3. The compiled intermediate version is then executed by the interpreter.

When this program is rerun the Python interpreter checks to see if a `.pyc` file is present. If no changes have been made to the source file since the `.pyc` was created, then the interpreter can skip steps 1 and 2 and immediately run the `.pyc` version of the program.

One interesting aspect of Python's usage is that it can be (and often is) used in an interactive fashion, with individual commands being entered and executed one at a time, with context information being built up. This can be useful in debugging situations.

1.7 Running Python Programs

There are several ways in which you can run a Python program, including


- Interactively using the Python interpreter.
- Stored in a file and run using the Python command.
- Run as a script file specifying the Python interpreter to use within the script file.
- From within a Python IDE (Integrated Development Environment) such as PyCharm.
- Using Jupyter Notebooks in a web browser.

1.7.1 *Interactively Using the Python Interpreter*

It is quite common to find that people will use Python in interactive mode. This uses the Python REPL (named after **R**ead **E**valuate **P**rint **L**oop style of operation).

Using the REPL, Python statements and expressions can be typed into the Python prompt and will then be executed directly. The values of variables will be remembered and may be used later in the session.

To run the Python REPL, Python must have been installed onto the computer system you are using. Once installed you can open a Command Prompt window (Windows) or a Terminal window (Mac) and type `Python` into the prompt. This is shown for a Windows machine below:

A screenshot of a terminal window titled "jeh — Python — 87x14". The terminal shows the following text:

```
jeh@Johns-iMac ~ % python
Python 3.11.2 (main, Feb 16 2023, 03:07:35) [Clang 14.0.0 (clang-1400.0.29.202)] on dar
win
Type "help", "copyright", "credits" or "license" for more information.
[>>> print('Hello World')]
Hello World
[>>> 5 + 4]
9
[>>> name = 'John']
[>>> print(name)]
John
[>>> ]
```

In the above example, we interactively typed in several Python commands and the Python interpreter ‘Read’ what we have typed in, ‘Evaluated’ it (worked out what it should do), ‘Printed’ the result and then ‘Looped’ back ready for further input. In this case we

- Printed out the string ‘Hello World’.
- Added 5 and 4 together and got the result 9.
- Stored the string ‘John’ in a variable called `name`.
- Printed out the contents of the variable `name`.


To leave the interactive shell (the REPL) and go back to the console (the system shell), press `Ctrl-Z` and then `Enter` on Windows, or `Ctrl-D` on OS X or Linux. Alternatively, you could also issue the command `quit()`.

You should note that once you have left the session anything you created within the session will be lost.

1.7.2 Running a Python File

We can of course store the Python commands into a file. This creates a program file that can then be run as an argument to the `python` command. It also allows the program held within the file to be run many times.


For example, given a file containing the following file (called `hello.py`) with the 4 commands in it:



```
hello.py
1 print('Hello World')
2 print(5 + 4)
3 name = 'John'
4 print(name)
5
```


Line 1, Column 1 Tab Size: 4 Python

To run the `hello.py` program on a PC using Windows we can use the `Python` command followed by the name of the file:



```
C:\Users\john>python hello.py
Hello World
9
John
C:\Users\john>
```

We can also run the same program on a Apple Mac using MacOS via the `Python` interpreter. For example on a Mac we can do the following:



```
temp — -bash — 52x5
Johns-iMac:temp jeh$ python hello.py
Hello World
9
John
Johns-iMac:temp jeh$
```

This makes it very easy to create Python programs that can be stored in files and run when needed on whatever platform is required (Windows, Linux or Mac). This illustrates the cross platform nature of Python and is just one of the reasons why Python is so popular.

1.7.3 Executing a Python Script

It is also possible to transform a file containing a stored Python program into a Script. A script is a stand-alone file that can be run directly without the need to (explicitly) use the `python` command.

This is done by adding a special line to the start of the Python file that indicates the Python command (or interpreter) to use with the rest of the file. This line must start with `#!` and must come at the start of the file.

To convert the previous sections file into a Script we would need to add the path to the python interpreter. Here path refers to the route that the computer must take to find the specified Python interpreter (or executable).


The exact location of the Python interpreter on your computer depends on what options upon selected when you (or whoever installed Python) set it up. Typically on a Windows PC Python will be found in the 'Program Files' directory or it might be installed in its own 'Python' directory.

Whatever the location of the Python interpreter to create a script we will need to add a first line to our `hello.py` file. This line must start with a `#!`. This combination of characters is known as a *shebang* and indicates to Linux and other Unix like operating systems (such as MacOS) how the remainder of the file should be executed.

For example, on a Apple Mac we might add:

```
/Library/Frameworks/Python.framework/Versions/3.11/bin/python3
```

When added to the `hello.py` file we now have:



```

1  #!/Library/Frameworks/Python.framework/Versions/3.7/bin/python3
2
3  print('Hello World')
4  print(5 + 4)
5  name = 'John'
6  print(name)
7

```

Line 7, Column 1 Tab Size: 4

However, we cannot just run the file as it stands. If we tried to run the file without any changes then we will get an error indicating that the permission to execute the file has been denied:

```

$ ./hello.py
-bash: ./hello.py: Permission denied
$

```


This is because by default you can't just run a file. We need to mark it as executable. There are several ways to do this, however one of the easiest on a Mac or Linux box is

to use the `chmod` command (which can be used to modify the permissions associated with the file). To make the file executable we can change the file permissions to include making it executable by using the following command from a terminal window when we are in the same directory as the `hello.py` file:

```
$ chmod +x hello.py
```

Where `+x` indicates that we want to add the executable permission to the file.

Now if we try to run the file directly it executes and the results of the commands within the file are printed out:

A screenshot of a terminal window on a Mac. The window title is "temp — -bash — 58x6". The prompt is "Johns-iMac:temp jeh\$". The user enters "chmod +x hello.py". The prompt is "Johns-iMac:temp jeh\$". The user enters "./hello.py". The output is "Hello World", "9", and "John". The prompt is "Johns-iMac:temp jeh\$".

```
Johns-iMac:temp jeh$ chmod +x hello.py
Johns-iMac:temp jeh$ ./hello.py
Hello World
9
John
Johns-iMac:temp jeh$
```

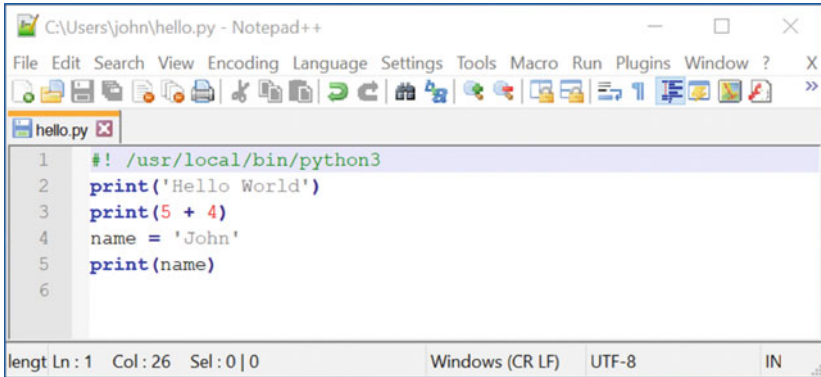
Note the use of the `./` preceding the file name in the above; this is used on Linux to tell the operating system to look in the current directory for the file to execute.

Different systems will store Python in different locations and thus might need different first lines, for example on a Linux we might write:

```
#!/usr/local/bin/python3
print('Hello, world')
print(5 + 4)
name = 'John'
print(name)
```

By default Windows does not have the same concept. However, to prompt cross platform portability the Python Launcher for Windows can also support this style of operation. It allows scripts to indicate a preference for a specific Python version using the same `#!` (Shebang) format as Unix style operating systems. We can now indicate that the rest of the file should be interpreted as a Python script; if multiple versions of Python are installed this may require Python 3 to be explicitly specified. The launcher also understands how to translate the Unix version into Windows versions so that `/usr/local/bin/python3` will be interpreted as indicating that python3 is required.

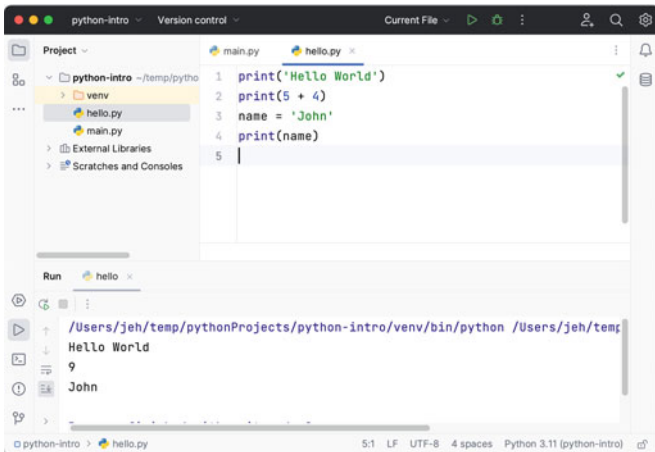
An example of the `hello.py` script for a Windows or Linux machine is given below using Notepad++ on a Windows box.



When the launcher was installed it should have been associated with Python files (i.e., files that have a `.py` extension). This means that if you double-click on one of these files from the Windows Explorer, then the Python launcher will be used to run the file.

1.7.4 Using Python in an IDE

We can also use an IDE such as PyCharm to writing and execute our Python program. The same program is shown using PyCharm below:



In the above figure the simple set of commands are again listed in a file called `hello.py`. However, the program has been run from within the IDE and the output is shown in an output console at the bottom of the display.

1.8 Jupyter Notebooks

An approach favored by many data analysts is to use the Jupyter Notebooks infrastructure to run Python commands and programs. To do this they use a web browser which will automatically display graphs and tabular format within the web page which is very useful in some domains.

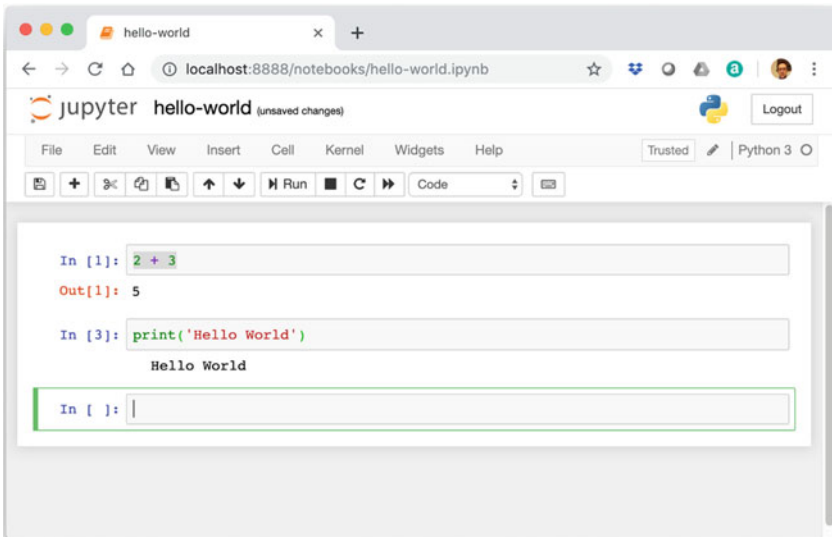
Jupyter notebooks have to be installed into your Python environment, for example using the command

```
pip install jupyter
```

Once installed it is possible to launch a new notebook using the command

```
jupyter notebook
```

This starts up a background Python compilation server and runtime as well as opening a Jupyter notebooks browser window such as that shown below:



You can now type Python statements and expressions into the browser, and the results generated will be displayed within the browser.

1.9 Useful Resources

There are a wide range of resources on the web for Python; we will highlight a few here that you should bookmark. We will not keep referring to these to avoid repetition but you can refer back to this section whenever you need to:

- https://en.wikipedia.org/wiki/Python_Software_Foundation Python Software Foundation.
- <https://docs.python.org/3/> The main Python 3 documentation site. It contains tutorials, library references, set up and installation guides as well as Python how-tos.
- <https://docs.python.org/3/library/index.html> A list of all the built-in features for the Python language—this is where you can find online documentation for the various class and functions that we will be using throughout this book.
- <https://pymotw.com/3/> The Python 3 Module of the week site. This site contains many, many Python modules with short examples and explanations of what the modules do. A python module is a library of features that build on and expand the core Python language. For example, if you are interested in building games using Python then pyjama is a module specifically designed to make this easier.
- <http://www.pythonweekly.com/> is a free weekly summary of the latest Python articles, projects, videos and upcoming events.

Chapter 2

Where is Python Used




2.1 Introduction

Python is one of the biggest and most popular programming languages in today's technology world; whatever index or survey you look at Python is usually in the top 3 (along with Java and JavaScript). However, you might be forgiven for thinking that Python is a language exclusively used for Data Analytics or Data Science given the number of times these terms turn up alongside Python. However, Python has a wide and diverse range of uses. In this chapter we will look at eight uses of Python, some of which might surprise you.

2.1.1 *Data Analytics*

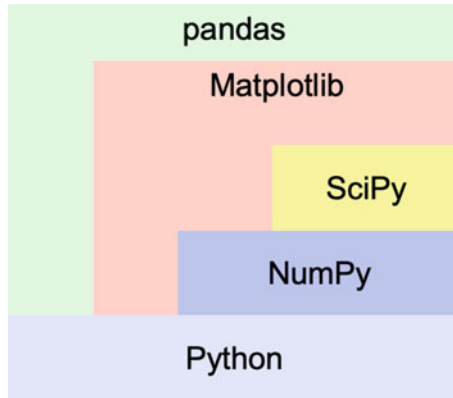
This is one of the two largest application areas for Python and may well be why you are reading this book in the first place. There are several widely used Data Analytics libraries available for Python including Pandas. Data Analytics (or analysis) is the practice of taking raw data and analyzing it to either understand that data and/or look for patterns in the data or to create new systems that can use knowledge gained from previous data to predict new trends or patterns in the future data. In both cases the aim is to provide value in some way to an organization or business as illustrated below.



The approaches to Data Analytics are not completely independent but illustrate the different purposes. For example, the “Discovery of Data Insight” side typically involves understanding and trends, patterns or relationships in the data to help an organization or business to improve their performance. For example, an online sales shop might use such insights to understand how and where to target specific promotions, a governmental health organization might use their data insights to determine where and when to target vaccinations, etc.

In contrast the ‘Development of Data Product(s)’ typically involves using information with the data to help develop new systems that will be applied to future, as yet unseen data. For example, analysis of the data might help to identify patterns of fraudulent behaviour which might be used to develop new monitoring systems to catch similar fraud earlier and faster.

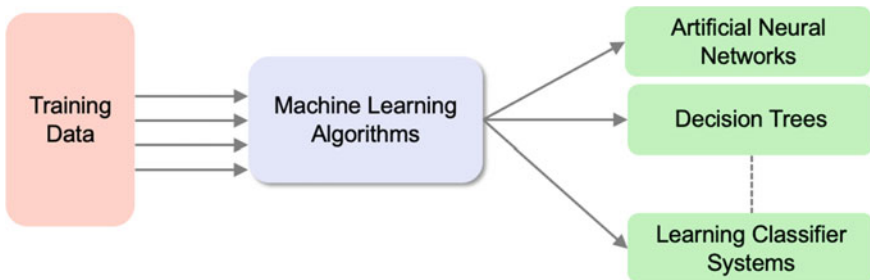
The techniques used for such analysis are typically oriented around statistical analysis and data exploration. These are the facilities provided by libraries such as Pandas (see <https://pandas.pydata.org/>) which is by far the most widely used Python data analysis library. However, Pandas does not provide a complete solution on its own. For example, Pandas builds on other Python libraries such as NumPy (which provides sophisticated facilities for handling numbers see <https://numpy.org/>). Pandas is also often used in combination with other libraries such as SciPy (which stands for Scientific Python see <https://scipy.org/>) and Matplotlib (see <https://matplotlib.org/>) which is a graphing library. The relationship between Pandas, Python and other commonly used data analysis-oriented libraries is shown below.



In addition several Data Analytics reporting systems and toolkits used outside the Python world also provide Python APIs allowing the results of any Python analysis to be incorporated into these systems. For example, the Tableau (<https://www.tableau.com/>) suite of tools provides a Python API allowing results created by Python to be incorporated into a Tableau dashboard for easy access and availability.

2.1.2 Machine Learning and AI

Following on from the previous section on data analysts it is not uncommon to apply some machine learning techniques when analyzing data, for example, to perform some classification of new data using previously seen data. That is, machine learning is a method of data analysis that automates analytical model building. Using algorithms that iteratively learn from data, machine learning allows computers to find hidden insights without being explicitly programmed. The general model is shown below.



This illustrates that training data (which may be what is called supervised or unsupervised) is fed into a machine learning algorithm and the output from this is a

system that might be used to classify new data, predict behaviour based on new data or make some decision based on new data.

Supervised learning tasks involve teaching a learning systems using data that is tagged or marked with known results. For example, loan data with applications marked as being either fraudulent or not, etc.

In contrast unsupervised data involves presenting a learning system with data—with no indications of expected results, and then identifying patterns or clusters within that data. This allows users to understand or learn more about the data. The term unsupervised is intended to indicate that there is no known (or at least provided) correct answers (and hence there is no teacher). In some situations there is a large amount of data, and only some of that data can be labeled with appropriate results or conclusions. Such problems are sometimes referred to as semi-structured machine learning. In many cases a combination of supervised and unsupervised techniques is required to analyze such data.

Each of the machine learning algorithms commonly used with Python has pros and cons and situations in which they are best applied. It is therefore useful to have a toolbox of such algorithms that can be applied as and when appropriate.

There are several commonly used libraries available in Python for this including SciKitLearn (aka SKLearn), TensorFlow and PyTorch. In general SciKitLearn and TensorFlow have about the same share of the Python market place with PyTorch being just behind, although all three are close to 30% in terms of Python developers' usage.

For further information on machine learning see:

- Harvard Business Review—“What every manager should know about Machine learning”
 - <http://hbr.org/2015/07/what-every-manager-should-know-about-machine-learning>
- Quick Reference cheat sheet on ML Algorithms/uses
 - <http://www.lauradhamilton.com/machine-learning-algorithm-cheat-sheet>
- 10 Machine Learning Methods that Every Data Scientist Should Know
 - <https://towardsdatascience.com/10-machine-learning-methods-that-every-data-scientist-should-know-3cc96e0eeee9>

2.1.3 Database Work

Most data is held within some form of database, and Python provides interfaces that allow a programmer to access that data. These interfaces make it easy to query a database for information and pull that information back into the Python program for further progressing. For example, a database of client information may be queried

to retrieve details held about clients. Python could then display this information in a tabular form via some form of table, graphically in some form of chart or on a map, etc.

2.1.4 Python for Animation


Python is one of the two programming languages that can be used in the Autodesk Maya toolkit. Maya is a 3D animation, modeling, simulation and rendering software systems used by games, film and TV companies. Within Maya, Python can be used for anything from creating scripts handling common tasks, through to develop complete plug-ins providing extensions to the core functionality of the tool. Maya currently supports Python 2 but will be moving toward Python 3 in the near future.

2.1.5 Python for Film Making

Industrial Light and Magic was formed in 1975 by George Lucas to help create special effects for the original Star Wars film. Since then it has gone on to provide special effects for a huge range of films and film makers. Back in 1996 it used Unix shell scripts to help it automate and control its production pipeline. However, they needed to develop their increasingly complex and computationally expensive production process and felt that their existing approach was not flexible enough. Although they looked at several alternatives (including TCL and Perl) they decided to adopt Python. This allowed for faster development times and greater flexibility. Over the years they have re-evaluated their decision and considered different alternatives, but Python remains key to their development processes.

2.1.6 Cross Platform UIs

Python is particularly well served in terms of Graphical User Interface frameworks (or GUI frameworks). Most of these are cross platform, although a few allow developers to exploit particular features of an underlying windowing system and are thus tied to a specific platform. Widely used GUI libraries include the cross platform TKinter, wxPython and pyQT and the platform specific PythonWin (for Windows) and PyObjc (for MacOS). As an example, the following 'Hello World' GUI application uses wxPython and is shown running on a Mac OS and Windows system:



2.1.7 Games Programming

Python has also been used for games development with several well-known games titles depending on it in one way or another. For example, Battlefield 2 by Digital Illusions CE is a military simulator first-person shooter game. Within this game, Python is used to handle portions of the game logic for Battlefield Heroes involving game modes and scoring. Other games that use Python include Civilisation IV, Pirates of the Caribbean Online and Overwatch. The pygame library is probably the most widely used library for creating games within the Python world. There are also many extensions available for pygame that help to create a wide range of different types of games. Pygame is built on top of the SDL library (or Simple Directmedia Layer). SDL is a cross platform development library designed to provide access to audio, keyboards, mouse, joystick and graphics hardware via OpenGL and Direct3D. To promote portability, pygame also supports a variety of additional backends including WinDIB, X11, Linux Frame Buffer, etc.

2.1.8 Integration Testing Framework

Morgan Stanley have developed a Python integration testing framework called Testplan which they have open sourced and made available through GitHub. Testplan is designed to simplify the process of configuring and driving integration tests for a range of programming languages and technologies. It supports driving integration tests that require messaging services, RESTful services, databases, files for code written in Python, C/C++, Java, etc.

2.1.9 Academic Research

Python is widely used within the academic community to support research work, not only within Computer Science departments but also across a range of different disciplines including Mechanical Engineering, Aerospace Engineering, Architecture, Pharmacology, Medicine, etc. In these settings it has been used to help develop distributed analysis systems, identify patterns in experiments of laboratory data,

evolve designs toward optimal solutions, provide natural language frontends to research applications (e.g., using the Python NLTK library), etc.

2.1.10 Web Services

Python is also widely used as a server-side language for the creation of web services, whether these are RESTful services or the newer GraphQL-based services. There are a range of frameworks available to help develop such services including Flask, Django and CherryPy. Flask and CherryPy are lightweight frameworks that can be used to create RESTful services, whereas Django is a full-stack web framework aimed at developing not just web services but full blown websites. A wide range of organizations used Python in this way for example Reddit, Spotify and Instagram.

2.1.10.1 DevOps

DevOps is another hot trend at the moment; it represents the combination of software developers and operations working together usually to automate operations processes that were previously either handled manually or were executed as individual steps. Python is one of the key programming languages used within the DevOps world. It can be used as a scripting language to help automate operations activities as well as a tool to analyze production data and for data visualizations.

2.2 Useful Resources

- <https://www.autodesk.co.uk/products/maya/overview> for information on Autodesk Maya computer animation software.
- <https://www.pygame.org> for information on pygame.
- <https://www.python.org/about/success/ilm/> ILM and their use of Python.
- <https://github.com/Morgan-Stanley/testplan> Testplan Python integration testing framework.
- <https://devops.com/how-python-is-transforming-the-devops-landscape/> Article on DevOps and Python.

Chapter 3

Setting up the Python Environment



3.1 Introduction


In this chapter we will check to see if you have Python installed on your computer. If you do not have Python installed then we will step through the process of installing Python. This is necessary because when you run a Python program it looks for the `python` interpreter that is used to execute your program or script. Without the `python` interpreter installed on your machine Python programs are just text files!

3.2 Check to See if Python is Installed

The first thing you should do is to see if Python 3 is already installed on your computer. First check to see that you don't have Python installed. If it is you don't need to do anything unless it is a very old version of Python 3 such as 3.1 or 3.2.

On a Windows 11 machine you can check the version installed by opening a Command Prompt window (this can be done by searching for `cmd` in the 'Search' box in Windows).

Once the Command window is open try typing in `python`. This is illustrated below:



```

Microsoft Windows [Version 10.0.22621.1555]
(c) Microsoft Corporation. All rights reserved.

C:\Users\phoeb>python

C:\Users\phoeb>python3


C:\Users\phoeb>

```

You may be prompted to look in the Microsoft store, you can ignore that and close the window.

Note in the above where we have tried both `python` and `python3` in case the latest version has been installed using that name. This is not an uncommon thing to do as on some machines you may have both Python 2 and Python 3 available.

On a system such as a Mac you can use the Terminal and do the same thing. You will probably find that at least Python (2) is pre-installed for you. For example, if you type in `python` on a Mac you will get something like this:



```

jeh — Python — 80x9
[Johns-iMac:~ jeh$ python
Python 2.7.15 (v2.7.15:ca079a3ea3, Apr 29 2018, 20:59:26)
[GCC 4.2.1 Compatible Apple LLVM 6.0 (clang-600.0.57)] on darwin
Type "help", "copyright", "credits" or "license" for more information.
>>> █

```

This indicates that the above user has version 2.7.15 installed (note you may have another 2.x version installed).


However, be careful if you find that you have Python 2 installed on your machine; this book is focused solely on Python 3.

If you have started a Python interpreter then

- Use `quit()` or `exit()` to exit the Python interpreter; `exit()` is an alias for `quit()` and is provided to make Python easier to use.

If Python 3 was not available, then the following steps will help you to install it.

If the correct version of Python is already available on your computer then you can skip to the next chapter. An example of what might happen if you type in Python and Python 3 is installed on a Mac is given below:



Here you can see that this machine has Python 3.11.3 installed.

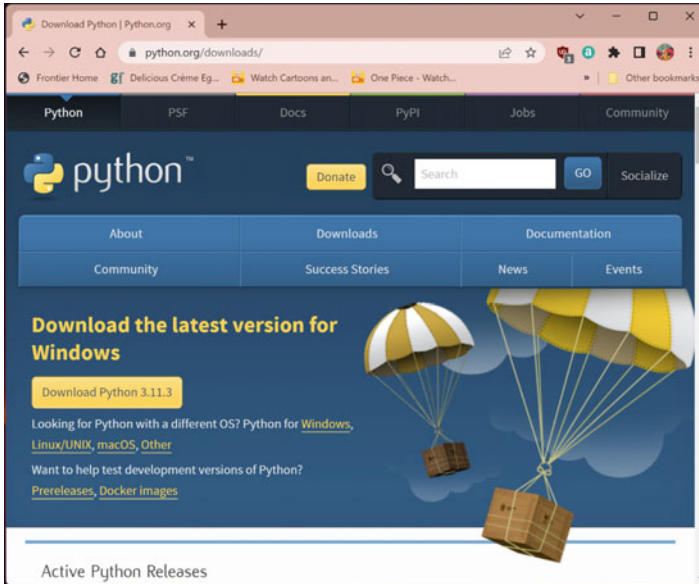
3.3 Installing Python on a Windows PC

3.3.1 Step 1: Downloading Python

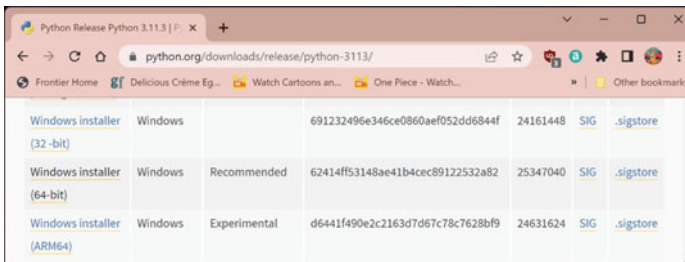
Python is available for a wide range of platforms from Windows, to Mac OS and Linux; you will need to ensure that you download the version for your operating system.

Python can be downloaded from the main Python web site which can be found at <http://www.python.org/>.

As you will see the 'Downloads' link is the second from the left on the large menu bar below the search field. Click on this and you will be taken to the download page; the current Python 3 version at the time of writing is Python 3.11 which is the version we will download (if you find that the version is 3.12 or later then that is fine as well). Click on the 'Download Python 3.x.x' link. For example, for a Windows 11 machine, using 3.11.3, you will see:



Click on the ‘Download’ button—this may take you to a page with several different options on it—look for the ‘Recommended’ option and download that installer.



Even if a newer version of Python is available (which is likely as the version is updated frequently) the steps should be fundamentally the same.

3.3.2 Step 2: Running the Installer

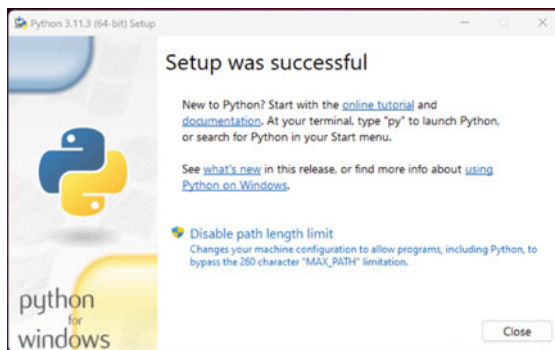
Once you have download the install *open* it. You will now be prompted for the location to install Python, for example:



Note that it is easiest if you click on the ‘Add python.exe to PATH’ option as this will make it available from the command line. If you don’t do this then don’t worry, we can add Python to the PATH later on (the PATH environment variable is used by Windows to find the location of programs such as the Python interpreter).

Next select the ‘Install Now’ option and follow the installation steps.

If everything went as expected you should now see a confirmatory dialog such as:



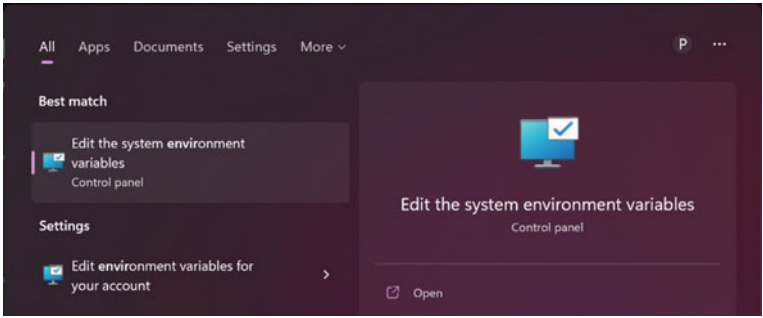
If you are on Windows now close any Command Prompt windows that you have open (the PATH variable is not updated for existing Commands). This can be done by typing in ‘exit’ or closing the window.

3.3.3 Step 3: Set up the PATH (Optional)

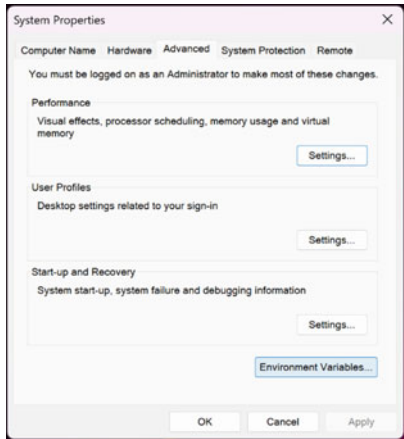
If you did not select ‘Add python.exe to PATH’ at the start of the installation you will now need to set up the PATH environment variable. If you did then skip onto the next step.

You can set the PATH environment variable using the system environment variables editor.

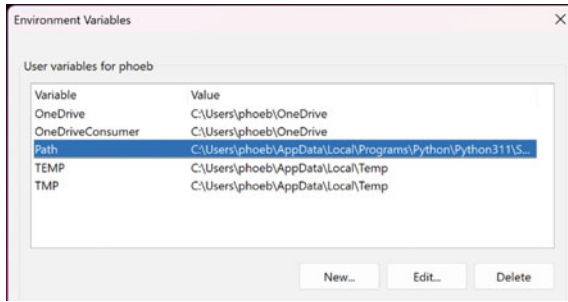
The easiest way to find this is to type 'envir' into the Windows search box, it will then list all applications which match this pattern including the 'Edit the system environment variables' editor. If you are logged into your machine as a user with 'admin' rights use the first one listed, if you are logged in as a user without admin rights select the 'Edit environment variables for your account' option.



One the resulting dialog select the 'Environment Variables ...' button at the bottom of the page:



One the next dialog select the PATH environment variable and select Edit:




Now add the location in which you installed Python, by default this will be something like.

```
C:\Users\\AppData\Local\Programs\Python\Python311
C:\Users\\AppData\Local\Programs\Python\Python311\Scripts
```

Note Python311 should be replaced by the version of Python you are installing if it is different.

The end result should look something like:



Now click on OK until all the windows are closed.

3.3.4 Step 4: Verify the Installation

Next open a new Command Prompt window and type `python` as shown below:

```
Command Prompt - python
Microsoft Windows [Version 10.0.22621.1555]
(c) Microsoft Corporation. All rights reserved.

C:\Users\phoeb>python
Python 3.11.3 (tags/v3.11.3:f3909b8, Apr 4 2023, 23:49:59) [MSC v.1934 64 bit (AMD64)] on win32
Type "help", "copyright", "credits" or "license" for more information.
>>>
```

Congratulations you have installed Python and run the `python` interpreter!


3.3.5 Step 5: Run Some Python

Now at the '`>>>`' prompt type in

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

Be careful to make sure you use all lowercase letters for the 'print' function; Python is very case sensitive which means that as far as Python is concerned `print('Hello World')` and `Print('Hello World')` are completely different things.

Also make sure that you have single quotes around the *Hello World* text this makes it a string. If you get this right then on Windows you should see:




```
Command Prompt - python
Microsoft Windows [Version 10.0.22621.1555]
(c) Microsoft Corporation. All rights reserved.

C:\Users\phoeb>python
Python 3.11.3 (tags/v3.11.3:f3909b8, Apr 4 2023, 23:49:59) [MSC v.1934 64 bit (AMD64)] on win32
Type "help", "copyright", "credits" or "license" for more information.
>>> print('Hello World')
Hello World
>>>
```

You have now run your first Python program. In this case your program printed out the message 'Hello World' (a traditional first program in most programming languages).

3.3.6 Step 6: Exit the Python Interpreter

To exit the Python interpreter use `exit()` or `Ctrl-Z` plus `Return`.



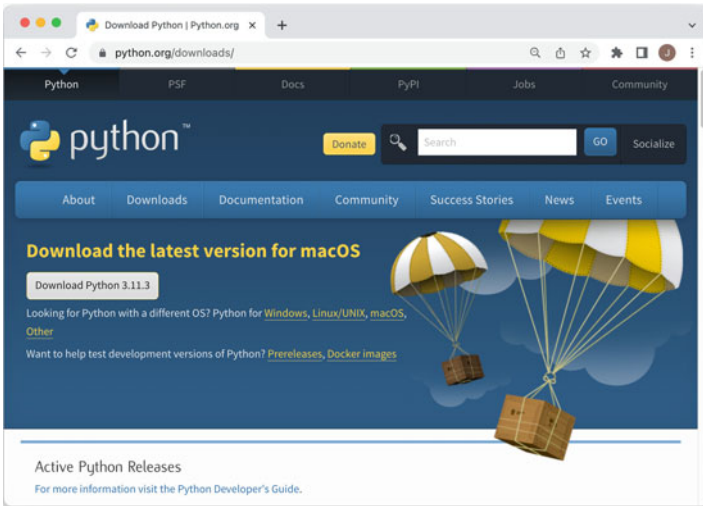
```
Command Prompt
Microsoft Windows [Version 10.0.22621.1555]
(c) Microsoft Corporation. All rights reserved.

C:\Users\phoeb>python
Python 3.11.3 (tags/v3.11.3:f3909b8, Apr 4 2023, 23:49:59) [MSC v.1934 64 bit (AMD64)] on win32
Type "help", "copyright", "credits" or "license" for more information.
>>> print('Hello World')
Hello World
>>> exit()

C:\Users\phoeb>
```

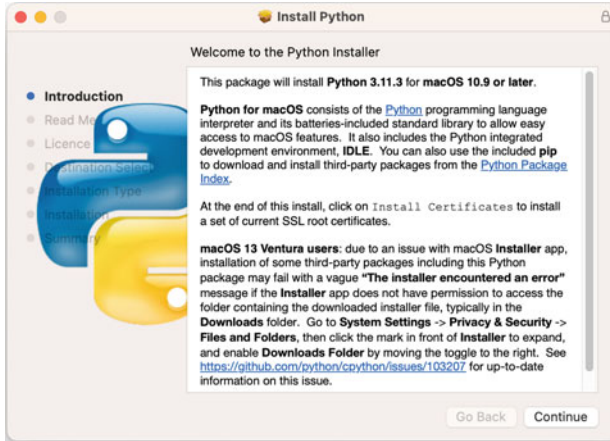
3.4 Setting up on a Mac

Installing Python on a Mac is similar to installing it on a Windows machine in that you can download the Python installer for the Apple Mac from the Python Software Foundation web site (<https://www.python.org>). Of course this time you need to make sure you select the macOS version of the download as shown below:



From here you can select the macOS 64-bit installer (make sure you select the appropriate one for your version of the operating system). This will download an Apple package that can be installed (it will have a name similar to `python-3.11.3-macos11.pkg` although the version number of Python and the macOS operating system may be different). You will need to run this installer.


When you run this installer the Python Installer wizard dialog will open as shown below:



Step through the dialogs presented to you by the installer wizard accepting each option until the installation starts. Once the installation has completed you will be presented with a summary screen confirming the installation was successful. You can now close the installer.

This will have created a new folder in your Applications folder for Python 3.11.

Note that on a Mac which already has Python 2 installed (which is installed by default), Python 3 can be installed along side it and can be accessible via the `python3` command (as shown below). You should confirm that Python has been installed successfully by opening a Terminal window and entering the Python 3 REPL:




Now at the '>>>' prompt type in

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

Be careful to make sure you use all lowercase letters for the 'print' function; as on a Windows machine Python is very case sensitive which means that as far as Python is concerned `print('Hello World')` and `Print('Hello World')` are completely different things.

The result should be as shown below:




You can now exit the REPL using the `exit()` or `quit()`.

3.5 Setting up an IDE

You can use whatever editor you want along with this book, and most chapters are IDE independent. However, if you wish to use the PyCharm editor to make life easier then this section outlines downloading, installing and getting started with this editor. Almost all other chapters are IDE independent apart from one section in the ‘Python Virtual Environments’ chapter.

3.5.1 Downloading the PyCharm IDE

PyCharm is provided by JetBrains who make tools for a variety of different languages. The PyCharm IDE can be downloaded from their site—see <https://www.jetbrains.com/>. Look for the menu heading ‘Developer Tools’ and select that. You will see a long list of tools, which should include PyCharm.



Select the PyCharm option.

On the next page displayed there should be a large button labelled ‘DOWNLOAD’ select this option:



On the next page there are then two download options available: *Professional* and *Community*.

- the Professional version is the charged for option but included several additional elements over the free community version,
- the Community version is free.

For most of the work I do in Python the Community version is more than adequate, and it is therefore the version we will download and install (note with the Professional version you do get a free trial but will need to either pay for the full version at the end of the trial or reinstall the Community version at that point).

Download PyCharm

Windows macOS Linux

Professional For both Scientific and Web Python development. With HTML, JS, and SQL support.	Community For pure Python development
Download .dmg ▼	Download .dmg ▼
Free 30-day trial available	Free, open-source

Assuming you selected the Community edition the installer will now download, and you will be prompted to run it. Note you can ignore the request to subscribe if you want). Note that the web browser should select the correct operating system for you however, if it does not then you can manually select it above the ‘Professional’ and ‘Community’ headings.

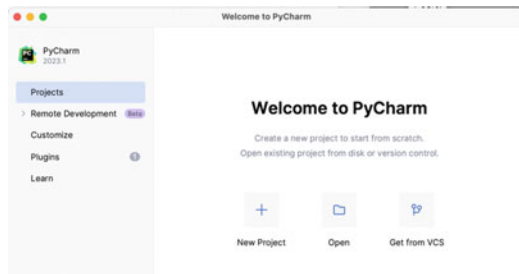
You can now run the installer and follow the instructions provided.

3.5.2 *Setting up the IDE*

You need to first start the PyCharm IDE. Once started the first dialog shown to you asks if you want to import any settings you may have had for another version of PyCharm. At this point select ‘Do not import settings’.

Step through the next set of dialogs selecting the look and feel (I like the light version of the IDE), whether you want to share data with JetBrains, etc. Once you have completed this click the ‘Start PyCharm’ option.

You should now be presented with the landing screen for PyCharm:




We will now create a project for you to work in. A Project in PyCharm is where you write your programs and how you configure what version of Python you are using and any libraries that you might need (such as graphics libraries etc.).

Click on the ‘New Project’ option in the landing dialog.

You will now be asked where you want to create this new project. Again you can use the default location but you will need to give it a name, we will call ours *python-intro*. This is done by changing the last element of the ‘Location’ field to read ‘python-intro’.

It is also worth at this stage making sure that the Python interpreter you installed has been picked up by the IDE. You can do this by opening the ‘Project Interpreter: New Virtualenv environment’ option and making sure that the Base interpreter field is populated with the installation of Python on your machine (this should be picked up automatically for you). If all is OK then select ‘Create’; if the Base Interpreter is not specified or is incorrect then click on the ‘...’ button to the right of the field and browse to the appropriate location.

The ‘New project’ dialog is shown below on the left for Windows 11 and on the right for an Apple Mac:




Next make sure that the ‘Create a main.py welcome script’ is **NOT** selected as this will create a sample program for you which contains far more than a simple ‘Hello World’ message.

On opening the PyCharm project you might see a Welcome message; click ‘Close’ and the project will be set up for you.

When you open the project, you will be shown a blank work space. On the left-hand side you should see at the top of the thin bar three icons:




Select the top one which looks a bit like a folder—this is the project view. This will look something like:



This illustrates the structure of your project. You should ignore ‘External Libraries’ and ‘Scratches and Consoles’ and open the top node. If you called your project something other than python-intro you will find that this node has the name you entered rather than python-intro but that is fine.

When you open the top node you will see that it currently contains a directory called venv—ignore this one as well and close this node so that you should now see:



We are now ready to get started.

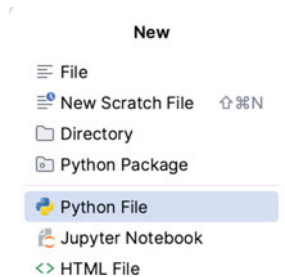
3.6 Writing Python Programs

In this section we will use PyCharm to write our first Python program. If you are using a different IDE then the basic ideas will be the same—create file with a.py extension and run it—however the menu options and buttons may be (will be) different.

The first step will be to create a new file within PyCharm to store our ‘Hello World’ program.

First make sure that the selected node in the project tree is the top level node with the name of your project.

Next select the PyCharm File->New menu option. This will display a ‘New’ menu with the option to create a new ‘Python File’. Make sure that this is the option you select, do not select the plain File option. If you don’t select the Python File option PyCharm will not realize that you want to run the associated file as a Python program. For example:




When you select the ‘Python File’ option PyCharm will display a new dialog allowing you to enter the name of the file you want to create. All Python files must end with a.py extension thus you could enter the name ‘main.py’ into this dialog, however the ‘.py’ is optional here as PyCharm will add that extension for you, thus you can merely enter ‘main’ if you wish:



Again make sure you have selected the 'Python file' option and neither of the other options as they will not work with the simple Hello World program we are creating. Now press enter.


Once you have entered a filename for your program, the enter will create the file for you, for example called `main.py`, and display it on the right-hand side of your editor as an empty file. Note that the editor is divided into two areas, one is the project tree on the left and one is the editor window on the right, for example:




Within the editor window enter the code to print out 'Hello World', for example:

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


Now save the file. For example,



You can now run the file. To do this look above the name of the file in the editor and check that the field next to the little green arrow head say 'Current File'. This indicates that the IDE will run the current file when you click on the green arrow, for example:



Now click on the green arrow head and your program should now run. Below the left hand screen shot is from running PyCharm on a Windows 11 machine and the right-hand one is from an Apple Mac. As you can see they are very similar:



The output from the program is displayed in a ‘Run’ window at the bottom of the IDE. This output shows you what was run (Python and the `main.py` file), the output from the program ‘Hello World’ and that the process (that ran the program) exited successfully with an exit code of ‘0’.

You have now run your first Python program!

3.7 A Note on Filenames

One point to note is that file names in Python need a little care.

Firstly, you should always name your files using all lowercase letters. For example the filename `utils.py` is fine but the file `Utils.py` should be avoided (by convention).

Secondly, you should try to avoid splitting file names up into constituent words, such as `my_utils.py`, this is a style thing mostly, and one that not everyone (including myself) adheres to. However, if you do need to split a filename up to make it more readable always use an underbar ‘_’, never use a hyphen ‘-’. If you do use a hyphen this will not cause a problem right now, but at some point in the future it may come back to bite you. This is because if you ever need to reference one Python file from another the hyphen ‘-’ will break such an operation as the ‘import’ command parses the ‘-’ as a minus sign and thus will try to subtract part of the file name from another part resulting in an error instead of just being the file name.

3.8 Online Resources

See the Python Standard Library documentation for:

- <https://docs.python.org/3/using/index.html> with documentation for Python setup and usage.
- <https://docs.python.org/3/faq/windows.html> Python on Windows FAQ.
- <https://www.jetbrains.com/pycharm/> The PyCharm IDE home page.

Chapter 4

A First Python Program



4.1 Introduction

In this chapter we will return to the Hello World program from the last chapter and look at what it is doing. We will also modify it to become more interactive and will explore the concept of Python variables. We will cover concepts such as naming conventions comments in code, and what as far as this book is concerned is the difference between a script and a program.

4.2 Hello World

As mentioned in the previous chapter, it is traditional to get started in a new programming language with writing a *Hello World* style program. This is very useful as it ensures that your environment, that is, the interpreter, any environmental settings, your editor (or IDE), etc. is all set up appropriately and can process (or compile) and execute (or run) your program. As the 'Hello World' program is about the simplest program in any language, you are doing this without the complexities of the actual language being used.

Our 'Hello World' program has already been presented in the previous chapter of this book, however we will return to it here and take a closer look at what is going on.

In Python the simplest version of the *Hello World* program merely prints out a string with the welcome message:

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

You can use any text editor or Integrated Development Editor (IDE) to create a Python file. Examples of editors commonly used with Python include Emacs, Vim, Notepad++, Sublime Text or Visual Studio Code; examples of IDEs for Python include PyCharm and Eclipse. Using any of these tools we can create file with a `py`

extension. Such a file can contain one or more Python statements that represent a Python program or script.

For example, we can create a file called `hello.py` containing the above `print()` function in it.

One question this raises is where does the `print()` function come from?

In fact, `print()` is a predefined function that can be used to *print things out*, for example to the user. The output is actually printed to what is known as the output stream. This handles a stream (sequence) of data such as letters and numbers. This output stream of data can be sent to an output window such as the terminal on a Mac or command window on a Windows PC. In this case we are printing the string 'Hello World'.

By predefined here we mean that it is *built into* the Python environment and is understood by the Python interpreter. This means that the interpreter knows where to find the definition of the `print()` function which tells it what to do when it encounters the `print()` function.

You can of course write your own functions, and we will be looking at how to do that later in this book.

The `print()` function actually tries to print whatever you give it,

- When it is given a string it will print a string.
- If it is given an integer such as 42 it will print 42.
- If it is given a real number such as 23.56 then it will print 23.56.

Thus, when we run this program, the string 'Hello World' is printed out to the console window.

Also note that the text forming the Hello World string is wrapped within two single quote characters; these characters delimit the start and end of the string; if you miss one of them out then there will be an error.

To run the program, if you are using an IDE such as PyCharm, then you can click on the green arrow on the top right of the editor window (as long as the 'Current File' option is selected next to the arrow).

If you are running it from the command line type in `python` followed by the name of the file, for example:

```
> python hello.py
```

This should be done in the directory where you created the file.

4.3 Interactive Hello World

Let us make our program a little more interesting; let's get it to ask us our name and say hello to us personally.

The updated program is now listed below:

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

```
user_name = input('Enter your name: ')\nprint('Hello ', user_name)
```

Now after printing the original 'Hello World' string, the program then has two additional statements.

The result of running this program is:

```
Hello World\nEnter your name: John\nHello John
```

We will look at each of the new statements separately.

The first statement is:


```
user_name = input('Enter your name: ')
```

This statement does several things. It first executes another function called `input()`. This function has passed a string—which is known as an argument—to use when it prompts the user for input.

This function `input()` is again a *built-in* function that is part of the Python language. In this case it will display the string you provide as a prompt to the user and wait until the user types something in followed by the return key.

Whatever the user types in is then returned as the result of executing the `input()` function. In this case that result is then stored in the *variable* `user_name`.

A *variable* is a named area of the computers' memory that can be used to hold things (often referred to as data) such as strings, numbers, Boolean values such as True/False. In this case the variable `user_name` is acting as a label for an area of memory which will hold the string entered by the user. The basic idea is illustrated in the following diagram:



This simplified picture illustrates how a variable, such as `user_name`, can reference an area of memory containing actual data. In this diagram the memory is shown as a two-dimensional grid of memory *locations*. Each location has an address associated with it. This address is unique within the memory and can be used to return

to the data held at that location. This address is often referred to as the memory address of the data. It is this memory address that is actually held in the variable `user_name`; this is why the `user_name` variable is shown as pointing to the area in memory containing the string 'John'.

Thus the variable `user_name` allows us to access this area of memory easily and conveniently.

For example, if we want to get hold of the name entered by the user in another statement, we can do so merely by referencing the variable name. Indeed, this is exactly what we do in the second statement we added to our program. This is shown below:

```
print('Hello', user_name)
```

The last statement again uses the built-in `print()` function, but this time it takes two arguments. This is because the `print()` function can actually take a variable number of arguments (data items that we pass into it). Each argument is separated by a comma. In this case we have passed in the string 'Hello' and whatever value is referenced by (present at the memory address indicated by) the variable `user_name`.

Something to note is that Python is very case sensitive; thus the variable `My_Name` is not the same as `my_name`! As far as Python is concerned these are two completely separate variables, and you should therefore be very careful with the case you use. In general all keywords are lowercase, all variables should be lowercase, while most built-in types use lowercase names, and there are some built-in types that may use an initial capital letter.

4.4 Variables

You may wonder why the element holding the user's name above is referred to as a *variable*. It is called a *variable* because the value it references in memory can vary during the lifetime of the program.

For example we can modify our *Hello World* program to ask the user for the name of their best friend and print out a welcome message to that best friend. If we want to, we can *reuse* the variable to hold the name of that best friend. For example:

```
print('Hello, world')
name = input('Enter your name: ')
print('Hello', name)
name = input('What is the name of your best friend: ')
print('Hello Best Friend', name)
```

When we run this version of the program and the user enters 'John' for their name and 'Denise' for their best friends' name we will see:

```
Hello, world
Enter your name: John
```

```

Hello John
What is the name of your best friend: Denise
Hello Best Friend Denise

```

As you can see from this when the string ‘Hello Best Friend’ is printed, it is the name ‘Denise’ that is printed alongside it.

This is because the area of memory that previously held the string ‘John’ now holds the string ‘Denise’.

In fact, in Python the variable name is not restricted to hold a string such as ‘John’ and ‘Denise’; it can also hold other types of data such as numbers or the values `True` and `False`. For example:

```

my_variable = 'John'
print(my_variable)
my_variable = 42
print(my_variable)
my_variable = True
print(my_variable)

```

The result of running the above example is

```

John
42
True

```

As you can see `my_variable` first holds (or references the area of memory containing) the string ‘John’, it then holds the number 42, and finally it holds the Boolean value `True` (Boolean values can only be `True` or `False`).

This is referred to in Python as *Dynamic Typing*. That is the *type* of the data held by a variable can *dynamically* change as the program executes. Although this may seem like the obvious way to do things, it is not the approach used by many programming languages such as Java and C# where variables are instead *statically typed*. The word static is used here to indicate that the type of data that a variable can hold will be determined when the program is first processed (or compiled). Later on it will not be possible to change the type of data it can hold; thus if a variable is to hold a number it cannot later on hold a string.

Both approaches have their pros and cons; but for many people the flexibility of Python’s variables is one of its major advantages.

4.5 Naming Conventions

You may have noticed something about some of the variable names we have introduced above such as `user_name` and `my_variable`. Both these variable names are formed of a set of characters with an underbar between the ‘words’ in the variable name.

Both these variable names highlight a very widely used naming convention in Python, which is that variable names should:

- be all lowercase,
- be in general more descriptive than variable names such as *a* or *b* (although there are some exceptions such as the use of variables *i* and *j* in looping constructs),
- with individual words separated by underscores as necessary to improve readability.

This last point is very important as in Python (and most computer programming languages) spaces are treated as separators which can be used to indicate where one thing ends and another starts. Thus it is not possible to define a variable name such as:

- user name

As the space is treated by Python as a separator and thus Python thinks you are defining two things ‘user’ and ‘name’.

When you create your own variables, you should try to name them following the Python accepted style thus name such as:

- `my_name, your_name, user_name, account_name`
- `count, total_number_of_users, percentage_passed, pass_rate`
- `where_we_live, house_number,`
- `is_okay, is_correct, status_flag`

are all acceptable but

- `A, Aaaaa, aaAAAAa`
- `Myname, myName, MyName` or `MYName`
- `WHEREWELIVE`

do not meet the accepted conventions.

However, it is worth mentioning that there are merely commonly adhered to *conventions* and even Python itself does not always comply with these conventions. Thus if you define a variable name that does not conform to the convention Python will not complain.

4.6 Assignment Operator

One final aspect of the statement shown below has yet to be considered

```
user_name = input('Enter your name: ')
```

What does exactly is this ‘ = ’ between `user_name` variable and the `input()` function?

It is called the *assignment* operator. It is used to assign the value returned by the function `input()` to the variable `user_name`. It is probably the most widely used operator in Python. Of course, it is not just used to assign values from functions as the earlier examples illustrated. For example, we also used it when we stored a string into a variable directly:

```
my_variable = 'Jason'
```

4.7 Python Statements

Throughout this chapter we have used the phrase *statement* to describe a portion of a Python program; for example the following line of code is a statement that prints out a string 'Hello' and the value held in `user_name`.

```
print('Hello', user_name)
```

So what do we mean by a statement? In Python a statement is an instruction that the Python interpreter can execute. This statement may be formed of a number of elements such as the ones above which includes a call to a function and an assignment of a value to a variable. In many cases a statement is a single line in your program, but it is also possible for a statement to extend over several lines particularly if this helps the readability or layout of the code. For example, the following is a single statement, but it is laid out over 6 lines of code to make it easier to read:

```
print('The total population for',  
      city,  
      'was',  
      number_of_people_in_city,  
      'in',  
      year)
```

As well as statements there are also expressions. An expression is essentially a computation that generates a value, for example:

```
4 + 5
```

This is an expression that adds 4 and 5 together and generates the value 9.

4.8 Comments in Code

It is common practice (although not universally so) to add comments to code to help anyone reading the code to understand what the code does, what its intent was, any design decisions the programmer made, etc.

Comments are sections of a program that are ignored by the Python interpreter—they are not executable code.

A comment is indicated by the ‘#’ character in Python. Anything following that character to the end of the line will be ignored by the interpreter as it will be assumed to be a comment, for example:

```
# This is a comment
name = input('Enter your name: ')
# This is another comment
print(name) # this is a comment to the end of the line
```

In the above, the two lines starting with a # are comments—they are for our human eyes only. Interestingly the line containing the `print()` function also has a comment—that is fine the command starts with the # and runs to the end of the line, anything *before* the # character is not part of the comment.

4.9 Scripts Versus Programs

Python can be run in several ways:

- via the Python interpreter by entering the REPL; the interactive Python session,
- by having the Python interpreter run a file containing stored Python commands,
- by setting up an association at the operating system level so that any files ending with `.py` are always run by the Python interpreter,
- by indicating the Python interpreter to use at the start of the Python file. This is done by including the first line in a file with something similar to `#!/usr/bin/env python`. This indicates that the rest of the file should be passed to the Python interpreter.

All of these can be defined as a way to run a Python program or indeed a Python script.

However, for the purposes of this book we will treat a file containing the first line specifying the Python interpreter to use as a *script*. All other Python code that represents some code to execute for a specific purpose will be called a Python program. However, this is really only a distinction being made here to simplify terminology.

4.10 Online Resources

See the Python Standard Library documentation for:

- https://docs.python.org/3/reference/simple_stmts.html For information on statements in Python.

4.11 Exercises

At this point you should try to write your own Python program. It is probably the easiest to start by modifying the *Hello World* program we have already been studying. The following steps take you through this:

1. If you have not yet run the *Hello World* program, then do so now. To run your program you have several options. The easiest if you have set up an IDE such as PyCharm is to use the ‘run’ menu option. Otherwise if you have set up the Python interpreter on your computer, you can run it from a command prompt (on Windows) or a terminal window (on a Mac/Linux box).
2. Now ensure that you are comfortable with what the program actually does. Try commenting out some lines—what happens; is that the behaviour you expected? Check that you are happy with what it does.
3. Once you have done that, modify the program with your own prompts to the user (the string argument given to the input function). Make sure that each string is surrounded by the single quote characters (’); remember these denote the start and end of a string.
4. Try creating your own variables and storing values into those instead of the variable `user_name`.
5. Add a `print()` statement to the program with your own prompt.
6. Include an assignment that will add two numbers together (e.g., $4 + 5$) and then assign the result to a variable.
7. Now print out that variables value once it has been assigned a value.
8. Make sure you can run the program after each of the above changes. If there is an error reported attempt to fix that issue before moving on.

You must also be careful with the indentation of your program—Python is very sensitive to how code is laid out, and at this point all statements should start at the beginning of the line.

Chapter 5

Python Strings



5.1 Introduction

In the previous chapter we used strings several times, both as prompts to the user and as output from the `print()` function. We even had the user type in their name and store it in a variable that could be used to access this name at a later point in time. In this chapter we will explore what a string is and how you can work with and manipulate them.

5.2 What Are Strings?

During the description of the *Hello World* program we referred to Python strings several times, but what is a string?

In Python a string is a series, or sequence, of characters in order. In this definition a *character* is anything you can type on the keyboard in one keystroke, such as a letter 'a', 'b', 'c' or a number '1', '2', '3' or a special character such as '\', '[', '\$', and a space is also a character ' ', although it does not have a visible representation.

It should also be noted that strings are *immutable*. Immutable means that once a string has been created it cannot be changed. If you try to change a string you will in fact create a new string, containing whatever modifications you made, you will not affect the original string in anyway. For the most part you can ignore this fact, but it means that if you try to get a sub string or split a string you must remember to store the result—we will see this later on in this chapter.

To define the start and end of a string we have used the single quote character ' '; thus all of the following are valid strings:

- 'Hello'
- 'Hello World'
- 'Hello Andrea2000'

- `'To be or not to be that is the question!'`

We can also define an empty string which has no characters in it (it is defined as a single quote followed immediately by a second single quote with not gap between them). This is often used to initialize or reset a variable holding a reference to a string, for example

- `some_string = ""`

5.3 Representing Strings

As stated above, we have used single quotes to define the start and end of a string, however in Python single or double quotes can be used to define a string, and thus both of the following are valid:

- `'Hello World'`
- `"Hello World"`

In Python these forms are exactly the same, although by convention we default to use single quotes. This approach is often referred to as being more Pythonic (which implies it is more the convention used by experienced Python programmers), but the language does not enforce it.

You should note, however, that you cannot mix the two styles of start and end strings; that is, you cannot start a string with a single quote and end a string with a double quote, and thus the following are both illegal in Python:

- `'Hello World" # This is illegal`
- `"Hello World' # So is this`

The ability to use both “ and ‘, however, comes in useful if your string needs to contain one of the other type of string delimiters. This is because a single quote can be embedded in a string defined using double quotes and vice versa; thus we can write the following:

```
print("It's the day")
print('She said "hello" to everyone')
```

The output of these two lines is:

```
It's the day
She said "hello" to everyone
```

A third alternative is the use of triple quotes, which might at first hand seem a bit unwieldy, but they allow a string to support multiline strings, for example:

```
z = """
Hello
World
"""
```

```
print(z)
```

which will print out

```
Hello  
World
```

5.4 What Type is String?

It is often said that Python is untyped, but this is not strictly true—as was stated in the first chapter it is a dynamically typed language with all data having an associated type.

The type of an item of data (such as a string) determines what it is legal to do with the data and what the effect of various actions will be. For example, what the effect of using a ‘+’ operator is will depend on the *types* involved; if they are numbers then the plus operator will add them together; if however strings are involved then the strings will be concatenated (combined) together, etc.

It is possible to find out what type a variable currently holds using the built-in `type()` function. This function takes a variable name and will return the type of the data held by that variable, for example:

```
my_variable = 'Bob'  
print(type(my_variable))
```

The result of executing these two lines of code is the output:

```
<class 'str'>
```

This is a shorthand for saying that what is held in `my_variable` is currently a class (type) of string (actually string is a class, Python supports ideas from object-oriented programming such as classes, and we will encounter them later in the book).

5.5 What Can You Do with Strings?

In Python terms this means what operations or functions are their available or built in that you can use to work with strings. The answer is that there are very many. Some of these are described in this section.

5.5.1 String Concatenation

You can concatenate two strings together using the '+' operator (an operator is an operation or behaviour that can be applied to r = types involved). That is you can take one string and add it to another string to create a new third string:

```
string_1 = 'Good'
string_2 = " day"
string_3 = string_1 + string_2
print(string_3)
print('Hello ' + 'World')
```

The output from this is

```
Good day
Hello World
```

Notice that the way in which the string is defined does not matter here, `string_1` used single quote, but `string_2` used double quotes; but they are both just strings.

5.5.2 Length of a String

It can sometimes be useful to know how long a string is; for example if you are putting a string into a user interface you might need to know how much of the string will be displayed within a field. To find out the length of a string in Python you use the `len()` function, for example:

```
print(len(string_3))
```

This will print out the length of the string currently being held by the variable `string_3` (in terms of the number of characters contained in the string).

5.5.3 Accessing a Character

As a string is a fixed sequence of letters, it is possible to use square brackets and an index (or position) to retrieve a specific character from within a string. For example:

```
my_string = 'Hello World'
print(my_string[4])
```

However, you should note that strings are indexed from *zero*! This means that the first character is in position 0, the second in position 1, etc. Thus stating `[4]` indicates that we want to obtain the fifth character in the string, which in this case is the letter 'o'. This form of indexing elements is actually quite common in programming languages and is referred to a *zero-based indexing*.

5.5.4 Accessing a Subset of Characters

It is also possible to obtain a subset of the original string, often referred to as a substring (of the original string). This can again be done using the square brackets notation but using a ':' to indicate the start and end points of the sub string. If one of the positions is omitted then the start or end of the string is assumed (depending upon the omission), for example:

```
my_string = 'Hello World'
print(my_string[4]) # characters at position 4
print(my_string[1:5]) # from position 1 to 5
print(my_string[:5]) # from start to position 5
print(my_string[2:]) # from position 2 to the end
```

will generate

```
o
ello
Hello
llo World
```

As such `my_string[1:5]` returns the substring containing the 2nd to 6th letters (that is 'ello'). In turn `my_string[:5]` returned the substring containing the 1st to 6th letters and `my_string[2:]` the sub string containing the 3rd to the last letters.

5.5.5 Repeating Strings

We can also use the '*' operator with strings. In the case of strings this means repeat the given string a certain number of times. This generates a new string containing the original string repeated *n* number of times. For example:

```
print('*' * 10)
print('Hi' * 10)
```

will generate

```
*****
HiHiHiHiHiHiHiHiHiHiHiHiHiHiHi
```

5.5.6 Splitting Strings

A very common requirement is the need to split a string up into multiple separate strings based on a specific character such as a space or a comma.

This can be done with the `split()` function that takes a string to be used in identifying how to split up the receiving string. For example:

```
title = 'The Good, The Bad, and the Ugly'
print('Source string:', title)
print('Split using a space')
print(title.split(' '))
print('Split using a comma')
print(title.split(','))
```

This produces as output

```
Source string: The Good, The Bad, and the Ugly
Split using a space
['The', 'Good', 'The', 'Bad', 'and', 'the', 'Ugly']
Split using a comma
['The Good', ' The Bad', ' and the Ugly']
```

As can be seen from this the result generated is either a list of each word in the string or three strings as defined by the comma.

You may have noticed something odd about the way in which we wrote the call to the split operation. We did not pass the string into `split()` rather we used the format of the variable containing the string followed by `'.'` and then `split()`.

This is because `split()` is actually what is referred to as a *method*. We will return to this concept more when we explore classes and objects. For the moment merely remember that methods are *applied* to things like string using the *dot* notation.

For example, given the following code

```
title = 'The Good, The Bad, and the Ugly'
print(title.split(' '))
```

This means take the string held by the variable `title` and split it based on the character space.

5.5.7 Counting Strings

It is possible to find out how many times a string is repeated in another string. This is done using the `count()` operation for example

```
my_string = 'Count, the number of spaces'
print("my_string.count(' '):", my_string.count(' '))
```

which has the output

```
my_string.count(' '): 8
```

indicating that there are 8 spaces in the original string.

5.5.8 Replacing Strings

One string can replace a substring in another string. This is done using the `replace()` method on a string. For example:

```
welcome_message = 'Hello World!'
print(welcome_message.replace("Hello", "Goodbye"))
```

The output produced by this is thus

```
Goodbye World!
```

5.5.9 Finding Sub Strings

You can find if one string is a substring of another string using the `find()` method. This method takes a second string as a parameter and checks to see if that string is in the string receiving the `find()` method, for example:

```
string.find(string_to_find)
```

The method returns `-1` if the string is not present. Otherwise it returns an index indicating the start of the substring. For example

```
print('Edward Alun Rawlings'.find('Alun'))
```

This prints out the value `7` (the index of the first letter of the substring 'Alun') note strings are indexed from zero; thus the first letter is at position zero, the second at position one, etc.

In contrast the following call to the `find` method prints out `-1` as 'Alun' is no longer part of the target string:

```
print('Edward John Rawlings'.find('Alun'))
```

5.5.10 Converting Other Types into Strings

If you try to use the `' + '` concatenation operator with a string and some other type such as a number then you will get an error. For example if you try the following:

```
msg = 'Hello Lloyd you are ' + 21
print(msg)
```

You will get an *error* message indicating that you can only concatenate string with strings not ints with strings. To concatenate a number such as `21` with a string you

must convert it to a string. This can be done using the `str()` function. This contests any type into a string representation of that type. For example:

```
msg = 'Hello Lloyd you are ' + str(21)
print(msg)
```

This code snippet will print out the message:

```
Hello Lloyd you are 21
```

5.5.11 Remove Prefix and Suffix

Two new string-related methods were added in Python 3.9. These methods are called `removePrefix()` and `removeSuffix()` and are used to remove a sub string from the front or the back of a string. For example:

```
info = '-John Hunt*'
print(info)
print(info.removeprefix('-'))
print(info.removesuffix('*'))
```

The output from this code snippet is:

```
-John Hunt*
John Hunt*
-John Hunt
```

5.5.12 Comparing Strings

To compare one string with another you can use the `'=='` equality and `'!='` not equals operators. These will compare two strings and return wither True or False indicating whether the strings are equal or not.

For example:

```
print('James' == 'James') # prints True
print('James' == 'John') # prints False
print('James' != 'John') # prints True
```

You should note that strings in Python are case sensitive; thus the string `'James'` does not equal the string `'james'`. Thus:

```
print('James' == 'james') # prints False
```

5.5.13 Other String Operations

There are in fact very many different operations available for strings, including checking that a string starts or ends with another string, that is, upper or lowercase tc. It is also possible to replace part of a string with another string, convert strings to upper, lower or title case, etc.

Example of these are given below (note all of these operations use the dot notation):

```
some_string = 'Hello World'
print('Testing a String')
print('-' * 20)
print('some_string', some_string)
print("some_string.startswith('H')", some_string.startswith('H'))
print("some_string.startswith('h')", some_string.startswith('h'))
print("some_string.endswith('d')", some_string.endswith('d'))
print('some_string.istitle()', some_string.istitle())
print('some_string.isupper()', some_string.isupper())
print('some_string.islower()', some_string.islower())
print('some_string.isalpha()', some_string.isalpha())

print('String conversions')
print('-' * 20)
print('some_string.upper()', some_string.upper())
print('some_string.lower()', some_string.lower())
print('some_string.title()', some_string.title())
print('some_string.swapcase()', some_string.swapcase())
print('String leading, trailing spaces', " xyz ".strip())
```

The output from this is

```
Testing a String
-----
some_string Hello World
some_string.startswith('H') True
some_string.startswith('h') False
some_string.endswith('d') True
some_string.istitle() True
some_string.isupper() False
some_string.islower() False
some_string.isalpha() False
String conversions
-----
some_string.upper() HELLO WORLD
some_string.lower() hello world
some_string.title() Hello World
some_string.swapcase() hELLO wORLD
String leading, trailing spaces xyz
```

5.6 Hints on Strings

5.6.1 Python Strings are Case Sensitive

In Python the string 'l' is not the same as the string 'L'; one contains the lowercase letter 'l' and one the uppercase letter 'L'. If case sensitive does not matter to you then you should convert any strings you want to compare into a common case before doing any testing; for example, using `lower()` as in

```
some_string.lower().startswith('h')
```

5.6.2 Function/Method Names

Be very careful with capitalization of function/method names; in Python `isupper()` is a completely different operation to `isUpper()`. If you use the wrong case Python will not be able to find the required function or method and will generate an error message.

5.6.3 Function/Method Invocations

Also be careful of always including the round brackets when you call a function or method; even if it takes no parameters/arguments. There is a significant difference between `isupper` and `isupper()`. The first one is the *name* of an operation on a string, while the second is a call to that operation so that the operation executes. Both formats are legal in Python, but the result is very different, for example:

```
print(some_string.isupper)
print(some_string.isupper())
```

produces the output:

```
<built-in method isupper of str object at 0x105eb19b0>
False
```

Notice that the first print out tells you that you are referring to the built-in method called `isupper` defined on the type string, while the second actually runs `isupper()` for you and returns either `True` or `False`.

5.7 String Formatting

Python provides a sophisticated formatting system for strings that can be useful for printing information out or logging information from a program.

The string formatting system uses a special string known as the *format* string that acts as a pattern defining how the final string will be laid out. This format string can contain placeholders that will be replaced with actual values when the final string is created. A set of values can be applied to the format string to fill the placeholders using the `format()` method.

The simplest example of a *format* string is one that provides a single placeholder indicated by two curly braces (e.g., `{}`). For example, the following is a *format* string with the pattern 'Hello' followed by a placeholder:

```
format_string = 'Hello {}!'
```

This can be used with the `format()` string method to provide a value (or populate) the placeholder, for example:

```
print(format_string.format('Phoebe'))
```

The output from this is:

```
Hello Phoebe!
```

A *format* string can have any number of placeholders that must be populated; for example, the next example has two placeholders that are populated by providing two values to the `format()` method:

```
# Allows multiple values to populate the string
name = "Adam"
age = 20
print("{} is {} years old".format(name, age))
```

In this case the output is:

```
Adam is 20 years old
```

It also illustrates that variables can be used to provide the values for the format method as well as literal values. A literal value is a fixed value such as 42 or the string 'John'.

By default the values are bound to the placeholders based on the order that they are provided to the `format()` method; however this can be overridden by providing an *index* to the placeholder to tell it which value should be bound, for example:

```
# Can specify an index for the substitution
format_string = "Hello {1} {0}, you got {2}%"
print(format_string.format('Smith', 'Carol', 75))
```

In this case the second string 'Carol' will be bound the first placeholder; note that the parameters are numbered from *zero* not one.

The output from the above example is:

```
Hello Carol Smith, you got 75%
```

Of course when ordering the values it is easy to get something wrong either because a developer might think the strings are indexed from 1 or just because they get the order wrong.

An alternative approach is to use *named* values for the placeholders. In this approach the curly brackets surround the name of the value to be substituted, for example, {artist}. Then in the `format()` method a `key = value` pair is provided where the key is the name in the *format* string; this is shown below:

```
# Can use named substitutions, order is not significant
format_string = "{artist} sang {song} in {year}"
print(format_string.format(artist='Paloma Faith',
song='Guilty', year=2017))
```

In this example the order no longer matters as the name associated with the parameter passed into the `format()` method is used to obtain the value to be substituted. In this case the output is:

```
Paloma Faith sang Guilty in 2017
```

It is also possible to indicate alignment and width within the format string. For example, if you wish to indicate a width to be left for a placeholder whatever the actual value supplied, you can do this using a colon (':') followed by the width to use. For example to specify a gap of 25 characters which can be filled with a substitute value you can use `{:25}` as shown below:

```
print('|{:25}|'.format('25 characters width'))
```

In the above the vertical bars are merely being used to indicate where the string starts and ends for reference, and they have no meaning within the format method. This produces the output:

```
|25 characters width |
```

Within this gap you can also indicate an alignment where:

- < indicates left alignment (the default).
- > indicates right alignment.
- ^ indicates centered.

These follow the colon (':') and come before the size of the gap to use, for example:

```
print('|{:<25}|'.format('left aligned')) # The default
print('|{:>25}|'.format('right aligned'))
print('|{: ^25}|'.format('centered'))
```

which produces:

```
|left aligned      |
|      right aligned|
|      centered   |
```

Another formatting option is to indicate that a number should be formatted with separators (such as a comma) to indicate thousands:

```
# Can format numbers with comma as thousands separator
print('{:,}'.format(1234567890))
print('{:,}'.format(1234567890.0))
```

Which generates the output:

```
1,234,567,890
1,234,567,890.0
```

There are in fact numerous options available to control the layout of a value within the format string, and the Python documentation should be referenced for further information.

5.8 String Templates

An alternative to using string formatting is to use string *templates*. These were introduced into Python 2.4 as a simpler, less error prone, solution to most string formatting requirements.

A string template is a class (type of thing) that is created via the `string.Template()` function. The template contains one or more *named* variables preceded with a `$` symbol. The template can then be used with a set of values that replace the template variables with actual values.

For example:

```
import string

# Initialise the template with variables that
# will be substitute with actual values
template = string.Template('$artist sang $song in $year')
```

Note that it is necessary to include an `import` statement at the start of the program as templates are not provided by default in Python; they must be loaded from a library of additional string features. This library is part of Python, but you need to tell Python that you want to access these extra string facilities.

We will return to the `import` statement later in the book; for now just except that it is needed to access the template functionality.

The template itself is created via the `string.Template()` function. The string passed into the `Template()` function can contain any characters plus the template variables (which are indicated by the `$` character followed by the name of the variable such as `$artist` above).

The above is thus a template for the pattern ‘*some-artist sang some-song in some-year*’.

The actual values can be substituted into the template using the `substitute()` function. The `substitute` function takes a set of `key = value` pairs, in which the `key` is the name of the template variable (minus the leading `$` character) and the `value` is the value to use in the string.

```
print(template.substitute(artist='Freddie Mercury', song='The
Great Pretender', year=1987))
```

In this example `$artist` will be replaced by ‘Freddie Mercury’, `$song` by ‘The Great Pretender’ and `$year` by 1987. The `substitute` function will then return a new string containing

```
'Freddie Mercury sang The Great Pretender in 1987'
```

This is illustrated in the following code:

```
import string

# Initialise the template with $variables that
# will be substitute with actual values
template = string.Template('$artist sang $song in $year')

# Replace / substitute template variables with actual values
# Can use a key = value pairs where the key is the name of
# the template Variable and the value is the value to use
# in the string
print(template.substitute(artist='Freddie Mercury', song='The
Great Pretender', year=1987))
```

This produces:

```
Freddie Mercury sang The Great Pretender in 1987
```

We can of course reuse the template by substituting other values for the template variables; each time we call the `substitute()` method it will generate a new string with the template variables replaced with the appropriate values:

```
print(template.substitute(artist='Ed Sheeran', song='Galway
Girl', year=2017))
print(template.substitute(artist='Camila Cabello',
song='Havana', year=2018))
```

With the above producing:

```
Ed Sheeran sang Galway Girl in 2017
Camila Cabello sang Havana in 2018
```

Alternatively you can create what is known as a dictionary. A dictionary is a structure comprised of `key:value` pairs in which the key is unique. This allows a *data structure* to be created containing the values to use and then applied to the `substitute` function:

```
d = dict(artist = 'Billy Idol', song='Eyes Without a Face', year =
1984)
```



```
print(template.substitute(d))
```

This produces a new string:

Billy Idol sang Eyes Without a Face in 1984

We will discuss dictionaries in greater detail later in the book.

Template strings can contain template variables using the format \$name-of-variable; however there are a few variations that are worth noting:

- \$\$ allows you to include a '\$' character in the string without Python interpreting it as the start of a template variable the double '\$\$' is replaced with a single \$. This is known as escape a control character.
- \${template_variable} is equivalent to \$template_variable. It is required when valid identifier characters follow the placeholder but are not part of the placeholder, such as "\${noun}ification".

Another point to note about the template.substitute() function is that if you fail to provide all the template variables with a value then an error will be generated. For example:

```
print(template.substitute(artist='David Bowie',      song='Rebel
Rebel'))
```

will result in the program failing to execute and an error message being generated:

```
Traceback (most recent call last):
  File "template_examples.py", line 18, in <module>
    print(template.substitute(artist='David Bowie', song='Rebel
Rebel'))
^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
  File    "/usr/local/Cellar/python@3.11/3.11.3/Frameworks/
Python.framework/Versions/3.11/lib/python3.11/string.py",  line
121, in substitute
    return self.pattern.sub(convert, self.template)
^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
  File    "/usr/local/Cellar/python@3.11/3.11.3/Frameworks/
Python.framework/Versions/3.11/lib/python3.11/string.py",  line
114, in convert
    return str(mapping[named])
~~~~~
KeyError: 'year'
```

This is because the template variable \$year has not been provided with a value. If you do not want to have to worry about providing all the variables in a template with a value then you should use the safe_substitute() function:

```
print(template.safe_substitute(artist='David Bowie', song='Rebel
Rebel'))
```

This will populate the template variables provided and leave any other template variables to be incorporated into the string as they are, for example:

```
David Bowie sang Rebel Rebel in $year
```

5.9 Formatting Using f-strings

The simplest approach to string formatting was introduced in Python 3.6. For this version onward, a formatted string could be indicated by prefixing the string with the letter 'f'. This so called f-string is then parsed by the Python runtime and allows data to be inserted into the string template. The values to be inserted are indicated using curly brackets {}. Within the curly brackets you can enter any valid Python expression; thus you can enter the name of a variable, an expression such as $3 + 5$ or indeed call a Python function such as `len(some_string)`. For example:

```
name = "Natalia"
age = 23

message = f"Hello {name}, you are {age}"
print(message)
```

This code, when run, generates the following output:

```
Natalia is 23 years old
```

Another example illustrates the use of expressions within the curly brackets:

```
welcome = 'Hello'
message = f'The value of 2 + 3 is: {2 + 3}, the len of {welcome} is
{len(welcome)}'
print(message)
```

The output of this is:

```
The value of 2 + 3 is: 5, the len of Hello is 5
```

A variation on this theme can be used to format floating point numbers. This approach uses format specifiers to indicate the width to use for the number and the number of decimal places (or precision) to include and a type specifier to indicate that you are using floating point numbers. This format is:

```
f' {value: {width} . {precision} f }'
```

where:

- value is any expression that evaluates to a number.
- width specifies the number of characters used in total to display, but if value needs more space than the width specifies then the additional space is used. This is an optional field as leaving it blank indicates that you want whatever space is required to be used.

- precision indicates the number of characters used after the decimal point.
- The `f` after the precision indicates that you are using fixed point presentation.

For example,

```
a = 10.1234
message = f'Result: {a:.2f}'
print(message)
```

The output from this is:

```
Result: 10.12
```

5.10 Online Resources

There is a great deal of online documentation available on strings in Python including:

- <https://docs.python.org/3/library/string.html> which presents common string operations.
- <https://docs.python.org/3/library/stdtypes.html#text-sequence-type-str> this provides information on strings and the `str` class in Python.
- <https://pyformat.info> has a simple introduction to Python string formatting.
- <https://docs.python.org/3/library/string.html#format-string-syntax> which presents detailed documentation on Python string formatting.
- <https://docs.python.org/3/library/string.html#template-strings> for documentation on string templates.

5.11 Exercises

We are going to try out some of the string-related operations.

1. Explore replacing a string
Create a string with words separated by ',' and replace the commas with spaces; for example replace all the commas in 'Denyse,Marie,Smith,21,London,UK' with spaces. Now print out the resulting string.
2. Handle user input
The aim of this exercise is to write a program to ask the user for two strings and concatenate them together, with a space between them and store them into a new variable called `new_string`.

Next:

- Print out the value of `new_string`.
- Print out how long the contents of `new_string` is.
- Now convert the contents of `new_string` to all uppercase.
- Now check to see if `new_string` contains the string 'Albus' as a substring.

Chapter 6

Numbers, Booleans and None



6.1 Introduction

In this chapter we will explore the different ways that numbers can be represented by the built-in types in Python. We will also introduce the Boolean type used to represent True and False. As part of this discussion we will also look at both numeric and assignment operators in Python. We will conclude by introducing the special value known as None.

6.2 Types of Numbers

There are three types used to represent numbers in Python; these are integers (or integral) types, floating point numbers and complex numbers.

This begs the question why? Why have different ways of representing numbers; after all humans can easily work with the number 4 and the number 4.0 and don't need completely different approaches to writing them (apart from the '.' of course).

This actually comes down to efficiency in terms of both the amount of memory needed to represent a number and the amount of processing power needed to work with that number. In essence integers are simpler to work with and can take up less memory than real numbers. Integers are whole numbers that do not need to have a fractional element. When two integers are added, multiplied or subtracted they will always generate another integer number.

In Python real numbers are represented as floating point numbers (or floats). These can contain a fractional part (the bit after the decimal point). Computers can best work with integers (actually of course only really 1 s and 0 s). They therefore need a way to represent a floating point or real number. Typically this involves representing the digits before and after the decimal point.

The term floating point is derived from the fact that there is no fixed number of digits before or after the decimal point; that is, the decimal point can float.

6.3.1 Converting to Ints

It is possible to convert another type into an integer using the `int()` function. For example, if we want to convert a string into an int (assuming the string contains a integer number) then we can do this using the `int()` function. For example

```
total = int('100')
```

This can be useful when used with the `input()` function.

The `input()` function always returns a string. If we want to ask the user to input an integer number, then we will need to convert the string returned from the `input()` function into an int. We can do this by wrapping the call to the `input` function in a call to the `int()` function, for example:

```
age = int(input('Please enter your age:'))
print(type(age))
print(age)
```

Running this gives:

```
Please enter your age: 21
<class 'int'>
21
```

The `int()` function can also be used to convert a floating point number into an int, for example:

```
i = int(1.0)
```

6.4 Floating Point Numbers

Real numbers, or floating point numbers, are represented in Python using the IEEE 754 double-precision binary floating point number format; for the most part you do not need to know this but it is something you can look up and read about if you wish.

The type used to represent a floating point number is called `float`.

Python represents floating point numbers using a decimal point to separate the whole part from the fractional part of the number, for example:

```
exchange_rate = 1.83
print(exchange_rate)
print(type(exchange_rate))
```

This produces output indicating that we are storing the number 1.83 as a floating point number:

```
1.83
<class 'float'>
```

The largest floating point number that you can represent using the default built-in floating point type is $1.7976931348623157e+308$ which is quite large, however if you need to represent larger numbers then there are other Python libraries (aka modules) that can be used to do this.

6.4.1 Converting to Floats

As with integers it is possible to convert other types such as an int or a string into a float. This is done using the `float()` function:

```
int_value = 1
string_value = '1.5'
float_value = float(int_value)
print('int value as a float:', float_value)
print(type(float_value))
float_value = float(string_value)
print('string value as a float:', float_value)
print(type(float_value))
```

The output from this code snippet is:

```
int value as a float: 1.0
<class 'float'>
string value as a float: 1.5
<class 'float'>
```

6.4.2 Converting an Input String into a Floating Point Number

As we have seen the `input()` function returns a string; what happens if we want the user to input a floating point (or real) number? As we have seen above, a string can be converted into a floating point number using the `float()` function and therefore we can use this approach to convert an input from the user into a float:


```
exchange_rate = float(input("Please enter the exchange rate to use: "))
print(exchange_rate)
print(type(exchange_rate))
```

Using this we can input the string 1.83 and convert it to a floating point number:

```
Please enter the exchange rate to use: 1.83
1.83
<class 'float'>
```


6.5 Complex Numbers

Complex numbers are Python's third type of built-in numeric type. A complex number is defined by a real part and an imaginary part and has the form $a + bi$ (where i is the imaginary part and a and b are real numbers):



The real part of the number (a) is the real number that is being added to the pure imaginary number.

The imaginary part of the number, or b , is the real number coefficient of the pure imaginary number.

The letter 'j' is used in Python to represent the imaginary part of the number, for example:

```

c1 = 1j
c2 = 2j
print('c1:', c1, ', c2:', c2)
print(type(c1))
print(c1.real)
print(c1.imag)

```

We can run this code and the output will be:

```

c1: 1j , c2: 2j
<class 'complex'>
0.0
1.0

```

As you can see the type of the number is 'complex' and when the number is printed directly it is done so by printing both the real and imaginary parts together.

Don't worry if this is confusing; it is unlikely that you will need to use complex numbers unless you are doing some very specific coding, for example within a scientific field.

6.6 Boolean Values

Python supports another type called Boolean; a Boolean type can only be one of True or False (and nothing else). Note that these values are True (with a capital T) and False (with a capital F); true and false in Python are not the same thing and have no meaning on their own.

The equivalent of the int or float class for Booleans is bool.

The following example illustrates storing the two Boolean values into a variable `all_ok`:

```
all_ok = True
print(all_ok)
all_ok = False
print(all_ok)
print(type(all_ok))
```

The output of this is

```
True
False
<class 'bool'>
```

The Boolean type is actually a sub type of integer (but with only the values `True` and `False`) so it is easy to translate between the two, using the functions `int()` and `bool()` to convert from Booleans to Integers and vice versa. For example:

```
print(int(True))
print(int(False))
print(bool(1))
print(bool(0))
```

Which produces

```
1
0
True
False
```

You can also convert strings into Booleans as long as the strings contain either `True` or `False` (and nothing else). For example:

```
status = bool(input('OK to proceed: '))
print(status)
print(type(status))
```

When we run this

```
OK to proceed: True
True
<class 'bool'>
```

6.7 Arithmetic Operators

Arithmetic operators are used to perform some form of mathematical operation such as addition, subtraction, multiplication and division, etc. In Python they are represented by one or two characters. The following table summarizes the Python arithmetic operators:

Operator	Description	Example
+	Add the left and right values together	1 + 2
-	Subtract the right value from the left value	3 - 2
*	Multiply the left and right values	3 * 4
/	Divide the left value by the right value	12/3
//	Integer division (ignore any remainder)	12//3
%	Modulus (aka the remainder operator)—only return any remainder	13%3
**	Exponent (or power of) operator—with the left value raised to the power of the right	3 ** 4

6.7.1 Integer Operations

Two integers can be added together using `+`, for example `10 + 5`. In turn two integers can be subtracted (`10 - 5`) and multiplied (`10 * 4`). Operations such as `+`, `-` and `*` between integers always produce integer results.

This is illustrated below:

```
home = 10
away = 15
print(home + away)
print(type(home + away))

print(10 * 4)
print(type(10*4))

goals_for = 10
goals_against = 7
print(goals_for - goals_against)
print(type(goals_for - goals_against))
```

The output from this is

```
25
<class 'int'>
40
<class 'int'>
3
<class 'int'>
```

However, you may notice that we have missed out division with respect to integers, why is this? It is because it depends on which division operator you use as to what the returned type actually is.

For example, if we divide the integer 100 by 20 then the result you might reasonably expect to produce might be 5; but it is not, it is actually 5.0:

```
print(100 / 20)
print(type(100 / 20))
```

The output is

```
5.0
<class 'float'>
```

And as you can see from this the type of the result is `float` (that is a floating point number). So why is this the case?

The answer is that division does not know whether two the integers involved divide into one another exactly or not (i.e., is there a remainder). It therefore defaults to producing a floating point (or real) number which can have a fractional part. This is of course necessary in some situations, for example if we divide 3 by 2:

```
res1 = 3/2
print(res1)
print(type(res1))
```

In this case 3 cannot be exactly divided by 2, we might say that 2 goes into 3 once with a remainder. This is what is shown by Python:

```
1.5
<class 'float'>
```

The result is that 2 goes into 3, 1.5 times with the type of the result being a `float`.

If you are only interested in the number of times 2 does go into 3 and are happy to ignore the fractional part then there is an alternative version of the divide operator `//`. This operator is referred to as the integer division operator:

```
res1 = 3//2
print(res1)
print(type(res1))
```

which produces

```
1
<class 'int'>
```

But what if you are only interested in the remainder part of a division, the integer division operator has lost that? Well in that case you can use the modulus operation (`%`). This operator returns the remainder of a division operation: for example:

```
print('Modulus division 4 % 2:', 4 % 2)
print('Modulus division 3 % 2:', 3 % 2)
```

Which produces:

```
Modulus division 4 % 2: 0
Modulus division 3 % 2: 1
```

A final integer operator we will look at is the power operator that can be used to raise a n integer by a given power, for example 5 to the power of 3. The power operator is `**`, this is illustrated below:

```
a = 5
b = 3
print(a ** b)
```

Which generates the number 125.

6.7.2 Negative Number Integer Division

It is also worth just exploring what happens in integer and true division when negative numbers are involved. For example,

```
print('True division 3/2:', 3 / 2)
print('True division 3//2:', -3 / 2)
print('Integer division 3//2:', 3 // 2)
print('Integer division 3//2:', -3 // 2)
```

The output from this is:

```
True division 3/2: 1.5
True division 3//2: -1.5
Integer division 3//2: 1
Integer division 3//2: -2
```

The first three of these might be exactly what you expect given our earlier discussion; however, the output of the last example may seem a bit surprising, why does `3//2` generate 1 but `-3//2` generates `-2`?

The answer is that Python always rounds the result of *integer* division towards minus infinity (which is the smallest negative number possible). This means it pulls the result of the integer division to the smallest possible number, 1 is smaller than 1.5 but `-2` is smaller than `-1.5`.

6.7.3 Floating Point Number Operators

We also have the multiple, subtract, add and divide operations available for floating point numbers. All of these operators produce new floating point numbers:

```
print(2.3 + 1.5)
print(1.5 / 2.3)
print(1.5 * 2.3)
```

```
print(2.3 - 1.5)
print(1.5 - 2.3)
```

These statements produce the output given below:

```
3.8
0.6521739130434783
3.4499999999999997
0.7999999999999998
-0.7999999999999998
```

6.7.4 *Integers and Floating Point Operations*

Any operation that involves both integers and floating point numbers will always produce a floating point number. That is, if one of the sides of an operation such as add, subtract, divide or multiple is a floating point number then the result will be a floating point number. For example, give the integer 3 and the floating point number 0.1 if we multiple them together then we get a floating point number:

```
i = 3 * 0.1
print(i)
```

Executing this we get

```
0.30000000000000004
```

Which may or may not have been what you expected (you might have expected 0.3); however this highlights the comment at the start of this chapter relating to floating point (or real) numbers being represented as an approximation within a computer system. If this was part of a larger calculation (such as the calculation of the amount of interest to be paid on a very large loan over a 10 year period) then the end result might well be out by a significant amount.

It is possible to overcome this issue using one of Python's modules (or libraries). For example, the `decimal` module provides the `Decimal` class that will appropriately handle multiplying 3 and 0.1.

6.7.5 *Complex Number Operators*

Of course you can use operators such as multiply, add, subtract and divide with complex numbers. For example:

```
c1 = 1j
c2 = 2j
c3 = c1 * c2
print(c3)
```

We can run this code and the output will be:

```
(-2+0j)
```

You can also convert another number or a string into a complex number using the `complex()` function. For example:

```
complex(1) # generates (1+0j)
```

In addition the `math` module provides mathematical functions for complex numbers.

6.8 Assignment Operators

In the Chap. 3 we briefly introduced the assignment operator (`=`) which was used to assign a value to a variable. There are in fact several different assignment operators that could be used with numeric values.

These assignment operators are actually referred to as *compound operators* as they combine together a numeric operation (such as add) with the assignment operator. For example, the `+=` compound operator is a combination of the add operator and the `=` operator such that

```
x = 0
x += 1 # has the same behaviour as x = x + 1
```

Some developers like to use these compound operators as they are more concise to write and can be interpreted more efficiently by the Python interpreter.

The following table provides a list of the available compound operators.

Operator	Description	Example	Equivalent
<code>+=</code>	Add the value to the left-hand variable	<code>x += 2</code>	<code>x = x + 2</code>
<code>-=</code>	Subtract the value from the left-hand variable	<code>x -= 2</code>	<code>x = x - 2</code>
<code>*=</code>	Multiple the left-hand variable by the value	<code>x *= 2</code>	<code>x = x * 2</code>
<code>/=</code>	Divide the variable value by the right-hand value	<code>x /= 2</code>	<code>x = x / 2</code>
<code>//=</code>	Use integer division to divide the variable's value by the right-hand value	<code>x //= 2</code>	<code>x = x // 2</code>
<code>%=</code>	Use the modulus (remainder) operator to apply the right-hand value to the variable	<code>x %= 2</code>	<code>x = x % 2</code>
<code>**=</code>	Apply the power of operator to raise the variable's value by the value supplied	<code>x **= 3</code>	<code>x = x ** 3</code>

6.9 None Value

Python has a special type, the `NoneType`, with a single value, `None`.

This is used to represent null values or *nothingness*.

It is not the same as `False`, or an empty string or `0`; it is a *non-value*. It can be used when you need to create a variable but don't have an initial value for it. For example:

```
winner = None
```

You can then test for the presence of `None` using `'is'` and `'is not'`, for example:

```
print(winner is None)
```

This will print out `true` if and only if the variable `winner` is currently set to `None`. Alternatively you can also write:

```
print(winner is not None)
```

Which will print out `True` only if the value of `winner` is *not* `None`.

Several example software using the value `None` and the `'is'` and `'is not'` operators are given below:

```
winner = None
print('winner:', winner)
print('winner is None:', winner is None)
print('winner is not None:', winner is not None)
print(type(winner))
print('Set winner to True')
winner = True
print('winner:', winner)
print('winner is None:', winner is None)
print('winner is not None:', winner is not None)
print(type(winner))
```

The output of this code snippet is:

```
winner: None
winner is None: True
winner is not None: False
<class 'NoneType'>
Set winner to True
winner: True
winner is None: False
winner is not None: True
<class 'bool'>
```


6.10 Online Resources

See the Python Standard Library documentation for:

- <https://docs.python.org/3/library/stdtypes.html#numeric-types-int-float-complex> Numeric Types.
- <https://docs.python.org/3/library/stdtypes.html#truth-value-testing> Boolean values and Boolean testing.
- <https://docs.python.org/3/library/decimal.html> which provides information on the Python decimal module.
- <https://docs.python.org/3/library/cmath.html> which discusses mathematical functions for complex numbers.

If you are interested in how floating point numbers are represented then a good starting points are:

- https://en.wikipedia.org/wiki/Double-precision_floating-point_format which provides an overview of floating point representation.
- https://en.wikipedia.org/wiki/IEEE_754 which is the Wikipedia page on the IEEE 754 Double-precision floating point number format.

6.11 Exercises

The aim of the exercises in this chapter is to explore the numeric types we have been looking at.

6.11.1 General Exercise

Try to explore the different number types available in Python.

You should try out the different numeric operators available and mix up the numbers being used, for example, 1 and well as 1.0, etc.

Check to see the results you get are what you expect.

6.11.2 Convert Kilometers to Miles

The aim of this exercise is to write a program to convert a distance in kilometers into a distance in miles.

1. Take input from the user for a given distance in kilometers. This can be done using the `input()` function.

2. Convert the value returned by the `input()` function from a string into an integer using the `int()` function.
3. Now convert this value into miles—this can be done by dividing the kilometres by 0.6214
4. Print out a message telling the user what the kilometers are in miles.

Chapter 7

Flow of Control Using if Statements



7.1 Introduction

In this chapter we are going to look at the if statement in Python. This statement is used to control the flow of execution within a program based on some condition. These conditions represent some choice point that will be evaluated too True or False. To perform this evaluation it is common to use a comparison operator (for example to check to see if the temperature is greater than some threshold). In many cases these comparisons need to take into account several values and in these situations logical operators can be used to combine two or more comparison expressions together. This chapter first introduces comparison and logical operators before discussing the if statement itself.

7.2 Comparison Operators

Before exploring if statements and various types of looping mechanism we need to discuss *comparison operators*. These are operators that return Boolean values. They are key to the conditional elements of flow of control statements such as if.

A comparison operator is an operator that performs some form of test and returns True or False.

These are operators that we use in everyday life all the time. For example, do I have enough money to buy lunch, or is this shoe in my size, etc.

In Python there are a range of comparison operators represented by typically one or two characters. These are:

Operator	Description	Example
==	Tests if two values are equal	3 == 3
!=	Tests that two values are <i>not</i> equal to each other	2 != 3
<	Tests to see if the left-hand value is less than the right-hand value	2 < 3
>	Tests if the left-hand value is greater than the right-hand value	3 > 2
<=	Tests if the left-hand value is less than <i>or</i> equal to the right-hand value	3 <= 4
>=	Tests if the left-hand value is greater than or equal to the right-hand value	5 >= 4

7.3 Logical Operators

In addition to comparison operators, Python also has logical operators.

Logical operators can be used to combined Boolean expressions together. Typically, they are used with comparison operators to create more complex conditions. Again, we use these every day for example we might consider whether we can afford an ice cream *and* whether we will be having our dinner soon, etc.

There are three logical operators in Python these are listed below:

Operator	Description	Example
and	Returns True if both left and right are true	(3 < 4) and (5 > 4)
or	Returns true if either the left or the right is true	(3 < 4) or (3 > 5)
not	Returns true if the value being tested is False	not 3 < 2

7.4 The If Statement

An if statement is used as a form of conditional programming; something you probably do every day in the real world. That is, you need to decide whether you are going to have tea or coffee or to decide if you will have toast or a muffin for breakfast, etc. In each of these cases you are making a choice, usually based on some information such as I had coffee yesterday, so I will have tea today.

In Python such choices are represented programmatically by the `if` condition statement.

In this construct if some condition is true some action is performed, optionally if it is not true some other action may be performed instead.

7.4.1 Working with an If Statement

In its most basic form, the if statement is

```
if <condition-evaluating-to-boolean>:
    statement
```

Note that the condition must evaluate to `True` or `False`. If the condition is `True` then we will execute the indented statement.

Note that indentation, this is *very* important in Python; indeed, layout of the code is *very, very* important in Python. Indentation is used to determine how one piece of code should be associated with another piece of the code.

Let us look at a simple example,

```
num = int(input('Enter a number: '))
if num < 0:
    print(num, 'is negative')
```

In this example, the user has input a number; if it is less than zero a message noting this will be printed to the user. If the number is positive; then nothing will be output.

For example,

```
Enter a number: -1
-1 is negative
```

If we wish to execute multiple statements when our condition is true we can indent several lines; in fact all lines indented to the same level after the if statement will automatically be part of it. For example,

```
num = int(input('Enter another number: '))
if num > 0:
    print(num, 'is positive')
    print(num, 'squared is ', num * num)
print('Bye')
```

If we now run this program and input 2 then we will see.

```
Enter another number: 2
2 is positive
2 squared is 4
Bye
```

However, if we enter the value `-1` then we get.

```
Enter another number: -1
Bye
```

Note that neither of the indented lines was executed.

This is because the two indented lines are associated with the if statement and will only be executed if the Boolean condition evaluates (returns) `True`. However, the statement `print('Bye')` is not part of the if statement; it is merely the next statement to be executed after the if statement (and its associated `print()` statements) have finished.

7.4.2 *Else in an If Statement*

We can also define an *else* part of an `if` statement; this is an optional element that can be run if the conditional part of the `if` statement returns `False`. For example:

```
num = int(input('Enter yet another number: '))
if num < 0:
    print('Its negative')
else:
    print('Its not negative')
```

Now when this code is executed, if the number entered is less than zero then the first `print()` statement will be run otherwise (else) the second `print()` statement will be run. However, we are *guaranteed* that at least one (and at most one) of the `print()` statements will execute.

For example, in run 1 if we enter the value 1:

```
Enter yet another number: 1
Its not negative
```

And in run 2 if we enter the value `-1`:

```
Enter yet another number: -1
Its negative
```

7.4.3 *The Use of elif*

In some cases there may be several conditions you want to test, with each condition being tested if the previous one failed. This *else-if* scenario is supported in Python by the `elif` element of an `if` statement.

The `elif` element of an `if` statement follows the `if` part and comes before any (optional) `else` part. It has the format:

```
elif <condition-evaluating-to-boolean>:
    statement
```

For example

```
savings = float(input("Enter how much you have in savings: "))
if savings == 0:
    print("Sorry no savings")
elif savings < 500:
    print('Well done')
elif savings < 1000:
    print('Thats a tidy sum')
elif savings < 10000:
```

```

    print('Welcome Sir!')
else:
    print('Thank you')

```

If we run this:

```

Enter how much you have in savings: 500
Thats a tidy sum

```

Here we can see that the first if condition failed (as savings is not equal to 0). However, the next `elif` also must have returned `False` as savings were greater than 500. In fact it was the second `elif` statement that returned `True` and thus the associated `print('Thats a tidy sum')` statement was executed. Having executed this statement the if statement then terminated (the remaining `elif` and `else` parts were ignored).

7.5 Nesting if Statements

It is possible to *nest* one `if` statement inside another. This term nesting indicates that one `if` statement is located within part of the another `if` statement and can be used to refine the conditional behaviour of the program.

An example is given below. Note that it allows some behaviour to be performed before and after the nested `if` statement is executed / run. Also note that indentation is key here as it is how Python works out whether the `if` statements are nested or not.

```

snowing = True
temp = -1
if temp < 0:
    print('It is freezing')
    if snowing:
        print('Put on boots')
    print('Time for Hot Chocolate')
print('Bye')

```

In this example, if the temperature is less than Zero then we will enter the if block of code. If it is not less than zero we will skip the whole If statement and jump to the `print('Bye')` statement which is after both If statements.

In this case the temperature is set to `-1` and so we will enter the `if` statement. We will then print out the 'It is freezing' string. At this point another `if` statement is *nested* within the first `if` statement. A check will now be made to see if it is snowing. Notice that `snowing` is already a Boolean value and so will be either `True` or `False` and illustrates that a Boolean value on its own can be used here.

As it is snowing, we will print out 'Put on boots'.

However, statement printing out 'Time for Hot Chocolate' is not part of the nested `if`. It is part of the outer if (this is where indentation is important). If you wanted it

to only be printed out if it was snowing then it must be indented to the same level as the first statement in the nested if block, for example:

```

snowing = True
temp = -1
if temp < 0:
    print('It is freezing')
    if snowing:
        print('Put on boots')
        print('Time for Hot Chocolate')
print('Bye')
```

This now changes the inner (or nested) if to have two print statements associated with it.

This may seem subtle, but it is *key* to how Python uses layout to link together individual statements.

7.6 If Expressions

An if expression is a short-hand form of an if statement that returns a value. In fact, the difference between an expression and a statement in a programming language is just that; expressions return a value; statements do not.

It is quite common to want to assign a specific value to a variable dependent on some condition. For example, if we wish to decide if someone is a teenager or not then we might check to see if they are over 12 and under 20. We could write this as:

```

age = 15
status = None
if (age > 12) and age < 20:
    status = 'teenager'
else:
    status = 'not teenager'
print(status)
```

If we run this, we get the string 'teenager' printed out.

However, this is quite long and it may not be obvious that the real intent of this code was to assign an appropriate value to `status`.

An alternative is an *if expression*. The format of an if expression is

```
<result1> if <condition-is-met> else <result2>
```

That is the result returned from the if expression is the first value unless the condition fails in which case the result returned will be the value after the else. It may seem confusing at first, but it becomes easier when you see an example.

For example, using the `if` expression we can perform a test to determine the value to assign to `status` and return it as the result of the `if` expression. For example:

```
status = 'teenager' if age > 12 and age < 20 else 'not teenager'
print(status)
```

Again, the result printed out is 'teenager' however now the code is much more concise, and it is clear that the purpose of the test is to determine the result to assign to `status`.

7.7 A Note on True and False

Python is actually quite flexible when it comes to what actually is used to represent `True` and `False`, in fact the following rules apply.

- 0, '' (empty strings), `None` equate to `False`.
- Non-zero, non-empty strings, any object equate to `True`.

However, we would recommend sticking to just `True` and `False` as it is often a cleaner and safer approach.

7.8 Hints

One thing to be very careful of in Python is layout.

Unlike language such as Java and C# the layout of your program is *part of your program*. It determines how statements are associated together and how flow of control elements such as `if` statements effect which statements are executed.

Also, be careful with the `if` statement and its use of the ':' character. This character is key in separating out the conditional part of the `if` statement from the statements that will be executed depending upon whether the condition is `True` or `False`.

7.9 Online Resources

See the Python Standard Library documentation for:

- <https://docs.python.org/3/library/stdtypes.html#boolean-operations-and-or-not> Boolean Operations.
- <https://docs.python.org/3/library/stdtypes.html#comparisons> Comparison operators.
- <https://docs.python.org/3/tutorial/controlflow.html> the online Python flow of control tutorial.

7.10 Exercises

There are three different exercises in this section, you can select which you are interested in or do all three.

7.10.1 *Check Input is Positive or Negative*

The aim of this exercise is to write a small program to test if an integer is positive or negative.

Your program should:

1. Prompt the user to input a number (use the `input()` function). You can assume that the input will be some sort of number.
2. Convert the string into an integer using the `int()` function.
3. Now check whether the integer is a positive number or a negative number.
4. You could also add a test to see if the number is Zero.

7.10.2 *Test if a Number is Odd or Even*

The exercises require you to write a program to take input from the user and determine if the number is odd or even. Again you can assume that the user will enter a valid integer number.

Print out a message to the user to let them know the result.

To test if a number is even you can use

```
(num % 2) == 0
```

which will return `True` if the number is even (note the brackets are optional but make it easier to read).

7.10.3 *Kilometers to Miles Converter*

In this exercise you should return to the kilometers to miles converter you wrote in the last chapter.

We will add several new tests to your program:

1. Modify your program such that it verify that the user has entered a positive distance (i.e., they cannot enter a negative number).
2. Now modify your program to verify that the input is a number; if it is not a number then do nothing; otherwise convert the distance to miles

To check to see if a string contains only digits use the method `isnumeric()` for example `'42'.isnumeric()`; which returns `True` if the string only contains numbers. Note this method only works for positive integers; but this is sufficient for this example.

Chapter 8

Iteration/Looping



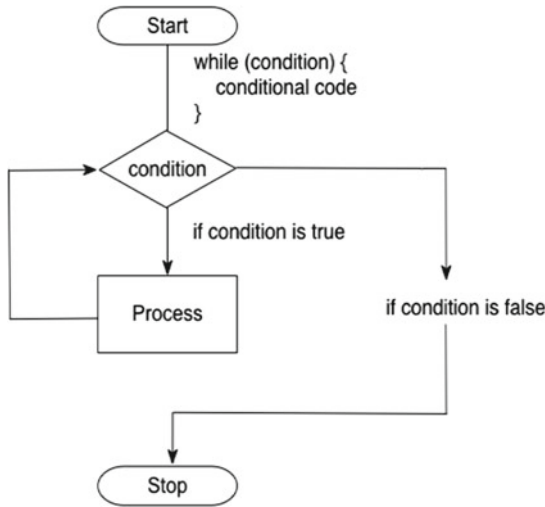
8.1 Introduction

In this section we will look at the `while` loop and the `for` loop available in Python. These loops are used to control the repeated execution of selected statements. A while loop is often referred to as a conditional loop as it will continue to loop while some condition is true. A for loop is often referred to as a counted loop as it will loop a specific number of times.

8.2 While Loop

The `while` loop exists in almost all programming languages and is used to iterative (or repeat) one or more code statements as long as the test condition (expression) is `True`. This iteration construct is usually used when the number of times when we need to repeat the block of code to execute it not know. For example, it may need to repeat until some solution is found or the user enters a particular value.

The behaviour of the while loop is illustrated in below.



The Python while loop has the basic form:

```
while <test-condition-is-true>:  
    statement or statements
```

As shown both in the diagram and can be inferred from the code; *while* the test condition/expression is True then the statement or block of statements in the while loop will be executed.

Note that the test is performed *before* each iteration, including the first iteration; therefore, if the condition fails the first time around the loop the statement or block of statements may never be executed at all.

As with the *if* statement, indentation is key here as the statements included in the while statement are determined by indentation. Each statement that is indented to the same level after the while condition is part of the *while* loop.

However, as soon as a statement is no longer following the *while* block, then it is no longer part of the *while* loop, and its execution is no longer under the control of the test condition.

The following illustrates an example *while* loop in Python:

```
count = 0  
print('Starting')  
while count < 10:  
    print(count, ' ', end=") # part of the while loop  
    count += 1             # also part of the while loop  
print() # not part of the while loop  
print('Done')
```

In this example while some variable `count` is less than the value 10 the `while` loop will continue to iterate (will be repeated). The `while` loop itself contains two statements: one prints out the value of `count` variable, while the other increments `count` (remember `count += 1` is equivalent to `Count = count + 1`).

We have used the version of `print()` that does not print a carriage return when it prints out a value (this is indicated by the `end = ''` option passed to the `print()` function).

The result of running this example is:

```
Starting
0 1 2 3 4 5 6 7 8 9
Done
```

As you can see the statements printing the starting and done messages are only run once. However, the statement printing out the `count` variable is run 10 times (printing out the values 0 to 9).


Once the value of `count` is equal to 10 then the loop finishes (or terminates).

Note that we needed to initialize the `count` variable before the loop. This is because it needs to have a value for the first iteration of the `while` loop. That is before the `while` loop does anything the program needs to already know the first value of `count` so that it can perform that very first test. This is a feature of the `while` loops' behaviour.

8.3 For Loop

In many cases we know how many times we want to iterate over one or more statements (as indeed we did in the previous section). Although the `while` loop can be used for such situations, the `for` loop is a far more concise way to do this. It is typically also clearer to another programmer that the loop must iterate for a specific number of iterations.

The `for` loop is used to step a *variable* through a series of values until a given test is met. The behaviour of the `for` loop is illustrated in below.



This flowchart shows that some sequence of values (e.g., all integer values between 0 and 9) will be used to iterate over a block of code to process. When the last item in the sequence has been reached, the loop will terminate.

Many languages have a `for` loop of the form:

```

for i = from 0 to 10
  statement or statements
  
```

In this case a variable 'i' would take the values 0, 1, 2, 3, etc. up to 10.

8.3.1 Python for Loops

In Python the approach is slightly different as the values from 0 up to 10 are represented by a *range*. This is actually a function that will generate the *range* of values to be used as the sequence in the `for` loop. This is because the Python `for` loop is very flexible and can loop over not only a range of integer values but also a set of values held in data structures such as a list of integers or strings. We will return to this feature of the `for` loop when we look at collections/containers of data in a later chapter.

The format of the Python `for` loop when using a range of values is

```
for <variable-name> in range(...):
    statement
    statement
```

An example is shown below which equates to the `while` loop we looked at earlier:

```
# Loop over a set of values in a range
print('Print out values in a range')
for i in range(0, 10):
    print(i, ' ', end=")
print()
print('Done')
```

When we run this code, the output is:

```
Print out values in a range
0 1 2 3 4 5 6 7 8 9
Done
```

As can be seen from the above; the end result is that we have generated a `for` loop that produces the same set of values as the earlier `while` loop. However,

- The code is more concise.
- It is clear we are processing a range of values from 0 to 9 (note that it is up to but not including the final value).
- We did not need to define the loop variable first.

For these reasons *for* loops are more common in programs in general than *while* loops.

Note that the Python `for` loop can be described as open-closed. That is the first number is included in the range (it is open), but the last number is not included in the range (it is closed), and thus specifying a `range(0, 10)` actually creates a range of the numbers 0 through to 9—it does not include the number 10.

8.3.2 Looping by Non-default Increments

One thing you might notice though is that in the `while` loop we are not constrained to incrementing the count variable by one (we just happened to do this). For example, we could have decided to increment count by 2 each time round the loop (a very common idea). In fact, `range` allows us to do exactly this; a third argument that can be provided to the `range` function is the value to increment the loop variable by each time round the loop, for example:

```
# Now use values in a range but increment by 2
print('Print out values in a range with an increment of 2')
for i in range(0, 10, 2):
    print(i, ' ', end=")
```



```
print()
print('Done')
```

When we run this code, the output is

```
Print out values in a range with an increment of 2
0 2 4 6 8
Done
```

Thus the value of the loop variable has jumped by 2 starting at 0. Once its value was 10 or more then the loop terminates. Of course, it is not only the value 2 we could use; we could increment by any meaningful integer 3, 4, 5, etc.

8.3.3 Anonymous Loop Variables

One interesting variation on the for loop is the use of a wild card (a `'_'`) instead of a looping variable; this can be useful if you are only interested in looping a certain number of times and not in the value of the loop counter itself, for example:

```
# Now use an 'anonymous' loop variable
for _ in range(0,10):
    print('.', end=" ")
print()
```

In this case we are not interested in the values generated by the range per se only in looping 10 times; thus there is no benefit in recording the loop variable. The loop variable is represented by the underbar character (`'_'`).

Note that in actual fact this is a valid variable and can be referenced within the loop; however by convention it is treated as being anonymous.

8.4 Negative Loops

It is also possible to loop from a larger number down to a smaller number by setting the loop increment to be a negative number. For example:

```
for i in range(10, 0, -1):
    print(f'{i}', ' ', end=" ")
print()
```

This loop will count down from 10 to zero. However, because the Python for loop is *open-closed*, this means that the number 10 will be included, but the number zero will not be included. The expression `range(10, 0)` actually creates a range of the numbers 10 through to 1—it does not include the zero. The output from this code snippet is therefore:

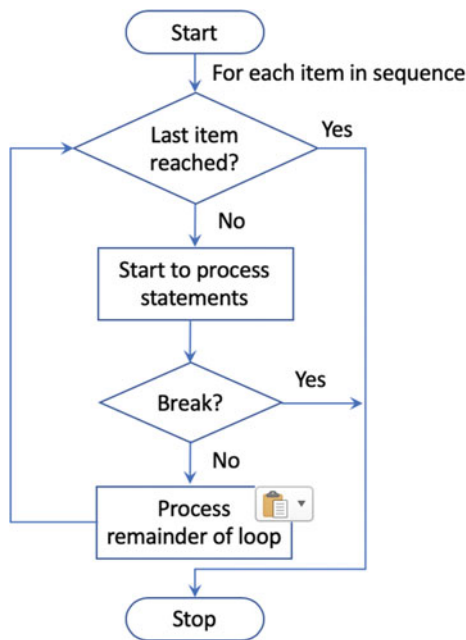
```
10, 9, 8, 7, 6, 5, 4, 3, 2, 1,
```

8.5 Break Loop Statement

Python allows programmers to decide whether they want to break out of a loop early or not (whether we are using a `for` loop or a `while` loop). This is done using the `break` statement.

The `break` statement allows a developer to alter the normal cycle of the loop based on some criteria which may not be predictable in advance (e.g., it may be based on some user input).

The `break` statement, when executed, will terminate the current loop and jump the program to the first line after the loop. The following diagram shows how this works for a `for` loop:



Typically, a guard statement (`if` statement) is placed on the `break` so that the `break` statement is conditionally applied when appropriate.

This is shown below for a simple application where the user is asked to enter a number which may or may not be in the range defined by the `for` loop. If the number is in the range, then we will loop until this value is reached and then break the loop (that is terminate early without processing the rest of the values in the loop):

```
print('Only print code if all iterations completed')
num = int(input('Enter a number to check for: '))
for i in range(0, 6):
    if i == num:
        break
    print(i, ' ', end="")
print('Done')
```

If we run this and enter the value 7 (which is outside of the range) then all the values in the loop should be printed out:

```
Enter a number to check for: 7
0 1 2 3 4 5 Done
```

However, if we enter the value 3 then only the values 0, 1 and 2 will be printed out before the loop breaks (terminates) early:

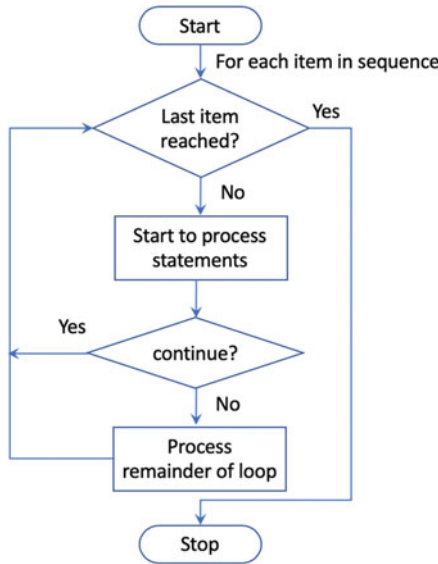
```
Enter a number to check for: 3
0 1 2 Done
```

Note that the string ‘Done’ is printed out in both cases as it is the first line after the `for` loop that is not indented and is thus not part of the `for` loop.

Note that the `break` statement can come anywhere within the block of code associated with the loop construct (whether that is a `for` loop or a `while` loop). This means that there can be statements before it and after it.

8.6 Continue Loop Statement

The `continue` statement also affects the flow of control within the looping constructs `for` and `while`. However, it does not terminate the whole loop; rather it only terminates the current iteration round the loop. This allows you to skip over part of a loop’s iteration for a particular value, but then to continue with the remaining values in the sequence.



A guard (if statement) can be used to determine when the continue statement should be executed.

As with the break statement, the continue can come anywhere within the body of the looping construct. This means that you can have some statements that will be executed for every value in the sequence and some that are only executed when the continue statement is not run.

This is shown below. In this program the continue statement is executed only for odd numbers, and thus the two print() statements are only run if the value of i is even:

```
for i in range(0, 10):
    print(i, ' ', end=" ")
    if i % 2 == 1:
        continue
    print('hey its an even number')
    print('we love even numbers')
print('Done')
```

When we run this code we get

```
0 hey its an even number
we love even numbers
1 2 hey its an even number
we love even numbers
3 4 hey its an even number
we love even numbers
5 6 hey its an even number
we love even numbers
```

```

7 8 hey its an even number
we love even numbers
9 Done

```

As you can see, we only print out the messages about a number being even when the values are 0, 2, 4, 6 and 8.

8.7 For Loop with Else

A for loop can have an *optional* else block at the end of the loop. The else part is executed if and only if all items in the sequence are processed. The for loop may fail to process all elements in the loop if for some reason an error occurs in your program (e.g., if you have a syntax error) or if you break the loop.

Here is an example of a for loop with an else part:

```

# Only print code if all iterations completed over a list
print('Only print code if all iterations completed')
num = int(input('Enter a number to check for: '))
for i in range(0, 6):
    if i == num:
        break
    print(i, ' ', end=")
else:
    print()
    print('All iterations successful')

```

If we run this code and enter the integer value 7 as the number to check for, then the else block executes as the if test within the for statement is never True, thus the loop is never broken, and all values are processed:

```

Only print code if all iterations completed
Enter a number to check for: 7
0 1 2 3 4 5
All iterations successful

```

However, if we enter the value 3 as the number to check for; then the if statement will be True when the loop variable 'i' has the value 3; thus only the values 0, 1 and 2 will be processed by the loop. In this situation the else part will *not* be executed because not all the values in the sequence were processed:

```

Only print code if all iterations completed
Enter a number to check for: 3
0 1 2

```

8.8 A Note on Loop Variable Naming

Earlier in the book we said that variable names should be meaningful and that names such as ‘a’ and ‘b’ were not in general a good idea. The one exception to this rule relates to loop variable names used with for loops over ranges. It is very common to find that these loop variables are called ‘i’, ‘j’, etc.

It is such a common convention that if a variable is called ‘i’ or ‘j’ people expect it to be a loop variable. As such.

- you should consider using these variable names in looping constructs,
- avoid using them elsewhere.

But this does raise the question why ‘i’ and ‘j’; the answer is that it all goes back to a programming language called Fortran which was first developed in the 1950s. In this programming language loop variables had to be called ‘i’, ‘j’, etc. Fortran was so ubiquitous for mathematical and scientific programming, where loops are almost *di rigour*, that this has become the convention in other languages which do not have this restriction.

8.9 Dice Roll Game

The following short program illustrates how a `while` loop can be used to control the execution of the main body of code. In this game we will continue to roll a pair of dice until the user indicates that they do not want to roll again. When this occurs the `while` loop will terminate:

```
import random

MIN = 1
MAX = 6

roll_again = 'y'

while roll_again == 'y':
    print('Rolling the dices...')
    print('The values are....')
    dice1 = random.randint(MIN, MAX)
    print(dice1)
    dice2 = random.randint(MIN, MAX)
    print(dice2)
    roll_again = input('Roll the dices again? (y / n): ')
```

When we run this program the results of rolling two dice are shown. The program will keep looping and printing out the two dice values until the user indicates that they no longer want to roll the dice:

```
Rolling the dices...
The values are....
```

```

2
6
Roll the dices again? (y / n): y
Rolling the dices...
The values are....
4
1
Roll the dices again? (y / n): y
Rolling the dices...
The values are....
3
6
Roll the dices again? (y / n): n

```

8.10 Online Resources

See the Python Standard Library documentation for:

- <https://docs.python.org/3/tutorial/controlflow.html> the online Python flow of control tutorial.

8.11 Exercises

There are two exercises for this chapter. The first exercise will require a simple for loop, while the second is more complicated requiring nested for loops and a break statement.

8.11.1 Calculate the Factorial of a Number

Write a program that can find the factorial of any given number. For example, find the factorial of the number 5 (often written as 5!) which is $1 * 2 * 3 * 4 * 5$ and equals 120.

The factorial is not defined for negative numbers, and the factorial of zero is 1; that is $0! = 1$.

Your program should take as input an integer from the user (you can reuse your logic from the last chapter to verify that they have entered a positive integer value using `isnumeric()`).

You should

1. If the number is less than zero return with an error message.
2. Check to see if the number is zero—if it is then the answer is 1—print this out.
3. Otherwise use a loop to generate the result and print it out.

8.11.2 Print All the Prime Numbers in a Range

A prime number is a positive whole number, greater than 1, that has no other divisors except the number 1 and the number itself.

That is, it can only be divided by itself and the number 1; for example the numbers 2, 3, 5 and 7 are prime numbers as they cannot be divided by any other whole number. However, the numbers 4 and 6 are not because they can both be divided by the number 2; in addition the number 6 can also be divided by the number 3.

You should write a program to calculate prime number starting from 1 up to the value input by the user.

If the user inputs a number below 2, print an error message.

For any number greater than 2 loop for each integer from 2 to that number and determine if it can be divided by another number (you will probably need two for loops for this; one nested inside the other).

For each number that cannot be divided by any other number (that is it's a prime number) print it out.

Chapter 9

Number Guessing Game



9.1 Introduction

In this chapter we are going to bring everything we have learned so far together to create a simple number guessing game.

This will involve creating a new Python program, handling user input, using the if statement as well as using looping constructs.

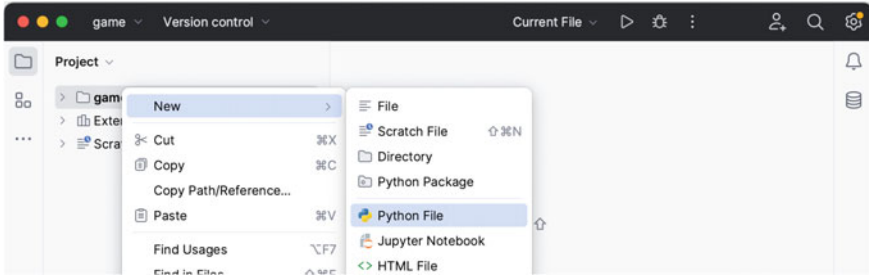
We will also, for the first time, use an additional library or module, that is not by default available by default to your program; this will be the random number generator module.

9.2 Setting up the Program

We want to make sure that we don't overwrite whatever you have done so far, and we would like this Python program to be separate from your other work. As such we will create a new Python file to hold this program. Of course, you don't have to use PyCharm for this, but we will assume you are; if you are using a different editor or IDE then you may have to look up the equivalent steps.

9.2.1 Create a New Python File

We will create a new Python file to hold your program. To do this select the project node in the left-hand view. Once you have selected this use the right mouse button menu. On this menu select 'New' followed by 'Python File', see below:



When you do this, you will be prompted to provide a name for the program file; you can call the file whatever you like, but I am using the name 'number_guess_game' as that is descriptive and will help me find this file again later:



Now click OK, notice I did not need to add .py to the end of the filename; PyCharm will do that for me (if you did add .py then don't worry PyCharm will be fine with that as well).

9.2.2 Add a Welcome Message


To make sure everything is working correctly we will add a simple `print()` statement to the file so that we can test out running it.

You can add anything you like, including printing out 'Hello World', however the example shown below shows a welcome message that can be used at the start of the game:

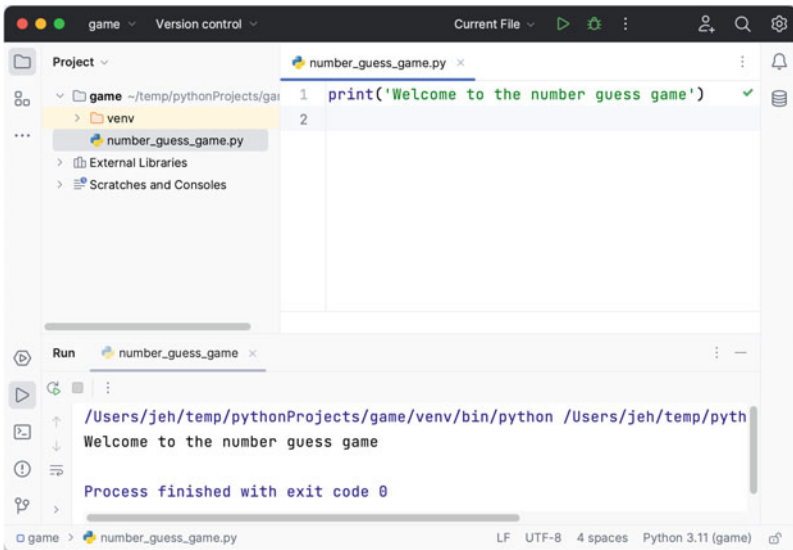


9.2.3 Running the Program

We can now run our embryonic program. To do this we can select the file in the left-hand view and use the right mouse menu to select the ‘Run’ option. Alternatively you can click on the green arrow bottom on the top right bar above the editor window as long as the ‘Current file’ option is displayed next to it:



When we do this the Python console will be opened at the bottom of the IDE and the output displayed:

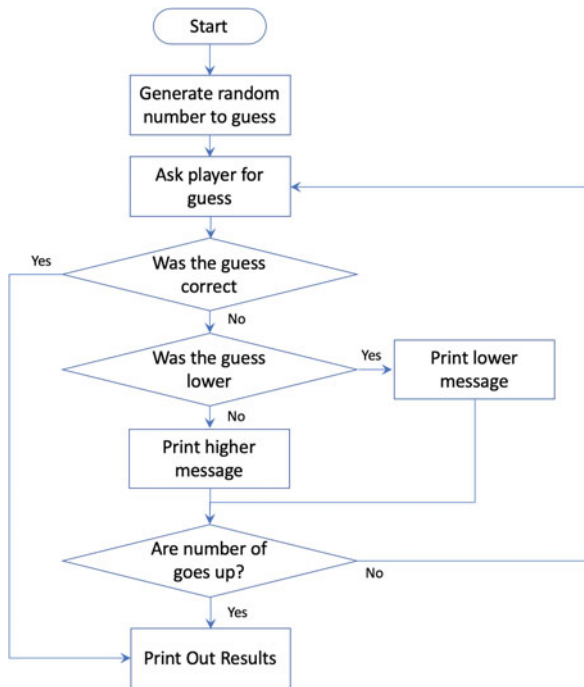


From now on you can rerun the `number_guess_game` program merely by clicking on the little green arrow at the top right-hand side of PyCharm (it allows you to rerun the last program that PyCharm executed).

9.3 What Will the Program Do?

The aim of our number guess game is to guess the number that the program has come up with.

The main flow of the program is shown in the following diagram:



Essentially the program logic is

- The program randomly selects a number between 1 and 10.
- It will then ask the player to enter their guess.
- It will then check to see if that number is the same as the one the computer randomly generated; if it is then the player has won.
- If the player's guess is not the same, then it will check to see if the number is higher or lower than the guess and tell the player.
- The player will have 4 goes to guess the number correctly; if they don't guess the number within this number of attempts, then they will be informed that they have lost the game and will be told what the actual number was.

9.4 Creating the Game

9.4.1 *Generate the Random Number*

We will start off by looking at how we can generate a random number. Up to this point we have only used the built-in functions that are provided by Python automatically. In actual fact Python comes with very many modules provided by the Python organization itself, by third-party vendors and by the open source (typically free) software community.

The Python `random` module or library is one that is provided with Python as part of the default environment; but the functions within it are not automatically loaded or made available to the programmer. This is partly because there are so many facilities available with Python that it could become overwhelming; so for the most part Python only makes available by default the most commonly used facilities. Programmers can then explicitly specify when they want to use some facilities from one of the other libraries or modules.

The `random` module provides implementations of pseudo-random number generators for use in application programs. These random number generators are referred to as *pseudo* because it is very hard for a computer to truly generate a series of random numbers; instead it does its best to mimic this using an algorithm; which by its very nature will be based on some logic which will mean that it is not impossible to predict the next number—hence it is not truly random. For our purposes this is fine but there are applications, such as security and encryption, where this can be a real problem.

To access the `random` module in Python you need to *import* it; this makes the module visible in the rest of the Python file (in our case to our program). This is done using

```
import random
number_to_guess = random.randint(1,10)
```

Once we have imported it, we can use the functions within this module, such as `randint`. This function returns a random integer between the first and second value parameters. In the above example it means that the random number generated will be between 1 and 10 inclusive.

The variable `number_to_guess` will now hold an integer which the player of the game must guess.

9.4.2 *Obtain an Input From the User*

We now need to obtain input from the user representing their guess. We have already seen how to do this in previous chapters; we can use the `input()` function which returns a string and then the integer function that will convert that string into an

integer (we will ignore error checking to make sure they have typed in a number at this point). We can therefore write:

```
guess = int(input('Please guess a number between 1 and 10: '))
```

This will store the number they guessed into the variable `guess`.

9.4.3 *Check to See if the Player Has Guessed the Number*

We now need to check to see whether the player has guessed the correct number.

We could use an `if` statement for this, however we are going to want to repeat this test until either they have guessed correctly or until they have run out of goes.

We will therefore use a `while` loop and check to see if their guess equals the number to be guessed:

```
guess = int(input('Please guess a number between 1 and 10: '))
while number_to_guess != guess:
```

The loop above will be repeated if the number they entered did not match the number to be guessed.

We can therefore print a message telling the player that their guess was wrong:

```
guess = int(input('Please guess a number between 1 and 10: '))
while number_to_guess != guess:
    print('Sorry wrong number')
```

Now we need the player to make another guess otherwise the program will never terminate, so we can again prompt them to enter a number:

```
guess = int(input('Please guess a number between 1 and 10: '))
while number_to_guess != guess:
    print('Sorry wrong number')
    # TBD ...
    guess = int(input('Please guess again: '))
```

9.4.4 *Check They Haven't Exceeded Their Maximum Number of Guess*

We also said above that the player can't play forever; they have to guess the correct number within 4 goes. We therefore need to add some logic which will stop the game once they exceed this number.

We will therefore need a variable to keep track of the number of attempts they have made.

We should call this variable something meaningful so that we know what it represents, for example we could call it `count_number_of_tries` and initialize it with the value 1:

```
count_number_of_tries = 1
```

This needs to happen before we enter the `while` loop.

Inside the `while` loop we need to do two things

- check to see if the number of tries has been exceeded,
- increment the number of tries if they are still allowed to play the game.

We will use an `if` statement to check to see if the number of tries has been met; if it has we want to terminate the loop; the easiest way to this is via a `break` statement.

```
if count_number_of_tries == 4:
    break
```

If we don't break out of the loop we can increment the count using the `+=` operator. We thus now have:

```
count_number_of_tries = 1
guess = int(input('Please guess a number between 1 and 10: '))
while number_to_guess != guess:
    print('Sorry wrong number')
    if count_number_of_tries == 4:
        break
    # TBD ...
    guess = int(input('Please guess again: '))
    count_number_of_tries += 1
```

9.4.5 Notify the Player Whether Higher or Lower

We also said at the beginning that to make it easier for the player to guess the number; we should indicate whether their guess was higher or lower than the actual number. To do this we can again use the `if` statement; if the guess is lower we print one message but if it was higher we print another.

At this point we have a choice regarding whether to have a separate `if` statement to that used to decide if the maximum goes has been reached or to extend that one with an `elif`. Each approach can work but the latter indicates that these conditions are all related so that is the one we will use.

The `while` loop now looks like:

```
count_number_of_tries = 1
guess = int(input('Please guess a number between 1 and 10: '))
while number_to_guess != guess:
    print('Sorry wrong number')
    if count_number_of_tries == 4:
```

```

    break
elif guess < number_to_guess:
    print('Your guess was lower than the number')
else:
    print('Your guess was higher than the number')
guess = int(input('Please guess again: '))
count_number_of_tries += 1

```

Notice that the `if` statement has a final `else` which indicates that the guess was higher; this is fine as by this point it is the only option left.

9.4.6 End of Game Status

We have now covered all the situations that can occur while the game is being played; all that is left for us to do is to handle the end of game messages.

If the player has guessed the number correctly, we want to congratulate them; if they did not guess the number, we want to let them know what the actual number was. We will do this using another `if` statement which checks to see if the player guessed the number or not. After this we will print an end of game message:

```

if number_to_guess == guess:
    print('Well done you won!')
    print('You took', count_number_of_tries ,
          'goes to complete the game')
else:
    print('Sorry - you loose')
    print('The number you needed to guess was',
          number_to_guess)
print('Game Over')

```

9.5 The Complete Listing

For ease of reference the complete listing is provided below:

```

import random

print('Welcome to the number guess game')

# Initialise the number to be guessed
number_to_guess = random.randint(1,10)

# Initialise the number of tries the player has made
count_number_of_tries = 1

# Obtain their initial guess
guess = int(input('Please guess a number between 1 and 10: '))
while number_to_guess != guess:

```



```

print('Sorry wrong number')

# Check to see they have not exceeded the maximum
# number of attempts if so break out of loop otherwise
# give the user some feedback
if count_number_of_tries == 4:
    break
elif guess < number_to_guess:
    print('Your guess was lower than the number')
else:
    print('Your guess was higher than the number')

# Obtain their next guess and increment number of attempts
guess = int(input('Please guess again: '))
count_number_of_tries += 1

# Check to see if they did guess the correct number
if number_to_guess == guess:
    print('Well done you won!')
    print('You took', count_number_of_tries, 'goes to complete
the game')
else:
    print("Sorry - you loose")
    print('The number you needed to guess was',
        number_to_guess)

print('Game Over')

```

And a sample run of the program is shown here:

```

Welcome to the number guess game
Please guess a number between 1 and 10: 5
Sorry wrong number
Your guess was higher than the number
Please guess again: 3
Sorry wrong number
Your guess was lower than the number
Please guess again: 4
Well done you won!
You took 3 goes to complete the game
Game Over

```

9.6 Hints

9.6.1 *Initializing Variables*

In Python it is not necessary to declare a variable before you assign to it; however, it is necessary to give it an initial value before you reference it. Here *reference* it refers to obtaining the value held by a variable. For example

```
count = count + 1
```

what this says is obtain the value held by `count`, add 1 to it and then store the new value back into `count`.

If `count` does not have a value before you try to do this then you are trying to get hold of nothing and add 1 to it; which you can't do and thus an error will be generated in your program, such as:

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

This is also true for

```
count += 1
```

Remember this is just a shorthand form of `count = count + 1` and thus still relies on `count` having a value before this statement.

This is why we needed to initialize the `count_number_of_tries` variable before we used it.

9.6.2 *Blank Lines Within a Block of Code*

You may have noticed that we have used blank lines to group together certain lines of code in this example. This is intended to make it easier to read the code and are perfectly allowable in Python. Indeed, the Python layout guidelines encourage it.

It is also possible to have a blank line within an indented block of code; the block of statements is not terminated until another line with valid Python on it is encountered which is not indented to the same level.

9.7 Exercises

For this chapter the exercises all relate to adding additional features to the game:

1. Provide a cheat mode, for example if the user enters `-1` print out the number they need to guess and then loop again. This does not count as one of their goes.
2. If their guess is within 1 of the actual number tell the player this.
3. At the end of the game, before printing 'Game Over', modify your program so that it asks the user if they want to play again; if they say yes then restart the whole thing.

Chapter 10

Recursion



10.1 Introduction

Recursion is a very powerful way to implement solutions to a certain class of problems. This class of problems is one where the overall solution to a problem can be generated by breaking that overall problem down into smaller instances of the same problem. The overall result is then generated by combining together the results obtained for the smaller problems.

10.2 Recursive Behaviour

A recursive solution in a programming language such as Python is one in which a function calls itself one or more times in order to solve a particular problem. In many cases the result of calling itself is combined with the functions current state to return a result.

In most cases the recursive call involves calling the function but with a smaller problem to solve. For example, a function to traverse a tree data structure might call itself passing in a sub-tree to process. Alternatively a function to generate a factorial number might call itself passing in a smaller number to process etc.

The key here is that an overall problem can be solved by breaking it down into smaller examples of the same problem.

Functions that solve problems by calling themselves are referred to as *recursive* functions.

However, if such a function does not have a termination point then the function will go on calling itself to infinity (at least in theory). In most languages such a situation will (eventually) result in an error being generated.

For a recursive function to be useful it must therefore have a *termination* condition. That is a condition under which they do not call themselves and instead just return (often with some result). The termination condition may be because:

- A solution has been found (some data of interest in a tree structure).
- The problem has become so small that it can be solved without further recursion. This is often referred to as a base case. That is, a base case is a problem that can be solved without further recursion.
- Some maximum level of recursion has been reached, possibly without a result being found/generated.

We can therefore say that a recursive function is a function defined in terms of itself via self-referential expressions. The function will continue to call itself with smaller variations of the overall problem until some termination condition is met to stop the recursion. All recursive functions thus share a common format; they have a recursive part and a termination point which represents the base case part.

10.3 Benefits of Recursion

The key benefit of recursion is that some algorithms (solutions to computer problems) are expressed far more elegantly and with a great deal less code when implemented recursively than when using an iterative approach.

This means that the resulting code can be easier to write and easier to read.

These twin ideas are important both for the developers who initially create the software but also for those developers who must maintain that software (potentially a significant amount of time later).

Code that is easier to write tends to be less error prone. Similarly code that is easier to read tends to be easier to maintain, debug, modify and extend.

Recursion is also well suited to producing functional solutions to a problem as by its very nature a recursive function relies on its inputs and outputs and does not hold any hidden state. We will return to this in a later chapter on functional programming.


10.4 Recursively Searching a Tree

As an example of a problem that is well suited to a recursive solution, we will example how we might implement a program to traverse a binary tree data structure.

A binary tree is a tree data structure made up of nodes in which each node has a value, a left pointer and a right pointer.

The root node is the top most node in the tree. This root node then references a left and right subtree. This structure is repeated until a leaf node. A leaf node is a node in which both the right and left pointers are empty (that is they have the value `None`).

This is shown below for a simple binary tree:



Thus a binary tree is either empty (represented by a null pointer), or is made of a single node, where the left and right pointers each point to a binary tree.

If we now want to find out if a particular value is in the tree then we can start at the root node.

If the root node holds the value we print it; otherwise we can call the `search` function on the child nodes of the current node. If the current node has no children we just return without a result.

The pseudo code for this might look like:

```

search(value_to_find, current_node):
    If current_node.value == value_to_find:
        print('value found:', current_node.value)
    Else If current_node.has_children:
        search(value, current_node.left)
        search(value, current_node.right)
  
```

This illustrates how easy it is to write a recursive function that can solve what might at first appear to be a complex problem.

10.5 Recursion in Python

Most computer programs support the idea of recursion and Python is no exception. In Python it is perfectly legal to have a function that calls itself (that is within the body of the function a call is made to the same function). When the function is executed it will therefore call itself.

In Python we can write a recursive function such as:

```
def recursive_function():
    print('calling recursive_function')
    recursive_function()
```

Here the function `recursive_function()` is defined such that it prints out a message and then calls itself. Note that no special syntax is required for this as a function does not need to have been completely defined before it is used.

Of course in the case of `recursive_function()` this will result in infinite recursion as there is no termination condition. However, this will only become apparent at runtime when Python will eventually generate an error:

```
Traceback (most recent call last):
  File "recursion_example.py", line 5, in <module>
    RecursionError: maximum recursion depth exceeded while calling a
    Python object
```

However, as already mentioned a recursive function should have a *recursive* part and a *termination* or base case part. The termination condition is used to identify when the base case applies. We should therefore add a condition to identify the termination scenario and what the base case behaviour is. We will do this in the following section as we look at how a factorial for a number can be generated recursively.

10.6 Calculating Factorial Recursively

We have already seen how to create a program that can calculate the factorial of a number using iteration as one of the exercises in the last chapter; now we will create an alternative implementation that uses recursion instead.

Recall that the factorial of a number is the result of multiplying that number by each of the integer values up to that number, for example, to find the factorial of the number 5 (written as 5!) we can multiply $1 * 2 * 3 * 4 * 5$ which will generate the number 120.

We can create a recursive solution to the Factorial problem by defining a function that takes an integer to generate the factorial number for. This function will return the value 1 if the number passed in is 1—this is the base case. Otherwise it will multiply the value passed in to it with the result of calling itself (the `factorial()` function) with $n - 1$ which is the recursive part.

The function is given below:

```
def factorial(n):
    if n == 1: # The termination condition
        return 1 # The base case
    else:
        res = n * factorial(n-1) # The recursive call
```

```

    return res

print(factorial(5))

```

The key to understanding this function is that it has:

1. A termination condition that is guaranteed to execute when the value of n is 1. This is the base case; we cannot reduce the problem down any further as the factorial of 1 is 1!
2. The function recursively calls itself but with $n - 1$ as the argument; this means each time it calls itself the value of n is *smaller*. Thus the value returned from this call is the result of a smaller computation.

To clarify how this works we can add some print statements (and a depth indicator) to the function to indicate its behaviour:

```

def factorial(n, depth = 1):
    if n == 1:
        print('\t' * depth, 'Returning 1')
        return 1
    else:
        print('\t'*depth, 'Recursively calling factorial(', n-
1, ')')
        result = n * factorial(n-1, depth + 1)
        print('\t' * depth, 'Returning:', result)
        return result

print('Calling factorial( 5 )')
print(factorial(5))

```

When we run this version of the program then the output is:

```

Calling factorial( 5 )
  Recursively calling factorial( 4 )
    Recursively calling factorial( 3 )
      Recursively calling factorial( 2 )
        Recursively calling factorial( 1 )
          Returning 1
        Returning: 2
      Returning: 6
    Returning: 24
  Returning: 120
120

```

Note that the depth parameter is used merely to provide some indentation to the print statements.

From the output we can see that each call to the factorial program results in a simpler calculation until the point where we are asking for the value of 1! which is 1. This is returned as the result of calling `factorial(1)`. This result is multiplied with the value of n prior to that; which was 2. This causes `factorial(2)` to return the value 2 and so on.

10.7 Disadvantages of Recursion

Although recursion can be a very expressive way to define how a problem can be solved, it is not as efficient as iteration. This is because a function call is more expensive for Python to process than a `for` loop. In part this is because of the infrastructure that goes along with a function call; that is the need to set up the stack for each separate function invocation so that all local variables are independent of any other call to that function. It is also related to associated unwinding the stack when a function returns. However, it is also affected by the increasing amount of memory each recursive call must use to store all the data on the stack.

In some languages optimizations are possible to improve the performance of a recursive solution. One typical example relates to a type of recursion known as *tail recursion*. A tail recursive solution is one in which the calculation is performed before the recursive call. The result is then passed to the recursive step, which results in the last statement in the function just calling the recursive function.

In such situations the recursive solution can be expressed (internally to the computer system) as an iterative problem. That is the programmer writes the solution as a recursive algorithm but the interpreter or compiler converts it into an iterative solution. This allows programmers to benefit from the expressive nature of recursion while also benefiting from the performance of an iterative solution.

You might think that the `factorial` function presented earlier is tail recursive; however it is not because the last statement in the function performs a calculation that multiplies `n` by the result of recursive call.

However, we can refactor the `factorial` function to be tail recursive. This version of the `factorial` function passes the evolving result along via the accumulator parameter. It is given for reference here:

```
def tail_factorial(n, accumulator=1):
    if n == 0:
        return accumulator
    else:
        return tail_factorial(n - 1, accumulator * n)

print(tail_factorial(5))
```

However, it should be noted that Python currently *does not* perform tail recursion optimization; so this is a purely a theoretical exercise.

10.8 Online Resources

The following provides some references on recursion available online:

- [https://en.wikipedia.org/wiki/Recursion_\(computer_science\)](https://en.wikipedia.org/wiki/Recursion_(computer_science)) Provides Wikipedias introduction to recursion.
- <https://www.sparknotes.com/cs/recursion/whatisrecursion/section1/> Provides an introduction to the concept of recursion.


10.9 Exercises

In this set of exercises you will get the chance to explore how to solve problems using recursion in Python.

1. Write a program to determine if a given number is a *Prime Number* or not. Use recursion to implement the solution. The following code snippet illustrates how this might work:

```
print('is_prime(3):', is_prime(3)) # True
print('is_prime(7):', is_prime(7)) # True
print('is_prime(9):', is_prime(9)) # False
print('is_prime(31):', is_prime(31)) # True
```

1. Write a function which implements Pascal's triangle for a specified number of rows. Pascals triangle is a triangle of the binomial coefficients. The values held in the triangle are generated as follows: In row 0 (the topmost row), there is a unique nonzero entry 1. Each entry of each subsequent row is constructed by adding the number above and to the left with the number above and to the right, treating blank entries as 0. For example, the initial number in the first (or any other) row is 1 (the sum of 0 and 1), whereas the numbers 1 and 3 in the third row are added together to generate the number 4 in the fourth row. An example of Pascals triangle for 4 rows is given below:



For example, your function might be called `pascals_triangle()` in which case the following application illustrates how you might use it:

```
triangle = pascals_triangle(5)
for row in triangle:
    print(row)
```

The output from this might be:

```
[1]
[1, 1]
[1, 2, 1]
[1, 3, 3, 1]
[1, 4, 6, 4, 1]
```

Chapter 11

Introduction to Structured Analysis



11.1 Introduction

In the preceding chapters what we have seen is typical of the procedural approach to programming. In the next chapter we will begin to explore the definition of functions which allow a more modular style of programming.

In this chapter we will introduce an approach to the analysis and design of software systems called structured analysis/design. Within this area there are many specific and well-documented methods including Structured Systems Analysis and Design Method (SSADM) and the Yourden structured method. However, we will not focus on any one specific approach; instead we will outline the key ideas and the two basic elements of most structured analysis methods: functional decomposition and data flow analysis. We will then present flowcharts for designing algorithms.

11.2 Structured Analysis and Function Identification

The structured analysis methods typically employ a process-driven approach (with a set of prescribed steps or stages) which in one way or another consider what the inputs and outputs of the system are and how those inputs are transformed into the outputs. This transformation involves the application of one or more functions. The steps involved in structured analysis identify these functions and will typically iteratively break them down into smaller and smaller functions until an appropriate level of detail has been reached. This process is known as functional decomposition.

Although simple Python programs may only contain a sequence of statements and expressions, any program of a significant size will need to be structured such that it can be:

- understood easily by other developers,
- tested to ensure that it does what is intended,
- maintained as new and existing requirements evolve,

- debugged to resolve unexpected or undesirable behaviour.

Given these requirements it is common to want to organize your Python program in terms of functions and sub functions.

Functional decomposition supports the analysis and identification of these functions.

11.3 Functional Decomposition

Functional decomposition is one way in which a system can be broken down into its constituent parts. For example, for a computer payroll system to calculate how much an hourly paid employee should receive it might be necessary to:

1. load the employee's details from some form of permanent storage (such as a file or a database),
2. load how many hours the employee has worked for that week (possibly from another system that records the number of hours worked),
3. multiple the hours worked by the employee's hourly rate,
4. record how much the employee is to be paid in a payroll database or file,
5. print out the employee's pay slip,
6. transfer the appropriate funds from the company's bank account to the employees bank account,
7. record in the payroll database the everything has been completed.

Each of the above steps could represent a function performed by the system.

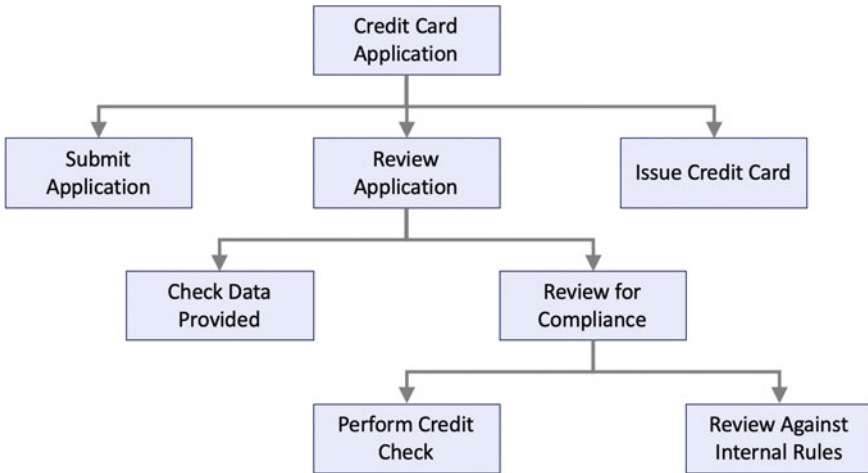
These top-level functions could themselves be broken down into lower-level functions. For example, printing out the employees payroll slip may involve printing their name and address in a particular format, printing the employee number, social security number, etc., as well as printing historical information such as how much they have been paid in the current financial year, how much tax they have paid, all in addition to printing the actual amount they are being paid.

This process of breaking down higher-level functions into lower-level functions helps with:

- testing the system (functions can be tested in isolation),
- understanding the system as the organization of the functions can give meaning to the code, as well as allowing each function to be understood in isolation from the rest of the system,
- maintaining the system as only those functions that need to be changed may be affected by new or modified requirements,
- debugging the system as issues can be isolated to specific functions which can be examined independently of the rest of the application.

It is also known as a top-down refinement approach. The term top-down refinement (also known as stepwise design) highlights the idea that we are breaking down a system into the sub-systems that make it up.

It is common to represent the functions identified in the form of a tree illustrating the relationships between the higher-level functions and lower-level functions. This is illustrated below:



This diagram illustrates how a credit card approval process can be broken down into sub functions.

11.3.1 Functional Decomposition Terminology

The key terms used within functional decomposition are:

- **Function.** This is a task that is performed by a device, system or process.
- **Decomposition.** This is the process by which higher-level functions are broken down into lower-level functions where each function represents part of the functionality of the higher-level function.
- **Higher-Level Function.** This is a function that has one or more sub functions. This higher-level function depends on the functionality of the lower-level functions for its behaviour.
- **Sub Function.** This is a function that provides some element of the behaviour of a higher-level function. A sub function can also be broken down into its own sub functions in a hierarchical manner. In the above diagram the *Review for Compliance* function is both a sub function and has its own sub functions.

- **Basic Function.** A basic function is a function that has no smaller sub functions. The *Perform Credit Check* and *Review Against internal Rules* functions are both basic functions.

11.3.2 Functional Decomposition Process


At a very high level, functional decomposition consists of a series of steps such as those outlined below:

1. Find/identify the inputs and outputs of the system.
2. Define how the inputs are converted to the outputs. This will help identify the top most, high-level function(s).
3. Look at the current function(s) and try to break them down into a list of sub functions. Identify what each sub function should do, and what its inputs and outputs are.
4. Repeat Step 2 for each function identified until the functions identified can't or should not be decomposed further.
5. Draw a diagram of the function hierarchy you have created. Viewing the functions and their relationships is a very useful thing to do as it allows developers to visualize the system functionally. There are many computer-aided software engineering (CASE) tools that help with this but any drawing tool (such as Visio) can be used.
6. Examine the diagram for any repeating functions. That is, functions that do the same thing but appear in different places in the design. These are probably more generic functions that can be reused. Also examine the diagram to see if you can identify any missing functions.
7. Refine/design the interfaces between one function and another. That is, what data/information is passed to and from a function to a sub function as well as between functions.

11.3.3 Calculator Functional Decomposition Example

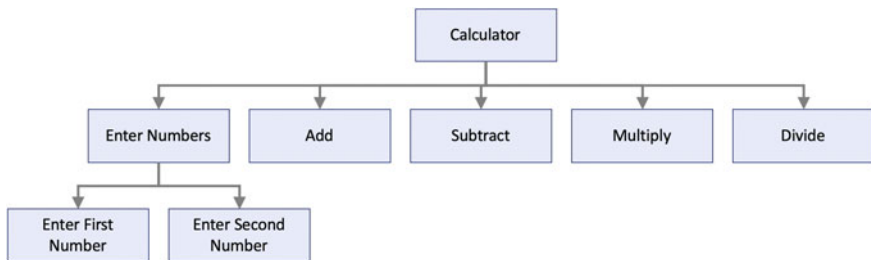
As an example of functional decomposition let us consider a simple calculator program.

We want this program to be able to perform a set of mathematical operations on two numbers, such as add, subtract, multiple and divide. We might therefore draw a functional decomposition diagram (or FDD) such as:




This illustrates that the calculator function can be decomposed into add, subtract, multiply and divide functions.

We might then identify the need to enter the two numbers to be operated on. This would result in one or more new functions being added:



We might also identify the need to determine which numerical operation should be performed based on user input. This function might sit above the numerical functions or alongside them. This is in fact an example of a design decision that the designer/ developer must make based on their understanding of the problem and how the software will be developed/tested/used, etc.

In the following version of the functional decomposition diagram the *operation selection* function is placed at the same level as the numerical operations as it provides information back to the top-level function.



11.4 Functional Flow

Although the decomposition hierarchy presented in the functional decomposition diagram illustrates the functions and their hierarchical relationships, it does not capture how the data flows between the functions or the order in which the functions are invoked.

There are several approaches to describing the interactions between the functions identified by functional decomposition including the use of pseudo-code, data flow diagrams and sequence diagrams:

- **Pseudo-code.** This is a form of structured English that is not tied to any particular programming language but which can be used to express simple ideas including conditional choices (similar to if statements) and iteration (as typified by looping constructs). However, as it is a pseudo-language, developers are not tied to a specific syntax and can include functions without the need to define those functions in detail.
- **Data Flow Diagrams.** These diagrams are used to chart the inputs, processes and outputs of the functions in a structured graphical form. A data flow diagram typically has no control flow; there are no decision rules and no loops. For each data flow, there must be at least one input and one end point. Each process (function) can be refined by another lower-level data flow diagram, which subdivides this process into sub-processes.
- **Sequence Diagrams.** These are used to represent interactions between different entities (or objects) in sequence. The functions invoked are represented as being called from one entity to another. Sequence diagrams are more typically used with object-oriented systems.

11.5 Data Flow Diagrams

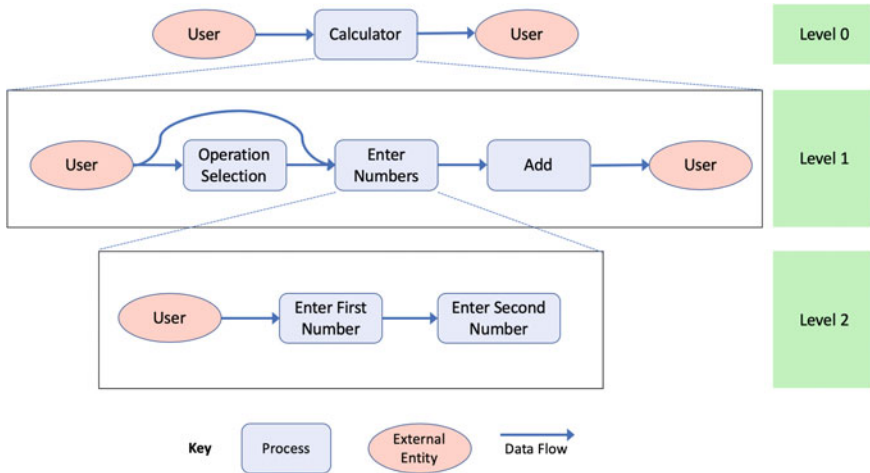
A data flow diagram consists of a set of inputs and outputs, processes (functions), flows, data stores (also known as warehouses) and terminators.

- **Process.** This is the process (or function or transformation) that converts inputs into outputs. The name of the process should be descriptive indicating what it does.
- **Data Flow.** The flow indicates the transfer of data/information from one element to another (that is a flow has a direction). The flow should have a name that suggests what information/data is being exchanged. Flows link processes, data stores and terminators.
- **Data Store/Warehouse.** A data store (which may be something such as a file, folder, database or other repository of data) is used to store data for later use. The name of the data store is a plural noun (e.g., employees). The flow from the data store usually represents the reading of the data stored in the data store, and the flow

to the waredata storehouse usually expresses data entry or updating (sometimes also deleting data).

- **Terminator.** The terminator represents an external entity (to the system) that communicates with the system. Examples of entities might be human users or other systems, etc.

An example of a data flow diagram is given below using the functions identified for the calculator:



In this diagram the hierarchy of DFDs is indicated by the levels which expand on how the function in the previous level is implemented by the functions at the next level. This DFD also only presents the data flow for the situation where the user selects the add function as the numerical operation to apply.

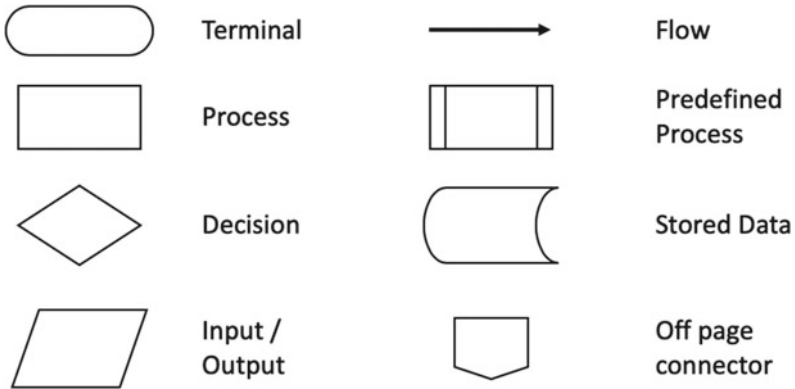
11.6 Flowcharts

A flowchart is a graphical representation of an algorithm, workflow or process for a given problem.

Flowcharts are used in the analysis, design and documentation of software systems. As with other forms of notation (such as DFDs) flowcharts help designers and developers to visualize the steps involved in a solution and thus aid in understanding the processes and algorithms involved.

The steps in the algorithm are represented as various types of boxes. The ordering of the steps is indicated by arrows between the boxes. The flow of control is represented by decision boxes.

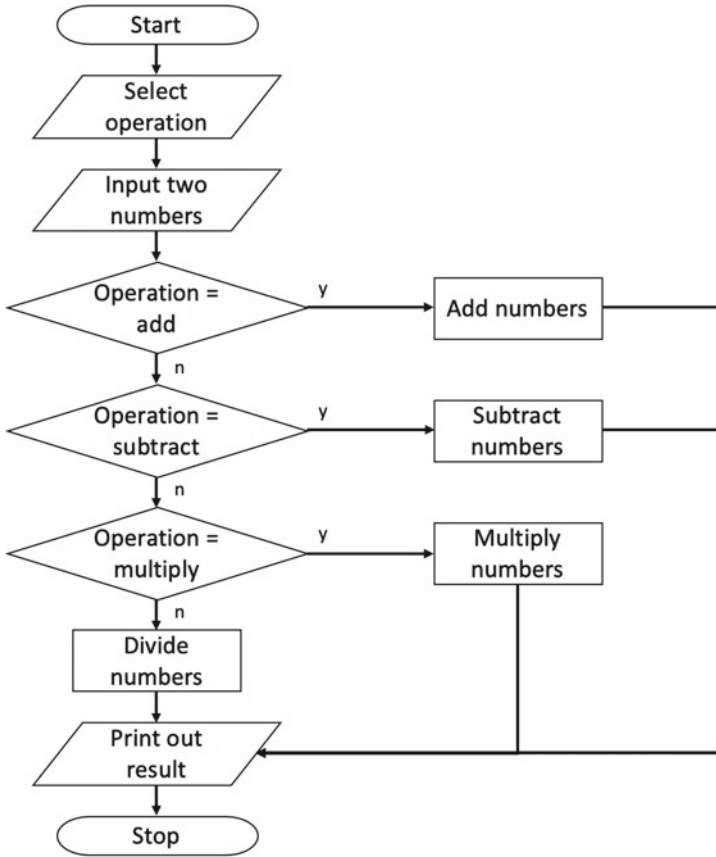
There are a number of common notations used with flowcharts, and most of those notations use the following symbols:



The meaning of these symbols is given below:

- **Terminal.** This symbol is used to indicate the start or end of a program or subprocess. They usually contain the words ‘Start’, ‘End’ or ‘Stop’ or a phrase indicating the start or end of some process such as ‘Starting Print Run’.
- **Process.** This symbol represents one or more operations (or programming statements/expressions) that in some way apply behaviour or change the state of the system. For example they may add two numbers together, run some form of calculation or change a Boolean flag, etc.
- **Decision.** This represents a decision point in the algorithm; that is it represents a decision point which will alter the flow of the program (typically between two different paths). The decision point is often represented as a question with a ‘yes’/ ‘no’ response, and this is indicated on the flowchart by the use of ‘yes’ (or ‘y’) and ‘no’ (or ‘n’) labels on the flowchart. In Python this decision point may be implemented using an if statement.
- **Input/Output.** This box indicates the input or output of data from the algorithm. This might represent obtaining input data from the user or printing out results to the user.
- **Flow.** These arrows are used to represent the algorithms order of execution of the boxes.
- **Predefined Process.** This represents a process that has been defined elsewhere.
- **Stored Data.** Indicates that data is stored in some form of persistent storage system.
- **Off page connector.** A labeled connector for use when the target is on another page (another flowchart).

Using the above symbols we can create a flowchart for our simple integer calculator program:



The above flowchart shows the basic operation of the calculator; the user selects which operation to perform, enters the two numbers and then depending upon the operation the selected operation is performed. The result is then printed out.

11.7 Data Dictionary

Another element commonly associated with structured analysis/design is the data dictionary. The data dictionary is a structured repository of data elements in the system. It stores the descriptions of all data flow diagram data elements. That is it records details and definitions of data flows, data stores, data stored in data stores and the processes. The format used for a data dictionary varies from method to method and project to project. It can be as simple as an Excel spreadsheet to an enterprise-wide software system such as Semanta (<https://www.semantacorp.com/data-dictionary>).

11.8 Online Resources

There are many online resources available that discuss functional decomposition, both from a theoretical and a practical Python-oriented point of view including:

- https://en.wikipedia.org/wiki/Structured_analysis Wikipedia Structured Analysis Page.
- https://en.wikipedia.org/wiki/Top-down_and_bottom-up_design Wikipedia page on Top Down and Bottom Up design.
- https://en.wikipedia.org/wiki/Edward_Yourdon#Yourdon_Structured_Method Wikipedia page on the Yourden method.
- https://en.wikipedia.org/wiki/Structured_systems_analysis_and_design_method Wikipedia page on SSADM.
- https://en.wikipedia.org/wiki/Functional_decomposition The Wikipedia page on Functional Decomposition.
- <https://docs.python.org/3/howto/functional.html> The Python standard documentation on functional decomposition.
- https://en.wikipedia.org/wiki/Data-flow_diagram Wikipedia page on Data Flow Diagrams (DFDs)
- https://en.wikipedia.org/wiki/Sequence_diagram Wikipedia page on Sequence Diagrams.
- https://en.wikipedia.org/wiki/Data_dictionary Wikipedia page on Data Dictionaries.
- <https://en.wikipedia.org/wiki/Flowchart> Wikipedia Flowchart page.

Chapter 12

Functions in Python



12.1 Introduction

As discussed in the last chapter, when you build an application of any size you will want to break it down into more manageable units, these units can then be worked on separately, tested and maintained separately. One way in which these units can be defined is as Python functions.

This chapter will introduce Python functions, how they are defined, how they can be referenced and executed. It also considers how parameters work in Python functions and how values can be returned from functions. It also introduces lambda or anonymous functions.

12.2 What are Functions?

In Python functions are groups of related statements that can be called together, that typically perform a specific task and which may or may not take a set of parameters or return a value.

Functions can be defined in one place and called or invoked in another. This helps to make code more modular and easier to understand.

It also means that the same function can be called multiple times or in multiple locations. This helps to ensure that although a piece of functionality is used in multiple places, it is only defined once and only needs to be maintained and tested in one location.

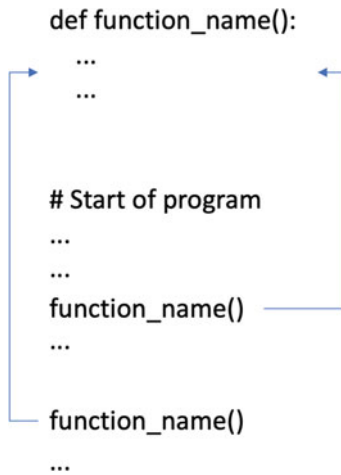
12.3 How Functions Work

We have said what they are and a little bit about why they might be good but not really how they work.

When a function is called (or invoked) the flow of control a program jumps from where the function was called to the point where the function was defined. The body of the function is then executed before control returns back to where it was called from.

As part of this process, all the values that were in place when the function was called are stored away (on something called the stack) so that if the function defines its own versions, they do not overwrite each other.

The invocation of a function is illustrated below:



Each time when the call is made to `function_name()` the program flow jumps to the body of the function and executes the statements there. Once the function finishes it returns to the point at which the function was called.

In the above this happens twice as the function is called at two separate points in the program.

12.4 Types of Functions

Technically speaking there are two types of functions in Python: *built-in* functions and *user-defined* functions.

Built-in functions are those provided by the language, and we have seen several of these already. For example, both `print()` and `input()` are built-in functions. We did not need to define them ourselves as they are provided by Python.

In contrast *user-defined* functions are those written by developers. We will be defining user-defined functions in the rest of this chapter, and it is likely that in many cases, most of the programs that you will write will include user-defined functions of one sort or another.

12.5 Defining Functions

The basic syntax of a function is illustrated below:

```
def function_name(parameter list):  
    """docstring"""  
    statement  
    statement(s)
```

This illustrates several things:

1. All (named) functions are defined using the *keyword* `def`; this indicates the start of a function definition. A keyword is a part of the syntax of the Python language and cannot be redefined and is not a function.
2. A function can have a name which uniquely identifies it; you can also have anonymous functions, but we will leave those until later in this chapter.
3. The naming conventions that we have been adopting for variables are also applied to functions, and they are all lowercase with the different elements of the function name separated by ‘_’.
4. A function can (optionally) have a list of parameters which allow data to be passed into the function. These are optional as not all functions need to be supplied with parameters.
5. A colon is used to mark the end of the *function header* and the start of the *function body*. The function header defines the signature of the function (what is called and the parameters it takes). The function body defines what the function does.
6. An optional documentation string (the *docstring*) can be provided that describes what the function does. We typically use the triple double quote string format as this allows the documentation string to go over multiple lines if required.
7. One or more Python statements make up the function body. These are indented relative to the function definition. All lines that are indented are part of the function until a line which is intended at the same level as the `def` line.
8. It is common to use 4 spaces (not a tab) to determine how much to indent the body of a function by.

12.5.1 An Example Function

The following is one of the simplest functions you can write; it takes no parameters and has only a single statement that prints out the message 'Hello World':

```
def print_msg():
    print('Hello World!')
```

This function is called `print_msg`, and when called (also known as invoked) it will run the body of the function which will print out the string, for example,

```
print_msg()
```

will generate the output

```
Hello World!
```

Be careful to include the round brackets () when you call the function. This is because if you just use the function's name then you are merely referring to the location in memory where the function is stored, and you are not invoking it.

We could modify the function to make it a little more general and reusable by providing a parameter. This parameter could be used to supply the message to be printed out, for example:

```
def print_my_msg(msg):
    print(msg)
```

Now the `print_my_msg` function takes a single parameter, and this parameter becomes a variable which is available within the body of the function. However, this parameter only exists within the body of the function; it is not available outside of the function.

This now means that we can call the `print_my_msg` function with a variety of different messages:

```
print_my_msg('Hello World')
print_my_msg('Good day')
print_my_msg('Welcome')
print_my_msg('Ola')
```

The output from calling this function with each of these strings being supplied as the parameter is:

```
Hello World
Good day
Welcome
Ola
```

12.6 Returning Values from Functions

It is very common to want to return a value from a function. In Python this can be done using the `return` statement. Whenever a `return` statement is encountered within a function then that function will terminate and return any values following the `return` keyword.

This means that if a value is provided, then it will be made available to any calling code.

For example, the following defines a simple function that squares whatever value has been passed to it:

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

When we call this function, it will multiply whatever it is given by itself and then return that value. The returned value can then be used at the point that the function was invoked, for example:

```
# Store result from square in a variable  
result = square(4)  
print(result)  
# Send the result from square immediately to another function  
print(square(5))  
# Use the result returned from square in a conditional expression  
if square(3) < 15:  
    print('Still less than 15')
```

When this code is run, we get:

```
16  
25  
Still less than 15
```

It is also possible to return multiple values from a function; for example, in this `swap` function the order in which the parameters are supplied is swapped when they are returned:

```
def swap(a, b):  
    return b, a
```

We can then assign the values returned to variables at the point when the function is called:

```
a = 2  
b = 3  
x, y = swap(a, b)  
print(x, ', ', y)
```

which produces

```
3 , 2
```


In actual fact the result returned from the swap function is what is called a *tuple* which is a simple way to grouping data together. This means that we could also have written:

```
z = swap(a, b)
print(z)
```

which would have printed the tuple out:

```
(3, 2)
```

We will look at tuples more when we consider collections of data.

12.7 Docstring

So far our example functions have not included any documentation strings (the *docstring*). This is because the docstring is *optional*, and the functions we have written have been very simple.

However, as functions become more complex and may have multiple parameters the documentation provided can become more important.

The *docstring* allows the function to provide some guidance on what is expected in terms of the data passed into the parameters, potentially what will happen if the data is incorrect, as well as what the purpose of the function is in the first place.

In the following example, the *docstring* is being used to explain the behaviour of the function and how the parameter is used.

```
def get_integer_input(message):
    """
    This function will display the message to the user
    and request that they input an integer.

    If the user enters something that is not a number
    then the input will be rejected
    and an error message will be displayed.

    The user will then be asked to try again. """
    value_as_string = input(message)
    while not value_as_string.isnumeric():
        print('The input must be an integer')
        value_as_string = input(message)

    return int(value_as_string)
```

When used, this method will guarantee that a valid integer will be returned to the calling code:

```
age = get_integer_input('Please input your age: ')
print('age is', age)
```

An example of what happens when this is run is given below:

```
Please input your age: John
The input must be an integer
Please input your age: 21
age is 21
```

The *docstring* can be read directly from the code but is also available to IDEs such as PyCharm so that they can provide information about the function. It is even available to the programmer via a very special property of the function called `__doc__` that is accessible via the name of the function using the dot notation:

```
print(get_integer_input.__doc__)
```

which generates

```
This function will display the message to the user
and request that they input an integer.
```

```
If the user enters something that is not a number
then the input will be rejected
and an error message will be displayed.
```

```
The user will then be asked to try again.
```

12.8 ReStructured Text

It is also possible to place *formatting* commands within a docstring that can be picked up by a tool such as PyCharm to improve the layout and information presented. There are several options available for this such as Google style docstrings and NumPy style docstrings (see <https://betterprogramming.pub/3-different-docstring-formats-for-python-d27be81e0d68>). However we are going to look at reStructured Text style docstrings as this is the style recommended by the Python organization itself.

ReStructured Text (aka ReST) is intended to be an easy-to-read markup syntax used to add additional meaning to the text within a docstring. Like many other simple markup languages it uses inline markup such as `*` or `#` to indicate emphasis or lists, etc. and as such is intended to be much simpler and easier to use than more complex markup languages such as HTML.

As a simple example, consider the following code:

```
def get_integer_input(message):
    """
    This function will display the message to the user
    and request that they input an integer.

    :param message: The message to print
    :returns: the integer entered by the user
    """
    value_as_string = input(message)
    while not value_as_string.isnumeric():
        print('The input must be an integer')
        value_as_string = input(message)
    return int(value_as_string)
```


This function has a docstring containing some reStructured Text formatting, namely the `:param` and `:returns` directives.

A directive is an item of *meta* information that will be used by the ReST parser to generate the formatted output.

- In this case the `:param` directive indicates that `message` is a parameter to the function and the information after the final `:` provides a description of that parameter.
- In turn the `:returns:` directive indicates that this function returns a value and the text after this described the meaning of the returned value.

PyCharm uses this information and generates reference documentation that can be viewed in the pop-up displayed when the programmer hovers over a call to the `get_integer_input` function. This is displayed below:

```
response = get_integer_input("Please input total: ")
print(f'The respons
```



```

/Users/Shared/workspaces/pycharm/cern-python-programming-
def get_integer_input(message: Any) -> int

This function will display the message to the user
and request that they input an integer.

Params: message – The message to print
Returns: the integer entered by the user

```

As you can see the `params` and `returns` directives are displayed in such a way as to make it easy to read and understand the information provided.

It is also possible to use ReST markup to emphasize text or to make the text bold or to indicate that a value is a literal:


- Emphasis is indicated using an asterisk `*`, for example: `*emphasis*`.
- Bold is indicated using a double asterisk `**`, for example: `**strong**`.
- Literal values, variables and small code samples can be indicated using two back quotes, for example: `“literal“`.

For example:

```
def get_input(prompt):
    """This function is primarily used to illustrate ReST, for
    example:
        This is used for *emphasis* while this is used for **bold**.
    Finally
        this is used for a literal “result“.
    """
    result = input(prompt)
    return result
```

This is rendered by PyCharm as shown below:

```
print(get_input("please enter a value: "))
```



The screenshot shows a tooltip from PyCharm. The tooltip text is: `/Users/Shared/workspaces/pycharm/cern-python-programming-samples`
`def get_input(prompt: Any) -> str`
 This function is primarily used to illustrate reST, for example: This is used for *emphasis* while this is used for **bold**. Finally this is used for a literal result.

Note that if asterisks or back quotes appear within the text and could be confused with inline markup delimiters, they should be escaped with a backslash, for example: `*`.

There are a few additional restrictions on this markup that you should be aware of, including:

- The markup cannot be nested; that is, you cannot nest a bold element within an emphasis element.
- The content may not start or end with whitespace: `* text*` is wrong, as is `*text *`.
- It must be separated from surrounding text by non-word characters. It is thus necessary to use a backslash escaped space to work around that for example: `thisis\ *one*\ word`.

It is also possible to create lists within ReST, bulleted lists are represented by placing an asterisk at the start of a line, and numbered lists are represented by a '#.'; it is also possible to explicitly number a list. Examples of these are shown below:

```
* This is a bulleted list.
* It has two items, the second
  item uses two lines.

1. This is a numbered list.
2. It has two items too.

#. This is a numbered list.
#. It has two items too.
```

For example, when used in a function docstring:

```
def get_more_input(prompt):
    """This function is used to illustrate lists, for example:

    * This is a bulleted list.
    * It has two items, the second item uses two lines.

    1. This is a numbered list.
    2. It has two items too.

    #. This is a numbered list.
    #. It has two items as well.

    """
    result = input(prompt)
    return result
```

Nested lists are possible, but be aware that they must be separated from the parent list items by blank lines, for example:

```
* this is
* a list

  * with a nested list
  * and some subitems

* and here the parent list continues
```

As an example, lists and sub lists are used below within a function docstring that also uses a couple of directives for the parameter and the return value.

```
def get_something(prompt):
    """
```

We can use lists:

```

* this is
* a list


    * with a nested list
    * and some subitems

* and here the parent list continues

:param prompt: the input prompt
:return: the value entered by the user
"""
result = input(prompt)
return result

```

This is illustrated below:



Section headers are created by underlining the section title with a punctuation character, at least as long as the text:

```

This is a heading
=====

```

Finally, if you want to embed some source code as an example within a docstring you can use the special marker `::`. This creates a literal block that will not be processed as formatted text but must be indented and with a gap following the `::`. For example:

```

def get_another_thing(prompt):
    """

```

To use this function see the code block::

```
result = get_another_thing("please input data: ")
print(result)
```

This is back to normal text
"""

```
result = input(prompt)
return result
```

PyCharm renders this as:



The screenshot shows a code editor window with the following content:

```
/Users/Shared/workspaces/pycharm/cern-python-programming
def get_another_thing(prompt: Any) -> str
```

To use this function see the code block:

```
result = get_another_thing("please input data: ")
print(result)
```

This is back to normal text

12.9 Function Parameters

Before we go any further it is worth clarifying some terminology associated with passing data into functions. This terminology relates to the parameters defined as part of the function header and the data passed into the function via these parameters:

- A *parameter* is a variable defined as part of the function header and is used to make data available within the function itself.
- An *argument* is the actual value or data passed into the function when it is called. The data will be held within the parameters.

Unfortunately many developers use these terms interchangeably, but it is worth being clear on the distinction.

12.9.1 Multiple Parameter Functions

So far the functions we have defined have only had zero or one parameter; however that was just a choice. We could easily have defined a function which defined two or more parameters. In these situations, the parameter list contains a list of parameter names separated by a comma.

For example,

```
def greeter(name, message):
    print('Welcome', name, '-', message)

greeter('Eloise', 'Hope you like Rugby')
```

Here the `greeter` function taken defines two parameters: `name` and `message`. These parameters (which are local to the function and cannot be seen outside of the function) are then used within the body of the function.

The output is

```
Welcome Eloise - Hope you like Rugby
```

You can have any number of parameters defined in a function (prior to Python 3.7 there was a limit of 256 parameters—although if you have this many then probably you have a major problem with the design of your function—however this limit has now gone).

12.9.2 Default Parameter Values

Once you have one or more parameters you may want to provide *default* values for some or all of those parameters, particularly for ones which might not be used in most cases.

This can be done very easily in Python; all that is required is that the default value must be declared in the function header along with the parameter name.

If a value is supplied for the parameter, then it will override the default. If no value is supplied when the function is called, then the default will be used.

For example, we can modify the `greeter()` function from the previous section to provide a default message such as 'Live Long and Prosper'.

```
def greeter(name, message = 'Live Long and Prosper'):
    print('Welcome', name, '-', message)

greeter('Eloise')
greeter('Eloise', 'Hope you like Rugby')
```

Now we can call the `greeter()` function with one or two arguments.

When we run this example, we will get:

```
Welcome Eloise - Live Long and Prosper
Welcome Eloise - Hope you like Rugby
```


As you can see from this in the first example (where only one argument was provided) the default message was used. However, in the second example where a message was provided, along with the name, then that message was used instead of the default.

Note we can use the terms *mandatory* and *optional* for the parameters in `greeter()`. In this case

- name is a mandatory field/parameter.
- message is an optional field/parameter (as it has a default value).

One subtle point to note is that any number of parameters in a function's parameter list can have a default value; however once one parameter has a default value all remaining parameters to the right of that parameter must also have default values. For example, we could *not* define the `greeter` function as

```
def greeter(message = 'Live Long and Prosper', name):
    print('Welcome', name, '-', message)
```

As this would generate an *error* indicating that `name` must have a default value as it comes after (to the right) of a parameter with a default value.

12.9.3 Named Arguments

So far we have relied on the position of a value to be used to determine which parameter that value is assigned to. In many cases this is the simplest and cleanest option.

However, if a function has several parameters, some of which have default values, it may become impossible to rely on using the position of a value to ensure it is given to the correct parameter (because we may want to use some of the default values instead).

For example, let us assume we have a function with four parameters

```
def greeter(name,
            title = 'Dr',
            prompt = 'Welcome',
            message = 'Live Long and Prosper'):
    print(prompt, title, name, '-', message)
```

This now raises the question how do we provide the name and the message arguments when we would like to have the default `title` and `prompt`?

The answer is to use *named* arguments (or *keyword* arguments). In this approach we provide the name of the parameter we want an argument/value to be assigned to; position is no longer relevant. For example:

```
greeter(message = 'We like Python', name = 'Lloyd')
```

In this example we are using the default values for `title` and `prompt` and have changed the order of `message` and `name`. This is completely legal and results in the following output:

```
Welcome Dr Lloyd - We like Python
```

We can actually mix *positional* and *named* arguments in Python, for example:

```
greeter('Lloyd', message = 'We like Python')
```

Here 'John' is bound to the `name` parameter as it is the first parameter, but 'We like Python' is bound to `message` parameter as it is a named argument.

However, you cannot place positional arguments after a named argument, so we cannot write:

```
greeter(name='John', 'We like Python')
```

As this will result in Python generating an error.

12.9.4 Arbitrary Arguments

In some cases, you do not know how many arguments will be supplied when a function is called. Python allows you to pass an arbitrary number of arguments into a function and then process those arguments inside the function.

To define a parameter list as being of arbitrary length, a parameter is marked with an asterisk (*). For example:

```
def greeter(*args):
    for name in args:
        print('Welcome', name)

greeter('John', 'Denise', 'Phoebe', 'Adam', 'Gryff', 'Jasmine')
```

This generates

```
Welcome John
Welcome Denise
Welcome Phoebe
Welcome Adam
Welcome Gryff
Welcome Jasmine
```

Note that this is another use of the `for` loop; but this time it is a sequence of strings rather than a sequence of integers that is being used.

12.9.5 *Positional and Keyword Arguments*

Some functions in Python are defined such that the arguments to the methods can either be provided using a variable number of positional or keyword arguments. Such functions have two arguments `*args` and `**kwargs` (for positional arguments and keyword arguments).

They are useful if you do not know exactly how many of either position or keyword arguments are going to be provided.

For example, the function `my_function` takes both a variable number of positional and keyword arguments:

```
def my_function(*args, **kwargs):
    for arg in args:
        print('arg:', arg)
    for key in kwargs.keys():
        print('key:', key, 'has value: ', kwargs[key])
```

This can be called with any number of arguments of either type:

```
my_function('John', 'Denise', daughter='Phoebe', son='Adam')
print('-' * 50)
my_function('Paul', 'Fiona', son_number_one='Andrew', son_
number_two='James', daughter='Joselyn')
```

which produces the output:

```
arg: John
arg: Denise
key: son has value: Adam
key: daughter has value: Phoebe
-----
arg: Paul
arg: Fiona
key: son_number_one has value: Andrew
key: son_number_two has value: James
key: daughter has value: Joselyn
```

Also note that the keywords used for the arguments are not fixed.

You can also define methods that only use one of the `*args` and `**kwargs` depending on your requirements (as we saw with the `greeter()` function above), for example:

```
def named(**kwargs):
    for key in kwargs.keys():
        print('arg:', key, 'has value:', kwargs[key])

named(a=1, b=2, c=3)
```

In this case, the `named` function only supports the provision of keyword arguments. Its output in the above case is:

```
arg: a has value: 1
arg: c has value: 3
arg: b has value: 2
```

In general, these facilities are most likely to be used by those creating libraries as they allow for great flexibility in how the library can be used.

12.10 Anonymous Functions

All the functions we have defined in this chapter have had a *name* that they can be referenced by, such as `greeter` or `get_integer_input`. This means that we can reference and reuse these functions as many times as we like.

However, in some cases we want to create a function and use it only once; giving it a name for this one time can *pollute* the namespace of the program (i.e., there are lots of names around) and also means that someone might call it when we don't expect them to.

Python therefore has another option when defining a function; it is possible to define an *anonymous* function. In Python an anonymous function is one that does not have a name and can only be used at the point that it is defined.

Anonymous functions are defined using the keyword `lambda`, and for this reason they are also known as lambda functions.

The syntax used to define an anonymous function is:

```
lambda arguments: expression
```

Anonymous functions can have any number of arguments but only one expression (that is a statement that returns a value) as their body. The expression is executed, and the value generated from it is returned as the result of the function.

As an example, let us define an anonymous function that will square a number:

```
double = lambda i : i * i
```

In this example the lambda definition indicates that there is one parameter to the anonymous function ('i') and that the body of the function is defined after the colon ':' which multiplies `i * i`; the value of which is returned as the result of the function. The whole anonymous function is then stored into a variable called `double`.

We can store the anonymous function into the variable as all functions are instances of the class function and can be referenced in this way (we just haven't done this so far).

To invoke the function, we can access the reference to the function held in the variable `double` and then use the round brackets to cause the function to be executed, passing in any values to be used for the parameters:

```
print(double(10))
```

When this is executed the value 100 is printed out.

Other examples of lambda/anonymous functions are given below (illustrating that an anonymous function can take any number of arguments):

```
func0 = lambda: print('no args')
func1 = lambda x: x * x
func2 = lambda x, y: x * y
func3 = lambda x, y, z: x + y + z
```

These can be used as shown below:

```
func0()
print(func1(4))
print(func2(3, 4))
print(func3(2, 3, 4))
```

The output from this code snippet is:

```
no args
16
12
9
```

12.11 Online Resources

See the Python Standard Library documentation for:

- <https://docs.python.org/3/library/functions.html> for a list of built-in functions in Python.
- https://www.w3schools.com/python/python_functions.asp the W3 Schools brief introduction to Python functions.
- https://www.w3schools.com/python/python_lambda.asp a short summary of lambda functions.
- <https://devguide.python.org/documentation/markup> for information on reStructured Text style docstrings.

12.12 Exercises

For this chapter the exercises involve the *number_guess_game* you created in the last chapter:

Take the number guess game and break it up into a number of functions. There is not necessarily a right or wrong way to do this; look for functions that are meaningful to you within the code, for example:

1. You could create a function to obtain input from the user.
2. You could create another function that will implement the main game playing loop.
3. You could also provide a function that will print out a message indicating if the player won or not.
4. You could create a function to print a welcome message when the game starts up.

Chapter 13

Scope and Lifetime of Variables



13.1 Introduction

We have already defined several variables in the examples we have been working within this book. In practice, most of these variables have been what are known as global variables. That is they are (potentially) accessible anywhere (or globally) in our programs.

In this chapter we will look again at local variable as defined within a function, at global variable and how they can be referenced within a function, and finally we will consider nonlocal variables.

13.2 Local Variables

In practice developers usually try to limit the number of global variables in their programs as global variables can be accessed anywhere and can be modified anywhere, and this can result in unexpected behaviours (and has been the cause of many, many bugs in all sorts of programs over the years).

However, not all variables are global. When we define a function, we can create variables which are scoped only to that function and are not accessible or visible outside of the function. These variables are referred to as local variables (as they are local to the function).

This is a great help in developing more modular code which has been proven to be easier to maintain and in fact develop and test.

In the following function local variable called `a_variable` has been created and initialized to hold the value 100.

```
def my_function():
    a_variable = 100
    print(a_variable)
my_function()
```

When this function is called `a_variable` will be initialized to 100 and will then be printed out to the console:

```
100
```

Thus when we ran the `my_function()` it successfully printed out the value 100 which was held in the local (to the function) variable `a_variable`.

However if we attempt to access `a_variable` outside the function, then it will not be defined and we will generate an error, for example:

```
my_function()
print(a_variable)
```

When we run this code, we get the number 100 printed out from the call `my_function()`. However, an error is then reported by Python:

```
100
Traceback (most recent call last):
  File "localvars.py", line 7, in <module>
    print(a_variable)
NameError: name 'a_variable' is not defined
```

This indicates that `a_variable` is *undefined* at the top level (which is the global scope). Thus, we can say that `a_variable` is not globally defined.

This is because `a_variable` only exists and only has meaning inside `my_function`; outside of that function it cannot be seen.

In fact, each time the function is called, `a_variable` comes back into existence as a *new* variable, so the value in `a_variable` is not even seen from one invocation of the function to another.

This raises the question what happens if a *global* variable called `a_variable` is defined? For example, if we have the following:

```
a_variable = 25
my_function()
print(a_variable)
```

Actually, this is fine and is supported by Python. There are now two versions of `a_variable` in the program: one of which is defined *globally* and one of which is defined within the *context* of the function.

Python does not get confused between these and treats them as completely separately. This is just like having two people called *John* in the same class in school. If they were only called John this might cause some confusion, but if they have different surnames then it is easy to distinguish between them via their full names such as *John Jones* and *John Smith*.

In this case we have *global* `a_variable` and `my_function` `a_variable`.
 Thus if we run the above code we get

```
100
25
```

The value 100 does not overwrite the value 25 as they are completely different variables.

13.3 The Global Keyword

But what happens if what you want is to reference the global variable within a function.

As long as Python does not think you have defined a local variable then all will be fine. For example

```
max = 100
def print_max():
    print(max)
print_max()
```

This prints out the value 100.

However, things go a bit astray if you try to modify the global variable inside the function. At this point Python thinks you are creating a local variable. If as part of the assignment you try to reference the current value of that (now) local variable you will get an error indicating that it currently does not have a value. For example, if we write:

```
def print_max():
    max = max + 1
    print(max)
print_max()
```

And then run this example, we will get

```
Traceback (most recent call last):
  File "localvars.py", line 17, in <module>
    print_max()
  File "localvars.py", line 14, in print_max
    max = max + 1
UnboundLocalError: local variable 'max' referenced before assignment
```

Indicating that we have referenced `max` before it was assigned a value—even though it was assigned a value globally before the function was called!

Why does it do this? To protect us from ourselves—Python is really saying ‘Do you really want to modify a global variable here?’. Instead it is treating `max` as a local variable, and as such it is being referenced before a value has been assigned to it.

To tell Python that we know what we are doing and that we want to reference the global variable at this point we need to use the keyword `global` with the name of the variable. For example:

```
max = 100
def print_max():
    global max
    max = max + 1
    print(max)

print_max()
print(max)
```

Now when we try to update the variable `max` inside the function `print_max()`, Python knows we mean the *global* version of the variable and uses that one. The result is that we now print out the value 101, and `max` is updated to 101 for everyone everywhere!

13.4 Nonlocal Variables

It is possible to define functions inside other functions, and this can be very useful when we are working with collections of data and operations such as `map()` (which map a function to all the elements of a collection of data).

However, local variables are local to a specific function; even functions defined within another function cannot modify the outer function’s local variable (as the inner function is a separate function). They can reference it, just as we could reference the global variable earlier; the issue is again modification.

The `global` keyword is no help here as the outer function’s variables are not global, they are local to a function.

For example, if we define a nested function (*inner*) inside the parent outer function (*outer*) and want the inner function to modify the local field we have a problem:

```
def outer():
    title = 'original title'

    def inner():
        title = 'another title'
```

```

        print('inner:', title)
    inner()
    print('outer:', title)
outer()

```

In this example both `outer()` and `inner()` functions modify the `title` variable. However, they are not the same `title` variable and as long as this is what we need then that is fine; both functions have their own version of a `title` local variable.

This can be seen in the output where the outer function maintains its own value for `title`:

```

inner: another title
outer: original title

```

However, if what we want is for the `inner()` function to modify the `outer()` function's `title` variable then we have a problem.

This problem can be solved using the `nonlocal` keyword. This indicates that a variable is not global but is also not local to the current function, and Python should look within the scope in which the function is defined to find a local variable with the same name:

If we now declare `title` as `nonlocal` in the `inner()` function, then it will use the `outer()` function's version of `title` (it will be shared between them), and thus when the `inner()` function changes the `title` it will change it for both functions:

```

def outer():
    title = 'original title'
    def inner():
        nonlocal title
        title = 'another title'
        print('inner:', title)
    inner()
    print('outer:', title)
outer()

```

The result of running this is

```

inner: another title
outer: another title

```

13.5 Hints

Points to note about the scope and lifetime of variables

1. The scope of a variable is the part of a program where the variable is known. Parameters and variables defined inside a function are not visible from outside. Hence, they have a local scope.
2. The lifetime of a variable is the period throughout which the variable exists in the memory of your Python program. The lifetime of variables inside a function is as long as the function executes. These local variables are destroyed as soon as the function returns or terminates. This means that the function does not store the values in a variable from one invocation to another.

13.6 Online Resources

See the Python Standard Library documentation for:

- <https://docs.python.org/3/faq/programming.html#what-are-the-rules-for-local-and-global-variables-in-python> which provides further information on the Python rules for local and global variables.

13.7 Exercise

Return to the number guess game—did you have to make any compromises with the variables to overcome the global variable issue? If so can you resolve them now with the use of the `global`?

Chapter 14

Implementing a Calculator Using Functions



14.1 Introduction

In this chapter we will step through the development of another Python program; this time the program will be to provide a simple calculator which can be used to add, subtract, multiple and divide numbers. The implementation of the calculator is based on the Function Decomposition performed earlier in the book in the Introduction to Structured Analysis chapter.

The calculator will be implemented using Python functions to help modularize the code.

14.2 What the Calculator Will Do

This will be a purely command-driven application that will allow the user to specify

- the operation to perform,
- the two numbers to use with that operation.

When the program starts up it can use a loop to keep processing operations until the user indicates that they wish to terminate the application.

We can also use an `if` statement to select the operation to perform, etc.

As such it will also build on several other features in Python that we have already been working with.

14.3 Getting Started

The first step will be to create a new Python file. If you are using the PyCharm IDE you can do it using the New > PythonFile menu option (look back at the Number Guess Game chapter if you can't remember how to do this). The file can be called anything you like, but *calculator* seems like a reasonable name.

In the newly created (and empty) `calculator.py` file type in a welcome print message such as:

```
print('Simple Calculator App')
```

Now run the `calculator.py` program (again if you don't remember how to do that look back at the Number Guess Game chapter).

You should see the message printed out in the Python console. This verifies that the file has been created properly and that you can run the Python code you will define in it.

14.4 The Calculator Operations

We are going to start by defining a set of functions that will implement the add, subtract, multiply and divide operations.

All of these functions take two numbers and return another number. We have also given each function a docstring to illustrate their use, although in practice the functions are so simple and self-describing that the docstring is probably redundant.

The functions are listed below; you can now add them to the `calculator.py` file:

```
def add(x, y):
    """ Adds two numbers """
    return x + y

def subtract(x, y):
    """ Subtracts two numbers """
    return x - y

def multiply(x, y):
    """ Multiplies two numbers """
    return x * y

def divide(x, y):
    """Divides two numbers"""
    return x / y
```

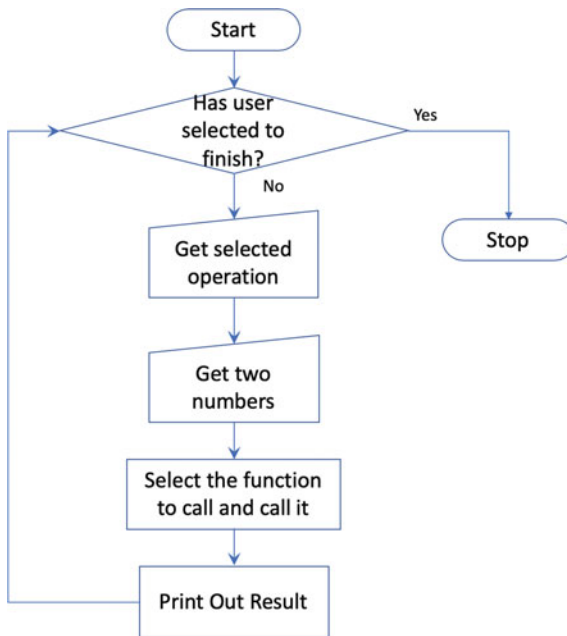
We now have the basic functions needed by the calculator.

14.5 Behaviour of the Calculator

We can now explore what the operation of the calculator program should be.

Essentially, we want to allow the user to be able to select the operation they want to perform and provide the two numbers to use with the operation and then for the program to call the appropriate function. The result of the operation should then be presented to the user.

We then want to ask the user whether they want to continue to use the calculator or to exit the program. This is illustrated below in flowchart form:



Based on this flowchart we can put in place the skeleton of the logic for the calculator's processing cycle.

We will need a while loop to determine whether the user has finished or not and a variable to hold the result and print it out.

The following code provides this skeleton:

```
finished = False
while not finished:
    result = 0
    # Get the operation from the user
    # Get the numbers from the user
    # Select the operation
    print('Result:', result)
    print('=====')
    # Determine if the user has finished

print('Bye')
```

If you try to run this right now then you will find that this code will loop forever as the user is not yet prompted to say if they wish to continue or not. However, it does provide the basic framework; we have

- a variable, `finished`, with a Boolean flag in it to indicate if the user has finished or not. This is referred to as a flag because it is a Boolean value and because it is being used to determine whether to terminate the main processing loop or not,
- a variable to hold the result of the operation and the two numbers,
- the `while` loop representing the main processing loop of the calculator.

14.6 Identifying Whether the User has Finished

We could address several of the remaining areas next; however, we will select the last step—that of determining if the user has finished or not. This will allow us to start to run the application so that we can test out the behaviour.

To do this we need to prompt the user to ask them if they want to continue using the calculator.

At one level this is very straight forward; we could ask the user to input ‘y’ or ‘n’ to indicate *yes* I have finished or *no* I want to keep going.

We could therefore use the input function as follows:

```
user_input = input('Do you want to finish (y/n):')
```

We could then check to see if they have entered a ‘y’ character and terminate the loop.

However, anytime if we take any input from something outside of our program (such as the user) we should verify the input. For example, what should the program do if the user enters ‘x’ or the number ‘1’? One option is to treat anything that is not a ‘y’ as being—I *want to keep going*. However, this is opening up our simple program to bad practices and (in a much larger system) to potential security issues and certainly to potential hacker attacks.

It is a much better idea to verify that the input is what is expected and to reject any input until it is either a ‘y’ or an ‘n’.

This means that the code is more complex than a single input statement; there is, for example, an implied loop here and as well as some idea of input validation.

This means that this is an ideal candidate for a function that will *encapsulate* this behaviour into a separate operation. We can then test this function which is always a good idea. It also means that where we use the function, we have a level of *abstraction*. That is we can name the function appropriately which will make it easier to see what we intended, instead of having a mass of code in one place.

We will call the function `check_if_user_has_finished`; this name makes it very clear what the purpose of the function is. It also means that when we use it in our main processing loop its role in that loop will be obvious.

The function is given below:

```
def check_if_user_has_finished():
    """
    Checks that the user wants to finish or not.
    Performs some verification of the input."""
    ok_to_finish = True
    user_input_accepted = False
    while not user_input_accepted:
        user_input = input('Do you want to finish (y/n): ')
        if user_input == 'y':
            user_input_accepted = True
        elif user_input == 'n':
            ok_to_finish = False
            user_input_accepted = True
        else:
            print('Response must be (y/n), please try again')
    return ok_to_finish
```

Notice the use of two variables that are local to the function:

- The first variable (`ok_to_finish`) holds the result of the function; whether it is OK to finish or not. It is given a default value of `True`; this follows the fail closed approach—which suggests that it is always better to fail by closing down an application or connection. In this case it means that if something goes wrong with the code (if it contains a software bug or logic error) the user will not keep looping forever.
- The second variable (`user_input_accepted`) is used to indicate whether the user has provided an acceptable input or not (i.e., have they entered 'y' or 'n') until they do the loop inside the function will repeat.

The loop itself is interesting as we are looping while the user input has not been accepted; note that we can (almost) read the while loop as plain English text. This is both a feature of Python (it is intended to be easily readable) and also of the use of a meaningful name for the variable itself.

Within the loop we obtain the input from the user; check to see if it is 'y' or 'n'. If it is either of these options, we set the `user_input_accepted` flag to `True`. Otherwise the code will print out a message indicating that the only acceptable input is a 'y' or 'n'.

Notice that we only set the `ok_to_finish` variable to `False` if the user inputs a 'n'; this is because the `ok_to_finish` variable by default has a value of `True` and thus there is no need to reassign `True` to it if the user select 'n'.

We can now add this function into our main processing loop in place of the last comment:

```
finished = False
while not finished:
    result = 0
    # Get the operation from the user
    # Get the numbers from the user
    # Select the operation
    print('Result:', result)
    print('=====')
    finished = check_if_user_has_finished()

print('Bye')
```

We can now run the application.

You may wonder why we would do this at this point as it does not yet do any calculations for us; the answer is that we can verify that the overall behaviour of the main loop works and that the `check_if_user_has_finished()` function operates correctly.

14.7 Selecting the Operation

Next let us implement the function used to obtain the operation to perform.

Again, we want to name this function in such a way as to help with the comprehensibility of our program. In this case we are asking the user to select which operation they want to perform, so let's call the function `get_operation_choice`.

This time we need to present a list of options to the user and then ask them to make a selection. Again, we want to write our function defensively, so that it makes sure the user only inputs a valid option; if they do not then the function prompts them for another input. This means our function will have a loop and some validation code.

There are four options available to the user: add, subtract, multiply and divide. We will therefore number them 1–4 and ask the user to select an option between 1 and 4.

There are several ways in which we can verify that they have entered a number in this range, including

- converting the string entered into a number and using numerical comparison (but then we need to check that they entered an integer),
- having multiple `if` and `elif` statements (but that seems a bit long winded),
- by checking that the entered character is one of a set of values (which is the approach we will use).

To check that a value is *in* a set of other value (that it is one of the values in the set) you can use the 'in' operator, for example:

```
user_selection in ('1', '2', '3', '4')
```

This will return True if (and only if) user_selection contains one of the strings '1', '2', '3' or '4'.

We can therefore use it in our function to verify that the user entered a valid input. The get_operation_choice function is shown below:

```
def get_operation_choice():
    input_ok = False
    while not input_ok:
        print('Menu Options are:')
        print('\t1. Add')
        print('\t2. Subtract')
        print('\t3. Multiply')
        print('\t4. Divide')
        print('-----')
        user_selection = input('Please make a selection: ')
        if user_selection in ('1', '2', '3', '4'):
            input_ok = True
        else:
            print('Invalid Input (must be 1 - 4)')
    print('-----')
    return user_selection
```

Work through this function and make sure you are comfortable with all its elements. The '\t' character is a special character denoting a Tab.

We can now update our main calculator loop with this function:

```
finished = False

while not finished:
    result = 0
    menu_choice = get_operation_choice()

    # Get the numbers from the user

    # Select the operation
    print('Result:', result)
    print('=====')
    finished = check_if_user_has_finished()

print('Bye')
```

14.8 Obtaining the Input Numbers

Next we need to obtain two numbers from the user to use with the selected operation.

In our introduction to Functions in Python chapter we looked at a function (the `get_integer_input()` function) that could be used to take input from the user and convert it (safely) into an integer; if the user entered a non-number then this function would prompt them to enter an actual number. We can reuse the function here.

However, we need to ask the user for two numbers; we will therefore create a function which uses the `get_integer_input()` function to prompt the user for two numbers and then return both numbers. Both functions are shown here:

```
def get_numbers_from_user():
    num1 = get_integer_input('Input the first number: ')
    num2 = get_integer_input('Input the second number: ')
    return num1, num2

def get_integer_input(message):
    value_as_string = input(message)
    while not value_as_string.isnumeric():
        print('The input must be an integer')
        value_as_string = input(message)
    return int(value_as_string)
```

Having one function call another function is very common, and indeed we have already been doing this; the `input()` function has been used several times; the only difference here is that we have written the `get_integer_input()` function ourselves.

When we can the `get_numbers_from_user()` function we can store the results returned into two variables: one for each result; for example:

```
n1, n2 = get_numbers_from_user()
```

We can now add this statement to the main calculator loop:

```
finished = False
while not finished:
    result = 0
    menu_choice = get_operation_choice()
    n1, n2 = get_numbers_from_user()
    # Select the operation
    print('Result:', result)
    print('=====')
    finished = check_if_user_has_finished()

print('Bye')
```

14.9 Determining the Operation to Execute

We are now almost there and can update our main calculation loop with some logic to determine the actual operation to invoke. To do this we will use an `if` statement with the optional `elif` parts. The `if` statement will be conditional on the operation selected and will then call the appropriate function (such as `add` and `subtract`) as shown here:

```

if menu_choice == '1':
    result = add(n1, n2)
elif menu_choice == '2':
    result = subtract(n1, n2)
elif menu_choice == '3':
    result = multiply(n1, n2)
elif menu_choice == '4':
    result = divide(n1, n2)

```

Each part of the `if` statement calls a different function, but they all store the value returned into the `result` variable.

We can now add this to the calculation loop to create our fully functional calculator loop:

```

finished = False
while not finished:
    result = 0
    menu_choice = get_operation_choice()
    n1, n2 = get_numbers_from_user()
    if menu_choice == '1':
        result = add(n1, n2)
    elif menu_choice == '2':
        result = subtract(n1, n2)
    elif menu_choice == '3':
        result = multiply(n1, n2)
    elif menu_choice == '4':
        result = divide(n1, n2)
    print('Result:', result)
    print('=====')
    finished = check_if_user_has_finished()

print('Bye')

```

14.10 Running the Calculator

If you now run the calculator you will be prompted as appropriate for input. You can try and break the calculator by entering characters when numbers are requested or values out of range for the operations, etc., and it should be resilient enough to handle these erroneous inputs, for example:

```
Simple Calculator App
Menu Options are:
  1. Add
  2. Subtract
  3. Multiply
  4. Divide
-----
Please make a selection: 5
Invalid Input (must be 1 - 4)
Menu Options are:
  1. Add
  2. Subtract
  3. Multiply
  4. Divide
-----
Please make a selection: 1
-----
Input the first number: 5
Input the second number: 4
Result: 9
=====
Do you want to finish (y/n): y
Bye
```

14.11 Exercises

For this chapter the exercises relate to extensions to the calculator:

1. Add an option to apply the modulus (%) operator to the two numbers input by the user. This will involve defining an appropriate function and adding this as an option to the menu. You will also need to extend the main calculator control loop to handle this option.
2. Add a power of (**) option to the calculator.
3. Modify the program to take floating point numbers instead of simple integers.
4. Allow the choice of division operator or integer division operator (this have both '/' and '//') available.

Chapter 15

Introduction to Functional Programming



15.1 Introduction

There has been much hype around Functional Programming in recent years. However, Functional Programming is not a new idea and indeed goes right back to the 1950s and the programming language LISP. However, many people are not clear as to what Functional Programming is and instead jump into code examples and never really understand some of the key ideas associated with Functional Programming such as Referential Transparency.

This chapter introduces Functional Programming (also known as FP) and the key concept of Referential Transparency (or RT).

One idea to be aware of is that Functional Programming is a software coding style or approach and is separate from the concept of a function in Python.

Python functions can be used to write functional programs but can also be used to write procedural style programs; so do not get too hung up on the syntax that might be used or the fact that Python has functions just yet. Instead explore the idea of defining a functional approach to your software design.

15.2 What is Functional Programming?

Wikipedia describes Functional Programming as:

... a programming paradigm, a style of building the structure and elements of computer programs, that treats computation as the evaluation of mathematical functions and avoids state and mutable data.

There are a number of points to note about this definition. The first is that it is focused on the *computational side* of computer programming. This might seem obvious but most of what we have looked at so far in Python would be considered procedural in nature.

Another thing to note is that the way in which the computations are represented emphasizes functions that generate results based purely on the data provided to them. That is these functions only rely on their inputs to generate a new output. They do not rely on any *side effects* and do not depend on the *current state* of the program. As an example of a side effect, if a function stored a running total in a global variable and another function used that total to perform some calculation, then the first function has a side effect of modifying a global variable and the second relies on some global state for its result.

Taking each of these in turn:

1. **Functional Programming aims to avoid side effects.** A function should be replaceable by taking the data it receives and in lining the result generated (this is referred to as Referential Transparency). This means that there should be no hidden side effects of the function. Hidden side effects make it harder to understand what a program is doing and thus make comprehension, development and maintenance harder. Pure functions have the following attributes:
 2. The only observable output is the return value.
 3. The only output dependency is the arguments.
 4. Arguments are fully determined before any output is generated.
5. **Functional Programming avoids concepts such as state.** Let us take these as separate issues. If some operation is dependent upon the (potentially hidden) state of the program or some element of a program, then its behaviour may differ depending upon that state. This may make it harder to comprehend, implement, test and debug. As all of these impact on the stability and probably reliability of a system, state-based operations may result in less reliable software being developed. As functions do not (should not) rely on any given state (only upon the data they are given) they should as a result be easier to understand, implement, test and debug.
6. **Functional Programming promotes immutable data.** Functional Programming also tends to avoid concepts such as mutable data. Mutable data is data that can change its state. By contrast *immutability* indicates that once created, data cannot be changed. In Python strings are immutable. Once you create a new string you cannot modify it. Any functions that apply to a string that might conceptually alter the contents of the string and result in a new string being generated. Many developers take this further by having a presumption of immutability in their code; that means that by default all data holding types are implemented as immutable. This ensures that functions cannot have hidden side effects and thus simplifies programming in general.
7. **Functional Programming promotes declarative programming** which means that programming is oriented around expressions that describe the solution rather than focus on the imperative approach of most procedural programming languages. Imperative languages emphasize aspects of how the solution is

derived. For example, an imperative approach to looping through some container and printing out each result in turn would look like this:

```
int sizeOfContainer = container.length
for (int i = 1 to sizeOfContainer) do
    element = container.get(i)
    print(element)
enddo
```

Whereas a Functional Programming approach would look like:

```
container.foreach(print)
```

Functional Programming has its roots in the lambda calculus, originally developed in the 1930s to explore computability. Many Functional Programming languages can thus be considered as elaborations on this lambda calculus. There have been numerous pure Functional Programming languages including Common Lisp, Clojure and Haskell. Python provides some support for writing in the functional style, particularly where the benefits of it are particularly strong (such as in processing various different types of data).

Indeed, when used judiciously, Functional Programming can be a huge benefit for, and an enhancement to, the toolkit available to developers.

To summarize then:

- **Imperative Programming** is what is currently perceived as traditional programming. That is, it is the style of programming used in languages such as C, C++, Java and C#. In these languages a programmer tells the computer what to do. It is thus oriented around control statements, looping constructs and assignments.
- **Functional Programming** aims to describe the solution, that is, *what* the program needs to do (rather than *how* it should be done).

15.3 Advantages to Functional Programming

There are a number of significant advantages to Functional Programming compared to imperative programming. These include:

1. **Less code.** Typically, a Functional Programming solution will require less code to write than an equivalent imperative solution. As there is less code to write, there is also less code to understand and to maintain. It is therefore possible that functional programs are not only more elegant to read but easier to update and maintain. This can also lead to enhanced programmer productivity as they spend less time writing reams of code as well as less time reading those reams of code.
2. **Lack of (hidden) side effects (Referential Transparency).** Programming without side effects is good as it makes it easier to reason about functions (that is a function is completely described by the data that goes in and the results that come back). This also means that it is safe to reuse these functions in different

situations (as they do not have unexpected side effects). It should also be easier to develop, test and maintain such functions.

3. **Recursion is a natural control structure.** Functional languages tend to emphasize recursion as a way of processing structures that would use some form of looping constructs in an imperative language. Although you can typically implement recursion in imperative languages, it is often easier to do in functional languages. It is also worth noting that although recursion is very expressive and a great way for a programmer to write a solution to a problem, it is not as efficient at runtime as iteration. However, any expression that can be written as a recursive routine can also be written using looping constructs. Functional Programming languages often incorporate *tail end recursive optimizations* to convert recursive routines into iterative ones at runtime. A util end recursive function is one in which the last thing a function does before it returns to call itself. This means that rather than actually invoking the function and having to set up the context for that function, it should be possible to reuse the current context and to treat it in an iterative manner as a loop around that routine. Thus the programmer benefits from the expressive recursive construct and the runtime benefits of an iterative solution using the same source code. This option is typically not available in imperative languages.
4. **Good for prototyping solutions.** Solutions can be created very quickly for algorithmic or behaviour problems in a functional language, thus allowing ideas and concepts to be explored in a rapid application development style.
5. **Modular functionality.** Functional Programming is modular in terms of functionality (where object-oriented languages are modular in the dimension of components). They are thus well suited to situations where it is natural to want to reuse or componentize the *behaviour* of a system.
6. **The avoidance of state-based behaviour.** As functions only rely on their inputs and outputs (and avoid accessing any other stored state) they exhibit a cleaner and simpler style of programming. This avoidance of state-based behaviour makes many difficult or challenging areas of programming simpler (such as those in concurrent applications).
7. **Additional control structures.** A strong emphasis is on additional control structures such as pattern matching, managing variable scope, tail recursion optimizations, etc.
8. **Concurrency and immutable data.** As Functional Programming systems advocate immutable data structures it is simpler to construct concurrent systems. This is because the data being exchanged and accessed is immutable. Therefore, multiple executing thread or processes cannot affect each other adversely. The Akka Actor model builds on this approach to provide a very clean model for multiple interacting concurrent systems.
9. **Partial evaluation.** Since functions do not have side effects, it also becomes practical to bind one or more parameters to a function at compile time and to reuse these functions with bound values as new functions that take fewer parameters.

15.4 Disadvantages of Functional Programming

If Functional Programming has all the advantages previously described, why isn't it the mainstream force that imperative programming languages are? The reality is that Functional Programming is not without its disadvantages, including:

- **Input–Output is harder in a purely functional language.** Input–Output flows naturally align with stream style processing, which does not neatly fit into the data in, results out, nature of functional systems.
- **Interactive applications are harder to develop.** Interactive applications are constructed via request response cycles initiated by a user action. Again, these do not naturally sit within the purely functional paradigm.
- **Continuously running programs** such as services or controllers may be more difficult to develop, as they are naturally based upon the idea of a continuous loop.
- **Functional Programming languages have tended to be less efficient on current hardware platforms.** This is partly because current hardware platforms are not designed with Functional Programming in mind and also because many of the systems previously available were focused on the academic community where out and out performance was not the primary focus. However, this has changed to a large extent with modern functional languages such as Scala and Haskell.
- **Not data oriented.** A pure functional language does not really align with the needs of the primarily data-oriented nature of many of today's systems. Many (most) commercial systems are oriented around the need to retrieve data from a database, manipulate it in some way and store that data back into a database. Such data can be naturally represented via objects in an object-oriented language.
- **Programmers are less familiar** with Functional Programming concepts and thus find it harder to pick up function-oriented languages.
- **Functional Programming idioms are often less intuitive** to (traditional) procedural programmers than imperative idioms which can make debugging and maintenance harder. Although with the use of a functional approach in many other languages now becoming more popular (including in Python) this trend is changing.
- Many Functional Programming languages have been viewed as **Ivory tower languages** that are only used by academics. This has been true of some older functional languages but is increasingly changing with the advent of languages such as Scala and with the facilities provided in more mainstream programming languages such as Python.

15.5 Referential Transparency

An important concept within the world of Functional Programming is that of Referential Transparency.

An operation is said to be *Referentially Transparent* if it can be replaced with its corresponding value, without changing the program's behaviour, for a given set of parameters.


For example, let us assume that we have defined the function `increment` as shown below.

```
def increment(num):
    return num + 1
```

If we use this simple example in an application to increment the value 5:

```
print(increment(5))
print(increment(5))
```

We can say that the function is Referentially Transparent (or RT) if it always returns the same result for the same value (i.e., that `increment(5)` always returns 6):



The screenshot shows a terminal window titled "referential_transparency_example". The terminal output is as follows:

```

/Users/Shared/workspaces/pycharm/beginnerspython3_2nd
6
6
Process finished with exit code 0

```

Any function that references a value which has been captured from its surrounding context and which can be modified cannot be guaranteed to be RT. This can have significant consequences for the maintainability of the resulting code. This can happen if for example the `increment` function did not add 1 to the parameter but added a global value. If this global value is changed then the function would suddenly start to return different values for the previously entered parameters. For example, the following code is no longer Referentially Transparent:

```
amount = 1
def increment(num):
    return num + amount

print(increment(5))
amount = 2
print(increment(5))
```

The output from this code is not 6 and 7—as the value of `amount` has changed between calls to the `increment()` function.

A closely related idea is that of *No Side Effects*. That is, a function should not have any side effects, it should base its operation purely on the values it receives, and its only impact should be the result returned. Any hidden side effects again makes software harder to maintain.

Of course, within most applications there is a significant need for side effects; for example, any logging of the actions performed by a program has a side effect of updating some logged information somewhere (typically in a file), and any database updates will have some side effect (i.e., that of updating the database). In addition some behaviour is inherently non-RT; for example, a function which returns the *current* time can never be Referentially Transparent.

However, for pure functions it is a useful consideration to follow.

15.6 Further Reading

There is a large amount of material on the web that can help you learn more about Functional Programming including:

- <https://codeburst.io/a-beginner-friendly-intro-to-functional-programming-4f69aa109569> intended as a friendly introduction to Functional Programming.
- <https://medium.freecodecamp.org/an-introduction-to-the-basic-principles-of-functional-programming-a2c2a15c84> which provides an introduction to the basic principles of Functional Programming.
- https://www.tutorialspoint.com/functional_programming which provides a good grounding in the basic concepts of Functional Programming.
- <https://docs.python.org/3/howto/functional.html> which is the Python standard library tutorial on Functional Programming.

Chapter 16

Higher-Order Functions



16.1 Introduction

In this chapter we will explore the concept of high-order functions. These are functions that take as a parameter, or return (or both), a function. To do this we will first look into how Python represents functions in memory and explore what actually happens when we execute a Python function.

16.2 Recap on Functions in Python

Let us first recap a few things regarding functions in Python:

Functions (mostly) have a name, and when invoked (or executed) the body of code associated with the function name is run.

There are some important ideas to remember when considering functions:

- Functions can be viewed as named blocks of code and are one of the main ways in which we can organize our programs in Python.
- Functions are defined using the keyword `def` and constitute a function header (the function name and the parameters, if any, defined for that function) and the function body (what is executed when the function is run).
- Functions are invoked or executed using their name followed by round brackets `()` with or without parameters depending on how the function has been defined.

This means we can write a function such as the following `get_msg` function:

```
def get_msg():  
    return 'Hello Python World!'
```

We can then call it by specifying its name and the round brackets:

```
message = get_msg()
print(message)
```

This of course prints out the string 'Hello Python World!' which is what you should expect by now.

16.3 Functions as Objects

A few chapters back we threw in something stating that if you forgot to include the round brackets then you were referencing the function itself rather than trying to execute it!

What exactly does that mean? Let's see what happens if we forgot to include the round brackets:

```
message = get_msg
print(message)
```

The output generated now is:

```
<function get_msg at 0x10ad961e0>
```

which might look very confusing at first sight.

What this is actually telling you is that you have referenced a function called `get_msg` that is located at a (hexidecimal) address in memory.

It is interesting to note that just as data has to be located in memory so does program code (so that it can be found and run), although typically data and code are located in separate areas of memory (as data tends to be short lived).

Another interesting thing to do is to find out what the *type* of `get_msg` is—hey it's a function—but what does that mean?

If we issue this statement and run it in Python:

```
print(type(get_msg))
```


Then we will get the following:

```
<class 'function'>
```

This means that it is of the class of things that are functions just as 1 is of the class of things called integers, 'John' is of the class of things called strings, and 42.6 is of the class of things called floating point numbers.

Taking this further it actually means that the *thing* being referenced by `get_msg` is a function object (an example or instance of the function class). This `get_msg` is really a type of variable that references (or points at) at the *function object* in memory which we can execute using the round brackets.

This is illustrated by the following diagram:



This means that when we run `get_msg()` what actually happens is we go to the `get_msg` variable and following the reference (or pointer) there to the function and then because we have the round brackets we run that function.

This has two implications:


1. We can pass the reference to a function around.
2. We can make `get_msg` reference (point) at a different function.

Let us look at the first of these implications. If we assign the reference represented by `get_msg` to something else, then in effect we have an alias for this function. This is because another variable now also references the same function. For example, if we write:

```
another_reference = get_msg
print(another_reference())
```

Then the result is that the string 'Hello Python World!' is again printed out.

What this has done is to copy the reference held in `get_msg` into `another_reference` (but it is a copy of that reference and that is the address of the function in memory). Thus, we now have in memory:



So just to emphasize this—we did not make a copy of the function; only its address in memory. Thus the same value is held in both `get_msg` and `another_reference`, and both these values are references to the same function object in memory.

What does this mean and why should we care? Well it means that we can pass references to functions around within our program which can be a very useful feature that we will look at later in this chapter.


Now let us go back to the second implication mentioned above; we can reassign another function to `get_msg`.

For example, let's say we wrote this next:


```
def get_some_other_msg():
    return 'Some other message!!!'

get_msg = get_some_other_msg
print(get_msg())
```

Now `get_msg` no longer references the original functions; it now references the new function defined by `get_some_other_msg`. It means that in memory we now have



which means the result of calling `print(get_msg())` will be that the string `'Some other message!!!'` is returned and printed out (rather than the `'Hello Python World!'`).

However, notice that we did not overwrite the original function; it is still being referenced by the `another_reference` variable and indeed can still be called via this variable. For example, the code:

```
print(get_msg())
print(another_reference())
```

now generates the output:

```
Some other message!!!
Hello Python World!
```

This illustrates some of the power but also the potential confusion that comes from how functions are represented and can be manipulated in Python.

16.4 Higher-Order Function Concepts

Given that we can assign a reference into a function to a variable, then this might imply that we can also use the same approach to pass a reference to a function as an argument to another function.

This means that one function can take another function as a parameter. Such functions are known as higher-order functions and are one of the key constructs in Functional Programming.

That is, a function that takes another function as a parameter is known as a *higher-order function*.

In fact, in Python, higher-order functions are functions that do at least one of the following (and may do both):

- Take one or more functions as a parameter.
- Return as a result a function.

All other functions in Python are *first-order* functions.

Many of the functions found in the Python libraries are higher-order functions. It is a common enough pattern that once you are aware of it you will recognize it in many different libraries.

16.4.1 Higher-Order Function Example

As an abstract example, consider the following higher-order function *apply*. This function (written in pseudo-code—not a real programming language) takes an integer and a function. Within the body of the function being defined, the function passed in as a parameter is applied to the integer parameter. The result of the function being defined is then returned:

```
def apply(x, function):  
    result = function(x)  
    return result
```

The function `apply` is a *higher-order* function because its behaviour (and its result) will depend on the behaviour defined by another function—the one passed into it.

We could also define a function that multiplies a number by 10.0, for example:

```
def mult(y):  
    return y * 10.0
```

Now we can use the function `mult` with the function `apply`, for example:

```
apply(5, mult)
```

This would return the value 50.0.

16.5 Python Higher-Order Functions

As we have already seen when we define a function it actually creates a function object that is referenced by the name of the function. For example, if we create the function `mult_by_two`:

```
def mult_by_two(num):
    return num * 2
```

Then this has created a function object referenced by the name `multi_by_two` that we can invoke (execute) using the round brackets `'()`'.

It is also a one parameter function that takes a number and returns a value which is twice that number.

Thus, a parameter that expects to be given a reference to a function that takes a number and returns a number can be given a reference to any function that meets this (implied) contract. This includes our `mult_by_two` function but also any of the following:

```
def mult_by_five(num):
    return num * 5

def square(num):
    return num * num

def add_one(num):
    return num + 1
```

All of the above could be used with the following higher-order function:

```
def apply(num, func):
    return func(num)
```

For example:

```
result = apply(10, mult_by_two)
print(result)
```

The output from this code is:

```
20
```

The following listing provides a complete set of the earlier sample functions and how they may be used with the `apply` function:

```
print(apply(10, mult_by_five))
print(apply(10, square))
print(apply(10, add_one))
print(apply(10, mult_by_two))
```

The output from this is:

```
50
100
11
20
```

16.5.1 Using Higher-Order Functions

Looking at the previous section you may be wondering why you would want to use a higher-order function or indeed why define one. After all, could you not have called one of the functions (`multi_by_five`, `square`, `add_one` or `mult_by_two`) directly by passing in the integer to be used? Yes, we could have, for example, we could have done:

```
square(10)
```

And this would have exactly the same effect as calling:

```
apply(10, square)
```

The first approach would seem to be both simpler and more efficient.

The key to why higher-order functions are so powerful is to consider what would happen if we know that some function should be applied to the value 10, but we do not yet know what it is. The actual function will be provided at some point in the future. Now we are creating a reusable piece of code that will be able to apply an appropriate function to the data we have when that function is known.

For example, let us assume that we want to calculate the amount of tax someone should pay based on their salary. However, we do not know how to calculate the tax that this person must pay as it is dependent on external factors. The `calculate_tax` function could take an appropriate function that performs that calculation and provides the appropriate tax value.

The following listing implements this approach. The function `calculate_tax` does not know how to calculate the actual tax to be paid, instead a function must be provided as a parameter to the `calculate_tax` function. The function passed in takes a number and returns the result of performing the calculation. It is used with the `salary` parameter also passed into the `calculate_tax` function.

```
import math

def simple_tax_calculator(amount):
    return math.ceil(amount * 0.3)

def calculate_tax(salary, func):
    return func(salary)

print(calculate_tax(45000.0, simple_tax_calculator))
```

The `simple_tax_calculator` function defines a function that takes a number and multiplies it by 0.3 and then uses the `math.ceil` function (imported from the `math` library/module) to round it up to a whole number. A call is then made to the `calculate_tax` function passing in the float `45000.0` as the `salary` and a *reference* to the `simple_tax_calculator` function. Finally, it prints out the tax calculated. The result of running this program is:

```
Run: higher_order_examples x
/Library/Frameworks/Python.framework/Versions/3.7/bin/python3.7
13500
Process finished with exit code 0
```

Thus, the function `calculate_tax` is a reusable function that can have different tax calculation strategies defined for it.

16.5.2 Functions Returning Functions

In Python as well as passing a function into another function, functions can be returned from a function. This can be used to select among a number of different options or to create a new function based on the parameters.

For example, the following code creates a function that can be used to check whether a number is even, odd or negative based on the string passed into it:

```
def make_checker(s):
    if s == 'even':
        return lambda n: n%2 == 0
    elif s == 'positive':
        return lambda n: n >= 0
    elif s == 'negative':
        return lambda n: n < 0
    else:
        raise ValueError('Unknown request')
```

Note the use of the `raise ValueError`; for the moment we will just say that this is a way of showing that there is a problem in the code which may occur if this function is called with an inappropriate parameter value for `s`.

This function is a *factory* for functions that can be created to perform specific operations. It is used below to create three functions that can be used to validate what type a number is:

```
f1 = make_checker('even')
f2 = make_checker('positive')
f3 = make_checker('negative')
print(f1(3))
print(f2(3))
print(f3(3))
```

Of course, it is not only anonymous functions that can be returned from a function; it is also possible to return a named function. This is done by returning just the name of the function (i.e., without the round brackets).

In the following example, a named function is defined within an outer function (although it could have been defined elsewhere in the code). It is then returned from the function:

```
def make_function():  
    def adder(x, y):  
        return x + y  
    return adder
```

We can then use this `make_function` to create the `adder` function and store it into another variable. We can now use this function in our code, for example:

```
f1 = make_function()  
print(f1(3, 2))  
print(f1(3, 3))  
print(f1(3, 1))
```

which produces the output

```
5  
6  
4
```

16.6 Online Resources

Further information on higher-order functions in Python can be found using the following online resources:

1. https://en.wikipedia.org/wiki/Higher-order_function Wikipedia page on higher-order functions.
2. <https://docs.python.org/3.1/library/functools.html> A module to support the creation and use of higher-order functions.
3. https://www.tutorialspoint.com/functional_programming/functional_programming_higher_order_functions.htm A tutorial on higher-order functions.

16.7 Exercises

The aim of this exercise is to explore higher-order functions.

You should write a higher-order function called `my_higher_order_function(i, func)`. This function takes a parameter and a second function to apply to the parameter.

Now you should write a sample program that uses the higher-order function you just created to perform. An example of the sort of thing you might implement is given below:

```
print(my_higher_order_function(2, double))
print(my_higher_order_function(2, triple))
print(my_higher_order_function(16, square_root))
print(my_higher_order_function(2, is_prime))
print(my_higher_order_function(4, is_prime))
print(my_higher_order_function('2', is_integer))
print(my_higher_order_function('A', is_integer))
print(my_higher_order_function('A', is_letter))
print(my_higher_order_function('1', is_letter))
```

If you are using the above code as your test application then you should write each of the supporting functions; each should take a single parameter.

Sample output from this code snippet is:

```
4
8
4.0
True
False
True
False
True
False
```

Note a simple way to find the square root of a number is to use the exponent (or power of) operator and multiply by 0.5.

Chapter 17

Curried Functions



17.1 Introduction

Currying is a technique which allows new functions to be created from existing functions by *binding* one or more parameters to a specific value. It is a major source of reuse of functions in Python which means that functionality can be written once, in one place and then reused in multiple other situations.

The name currying may seem obscure, but the technique is named after Haskell Curry (for whom the Haskell programming language is also named).

This chapter introduces the core ideas behind currying and explores how currying can be implemented in Python. The chapter also introduces the concept of closures and how they affect curried functions.

17.2 Currying Concepts

At an abstract level, consider having a function that takes two parameters. These two parameters x and y are used within the function body with the multiply operator in the form $x * y$. For example, we might have:

```
operation(x, y): return x * y
```

This function `operation()` might then be used as follows

```
total = operation(2, 5)
```

Which would result in 5 being multiplied by 2 to give 10. Or it could be used:

```
total = operation(10, 5)
```

Which would result in 5 being multiplied by 10 to give 50.

If we needed to double a number, we could thus reuse the `operation()` function many times, for example:

```
operation(2, 5)
operation(2, 10)
operation(2, 6)
operation(2, 151)
```

All of the above would double the second number. However, we have had to remember to provide the 2 so that the number can be doubled. However, the number 2 has not changed between any of the invocations of `operation()` function. What if we fixed the first parameter to always be 2, thus would mean that we could create a new function that apparently only takes one parameter (the number to double). For example, let us say we could write something like:

```
double = operation(2, *)
```

Such that we could now write:

```
double(5)
double(151)
```

In essence `double()` is an alias for `operation()`, but an alias that provides the value 2 for the first parameter and leaves the second parameter to be filled in by the future invocation of the `double` function.

17.3 Python and Curried Functions

A curried function in Python is a function where one or more of its parameters have been *applied or bound* to a value, resulting in the creation of a new function with one fewer parameters than the original. For example, let us create a function that multiplies two numbers together:

```
def multiply(a, b):
    return a * b
```

This is a general function that does exactly what it says; it multiplies any two numbers together. These numbers could be any two integers or floating-point numbers, etc.

We can thus invoke it in the normal manner:

```
print(multiply(2, 5))
```

The result of executing this statement is:

```
10
```

We could now define a new method that takes a function and a number and returns a new (anonymous) function that takes one *new* parameter and calls the function passed in with the number passed in and the new parameter:

```
def multby(func, num):
    return lambda y: func(num, y)
```

Look carefully at this function; it has used or *bound* the number passed into the `multby` function to the invocation of the function passed in, but it has also defined a new variable ‘y’ that will have to be provided when this new anonymous function is invoked. It then returns a reference to the anonymous function as the result of `multby`.

The `multby` function can now be used to bind the first parameter of the multiply function to anything we want. For example, we could bind it to 2 so that it will always double the second parameter and store the resulting function reference into a property `double`:

```
double = multby(multiply, 2)
```

We could also bind the value 3 to the first parameter of `multiply` to make a function that will triple any value:

```
triple = multby(multiply, 3)
```

Which means we can now write:

```
print(double(5))
print(triple(5))
```

which produces the output

```
10
15
```

You are not limited to just binding one parameter; you can bind any number of parameters in this way.

Curried functions are therefore very useful for creating new functions from existing functions.

17.4 Closures

One question that might well be on your mind now is what happens when a function references some data that is in scope where it is defined but is no longer available when it is evaluated? This question is answered by the implementation of a concept known as *closure*.

Within Computer Science (and programming languages in particular) a closure (or a lexical closure or function closure) is a function (or more strictly a reference to a function) together with a referencing environment. This referencing environment

records the context within which the function was originally defined and if necessary, a reference to each of the non-local variables used by that function. These non-local or free variables allow the function body to reference variables that are external to the function, but which are utilized by that function. This referencing environment is one of the distinguishing features between a functional language and a language that supports function pointers (such as C).

The general concept of a lexical closure was first developed during the 1960s but was first fully implemented in the language Scheme in the 1970s. It has since been used within many functional programming languages including LISP and Scala.

At the conceptual level, a closure allows a function to reference a variable available in the scope where the function was originally defined, but not available by default in the scope where it is executed.

For example, in the following simple programme, the variable `more` is defined outside the body of the function named `increase`. This is permissible as the variable is a *global* variable. Thus, the variable `more` is *within scope* at the point of definition.

```
more = 100

def increase(num):
    return num + more

print(increase(10))
more = 50
print(increase(10))
```

Within our program we then invoke the `increase` function by passing in the value 10. This is done twice with the variable `more` being reset to 50 between the two. The output from this program is shown below:

```
110
60
```

Note that it is the *current* value of `more` that is being used when the function executes and not the value of `more` present at the point that the function was defined. Hence the output is 110 and 60 that is $100 + 10$ and then $50 + 10$.

This might seem obvious as the variable `more` is still in scope within the same function as the invocations of the function referenced by `increase`.

However, consider the following example:

```
def increment(num):
    return num + 1

def reset_function():
    global increment
    addition = 50
    increment = lambda num: num + addition

print(increment(5))
reset_function()
print(increment(5))
```


In the above listing the function `increment` initially adds 1 to whatever value has been passed to it. Then in the program this function is called with the value 5 and the result returned by the function is printed. This will be the value 6.

However, after this a second function, `reset_function()` is invoked. This function has a variable that is *local* to the function. That is, *normally* it would only be available within the function `reset_function`. This variable is called `addition` and has the value 50.

The variable `addition` is, however, used within the function body of a new anonymous function definition. This function takes a number and adds the value of `addition` to that number and returns this as the result of the function. This new function is then assigned to the name `increment`. Note that to ensure we reference the *global* name `increment` we must use the keyword `global` (otherwise we will create a local variable that just happens to have the same name as the function).

Now, when the second invocation of `increment` occurs, the `reset_function()` method has terminated and *normally* the variable `addition` would no longer even be in existence. However, when this program runs the value 55 is printed out from the second invocation of `increment`. That is the function being referenced by the name `increment`, when it is called the second time, is the one defined within `reset_function()` and which uses the variable `addition`.

The actual output is shown below:



```
Run: closures x
/Library/Frameworks/Python.framework/Versions/3.7/bin/python3.7
6
55
Process finished with exit code 0
```

So, what has happened here? It should be noted that the value 50 was not copied into the second function body. Rather it is a concrete example of the use of a reference environment with the closure concept. Python ensures that the variable `addition` is available to the function, even if the invocation of the function is somewhere different to where it was defined by binding any free variables (those defined outside the scope of the function) and storing them so that they can be accessed by the function's context (in effect moving the variable from being a *local* variable to one which is available to the function anywhere; but only to the function).

17.5 Online Resources

Further information on currying see:

- <https://en.wikipedia.org/wiki/Currying> Wikipedia page on currying.
- <https://wiki.haskell.org/Currying> A page introducing currying (based on the Haskell language but still a useful reference).
- https://www.python-course.eu/currying_in_python.php A tutorial on currying in Python.

17.6 Exercises

This exercise is about creating a set of functions to perform currency conversions based on specified rates using currying to create those functions.

Write a function that will curry another function and a parameter in a similar manner to `multby` in this chapter—call this function `curry()`.

Now define a function that can be used to convert an amount into another amount based on a rate. The definition of this conversion function is very straight forward and just involves multiplying the number by the rate.

Now create a set of functions that can be used to convert a value in one currency into another currency based on a specific rate. We do not want to have to remember the rate, only the name of the function. For example:

```
dollars_to_sterling = curry(convert, 0.77)
print(dollars_to_sterling(5))

euro_to_sterling = curry(convert, 0.88)
print(euro_to_sterling(15))

sterling_to_dollars = curry(convert, 1.3)
print(sterling_to_dollars(7))

sterling_to_euro = curry(convert, 1.14)
print(sterling_to_euro(9))
```

If the above code is run the output would be:

```
3.85
13.2
9.1
10.26
```

Chapter 18

Introduction to Object Orientation



18.1 Introduction

This chapter introduces the core concepts in Object Orientation. It defines the terminology used and attempts to clarify issues associated with hierarchies. It also discusses some of the perceived strengths and weaknesses of the object-oriented approach. It then offers some guidance on the approach to take in learning about objects.

18.2 Classes

A class is one of the basic building blocks of Python. It is also a core concept in a style of programming known as Object-Oriented Programming (or OOP). OOP provides an approach to structuring programs/applications so that the data held, and the operations performed on that data are bundled together into classes and accessed via objects.

As an example, in an OOP style program, employees might be represented by a class `Employee` where each employee has an `id`, a `name`, a `department` and a `desk_number`, etc. They might also have operations associated with them such as `take_a_holiday()` or `get_paid()`.

In many cases classes are used to represent real-world entities (such as employees), but they do not need to, they can also represent more abstract concepts such as a transaction between one person and another (for example, an agreement to buy a meal).

Classes act as *templates* which are used to construct instances or examples of a class of things. Each example of the class `Person` has a `name`, an `age`, an `address`, etc., but they have their own values for their `name`, `age` and `address`. For example, to represent the people in a family we might create a class `Person` with the name `Paul`, the age `52` and the address set to `London`. We may also create another `Person` object

(instance) with the name Fiona, the age 48 and the address also of London and so on.

An instance or object is therefore an example of a class. All instances/objects of a class possess the same data variables but contain their own data. Each instance of a class responds to the same set of requests.

Classes allow programmers to specify the *structure* of an object (i.e., its attributes or fields, etc.) and the its behaviour separately from the objects themselves.

This is important, as it would be extremely time-consuming (as well as inefficient) for programmers to define each object individually. Instead, they define classes and create *instances* or *objects* of those classes.

They can then store related data together in a named concept which makes it much easier to structure and maintain code.

18.3 What Are Classes for?

We have already seen several types of data in Python such as integer, string, Boolean, etc. Each of these allowed us to hold a single item of data (such as the integer 42 or the string 'John' and the value `True`). However, how might we represent a Person, a Student or an Employee of a firm? One way we can do this is to use a class to represent them.

As indicated above, we might represent any type of (more complex) data item use a combination of attributes (or fields) and behaviours. These attributes will use existing data types, these might be integers, strings, Booleans, floating-point numbers or other classes.

For example, when defining the class `Person` we might give it:

- a field or attribute for the person's name,
- a field or attribute for their age,
- a field or attribute for their email,
- some behaviour to give them a birthday (which will increment their age),
- some behaviour to allow us to send them a message via their email,
- etc.

In Python classes are thus used:

- as a template to create instances (or objects) of that class,
- define instance methods or common behaviour for a class of objects,
- define attributes or fields to hold data within the objects,
- be sent messages.

Objects (or instances), on the other hand, can:

- be created from a class,
- hold their own values for instance variables,
- be sent messages,

- execute instance methods,
- may have many copies in the system (all with their own data).

18.3.1 *What Should a Class Do?*

A class should accomplish one specific purpose; it should capture only one idea. If more than one idea is encapsulated in a class, you may reduce the chances for reuse, as well as contravene the laws of encapsulation in object-oriented systems. For example, you may have merged two concepts together so that one can directly access the data of another. This is rarely desirable.

The following guidelines may help you to decide whether to split the class with which you are working. Look at the comment describing the class (if there is no class comment, this is a bad sign in itself). Consider the following points:

- Is the description of the class short and clear? If not, is this a reflection on the class? Consider how the comment can be broken down into a series of short clear comments. Base the new classes around those comments.
- If the comment is short and clear, do the class and instance variables make sense within the context of the comment? If they do not, then the class needs to be re-evaluated. It may be that the comment is inappropriate, or the class and instance variables inappropriate.
- Look at how and where the attributes of the class are used. Is their use in line with the class comment? If not, then you should take appropriate action.

18.3.2 *Class Terminology*

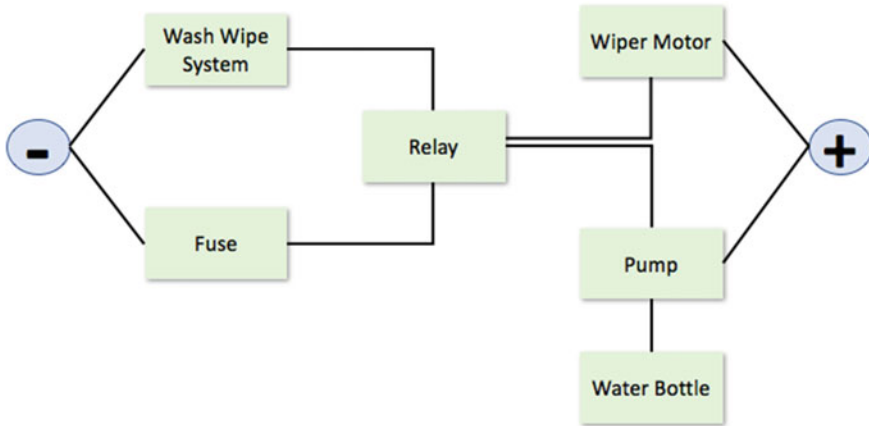
The following terms are used in Python (and other languages that support object orientation):

- *Class* A class defines a combination of data and behaviour that operates on that data. A class acts as a template when creating new instances.
- *Instance or object* An instance also known as an object is an example of a class. All instances of a class possess the same data fields/attributes but contain their own data values. Each instance of a class responds to the same set of requests.
- *Attribute/field/instance variable* The data held by an object is represented by its attributes (also sometimes known as a field or an instance variable). The “state” of an object at any particular moment relates to the current values held by its attributes.
- *Method* A method is a procedure defined within an object.
- *Message* A message is sent to an object requesting some operation to be performed or some attribute to be accessed. It is a request to the object to do something or return something. However, it is up to the object to determine how to execute that request. A message may be considered akin to a procedure call in other languages.

18.4 How is an OO System Constructed?

At this point you may be wondering how a system can be built from classes and objects instantiated from those classes? What would such an application look like? It is clear it is different to writing functions and free-standing application code that calls those functions?

Let's use a real-world (physical) system to explore what an OOP application might look like.



This system aims to provide a diagnosis tutor for the equipment illustrated above. Rather than use the wash-wipe system from a real car, students on a car mechanics diagnosis course use this software simulation. The software system mimics the actual system, so the behaviour of the pump depends on information provided by the relay and the water bottle.

The operation of the wash-wipe system is controlled by a switch which can be in one of five positions: off, intermittent, slow, fast and wash. Each of these settings places the system into a different state:

Switch setting	System state
Off	The system is inactive
Intermittent	The blades wipe the windscreen every few seconds
Slow	The wiper blades wipe the windscreen continuously
Fast	The wiper blades wipe the windscreen continuously and quickly
Wash	The pump draws water from the water bottle and sprays it onto the windscreen

For the pump and the wiper motor to work correctly, the relay must function correctly. In turn, the relay must be supplied with an electrical circuit. This electrical

circuit is negatively fused, and thus, the fuse must be intact for the circuit to be made. Cars are negatively switched as this reduces the chances of short circuits leading to unintentional switching of circuits.

18.4.1 *Where Do We Start?*

This is often a very difficult point for those new to object-oriented systems. That is, they have read the basics and understand simple diagrams, but do not know where to start. It is the old chestnut, “I understand the example but don’t know how to apply the concepts myself”. This is not unusual and, in the case of object orientation, is probably normal.

The answer to the question “where do I start?” may at first seem somewhat obscure; you should start with the data. Remember that objects are things that exchange messages with each other. The things possess the data that is held by the system and the messages request actions that relate to the data. Thus, an object-oriented system is fundamentally concerned with data items.

Before we go on to consider the object-oriented view of the system, let us stop and think for a while. Ask yourself where could I start; it might be that you think about starting “with some form of functional decomposition” (breaking the problem down in terms of the functions it provides) as this might well be the view the user has of the system. As a natural part of this exercise, you would identify the data required to support the desired functionality. Notice that the emphasis would be on the system functionality.

Let us take this further and consider the functions we might identify for the example presented above:

Function	Description
Wash	Pump water from the water bottle to the windscreen
Wipe	Move the windscreen wipers across the windscreen

We would then identify important system variables and sub-functions to support the above functions.

Now let us go back to the object-oriented view of the world. In this view, we place a great deal more emphasis on the data items involved and consider the operations associated with them (effectively, the reverse of the functional decomposition view). This means that we start by attempting to identify the primary data items in the system; next, we look to see what operations are applied to, or performed on, the data items; finally, we group the data items and operations together to form objects. In identifying the operations, we may well have to consider additional data items, which may be separate objects or attributes of the current object. Identifying them is mostly a matter of skill and experience.


The object-oriented design approach considers the operations far less important than the data and their relationships. In the next section we examine the objects that might exist in our simulation system.

18.4.2 Identifying the Objects

We look at the system as a whole and ask what indicates the state of the system. We might say that the position of the switch or the status of the pump is significant. This results in the data items shown below

Data item	States
switch setting	Is the switch set to off, intermittent, wipe, fast wipe or wash?
wiper motor	Is the motor working or not?
pump state	Is the pump working or not?
fuse condition	Has the fuse blown or not?
water bottle level	The current water level
relay status	Is current flowing or not?

The identification of the data items is considered in greater detail later. At this point, merely notice that we have not yet mentioned the functionality of the system or how it might fit together, we have only mentioned the significant items. As this is such a simple system, we can assume that each of these elements is an object and illustrate it in a simple object diagram:



Notice that we have named each object after the element associated with the data item (e.g., the element associated with the fuse condition is the fuse itself) and that the actual data (e.g., the condition of the fuse) is an instance variable of the object. This is a very common way of naming objects and their instance variables. We now have the basic objects required for our application.

18.4.3 Identifying the Services or Methods

At the moment, we have a set of objects each of which can hold some data. For example, the water bottle can hold an integer indicating the current water level. Although object-oriented systems are structured around the data, we still need some procedural content to change the state of an object or to make the system achieve some goal. Therefore, we also need to consider the operations a user of each object might require. Notice that the emphasis here is on the user of the object and what they require of the object, rather than what operations are performed on the data.

Let us start with the switch object. The switch state can take a number of values. As we do not want other objects to have direct access to this variable, we must identify the services that the switch should offer. As a user of a switch we want to be able to move it between its various settings. As these settings are essentially an enumerated type, we can have the concept of incrementing or decrementing the switch position. A switch must therefore provide a `move_up` and a `move_down` interface. Exactly how this is done depends on the programming language; for now, we concentrate on specifying the required facilities.

If we examine each object in our system and identify the required services, we may end up with the following table:

Object	Service	Description
switch	<code>move_up</code>	Increment switch value
	<code>move_down</code>	Decrement switch value
	State?	Return a value indicating the current switch state
fuse	<code>working?</code>	Indicate if the fuse has blown or not
wiper motor	<code>working?</code>	Indicate whether the wipers are working or not
relay	<code>working?</code>	Indicate whether the relay is active or not
pump	<code>working?</code>	Indicate whether the pump is active or not
water bottle	<code>fill</code>	Fill the water bottle with water
	<code>extract</code>	Remove some water from the water bottle
	<code>empty</code>	Empty the water bottle

We generated this table by examining each of the objects in isolation to identify the services that might reasonably be required. We may well identify further services when we attempt to put it all together.

Each of these services should relate to a method within the object. For example, the `moveUp` and `moveDown` services should relate to methods that change the state instance variable within the object. Using a generic pseudo-code, the `move_up` method, within the switch object, might contain the following code:


```
def move_up(self) :
    if self.state == "off" then
        self.tate = "wash"
    else if self.state == "wash" then
```

```
self.state = "wipe"
```

This method changes the value of the state variable in switch. The new value of the instance variable depends on its previous value. You can define `moveDown` in a similar manner. Notice that the reference to the instance variable illustrates that it is global to the object. The `moveUp` method requires no parameters. In object-oriented systems, it is common for few parameters to be passed between methods (particularly of the same object), as it is the object that holds the data.

18.4.4 Refining the Objects

If we look back to able table, we can see that fuse, wiper motor, relay and pump all possess a service called `working?`. This is a hint that these objects may have something in common. Each of them presents the same interface to the outside world. If we then consider their attributes, they all possess a common instance variable. At this point, it is too early to say whether fuse, wiper motor, relay and pump are all instances of the same class of object (e.g., a `Component` class) or whether they are all instances of classes which inherit from some common superclass (see below). However, this is something we must bear in mind later.



18.4.5 *Bringing It All Together*

So far, we have identified the primary objects in our system and the basic set of services they should present. These services were based solely on the data the objects hold. We must now consider how to make our system function. To do this, we need to consider how it might be used. The system is part of a very simple diagnosis tutor; a student uses the system to learn about the effects of various faults on the operation of a real wiper system, without the need for expensive electronics. We therefore wish to allow a user of the system to carry out the following operations:

- change the state of a component device
- ask the motor what its new state is

The `move_up` and `move_down` operations on the switch change the switch's state. Similar operations can be provided for the fuse, the water bottle and the relay. For the fuse and the relay, we might provide a `change_state` interface using the following algorithm:

```
define change_state(self)
  if self.state == "working" then
    self.tate = "notWorking"
  else
    self.state = "working"
```

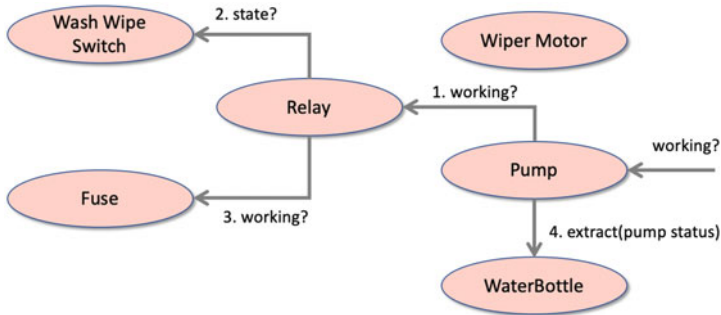
Discovering the state of the motor is more complicated. We have encountered a situation where one object's state (the value of its instance variable) is dependent on information provided by other objects. If we write down procedurally how the value of other objects affect the status of the pump, we might get the following pseudo-code:

```
if fuse is working then
  if switch is not off then
    if relay is working then
      pump status = "working"
```

This algorithm says that the pump status depends on the relay status, the switch setting and the fuse status. This is the sort of algorithm you might expect to find in your application. It links the sub-functions together and processes the data.

In an object-oriented system, well-mannered objects pass messages to one another. How then do we achieve the same effect as the above algorithm? The answer is that we must get the objects to pass messages requesting the appropriate information. One way to do that is to define a method in the pump object that gets the required information from the other objects and determines the motor's state. However, this requires the pump to have links to all the other objects so that it can send them messages. This is a little contrived and loses the structure of the underlying system. It also loses any modularity in the system. That is, if we want to add new components then we have to change the pump object, even if the new components only affect the switch. This approach also indicates that the developer is thinking too procedurally and not really in terms of objects.

In an object-oriented view of the system, the pump object only needs to know the state of the relay. It should therefore request this information from the relay. In turn, the relay must request information from the switches and the fuse.



The above illustrates the chain of messages initiated by the pump object:

1. pump sends a `working?` message to the relay,
2. relay sends a `state?` message to the switch, the switch replies to the relay,
3. relay sends a second `working?` message to the fuse:
4. The fuse replies to the relay
5. the relay replies to the motor
6. If the pump is working, then the pump object sends the final message to the water bottle
7. pump sends a message `extract` to the water bottle

In step four, a parameter is passed with the message because, unlike the previous messages that merely requested state information, this message requests a change in state. The parameter indicates the rate at which the pump draws water from the water bottle.

The water bottle should not record the value of the pump's status as it does not own this value. If it needs the motor's status in the future, it should request it from the pump rather than using the (potentially obsolete) value passed to it previously.

In the above figure we assumed that the pump provided the service `working?` which allows the process to start. For completeness, the pseudo-code of the `working?` method for the pump object is:

```

def working?(self)
    self.status = relay.working()
    if self.status == "working" then
        water_bottle.extract(self.status)
  
```


This method is a lot simpler than the procedural program presented earlier. At no point do we change the value of any variables that are not part of the pump, although they may have been changed as a result of the messages being sent. Also, it only shows us the part of the story that is directly relevant to the pump. This means that it can

be much more difficult to deduce the operation of an object-oriented system merely by reading the source code. Some Python environments (such as the PyCharm IDE) alleviate this problem, to some extent, through the use of sophisticated browsers.


18.5 Where Is the Structure in an OO Program?

People new to object orientation may be confused because they have lost one of the key elements that they use to help them understand and structure a software system: the main program body. This is because the objects and the interactions between them are the cornerstone of the system. In many ways, the following figure shows the object-oriented equivalent of a main program. This also highlights an important feature of most object-oriented approaches: graphical illustrations. Many aspects of object technology, for example object structure, class inheritance and message chains, are most easily explained graphically.

Let us now consider the structure of our object-oriented system. It is dictated by the messages that are sent between objects. That is, an object must possess a reference to another object in order to send it a message. The resulting system structure is illustrated below.



In Python, this structure is achieved by making instance variables reference the appropriate objects. This is the structure which exists between the instances in the system and does not relate to the classes, which act as templates for the instances.

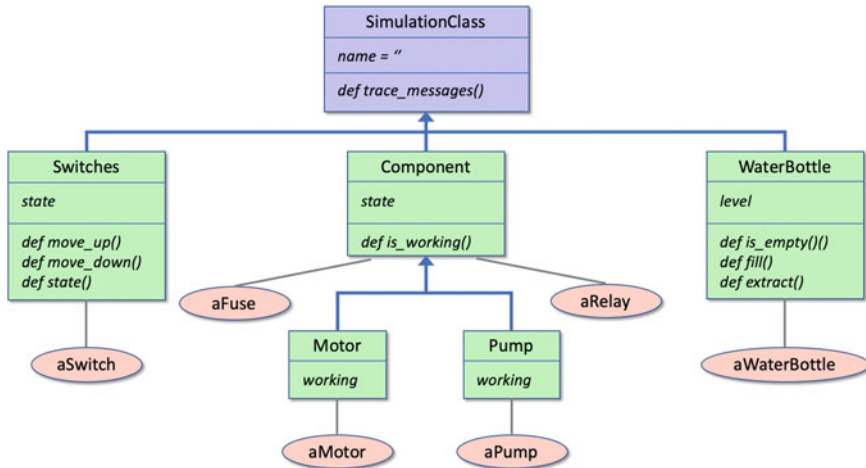


We now consider the classes that create the instances. We could assume that each object is an instance of an equivalent class (see above (a)). However, as has already been noted, some of the classes bear a very strong resemblance. In particular, the fuse, the relay, the motor and the pump share a number of common features. Table following table compares the features (instance variables and services) of these objects.

	fuse	relay	motor	pump
instance variable	state	state	state	state
services	working?	working?	working?	working?

From this table, the objects differ only in name. This suggests that they are all instances of a common class such as Component. This class would possess an additional instance variable, to simplify object identification.

If they are all instances of a common class, they must all behave in exactly the same way. However, we want the pump to start the analysis process when it receives the message working? so it must possess a different definition of working? from fuse and relay. In other ways it is very similar to fuse and relay, so they can be instances of a class (say Component) and pump and motor can be instances of classes that inherit from Component (but redefine working?). This is illustrated in the previous figure (c). The full class diagram is presented in the Figure below.



18.6 Further Readings

If you want to explore some of the ideas presented in this chapter in more detail here are some online references:

- https://en.wikipedia.org/wiki/Object-oriented_programming This is the Wikipedia entry for Object-Oriented Programming and thus provides a quick reference to much of the terminology and history of the subject and acts as a jumping off point for other references.
- <https://dev.to/charanrajgolla/beginners-guide---object-oriented-programming> which provides a light-hearted look at the four concepts within object orientations namely abstraction, inheritance, polymorphism and encapsulation.
- https://www.tutorialspoint.com/python/python_classes_objects.htm A Tutorialspoint course on Object-Oriented Programming and Python.

Chapter 19

Python Classes



19.1 Introduction

In Python everything is an object and as such is an example of a type or class of things. For example, integers are an example of the `int` class, real numbers are examples of the `float` class, etc. In this chapter we will look at what a class is in Python, how it is defined, how instances of a class can be created, how attributes (data) and methods (behaviour) can be defined for a class and outline how memory is managed within Python with respect to instance creation and deletion.

19.2 Python and Classes

As mentioned above everything in Python is an object! This is illustrated below for a number of different types within Python:

```
print(type(4))
print(type(5.6))
print(type(True))
print(type('Natalia'))
print(type([1, 2, 3, 4]))
```

This prints out a list of classes that define what it is to be an `int`, or a `float` or a `bool` etc. in Python:

```

<class 'int'>
<class 'float'>
<class 'bool'>
<class 'str'>
<class 'list'>

```

That is the number 4 is an example or instance of the class `int` representing how integers work in Python. In turn the number 5.6 is an instance of the class `float` (representing floating-point numbers). The value `True` is an instance of the class `bool` and the string `'Natalia'` is an instance of the class `str` (representing for strings in Python). Finally `[1, 2, 3, 4]` is an instance of a type of collection called a list (which is discussed later in this book).

However, you are not just restricted to the built-in types (aka classes); it is also possible to define user defined types (classes). These can be used to create your own data structures, your own data types, your own applications, etc.

Three remainder of this chapter considers the constructs in Python used to create user defined classes.

19.3 Class Definitions

In Python, a class definition has the following format

```

class nameOfClass (SuperClass) :
    __init__
    attributes
    methods

```

Although you should note that you can mix the order of the definition of attributes, and methods as required within a single class.

The following code is an example of a class definition:

```

class Person:
    def __init__(self, name, age):
        self.name = name
        self.age = age

```

Although this is not a hard and fast rule, it is common to define a class in a file named after that class. For example, the above code would be stored in a file called `Person.py`; this makes it easier to find the code associated with a class. This is shown below using the PyCharm IDE:

```

1 class Person:
2     """ An example class to hold a persons name and age"""
3
4     John Hunt
5     def __init__(self, name, age):
6         print('init called with', name, age)
7         self.name = name
8         self.age = age

```

The `Person` class possesses two *attributes* (or instance variables) called `name` and `age`.

There is also a special method defined called `__init__`. This is an initializer (also known as a constructor) for the class. It indicates what data must be supplied when an instance of the `Person` class is created and how that data is stored internally.

In this case a `name` and an `age` must be supplied when an instance of the `Person` class is created.

The values supplied will then be stored within an instance of the class (represented by the *special* variable `self`) in instance variables/attributes `self.name` and `self.age`. Note that the parameters to the `__init__` method are *local* variables and will disappear when the method terminates, but `self.name` and `self.age` are instance variables and will exist for as long as the object is available.

Let us look for a moment at the special variable `self`. This is the first parameter passed into any method. However, when a method is called we do not pass a value for this parameter ourselves; Python does. It is used to represent the object within which the method is executing. This provides the context within which the method runs and allows the method to access the data held by the object. Thus `self` is the object itself.

You may also be wondering about that term *method*. A method is the name given to behaviour that is linked directly to the `Person` class; it is not a free-standing function rather it is part of the definition of the class `Person`.

Historically, it comes from the language Smalltalk; this language was first used to simulate a production plant and a method represented some behaviour that could be used to simulate a change in the production line; it therefore represented a *method* for making a change.


19.4 Creating Examples of the Class Person

New instances/objects (examples) of the class `Person` can be created by using the name of the class and passing in the values to be used for the parameters of the initialization method (with the exception of the first parameter *self* which is provided automatically by Python).

For example, the following creates two instances of the class `Person`:

```
p1 = Person('John', 36)
p2 = Person('Phoebe', 21)
```

The variable `p1` holds a reference to the *instance* or *object* of the class `Person` whose attributes hold the values 'John' (for the name attribute) and 36 (for the age attribute). In turn the variable `p2` references an instance of the class `Person` whose name and age attributes hold the values 'Phoebe' and 21. Thus in memory we have:



The two variables reference separate *instances* or examples of the class `Person`. They therefore respond to the same set of methods/operations and have the same set of attributes (such as `name` and `age`); however, they have their own values for those attributes (such as 'John' and 'Phoebe').

Each instance also has its own unique identifier—that shows that even if the attribute values happen to be the same between two objects (for example, there happen to be two people called John who are both 36); they are still separate instances of the given class. This identifier can be accessed using the `id()` function, for example:

```
print('id(p1):', id(p1))
print('id(p2):', id(p2))
```

When this code is run `p1` and `p2` will generate different identifiers, for example:

```
id(p1): 4547191808
id(p2): 4547191864
```


Note that actual number generated may vary from that above but should still be unique (within your program).

19.5 Be Careful with Assignment

Given that in the above example, `p1` and `p2` reference different instances of the class `Person`; what happens when `p1` or `p2` are assigned to another variable? That is, what happens in this case:

```
p1 = Person('John', 36)
px = p1
```

What does `px` reference? Actually, it makes a complete copy of the value held by `p1`; however, `p1` does not hold the instance of the class `Person`; it holds the address of the object. It thus copies the address held in `p1` into the variable `px`. This means that both `p1` and `px` now reference (point at) the same instance in memory; we there have this:



This may not be obvious when you print `p1` and `px`:

```
print(p1)
print(px)
```

As this could just imply that the object has been copied:

```
John is 36
John is 36
```

However, if we print the unique identifier for what is referenced by `p1` and `px` then it becomes clear that it is the same instance of class `Person`:

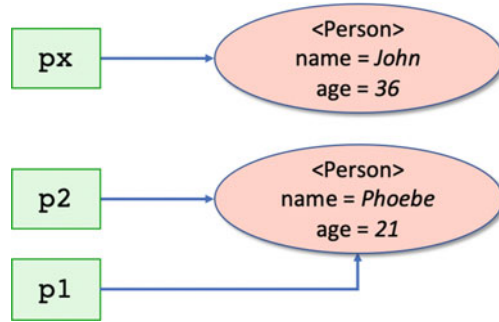
```
print('id(p1):', id(p1))
print('id(px):', id(px))
```

which prints out

```
id(p1): 4326491864
id(px): 4326491864
```

As can be seen the unique identifier is the same.

Of course, if `p1` is subsequently assigned a different object (for example, if we ran `p1 = p2`) then this would have no effect on the value held in `px`; indeed, we would now have:



19.6 Printing Out Objects

If we now use the `print()` function to print the objects held by `p1` and `p2`, we will get what might at first glance appear to be a slightly odd result:

```
print(p1)
print(p2)
```

The output generated is

```
<__main__.Person object at 0x10f08a400>
<__main__.Person object at 0x10f08a438>
```

What this is showing is the name of the class (in this case `Person`) and a hexadecimal number indicates where it is held in memory. Neither of which is particularly useful and certainly doesn't help us in knowing what information `p1` and `p2` are holding.

19.6.1 Accessing Object Attributes

We can access the attributes held by `p1` and `p2` using what is known as the *dot* notation. This notation allows us to follow the variable holding the object with a dot (`.`) and the attribute we are interested in access. For example, to access the name of a person object we can use `p1.name` or for their age we can use `p1.age`:


```
print(p1.name, 'is', p1.age)
print(p2.name, 'is', p2.age)
```

The result of this is that we output

```
John is 36
Phoebe is 21
```

Which is rather more meaningful.

In fact, we can also update the attributes of an object directly, for example we can write:

```
p1.name = 'Bob'
p1.age = 54
```

If we now run

```
print(p1.name, 'is', p1.age)
```

then we will get

```
Bob is 54
```

We will see in a later chapter (Python Properties) that we can restrict access to these attributes by making them into *properties*.

19.6.2 Defining a Default String Representation

In the previous section we printed out information from the instances of class `Person` by accessing the attributes `name` and `age`.

However, we now needed to know the internal structure of the class `Person` to print out its details. That is, we need to know that there are attributes called `name` and `age` available on this class.

It would be much more convenient if the object itself knew how to convert its self into a string to be printed out!

In fact we can make the class `Person` do this by defining a method that can be used to convert an object into a string for printing purposes.

This method is the `__str__` method. The method is expected to return a string which can be used to represent appropriate information about a class.

The signature of the method is

```
def __str__(self)
```

Methods that start with a double underbar ('`__`') are by convention considered special in Python and we will see several of these methods later on in the book. For the moment we will focus only on the `__str__()` method.

We can add this method to our class `Person` and see how that affects the output generated when using the `print()` function.

We will return a string from the `__str__` method that provides and the name and age of the person:

```
class Person:
    def __init__(self, name, age):
        self.name = name
        self.age = age

    def __str__(self):
        return self.name + ' is ' + str(self.age)
```

Note that in the `__str__` method we access the name and age attributes using the `self` parameter passed into the method by Python. Also note that it is necessary to convert the age number attribute into a string. This is because the `+` operator will do string *concatenation* unless one of the operands (one of the sides of the `+`) is a number; in which case it will try and do arithmetic addition which of course will not work if the other operand is a string!

If we now try to print out `p1` and `p2`:

```
print(p1)
print(p2)
```

The output generated is:

```
John is 36
Phoebe is 21
```

Which is much more useful.

19.6.3 *Defining a Default Storage Representation*

Somewhat confusingly Python provides two ways in which an instance of a class can be converted into a String. One approach is the `__str__()` method described in the previous section. The other is `__repr__()` method. An example of a `__repr__()` string representation function is given below:

```
class Person:
    def __init__(self, name, age):
        self.name = name
        self.age = age

    def __repr__(self):
        return f'Person(name = {self.name}, age = {self.age})'
```

If we now use the following code to print an instance of the class `Person`:

```
p1 = Person('John', 36)
print(p1)
```

The output that is generated is

```
Person(name=John, age=36)
```

So what's the difference between `__str__()` and `__repr__()`? The answer is that:

- `__str__()` is intended to be used to create a string version of an object that can be used for printing and logging purposes.
- `__repr__()` is intended to generate a string representation of an object that can be used to recreate the object - hence the format used above that would allow you to use `Person(name = John, age = 36)` to recreate an instance of the person John.

Which is used by Python depends on what you have define. If you print an object out then Python first tries to use the `__str__()` method, if that does not exist it uses the `__repr__()` method. You can therefore choose to just use the `__repr__()` method. However, if instances of a class are contained within for example a list, then when the list is printed out it will always use the `__repr__()` method. For this reason it is quite common to find that a class has both a `__str__()` and a `__repr__()` method using slightly different formats. Alternative, one of the two methods may call the other one.

For example we might define the class `Person` as:

```
class Person:
    """ An example class to hold a persons name and age """

    def __init__(self, name, age):
        print('init called with', name, age)
        self.name = name
        self.age = age

    def __str__(self):
        return self.name + ' is ' + str(self.age) .

    def __repr__(self):
        return f'Person(name = {self.name}, age = {self.age})'
```

19.7 Providing a Class Comment

It is common to provide a comment for a class defining what that class does, its purpose and any important points to note about the class.

This can be done by providing a *docstring* for the class just after the class declaration header; you can use the triple quotes string (" " " ... " " ") to create multiple line *docstrings*, for example:

```
class Person:
    """ An example class to hold a
        persons name and age"""

    def __init__(self, name, age):
        self.name = name
        self.age = age

    def __str__(self):
        return self.name + ' is ' + str(self.age)
```

The *docstring* is accessible through the `__doc__` attribute of the class. The intention is to make information available to users of the class, even at runtime. It can also be used by IDEs to provide information on a class.

Note that a class comment can also contain reStructured Text formatting commands and directives. This is a common practice as it provides improved information layout over a basic comment.

19.8 Adding a Birthday Method

Let us now add some behaviour to the class `Person`. In the following example, we define a *method* called `birthday()` that takes no parameters and increments the `age` attribute by 1:

```
class Person:
    """ An example class to hold a persons name and age"""

    def __init__(self, name, age):
        self.name = name
        self.age = age

    def __str__(self):
        return self.name + ' is ' + str(self.age)

    def birthday(self):
        print ('Happy birthday you were', self.age)
        self.age += 1
        print ('You are now', self.age)
```

Note that again the first parameter passed into the method `birthday` is `self`. This represents the instance (the example of the class `Person`) that this method will be used with.

If we now create an instance of the class `Person` and call `birthday()` on it, the age will be incremented by 1, for example:

```
p3 = Person('Adam', 19)
print(p3)
p3.birthday()
print(p3)
```

When we run this code, we get

```
Adam is 19
Happy birthday you were 19
You are now 20
Adam is 20
```

As you can see Adam is initially 19; but after his birthday he is now 20.

19.9 Defining Instance Methods

The `birthday()` method presented above is an example of what is known as an instance method; that is, it is tied to an instance of the class. In that case the method did not take any parameters, nor did it return any parameters; however, instance methods can do both.

For example, let us assume that the `Person` class will also be used to calculate how much someone should be paid. Let us also assume that the rate is £7.50 if you are under 21 but that there is a supplement of 2.50 if you are 21 or over.

We could define an instance method that will take as input the number of hours worked and return the amount someone should be paid:

```
class Person:
    """ An example class to hold a persons name and age """
    #...
    def calculate_pay(self, hours_worked):
        rate_of_pay = 7.50
        if self.age >= 21:
            rate_of_pay += 2.50
        return hours_worked * rate_of_pay
```

We can invoke this method again using the *dot* notation, for example:

```
pay = p2.calculate_pay(40)
print('Pay', p2.name, pay)
pay = p3.calculate_pay(40)
print('Pay', p3.name, pay)
```

Running this shows that Phoebe (who is 21) will be paid £400 while Adam who is only 19 will be paid only £300:

```
Pay Phoebe 400.0
Pay Adam 300.0
```

Another example of an instance method defined on the class `Person` is the `is_teenager()` method. This method does not take a parameter, but it does return a Boolean value depending upon the `age` attribute:

```
class Person:
    """ An example class to hold a persons name and age """
    #...
    def is_teenager(self):
        return self.age < 20
```

Note that the *implicitly* provided parameter `'self'` is still provided even when a method does not take a parameter.

19.10 Person Class Recap

Let us bring together the concepts that we have looked at so far in the final version of the class `Person`.

```
class Person:
    """ An example class to hold a persons name and age """

    def __init__(self, name, age):
        self.name = name
        self.age = age

    def __str__(self):
        return self.name + ' is ' + str(self.age)

    def birthday(self):
        print ('Happy birthday you were', self.age)
        self.age += 1
        print ('You are now', self.age)

    def calculate_pay(self, hours_worked):
        rate_of_pay = 7.50
        if self.age >= 21:
            rate_of_pay += 2.50
        return hours_worked * rate_of_pay

    def is_teenager(self):
        return self.age < 20
```

This class exhibits several features we have seen already and expands a few others:

- The class has a two-parameter initializer that takes a `String` and an `Integer`.
- It defines two attributes held by each of the instances of the class; `name` and `age`.
- It defines a `__str__` method so that the details of the `Person` object can be easily printed.

- It defines three methods `birthday()`, `calculate_pay()` and `is_teenager()`.
- The method `birthday()` does not return anything (i.e., it does not return a value) and is comprised of three statements, two print statements and an assignment.
- `is_teenager()` returns a Boolean value (i.e., one that returns True or False).

An example application using this class is given below:

```
p1 = Person('John', 36)
print(p1)
print(p1.name, 'is', p1.age)
print('p1.is_teenager', p1.is_teenager())
p1.birthday()
print(p1)
p1.age = 18
print(p1)
```

This application creates an instance of the `Person` class using the values 'John' and 36. It then prints out `p1` using `print` (which will automatically call the `__str__()` method on the instances passed to it). It then accesses the values of `name` and `age` properties and prints these. Following this it calls the `is_teenager()` method and prints the result returned. It then calls the `birthday()` method. Finally, it assigns a new value to the `age` property. The output from this application is given below:

```
John is 36
John is 36
p1.is_teenager False
Happy birthday you were 36
You are now 37
John is 37
John is 18
```

19.11 The Del Keyword

Having at one point created an object of some type (whether that is a `bool`, an `int` or a user defined type such as `Person`) it may later be necessary to delete that object. This can be done using the keyword `del`. This keyword is used to delete objects which allows the memory they are using to be reclaimed and used by other parts of your program.

For example, we can write

```
p1 = Person('John', 36)
print(p1)
del p1.
```

After the `del` statement the object held by `p1` will no longer be available and any attempt to reference it will generate an error.

You do not need to use `del` as setting `p1` above to the `None` value (representing nothingness) will have the same effect. In addition, if the above code was defined within a function or a method then `p1` will cease to exist once the function or method terminates and this will again have the same effect as deleting the object and freeing up the memory.

19.12 Automatic Memory Management

The creation and deletion of objects (and their associated memory) is managed by the Python Memory Manager. Indeed, the provision of a memory manager (also known as automatic memory management) is one of Python's advantages when compared to languages such as C and C++. It is not uncommon to hear C++ programmers complaining about spending many hours attempting to track down a particularly awkward bug only to find it was a problem associated with memory allocation or pointer manipulation. Similarly, a regular problem for C++ developers is that of memory creep, which occurs when memory is allocated but is not freed up. The application either uses all available memory or runs out of space and produces a run time error.

Most of the problems associated with memory allocation in languages such as C++ occur because programmers must not only concentrate on the (often complex) application logic but also on memory management. They must ensure that they allocate only the memory which is required and deallocate it when it is no longer required. This may sound simple, but it is no mean feat in a large complex application.

An interesting question to ask is "why do programmers have to manage memory allocation?". There are few programmers today who would expect to have to manage the registers being used by their programs, although 30 or 40 years ago the situation was very different. One answer to the memory management question, often cited by those who like to manage their own memory, is that "it is more efficient, you have more control, it is faster and leads to more compact code". Of course, if you wish to take these comments to their extreme, then we should all be programming in assembler. This would enable us all to produce faster, more efficient and more compact code than that produced by Python or languages such as Java.

The point about high level languages, however, is that they are more productive, introduce fewer errors, are more expressive and are efficient enough (given modern computers and compiler technology). The memory management issue is somewhat similar. If the system automatically handles the allocation and deallocation of memory, then the programmer can concentrate on the application logic.

This makes the programmer more productive, removes problems due to poor memory management and, when implemented efficiently, can still provide acceptable performance.

Python therefore provides automatic memory management. Essentially, it allocates a portion of memory as and when required. When memory is short, it looks for areas which are no longer referenced. These areas of memory are then freed up (deallocated) so that they can be reallocated. This process is often referred to as *Garbage Collection*.

19.13 Intrinsic Attributes

Every class (and every object) in Python has a set of *intrinsic* attributes set up by the Python runtime system. Some of these intrinsic attributes are given below for classes and objects.

Classes have the following intrinsic attributes:

- `__name__` the name of the class
- `__module__` the module (or library) from which it was loaded
- `__bases__` a collection of its base classes (see inheritance later in this book)
- `__dict__` a dictionary (a set of key-value pairs) containing all the attributes (including methods)
- `__doc__` the documentation string

For objects:

- `__class__` the name of the class of the object
- `__dict__` a dictionary containing all the object's attributes

Notice that these intrinsic attributes all start and end with a double underbar—this indicates their special status within Python.

An example of printing these attributes out for the class `Person` and a instance of the class are shown below:

```
print('Class attributes')
print(Person.__name__)
print(Person.__module__)
print(Person.__doc__)
print(Person.__dict__)
print('Object attributes')
print(p1.__class__)
print(p1.__dict__)
```

The output from this is:

```
Class attributes
Person
__main__
An example class to hold a persons name and age
```

```
{'__module__': '__main__', '__doc__': ' An example class to hold a
persons name and age', 'instance_count': 4, 'increment_instance_
count': <classmethod object at 0x105955588>, 'static_function':
<staticmethod object at 0x1059555c0>, '__init__': <function
Person.__init__ at 0x10595d268>, '__str__': <function Person._
__str__ at 0x10595d2f0>, 'birthday': <function Person.birthday
at 0x10595d378>, 'calculate_pay': <function Person.calculate_
pay at 0x10595d400>, 'is_teenager': <function Person.is_teenager
at 0x10595d488>, '__dict__': <attribute '__dict__' of 'Person'
objects>, '__weakref__': <attribute '__weakref__' of 'Person'
objects>}
Object attributes
<class '__main__.Person'>
{'name': 'John', 'age': 36}
```

19.14 Online Resources

See the following for further information on Python classes:

- <https://docs.python.org/3/tutorial/classes.html> the Python Standard library Class tutorial.
- https://www.tutorialspoint.com/python3/python_classes_objects.htm The tutorials point tutorial on Python 3 classes.

19.15 Exercises

The aim of this exercise is to create a new class called Account.

1. Define a new class to represent a type of bank account.
2. When the class is instantiated you should provide the account number, the name of the account holder, an opening balance and the type of account (which can be a string representing 'current', 'deposit' or 'investment', etc.). This means that there must be an `__init__` method and you will need to store the data within the object.
3. Provide three instance methods for the Account; `deposit(amount)`, `withdraw(amount)` and `get_balance()`. The behaviour of these methods should be as expected, deposit will increase the balance, withdraw will decrease the balance and `get_balance()` returns the current balance.
4. Define a simple test application to verify the behaviour of your Account class.

It can be helpful to see how your class Account is expected to be used. For this reason a simple test application for the Account is given below:

```
acc1 = Account('123', 'John', 10.05, 'current')
acc2 = Account('345', 'John', 23.55, 'savings')
```

```
acc3 = Account('567', 'Phoebe', 12.45, 'investment')

print(acc1)
print(acc2)
print(acc3)

acc1.deposit(23.45)
acc1.withdraw(12.33)
print('balance:', acc1.get_balance())
```

The following output illustrates what the result of running this test application might look like:

```
Account[123] - John, current account = 10.05
Account[345] - John, savings account = 23.55
Account[567] - Phoebe, investment account = 12.45
balance: 21.17
```

Chapter 20

Class Side and Static Behaviour



20.1 Introduction

Python classes can hold data and behaviour that is not part of an instance or object; instead they are part of the class.

This chapter introduces class-side data, behaviour and static behaviour. Note that in Python the concept of class-side data and behaviour is separate for the concept of static behaviour and so care should be taken when looking at these two aspects of a class.

20.2 Class-Side Data

In Python classes can also have attributes; these are referred to as *class* variables or attributes (as opposed to instance variables or attributes).

In Python variables defined within the scope of the class, but outside of any methods, are tied to the class rather than to any instance and are thus class variables.

For example, we can update the class `Person` to keep a count of how many instances of the class are created:

```
class Person:
    """ An example class to hold a persons name and age """

    instance_count = 0

    def __init__(self, name, age):
        Person.instance_count += 1
        self.name = name
        self.age = age
```

The variable `instance_count` is not part of an individual object, rather it is part of the class and all instances of the class can access that shared variable by prefixing it with the class name.

Now each time a new instance of the class is created, the `instance_count` is incremented, thus if we write:

```
p1 = Person('Jason', 36)
p2 = Person('Carol', 21)
p3 = Person('James', 19)
p4 = Person('Tom', 31)
print(Person.instance_count)
```

The output will be:

```
4
```

This is because 4 instances have been created and thus `__init__()` has been run 4 times and `instance_count` has been incremented four times.

20.3 Class-Side Methods

It is also possible to define behaviour that is linked to the *class* rather than an individual object; this behaviour is defined in a class method.

Class methods are written in a similar manner to any other method but are *decorated* with `@classmethod` and take a first parameter which represents the class rather than an individual instance. This decoration is written before the method declaration.

An example of a class method is shown below:

```
class Person:
    """ An example class to hold a persons name and age """
    instance_count = 0

    @classmethod
    def increment_instance_count(cls):
        cls.instance_count += 1

    def __init__(self, name, age):
        Person.increment_instance_count()
        self.name = name
        self.age = age
```

In this case the class method increments the `instance_count` variable; note that the `instance_count` variable is accessed via the `cls` parameter passed into the `increment_instance_count` method by Python. As this is a class method you do not need to prefix the class attribute with the name of class; instead the first parameter to the class method, `cls`, represents the class itself.

The class method can be accessed by *prefixing* it with the name of the class and using the *dot* notation to indicate which method to call. This is illustrated in the body of the `__init__()` method.

20.3.1 Why Class-Side Methods?

It may at first seem unclear what should normally go in an instance method as opposed to what should go in a class method. After all, they are both defined in the class. However, it is important to remember that

- Instance methods define the behaviour of the instance or object
- Class methods define the behaviour of the class.

Class-side methods should only perform one of the following roles:

- *Instance creation* This role is very important as it is how you can use a class as a factory for objects and can help hide a whole load of set up and instantiation work.
- *Answering enquiries about the class* This role can provide generally useful objects, frequently derived from class variables. For example, they may return the number of instances of this class that have been created.
- *Instance management* In this role, class-side methods control the number of instances created. For example, a class may only allow a single instance of the class to be created; this is termed a singleton class. Instance management methods may also be used to access an instance (e.g., randomly or in a given state).
- *Examples* Occasionally, class methods are used to provide helpful examples which explain the operation of a class. This can be very good practice.
- *Testing* Class-side methods can be used to support the testing of an instance of a class. You can use them to create an instance, perform an operation and compare the result with a known value. If the values are different, the method can report an error. This is a very useful way of providing regression tests.
- *Support* for one of the above roles

Any other tasks should be performed by an instance method.

20.4 Static Methods

There is one more type of method that can be defined on a class; these are *static* methods.

Static methods are defined within a class but are not tied to either the class nor any instance of the class; they do not receive the special first parameter representing either the class (`cls` for class methods) or the instances (`self` for instance methods).

They are in effect, the same as *free standing functions* but are defined within a class often for convenience or to provide a way to group such functions together.

A *static* method is a method that is decorated with the `@staticmethod` decorator.

An example of a static method is given below:

```
class Person:
    @staticmethod
    def static_function():
        print('Static method')
```

Static methods are invoked via the name of the class they are defined in, for example:

```
Person.static_function()
```

A note for Java and C# programmers; in both Java and C# the term class side and static are used interchangeably (not helped by the use of the keyword `static` for these methods). However, in both cases those methods are the equivalent of class-side methods in Python. In Python *class* methods and *static* methods are two very, very different things—do not use these terms interchangeably.

20.5 Hints

There are a range of special methods available in Python on a class.

All of these special methods start and end with a double underbars (`'__'`).

In general, in Python anything that starts and ends with these double underbars is considered special and so care should be taken when using them.

You should never name one of your own methods or functions `__<something>__` unless you intend to (re)define some default behaviour.

20.6 Online Resources

There is surprisingly little information available on Python's static and class methods (which in part explains why many developers are confused by them); however the following are available:

- <https://python-reference.readthedocs.io/en/latest/docs/functions/staticmethod.html> documentation on static methods.
- <https://python-reference.readthedocs.io/en/latest/docs/functions/classmethod.html?highlight=classmethod> documentation on class methods.
- <https://www.tutorialspoint.com/class-method-vs-static-method-in-python> tutorial on class methods versus static methods.

20.7 Exercises

The aim of this exercise is to add housekeeping style methods to the `Account` class.

You should follow these steps:

1. We want to allow the `Account` class from the last chapter to keep track of the number of instances of the class that have been created.
2. Print out a message each time a new instance of the `Account` class is created.
3. Print out the number of accounts created at the end of the previous test program.

For example add the following two statements to the end of the program:

```
print('Number of Account instances created:',  
      Account.instance_count)
```


Chapter 21

Class Inheritance



21.1 Introduction

Inheritance is a core feature of Object-Oriented Programming. It allows one class to *inherit* data or behaviour from another class and is one of the key ways in which reuse is enabled within classes.

This chapter introduces inheritance between classes in Python. It explains what inheritance is, how it is defined in Python, considers the purpose of inheritance and addresses the issues around single and multiple inheritance (both of which are supported by Python).

21.2 What is Inheritance?

Inheritance allows features defined in one class to be *inherited* and reused in the definition of another class. For example, a `Person` class might have the attributes `name` and `age`. It might also have behaviour associated with a `Person` such as `birthday()`.


We might then decide that we want to have another class `Employee` and that employees also have a `name` and an `age` and will have birthdays. However, in addition an `Employee` may have an `employee_id` attribute and a `calculate_pay()` behaviour.

At this point we could duplicate the definition of the `name` and `age` attributes and the `birthday()` behaviour in the class `Employee` (for example by cutting and pasting the code between the two classes).

However, this is not only inefficient; it may also cause problems in the future. For example we may realize that there is a problem or bug in the implementation of `birthday()` and may correct it in the class `Person`; however, we may forget to apply the same fix to the class `Employee`.

In general, in software design and development it is considered best practice to define something once and to reuse that something when required.

In an object-oriented system we can achieve the reuse of data or behaviour via inheritance. That is one class (in this case the `Employee` class) can *inherit* features from another class (in this case `Person`). This is shown pictorially below:



In this diagram the `Employee` class is shown as inheriting from the `Person` class. This means that the `Employee` class obtains all the data and behaviour of the `Person` class. It is therefore as though the `Employee` class has defined three attributes `name`, `age` and `id` and two methods `birthday()` and `calculate_pay()`.

A class that is defined as extending a parent class has the following syntax:

```
class SubClassName (BaseClassName) :
    class-body
```

Note that the parent class is specified by providing the name of the class in round brackets after the name of the new (child) class.

21.3 Defining Inheritance Between Classes

We can define the class `Person` in Python as before:

```
class Person:
    def __init__(self, name, age):
        self.name = name
        self.age = age

    def birthday(self):
        print('Happy birthday you were', self.age)
        self.age += 1
        print('You are now', self.age)
```

We could now define the class `Employee` as being a class whose definition builds on (or inherits from) the class `Person`:

```
class Employee(Person):
    def __init__(self, name, age, id):
        super().__init__(name, age)
        self.id = id

    def calculate_pay(self, hours_worked):
        rate_of_pay = 7.50
        if self.age >= 21:
            rate_of_pay += 2.50
        return hours_worked * rate_of_pay
```

Here we do several things:

1. The class is called `Employee` but it extends `Person`. This is indicated by including the name of the class being inherited in parentheses after the name of the class being defined (e.g., `Employee(Person)`) in the class declaration.
2. Inside the `__init__` method we reference the `__init__()` method defined in the class `Person` and used to initialize instances of that class (via the `super().__init__()` reference. This allows whatever initialization is required for `Person` to happen. This is called from within the `Employee` class's `__init__()` which then allows any initialization required by the `Employee` to occur. Note that the call to the `super().__init__()` initializer can come anywhere within the `Employee.__init__()` method; but by convention it comes first to ensure that whatever the `Person` class does during initialization does not over write what happens in the `Employee` class.
3. All instances of the class `Person` have a name, and age and have the behaviour `birthday()`.
4. All instances of the class `Employee` have a name, and age and an id and have the behaviours `birthday()` and `calculate_pay(hours_worked)`.
5. The method `calculate_pay()` defined in the `Employee` class can access the attributes name and age just as it can access the attribute id. In fact, it uses the employee's age to determine the rate of pay to apply.

We can go further, and we can subclass `Employee`, for example with the class `SalesPerson`:

```
class SalesPerson(Employee):
    def __init__(self, name, age, id, region, sales):
        super().__init__(name, age, id)
        self.region = region
        self.sales = sales

    def bonus(self):
        return self.sales * 0.5
```

Now we can say that the class `SalesPerson` has a name, an age and an `id` as well as a `region` and a `sales total`. It also has the methods `birthday()`, `calculate_pay(hourse_worked)` and `bonus()`.

In this case the `SalesPerson.__init__()` method calls the `Employee.__init__()` method as that is the next class up the hierarchy and thus we want to run that classes initialization behaviour before we set up the `SalesPerson` class (which of course in turn runs the `Person` classes initialization behaviour).

We can now write code such as:

```
print('Person')
p = Person('John', 54)
print(p)
print('--' * 25)

print('Employee')
e = Employee('Denise', 51, 7468)
e.birthday()
print('e.calculate_pay(40):', e.calculate_pay(40))
print('--' * 25)

print('SalesPerson')
s = SalesPerson('Phoebe', 21, 4712, 'UK', 30000.0)
s.birthday()
print('s.calculate_pay(40):', s.calculate_pay(40))
print('s.bonus():', s.bonus())
```

With the output being:

```
Person
John is 54
-----
Employee
Happy birthday you were 51
You are now 52
e.calculate_pay(40): 400.0
-----
SalesPerson
Happy birthday you were 21
You are now 22
s.calculate_pay(40): 400.0
s.bonus(): 15000.0
```

It is important to note that we have not done anything to the class `Person` by defining `Employee` and `SalesPerson`; that is it is not affected by those class definitions. Thus, a `Person` *does not* have an `employee id`. Similarly, neither an `Employee` nor a `Person` have a `region` or a `sales total`.

In terms of behaviour, instances of all three classes can run the method `birthday()`, but

- only `Employee` and `SalesPeron` objects can run the method `calculcate_pay()` and
- only `SalesPerson` objects can run the method `bonus()`.

21.4 Terminology Around Inheritance

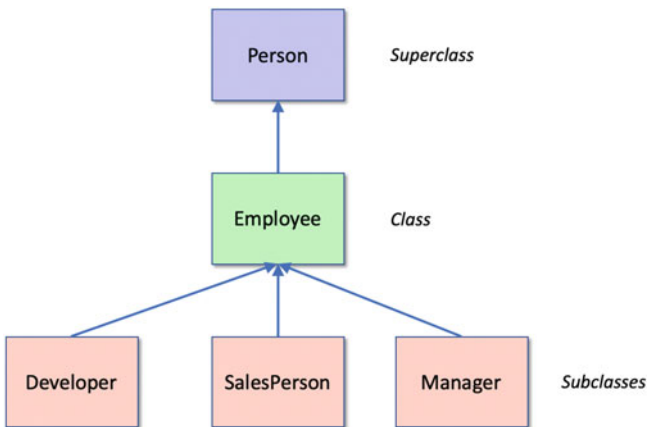
The following terminology is commonly used with inheritance in most object-oriented languages including Python:

Class A class defines a combination of data and procedures that operate on that data.

Subclass A subclass is a class that inherits from another class. For example, an `Employee` might inherit from a class `Person`. Subclasses are, of course, classes in their own right. Any class can have any number of subclasses.

Superclass A superclass is the parent of a class. It is the class from which the current class inherits. For example, `Person` might be the superclass of `Employee`. In Python, a class can have any number of superclasses.

Single or multiple inheritance Single and multiple inheritance refer to the number of super classes from which a class can inherit. For example, Java is a single inheritance system, in which a class can only inherit from one class. Python by contrast is a multiple inheritance system in which a class can inherit from one or more classes.



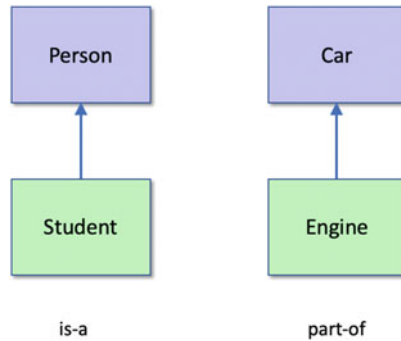
Note that a set of classes, involved in an inheritance hierarchy, such as those shown above is often named after the class at the root (top) of the hierarchy; in this case it would make these classes part of the `Person` class hierarchy.

21.4.1 Types of Hierarchy

In most object-oriented systems there are two types of hierarchy; one refers to *inheritance* (whether single or multiple) and the other refers to *instantiation*. The inheritance hierarchy has already been described. It is the way in which one class inherits features from a superclass.

The instantiation hierarchy relates to instances or objects rather than classes and is important during the execution of the object.

There are two types of instance relationships: one indicates a *part-of* relationship, while the other relates to a *using* relationship (it is referred to as an *is-a* relationship). This is illustrated below:

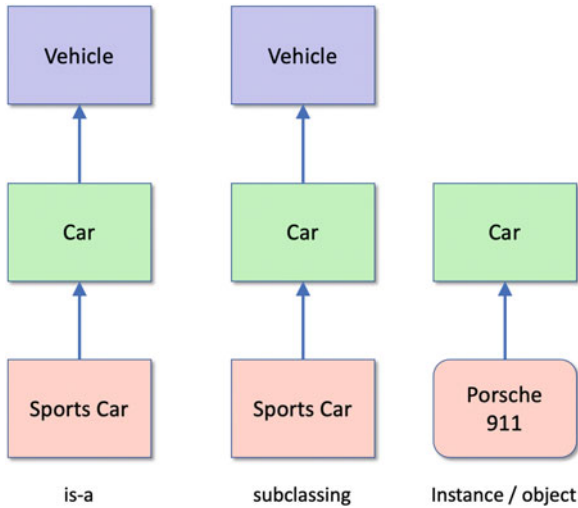


The difference between an *is-a* relationship and a *part-of* relationship is often confusing for new programmers (and sometimes for those who are experienced in non object-oriented languages). The above figure illustrates that a Student *is-a* type of Person whereas an Engine is *part-of* a Car. It does not make sense to say that a student is part-of a person or that an engine is-a type of car!

In Python, *inheritance* relationships are implemented by the sub-classing mechanism. In contrast, part-of relationships are implemented using instance attributes in Python.

The problem with classes, inheritance and is-a relationships is that on the surface they appear to capture a similar concept. In the following figure the hierarchies all capture some aspect of the use of the phrase *is-a*. However, they are all intended to capture a different relationship.

The confusion is due to the fact that in modern English we tend to overuse the term *is-a*. For example, in English we can say that an Employee is a type of Person or that Andrew is a Person; both are semantically correct. However, in Python classes such as Employee and Person and an object such as Andrew are different things. We can distinguish between the different types of relationship by being more precise about our definitions in terms of a programming language, such as Python.



21.5 The Class Object and Inheritance

Every class in Python extends one or more superclasses. This is true even of the class `Person` shown below:

```
class Person:
    def __init__(self, name, age):
        self.name = name
        self.age = age
```

This is because if you do not specify a superclass explicitly Python automatically adds in the class object as a parent class. Thus the above is exactly the same as the following listing which explicitly lists the class object as the superclass of `Person`:

```
class Person(object):
    def __init__(self, name, age):
        self.name = name
        self.age = age
```

Both listings above defines a class called `Person` that extends the class `object`.

In fact, between Python 2.2 and Python 3 it was required to use the long hand form to ensure that the *new style* classes were being used (as opposed to an older way in which classes were defined pre Python 2.2). As such it is common to find that Python developers still use the long hand (explicit) form when defining classes that directly extend `object`.

The fact that all class eventually inherit from the class `object` means that behaviour defined in `object` is available for all classes everywhere.

21.6 The Built-in Object Class

The class `object` is the base (root) class for all classes in Python. It has methods that are therefore available in all Python objects. It defines a common set of *special* methods and *intrinsic* attributes. The methods include the special methods `__str__()`, `__jnit__()`, `__eq__()` (equals) and `__hash__()` (hash method). It also defines attributes such as `__class__`, `__disct__`, `__doc__` and `__module__`.

21.7 Purpose of Subclasses

Subclasses are used to refine the behaviour and data structures of a superclass.

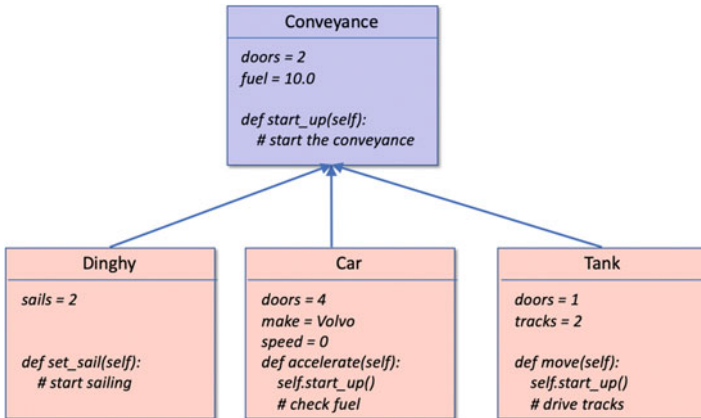
A parent class may define some generic/shared attributes and methods; these can then be inherited and reused by several other (sub) classes which add subclass specific attributes and behaviour.

In fact, there are only a small number of things that a subclass should do relative to its parent or super class. If a proposed subclass does not do any of these then your selected parent class is not the most appropriate super class to use.

A subclass should modify the behaviour of its parent class or extend the data held by its parent class. This modification should refine the class in one or more of these ways:

- Changes to the external protocol or interface of the class, that is it should extend the set of methods or attributes provided by the class.
- Changes in the implementation of the methods; i.e., the way in which the behaviour provided by the class are implemented.
- Additional behaviour that references inherited behaviour.

If a subclass does not provide one or more of the above, then it is incorrectly placed. For example, if a subclass implements a set of new methods, but does not refer to the attributes or methods of the parent class, then the class is not really a subclass of the parent (it does not extend it).



As an example, consider the class hierarchy illustrated above. A generic root class has been defined. This class defines a `Conveyance` which has `doors`, `fuel` (both with default values) and a method, `start_up()`, that starts the engine of the conveyance. Three subclasses of `Conveyance` have also been defined: `Dinghy`, `Car` and `Tank`. Two of these subclasses are appropriate, but one should probably not inherit from `Conveyance`. We shall consider each in turn to determine their suitability.

- The class `Tank` overrides the number of doors inherited, uses the `start_up` method within the method `move`, and provides a new attribute. It therefore matches all three of our criteria.
- Similarly, the class `Car` overrides the number of doors and uses the method `start_up()`. It also uses the instance variable `fuel` within a new method `accelerate()`. It also, therefore, matches our criteria.
- The class `Dinghy` defines a new attribute `sails` and a new method `set_sail()`. As such, it does not use any of the features inherited from `Conveyance`. However, we might say that it has extended `Conveyance` by providing this attribute and method. We must then consider the features provided by `Conveyance`. We can ask ourselves whether they make sense within the context of `Dinghy`. If we assume that a dinghy is a small sail-powered boat, with no cabin and no engine, then nothing inherited from `Conveyance` is useful. In this case, it is likely that `Conveyance` is misnamed, as it defines some sort of a *motor vehicle*, and the `Dinghy` class should not have extended it.

21.8 Overriding Methods

Overriding occurs when a method is defined in a class (for example, `Person`) and also in one of its subclasses (for example, `Employee`). It means that instances of `Person` and `Employee` both respond to requests for this method to be run but each has their own implementation of the method.

For example, let us assume that we define the method `__str__()` in these classes (so that we have a string representation of these objects to use with the `print` function). The pseudo code definition of this in `Person` might be:

```
def __str__(self):
    return 'Person ' + self.name + ' is ' + str(self.age)
```

In `Employee`, it might be defined as:

```
def __str__(self):
    return 'Employee(' + str(self.id) + ')'
```

The method in `Employee` replaces the version in `Person` for all instances of `Employee`. If we ask an instance of `Employee` for the result of `__str__()`, we get the string `'Employee(<some_id>')`. If you are confused, think of it this way:

If you ask an object to perform some operation, then, to determine which version of the method is run, look in the class used to create the instance. If the method is not defined there, look in the class's parent. Keep doing this until you find a method which implements the operation requested. This is the version which is used.

As a concrete example, see the classes `Person` and `Employee` below; in which the `__str__()` method in `Person` is overridden in `Employee`.

```
class Person:
    def __init__(self, name, age):
        self.name = name
        self.age = age
    def __str__(self):
        return self.name + ' is ' + str(self.age)

class Employee(Person):
    def __init__(self, name, age, id):
        super().__init__(name, age)
        self.id = id
    def __str__(self):
        return self.name + ' is ' + str(self.age) + ' - id(' +
str(self.id) + ')'
```

Instances of these classes will both be convertible to a string using `__str__()` but the version used by instances of `Employee` will differ from that used with instances of `Person`, for example:

```
p = Person('John', 54)
print(p)
e = Employee('Denise', 51, 1234)
print(e)
```

Generates as output:

```
John is 54
Denise is 51 - id(1234)
```

As can be seen from this the `Employee` class prints the name, age and `id` of the `Employee` while the `Person` class only prints the name and age.

21.9 Extending Superclass Methods

However, in the previous section we had to duplicated the code in `Person` down in `Employee` so that we could convert the `name` and `age` attributes into strings.

However we can avoid this duplication by invoking the parent class's method from within the child class version (as we in fact did for the `__init__()` initializer).

For example:

```
class Person:
    def __init__(self, name, age):
        self.name = name
        self.age = age

    def __str__(self):
        return self.name + ' is ' + str(self.age)

class Employee(Person):
    def __init__(self, name, age, id):
        super().__init__(name, age)
        self.id = id

    def __str__(self):
        return super().__str__() + ' - id(' + str(self.id) + ')'
```

In this version of the code the `Employee` classes version of the `__str__()` method first calls the parent classes version of this method and then adds the location information to the string returned from that. This means that we only have one location that converts `name` and `age` into a string.

The output from the code

```
p = Person('John', 54)
print(p)
e = Employee('Denise', 51, 1234)
print(e)
```

remains exactly the same:

```
John is 54
Denise is 51 - id(1234)
```

21.10 Inheritance-Oriented Naming Conventions

There are two naming conventions to be aware of with respect to Python classes and inheritance. These are that

- *Single underbar convention.* Methods or instance variables/attributes (those accessed via `self`) whose names start with a single under bar are considered to be *protected* that is they are private to the class but can be accessed from any subclass. Their scope is thus the class and any subclasses (either direct subclasses or any level of sub subclass).

- *Double underbar convention.* Method or instance variables/attributes (those accessed via self) whose names start with a double under bar should be considered *private* to that class and should not be called from outside of the class. This includes any subclasses; private means private to the class and only to that class.

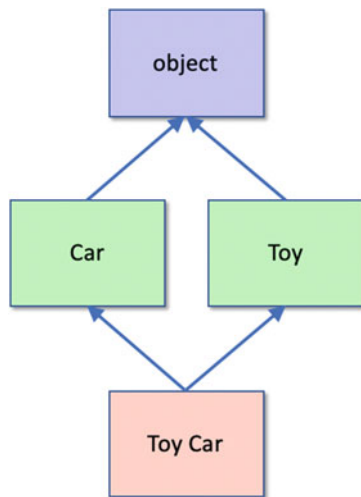
Any identifier of the form `__somename` (at least two leading underscores and at most one trailing underscore) is textually replaced with `__classname__somename`, where `classname` is the current class name with leading underscore(s) stripped.

Python does what is called name mangling to provide some support for methods that start with a double under bar. This mangling is done without regard to the syntactic position of the identifier, so it can be used to define class-private instance and class variables, methods, variables stored in globals, and even variables stored in instances.

21.11 Python and Multiple Inheritance

Python supports the idea of multiple inheritance; that is a class can inherit from one or more other classes (many object-oriented languages limit inheritance to a single class such as Java and C#).

This idea is illustrated by the following diagram:



In this case the class `ToyCar` inherits from the class `Car` and the class `Toy`. In turn the `Car` and `Toy` classes inherit from the (default) base class `object`.

The syntax for defining multiple inheritance in Python allows multiple super-classes to be listed in the parent class list (defined by the brackets following the class name). Each parent class is separated by a comma. The syntax is thus:


```
class SubClassName(BaseClassName1, BaseClassName2, ...  
BaseClassNameN):  
    class-body
```

For example:

```
class Car:  
    """ Car """  
  
class Toy:  
    """ Toy """  
  
class ToyCar(Car, Toy):  
    """ A Toy Car """
```

We can say that the class `ToyCar` inherits all the attributes (data) and methods (behaviour) defined in classes `Car`, `Toy` and `object`.

One of the fundamental questions that this raises is how is inheritance of behaviour managed within a multiple inheritance hierarchy. The challenge that multiple inheritance possesses is illustrated by adding a couple of methods to the class hierarchy we are looking at. In this example we have added the method `move()` to both the class `Car` and the class `Toy`:



The question here is which version of the method `move()` will be run when an instance of the `ToyCar` class is instantiated and we call `toy_car.move()`?

This illustrates (a simple version of) the so-called “diamond inheritance” problem.

The issue is that with multiple base classes from which attributes or methods may be inherited, there is often ambiguity that must be resolved. Here, when we create an instance of the class `ToyCar`, and call the `move()` method, does this invoke the one inherited from the `Car` base class or from the `Toy` base class?

The answer is that in Python 3, a *breadth first* search is used to find methods defined in parent classes; this means that when the method `move()` is called on `ToyCar`, it would first look in `Car`; it would then only look in `Toy` if it could not find a method `move()` in `Car`. If it cannot find the method in either `Car` or `Toy` it would then look in the class object.

As a result, it will find the version in `Car` first and use that version.

This is shown below:

```
class Car:
    def move(self):
        print('Car - move()')

class Toy:
    def move(self):
        print('Toy - move()')

class ToyCar(Car, Toy):
    """ A Toy Car """

tc = ToyCar()
tc.move()
```

The output of this is

```
Car - move()
```

However, if we alter the order in which the `ToyCar` inherits from the parent classes such that we swap `Toy` and `Car` around:

```
class ToyCar(Toy, Car):
    """ A Toy Car """
```

Then the `Toy` class is searched first and the output is changed to `Toy—move()`.

This shows that the order in which a class inherits from multiple classes is *significant* In Python.

21.12 Multiple Inheritance Considered Harmful

At first sight multiple inheritance in Python might appear to be particularly useful; after all it allows you to mix together multiple concepts into a single class very easily and quickly. This is certainly true, and it can be a very flexible feature if used with care. However, the word *care* is used here and should be noted.


Multiple inheritance can also be very dangerous and is quite a contentious topic for programmers and for those designing programming languages. Few things in programmer are inherently bad but multiple inheritance can result in a level of complexity (and unexpected behaviour) that can tie developers in knots.

Part of the problem highlighted by those protesting against multiple inheritance is down to the increased complexity and ambiguity that can occur with multiple inheritance trees that may interconnect between the different classes. One way to think of this is that if a class inherits from multiple classes, then that class may have the same classes in the class hierarchy multiple times, this can make it hard to determine which version of a method may execute and this may allow bugs to go untouched or indeed introduce expected issues due to different interactions between methods. This is exacerbated when inherited methods call `super()` using the same method name such as:

```
def get_data(self):  
    return super().get_data() + 'FData'
```

The following diagram presents a somewhat convoluted multiple inheritance examples where the class names A–X have been used so that there is no semantic meaning attributable to the inherited classes. Different classes define several common methods (`print_info()` and `get_data()`).

All the classes in the hierarchy define a `__str__()` method that returns the class name; if the class extends a class other than object, then the super version of `__str__()` is also invoked:



The code for this class hierarchy is given at the end of the section to avoid breaking up the flow.

We can now use the X class in a simple Python program:

```
x = X()
print('print(x):', x)
print('-' * 25)
x.print_info()
```

The question now is what is the output from this program?

What is the string printed to represent X? What is printed out as a result of calling the method `print_info()`?

The output from this simple code is:

```
print(x): CGFHJX
-----
HCDatAGDataFData
```

However, if we change the order of inheritance for the class 'H' from (F, G) to (G, F) then the output changes:

```
print(x): CFGHJX
-----
HCDatAFDataGData
```

This is of course because the search order, back up through the class hierarchy, is now different.

Note that this change came about not because of a modification we made to the class we instantiated (that is the class X), but from the order of the classes that one of its parents inherited from. This can be one of the unintended consequences of multiple inheritance; changing something in the multiple class hierarchy at one level can break some behaviour further down the hierarchy in a class that is unknown to the developer. Making the change.

Also note that the class inheritance diagram we presented earlier did not state what order the parent classes were listed in for any specific class (this was left to the discretion of the programmer).

Of course Python is not ambiguous nor does it get confused; it is the human developer that can get confused and be surprised with the behaviour that is then presented. Indeed, if you try and define a class hierarchy which Python cannot resolve into a consistent structure it will tell you so, for example:

```
Traceback (most recent call last):
  File "multiple_inheritance_example.py", line 65, in <module>
    class Z(H, J):
TypeError: Cannot create a consistent method resolution
order (MRO) for bases F, G
```

What can be confusing is that Python's ability to produce a consistent structure can also be dependent on the order of inheritance. For example, if we modify the classes that 'X' inherits from such that the order is I and J:

```
class X(I, J):
```



```
def __str__(self):
    return super().__str__() + 'X'
```

Then this compiles and can be used with the previous code (albeit with different output):

```
print(x): AIX
-----
A
```

However, if we change the order of the parent classes such that we swap I and J:

```
class X(J, I):
    def __str__(self):
        return super().__str__() + 'X'
```

We now get a `TypeError` exception raised:

```
Traceback (most recent call last):
  File "multiple_inheritance_example.py", line 73, in <module>
    class X(J, I):
TypeError: Cannot create a consistent method resolution
order (MRO) for bases J, I
```

Therefore, in general care needs to be taken when utilising multiple inheritance; but that is not to say that such situations are not useful. In some cases you want a class to inherit from parents that have complete different hierarchies and are completely separate from each other; in such situations multiple inheritance can be very useful—these so called *orthogonal behaviours* are one of the best uses of multiple inheritance and should not be ignored merely due to concerns of increased complexity.

The class definitions used for the multiple inheritance class hierarchy are given below:

```
class A:
    def __str__(self):
        return 'A'
    def print_info(self):
        print('A')

class B:
    def __str__(self):
        return 'B'

class C:
    def __str__(self):
        return 'C'
    def get_data(self):
        return 'CData'

class D:
    def __str__(self):
        return 'D'
    def print_info(self):
        print('D')
```

```

class E:
    def __str__(self):
        return 'E'
    def print_info(self):
        print('E')

class F(C, D, E):
    def __str__(self):
        return super().__str__() + 'F'
    def get_data(self):
        return super().get_data() + 'FData'
    def print_info(self):
        print('F' + self.get_data())

class G(C, D, E):
    def __str__(self):
        return super().__str__() + 'G'
    def get_data(self):
        return super().get_data() + 'GData'

class H(F, G):
    def __str__(self):
        return super().__str__() + 'H'
    def print_info(self):
        print('H' + self.get_data())

class J(H):
    def __str__(self):
        return super().__str__() + 'J'

class I(A, J):
    def __str__(self):
        return super().__str__() + 'I'

class X(J, H, B):
    def __str__(self):
        return super().__str__() + 'X'

```

21.13 Summary

To recap on the concept of inheritance. Inheritance is supported between classes in Python. For example, a class can extend (subclass) another class or a set of classes. A subclass inherits all the methods and attributes defined for the parent class(es) but may override these in the subclass.

In terms of the inheritance we say:

- A subclass inherits from a super class.
- A subclass obtains all code and attributes from the super class.
- A subclass can add new code and attributes.
- A subclass can override inherited code and attributes.
- A subclass can invoke inherited behaviour or access inherited attributes.

21.14 Online Resources

There are many resources available online relating to class inheritance including:

- <https://docs.python.org/3/tutorial/classes.html#inheritance> The Python software foundation tutorial on class inheritance.
- https://en.wikipedia.org/wiki/Multiple_inheritance which provides a discussion on multiple inheritance and the potential challenges it can introduce.

21.15 Exercises

The aim of these exercises is to extend the `Account` class you have been developing from the last two chapters by providing `DepositAccount`, `CurrentAccount` and `InvestmentAccount` subclasses.

Each of the classes should extend the `Account` class by:

1. `CurrentAccount` adding an overdraft limit as well as redefining the `withdraw` method.
2. `DepositAccount` by adding an interest rate.
3. `InvestmentAccount` by adding an investment type attribute.

These features are discussed below:

The `CurrentAccount` class can have an `overdraft_limit` attribute. This can be set when an instance of a class is created and altered during the lifetime of the object. The overdraft limit should be included in the `__str__()` method used to convert the account into a string.

The `CurrentAccount` `withdraw()` method should verify that the balance never goes below the overdraft limit. If it does then the `withdraw()` method should not reduce the balance instead it should print out a warning message.

The `DepositAccount` should have an interest rate associated with it which is included when the account is converted to a string.

The `InvestmentAccount` will have a `investment_type` attribute which can hold a string such as 'safe' or 'high risk'.

This also means that it is no longer necessary to pass the type of account as a parameter—it is implicit in the type of class being created.

For example, given this code snippet:

```
# CurrentAccount(account_number, account_holder,
#                 opening_balance, overdraft_limit)
acc1 = CurrentAccount('123', 'John', 10.05, 100.0)
# DepositAccount(account_number, account_holder, opening_
balance,
#                 interest_rate)
acc2 = DepositAccount('345', 'John', 23.55, 0.5)
# InvestmentAccount(account_number, account_holder, opening_
balance,
```

```
#             investment_type)
acc3 = InvestmentAccount('567', 'Phoebe', 12.45, 'high risk')

acc1.deposit(23.45)
acc1.withdraw(12.33)

print('balance:', acc1.get_balance())
acc1.withdraw(300.00)
print('balance:', acc1.get_balance())
```

Then the output might be:

```
balance: 21.17
Withdrawal would exceed your overdraft limit
balance: 21.17
```

Chapter 22

Why Bother with Object Orientation?



22.1 Introduction

The previous four chapters have introduced the basic concepts behind object orientation, the terminology and explored some of the motivation. This chapter looks at how object orientation addresses some of the issues that have been raised with procedural languages. To do this it looks at how a small extract of a program might be written in a language such as C, considers the problems faced by the C developer and then looks at how the same functionality might be achieved in an object-oriented language such as Python. Do not worry too much about the syntax you will be presented with; it is mostly a form of pseudo-code, and it should not detract from the legibility of the examples.

22.2 The Procedural Approach

Consider the following example:

```
record Date {  
    int day  
    int month  
    int year  
}
```

This defines a data structure for recording dates. There are similar structures in many procedural languages such as C, Ada and Pascal.

So, what is wrong with a structure such as this? Nothing, apart from the issue of visibility? That is, what can see this structure and what can update the contents of the structure? Any code can directly access and modify its contents. Is this problem? It could be, for example, some code could set the day to -1 , the month to 13 and the year to 9999.

As far as the structure is concerned the information it now holds is fine (that is day = 01, month = 13, year = 9999). This is because the structure only knows it is supposed to hold integers; it knows nothing about dates per se. This is not surprising, it is only data.

22.2.1 *Procedures for the Data Structure*

This data is associated with procedures that perform operations on it. These operations might be to

- test whether the date represents a date at a weekend or part of the working week.
- change the date (in which case the procedure may also check to see that the date is a valid one).

For example:

```
is_day_of_week(date)
in_month(date, 2)
next_day(date)
set_day(date, 9, 3, 1946)
```

How do we know that these procedures are related to the date structure we have just looked at? By the naming conventions of the procedures and by the fact that one of the parameters is a data (record).

The problem is that these procedures are not limited in what they can do to the data (for example the `setDay` procedure might have been implemented by a Brit who assumes that the data order is day, month and year. However, it may be used by an American who assumes that date order is month, day, year. Thus the mean of `set_day(date, 9, 3, 1946)` will be interpreted very differently. The American views this as the 3rd of September 1946, while the Brit views this as the 9th of March, 1946. In either case, there is nothing to stop the date record being updated with both versions. Obviously the `set_day()` procedure might check the new date to see it was legal, but then again it might not. The problem is that the data is naked and has no defense against what these procedures do to it. Indeed, it has no defense against what any procedures that can access it, may do to it.

22.2.2 *Packages*

One possibility is of course to use a package construct. In languages such as Ada, packages are common place and are used as a way of organizing code and restricting visibility. For example,

```

package Dates is
  type Date is ...
  function is_day_of_week(d: Date) return Boolean;
  function in_month(d: Date, m: Integer) return Boolean;
...

```

The package construct provides some ring fencing of the data structure and a grouping of the data structure with the associated procedures. In order to use this package a developer must import the package. They can then access the procedures and work with data of the specified type (in this case Date).

There can even be data that is hidden from the user within a private part. This therefore increases the ability to encapsulate the data (hide the data) from unwelcome attention.

22.3 Does Object Orientation Do Any Better?

This is an important question “*Does object orientation do any better?*” than the procedural approach described above? We will first consider classes then inheritance.

22.3.1 Packages Versus Classes

It has been argued (to me at least) that an Ada package is just like a class. It provides a template from which you can create executable code, it provides a wall around your data with well-defined gateways etc. However, there are a number of very significant differences between packages and classes.

Firstly, packages tend to be larger (at least conceptually) units than classes. For example, the TextIO package in Ada is essentially a library of textual IO facilities, rather than a single concept such as the class `string` in Python. Thus packages are not used to encapsulate a single small concept such as `string` or `Date`, but rather a whole set of related concepts (as indeed they are used in Python). Thus, a class is a finer level of granularity than a package.

Secondly, packages still provide a relatively loose association between the data and the procedures. An Ada package may actually deal with very many data structures with a wide range of methods. The data and the methods are related primarily via the related set of concepts represented by the package. In contrast a class tends to closely relate data and methods in a single concept. Indeed, one of the guidelines relating to good class design is that if a class represents more than one concept, then you should split it into two classes.

Thus, this close association between data and code means that the resulting concept is more than just a data structure (it is closer to a concrete realization of concept). For example:

```

class Date:

    def __init__(self, day, month, year):
        self.day = day
        self.month = month
        self.year = year

    def is_day_of_week(self):
        """Check if date is a week day"""
        # ... To be defined

    def in_month(self, month_index):
        """Check if month is in month_index"""
        return self.month == month_index

```

Anyone using an instance of `Date` now gets an object which can tell you whether it is a day of the week or not and can hold the appropriate data. Note that the `is_day_of_week()` method takes no parameters other than `self`, it doesn't need to as it and the date information are part of the same thing. This means that a user of a `Date` object will never need to get their hands on the actual data holding the date (i.e., the integers `day`, `month` and `year`). Instead, they should go via the methods. This may only seem a small step, but it is a significant one, nothing outside the object should need to access the data within the object. In contrast the data structure in the procedural version, is not only held separately to the procedures, the values for `day`, `month` or `year` it must also be modified directly.

For example, compare the differences between an excerpt from a program to manipulate dates (using a procedural programming language):

```

d: Date;
setDay(d, 28);
setMonth(d, 2);
setYear(d, 1998);
isDayOfWeek(d);
inMonth(d, 2);

```

Note that it was necessary to first create the data and then to set the fields in the data structure. Here we have been good and have used the interface procedures to do this. Once we had the data set up we could then call methods such as `IsDayOfWeek` and `InMonth` on that data.

In contrast the Python code uses a constructor to pass in the appropriate initialization information. How this is initialized internally is hidden from the user of the class `Date`. We then call method such as `is_day_of_week()` and `is_month(12)` directly on the object `date`.

```

date = Date(12, 2, 1998)
date.is_day_of_week()
date.in_month(12)

```

The thing to think about here is where would code be defined?

22.3.2 *Inheritance*

Inheritance is a key element in an object-oriented language allowing one class to inherit data and methods from another.

One of the most important features of inheritance (ironically) is that it allows the developer to get inside the encapsulation bubble in limited and controlled ways.

This allows the sub-class to take advantage of internal data structures and methods, without compromising the encapsulation afforded to objects. For example, let us define a subclass of the class `Date`:

```
class Birthday(Date):
    name = "
    age = 0
    def is_birthday():
        # ... Check to see if it is their birthday
```

The method `is_birthday()` could check to see if the current date, matched the birthday represented by an instance of `Birthday` and return true if it does and false if it does not.

Note however, that the interesting thing here is that not only have we not had to define integers to represent the date, nor have we had to define methods to access such dates. These have both been inherited from the parent class `Date`.

In addition, we can now treat an instance of `Birthday` as either a `Date` or as a `Birthday` depending on what we want to do!

What would you do in languages such as C, Pascal or Ada? One possibility is that you could define a new package `Birthday`, but that package would not extend `Dates`, it would have to import `Dates` and add interfaces to it, etc.? However, you certainly couldn't treat a `Birthday` package as a `Dates` package.

In languages such as Python, because of polymorphism, you can do exactly that. You can reuse existing code that only knew about `Date`, for example:

```
birthday = Birthday(12, 3, 1974)

def test(date):
    # Do something that works with a date

t.test(birthday)
```

This is because `birthday` is indeed a type of `Date` as well as being a type of `Birthday`.

You can also use all of the features defined for `Date` on `Birthdays`:

```
birthday.is_day_of_week()
```

Indeed, you don't actually know where the method is defined. This method could be defined in the class `Birthday` (where it would override that defined in the class `Date`). However, it could be defined in the class `Date` (if no such method is defined in `Birthday`); without looking at the source code there is no way of knowing!

Of course, you can also use the new methods defined in the class `Birthday` on instance (objects) of this class. For example:

```
birthday.is_birthday()
```

22.4 Summary

Classes in an object-oriented language provide a number of features that are not present in procedural languages. To summarize, the main points to be noted from this chapter on object orientation are:

1. Classes provide for inheritance.
2. Inheritance provides for reuse.
3. Inheritance provides for extension of a data type.
4. Inheritance allows for polymorphism.
5. Inheritance is a unique feature of object orientation.

Chapter 23

Operator Overloading



23.1 Introduction

It can be convenient when defining a user-defined class (type) to define operators that can make working with instance of those classes easier. Such operators might include the '+' plus operator or the '-' operator. This allows two instance of a class to be added together or subtracted. In this chapter we will explore operator overloading; what it is, how it works and why we want it.

23.2 Operator Overloading

23.2.1 Why Have Operator Overloading?

Operator overloading allows user defined *classes* to appear to have a natural way of using operators such as +, -, <, > or == as well as logical operators such as & (and) and | (or).

This leads to more succinct and readable code as it is possible to write code such as:

```
q1 = Quantity(5)
q2 = Quantity(10)
q3 = q1 + q2
```

It feels more natural for both developers and those reading the code. The alternative would be to create methods such as add and write code such as

```
q1 = Quantity(5)
q2 = Quantity(10)
q3 = q1.add(q2)
```

Which semantically might mean the same thing but feel less *natural* to most people.

23.2.2 Why not Have Operator Overloading?

If operator overloading is such a good idea, why don't all programming languages support it? Interestingly Java, a very widely used programming language, does not support operator overloading!

One answer is because it can be abused! For example, what is the meaning of the following code:

```
p1 = Person('John')
p2 = Person('Denise')
p3 = p1 + p2
```

It is not clear what '+' means in this context; in what way is Denise being added to John; does it imply they are getting married? If so, what is the result that is held in p3?

The problem here is that from a design perspective (which in this case may be purely intuitive but in other cases may relate to the intention of an application) the plus operator does not make sense for the type Person. However, there is nothing in the Python language to indicate this and thus anyone can code any operator into any class!

As a general design principle, developers should follow the semantics of built in types and thus should only implement those operators which are appropriate for the type being developed. For example, for arithmetic value types such as `Quantity` it makes perfect sense to provide a plus operator but for domain specific data-oriented types such as `Person` it does not.

23.2.3 Implementing Operator Overloading

To implement operators such as '+' in a user-defined class it is necessary to implement specific methods that are then mapped to the arithmetic or logical operators used by users of the class.

These methods are considered *special* in that they start with and end with a double underscore ('__'). Such methods are considered private and usually restricted for Python-oriented implementations (we have seen these already with methods such as `__init__()` and `__str__()`).

As an example, let us assume that we want to implement the '+' and '-' operators for our `Quantity` type. We also want our `Quantity` type to hold an actual value and be able to be converted into a string for printing purposes.

To implement the '+' and '-' operators we need to provide two special methods one will provide the implementation of the '+' operator and one will provide the implementation of the '-' operator:

- '+' operator is implemented by a method with the signature **def** `__add__(self, other):`
- '-' operator is implemented by a method with the signature **def** `__sub__(self, other):`

Where `other` represents another `Quantity` or other suitable type which will be either added to, or subtracted from, the current `Quantity` object.

The methods will be mapped by Python to the operators '+' and '-'; such that if someone attempts to add to quantities together then the `__add__()` method will be called, etc.

The definition of the class `Quantity` is given below; note that the class actually just wraps a number held in the attribute `value`.

```
class Quantity:
    def __init__(self, value=0):
        self.value = value
    def __add__(self, other):
        new_value = self.value + other.value
        return Quantity(new_value)
    def __sub__(self, other):
        new_value = self.value - other.value
        return Quantity(new_value)
    def __str__(self):
        return 'Quantity[' + str(self.value) + ']'
```

Using this class definition, we can create two instances of the type `Quantity` and add them together:

```
q1 = Quantity(5)
q2 = Quantity(10)
print('q1 =', q1, ', q2 =', q2)

q3 = q1 + q2
print('q3 =', q3)
```

If we run this code snippet we get:

```
q1 = Quantity[5] , q2 = Quantity[10]
q3 = Quantity[15]
```

Note that we have made the class `Quantity` *immutable*; that is once a `Quantity` instance has been created its value cannot be changed (it is fixed).

This means that when two quantities are added together a new instance of the class `Quantity` is created. This is analogous to how integers work, if you add together $2+3$ then you get 5; neither 2 or 3 are modified however; instead a new integer 5 is generated—this is an example of the general design principle; developers should follow the semantics of built in types; `Quantity` objects act like number objects.

23.3 Numerical Operators

There are nine different numerical operators that can be implemented by special methods; these operators are listed in the following table:

Operator	Expression	Method
Addition	$q1 + q2$	<code>__add__(self, q2)</code>
Subtraction	$q1 - q2$	<code>__sub__(self, q2)</code>
Multiplication	$q1 * q2$	<code>__mul__(self, q2)</code>
Power	$q1 ** q2$	<code>__pow__(self, q2)</code>
Division	$q1 / q2$	<code>__truediv__(self, q2)</code>
Floor Division	$q1 // q2$	<code>__floordiv__(self, q2)</code>
Modulo (Remainder)	$q1 \% q2$	<code>__mod__(self, q2)</code>
Bitwise Left Shift	$q1 \ll q2$	<code>__lshift__(self, q2)</code>
Bitwise Right Shift	$q1 \gg q2$	<code>__rshift__(self, q2)</code>

We have already seen examples of add and subtract; this table indicates how we can also provide operators for multiplication and division etc.

The above table also presents Bitwise shift operators (both left and right). These operate at the bit level used to represent numbers under the hood and can be a very efficient way of manipulating numeric values; however, we do not want to support these operators for our Quantity class therefore we will only implement the core numeric operators of multiplication, division and power.

Also note that the names of the division methods are not `div` but `__truediv__()` and `__floordiv__()` indicating the difference in behaviour between `/` and `//`.

The updated Quantity class is given below:

```
class Quantity:
    def __init__(self, value=0):
        self.value = value

    def __add__(self, other):
        new_value = self.value + other.value
        return Quantity(new_value)

    def __sub__(self, other):
        new_value = self.value - other.value
        return Quantity(new_value)

    def __mul__(self, other):
        new_value = self.value * other.value
        return Quantity(new_value)

    def __pow__(self, other):
        new_value = self.value ** other.value
```

```

    return Quantity(new_value)

def __truediv__(self, other):
    new_value = self.value / other.value
    return Quantity(new_value)

def __floordiv__(self, other):
    new_value = self.value // other.value
    return Quantity(new_value)

def __mod__(self, other):
    new_value = self.value % other.value
    return Quantity(new_value)

def __str__(self):
    return 'Quantity[' + str(self.value) + ']'

```

This means that we can now extend our simple application that uses the Quantity class to include some of these additional numerical operators:

```

q1 = Quantity(5)
q2 = Quantity(10)
print('q1 =', q1, ', q2 =', q2)

q3 = q1 + q2
print('q3 =', q3)
print('q2 - q1 =', q2 - q1)

print('q1 * q2 =', q1 * q2)
print('q1 / q2 =', q1 / q2)

```

The output from this is now:

```

q1 = Quantity[5] ,q2 = Quantity[10]
q3 = Quantity[15]
q2 - q1 Quantity[5]
q1 * q2 = Quantity[50]
q1 / q2 = Quantity[0.5]

```

One interesting point to note is that given the multiple and divide style methods, we might want to multiple a Quantity by an integer or divide a Quantity by an inter. There is nothing to stop us doing this and indeed this might be a very useful behaviour. This would allow a Quantity to be multiple by 2 or divided by 2, for example:

```

print('q1 * 2', q1 * 2)
print('q2 / 2', q2 / 2)

```

At the moment if we tried to run the above code we would generate an error telling us that an int does not have a value attribute. However, we can test to see if the argument passed into the `__mult__()` and `__truediv__()` methods is an int or not using the `is` instance function. This function takes a variable and the name of a class and returns a new Quantity. To do this it either adds the values of two Quantities together (if the other value is a Quantity) or an integer to the `self.value` integer (if the other value is an int), for example:

```

class Quantity:
    # Code omitted for brevity

    def __mul__(self, other):
        if isinstance(other, int):
            new_value = self.value * other
        else:
            new_value = self.value * other.value
        return Quantity(new_value)

    def __truediv__(self, other):
        if isinstance(other, int):
            new_value = self.value / other
        else:
            new_value = self.value / other.value
        return Quantity(new_value)

```

Now when we run the earlier print statements we generate the output:

```

q1 * 2 Quantity[10]
q2 / 2 Quantity[5.0]

```

23.4 Comparison Operators

Numerical types (such as integers and real numbers) also support comparison operators such as equals, not equals, greater than, less than as well as greater than or equal to and less than or equal to.

Python allows these comparison operators to be defined for user-defined types/classes as well.

Just as numerical operators such as '+' and '-' are implemented by special methods so are comparison operators. For example the '<' operator is implemented by a method called `__lt__(self, other)`.

The complete list of comparison operators and the associated special methods is given in the following table:

Operator	Expression	Method
Less than	<code>q1 < q2</code>	<code>__lt__(q1, q2)</code>
Less than or equal to	<code>q1 <= q2</code>	<code>__le__(q1, q2)</code>
Equal to	<code>q1 == q2</code>	<code>__eq__(q1, q2)</code>
Not Equal to	<code>q1 != q2</code>	<code>__ne__(q1, q2)</code>
Greater than	<code>q1 > q2</code>	<code>__gt__(q1, q2)</code>
Greater than or equal to	<code>q1 >= q2</code>	<code>__ge__(q1, q2)</code>

We can add these definitions to our `Quantity` class to provide a more complete type that can be used in comparison style tests (such as `if` statements).

The updated `Quantity` class is given below (with some of the numerical operators omitted for brevity):

```
class Quantity:
    def __init__(self, value=0):
        self.value = value
    def __add__(self, other):
        new_value = self.value + other.value
        return Quantity(new_value)

    # remaining numerical operators omitted for brevity ...

    def __eq__(self, other):
        return self.value == other.value

    def __ne__(self, other):
        return self.value != other.value

    def __ge__(self, other):
        return self.value >= other.value

    def __gt__(self, other):
        return self.value > other.value

    def __lt__(self, other):
        return self.value < other.value

    def __le__(self, other):
        return self.value <= other.value

    def __str__(self):
        return 'Quantity[' + str(self.value) + ']'
```

This now means that we can update our sample application to take advantage of these comparison operators:

```
q1 = Quantity(5)
q2 = Quantity(10)
print('q1 =', q1, ', q2 =', q2)
q3 = q1 + q2
print('q3 =', q3)
print('q1 < q2: ', q1 < q2)
print('q3 > q2: ', q3 > q2)
print('q3 == q1: ', q3 == q1)
```

The output from this is now:

```
q1 = Quantity[5] ,q2 = Quantity[10]
q3 = Quantity[15]
q1 < q2: True
q3 > q2: True
q3 == q1: False
```

23.5 Logical Operators

The final category of operators that can be defined on a class are logical operators. These are operators that can be used with and/or type tests; they typically return the values `True` or `False`.

As with the numerical operators and the comparison operators; the logical operators are implemented by a set of special methods.

The following table summarizes the logical operators and the methods used to implement them:

Operator	Expression	Method
And	<code>q1 & q2</code>	<code>__and__(q1, q2)</code>
OR	<code>q1 q2</code>	<code>__or__(q1, q2)</code>
XOR	<code>q1 ^ q2</code>	<code>__xor__(q1, q2)</code>
NOT	<code>~ q1</code>	<code>__invert__()</code>

As these operators do not really make sense for the `Quantity` type, we will not define them. For example, what would it mean to say:

```
q1 | q2
```

In what way is `q1` an alternative to `q2`?

23.6 Summary

Only use operators when they make sense and only implement those operators that work with the type you are defining. In general, this means

- Arithmetic operators should only be used for values types with a numeric property.
- Comparison operators typically only make sense for classes that can be ordered.
- Logical operators typically work for types that are similar in nature to Booleans.

23.7 Online Resources

Some online resources on operator overloading include:

- <https://docs.python.org/3/reference/datamodel.html> for information on operator overloading in Python.
- <https://pythonprogramming.net/operator-overloading-intermediate-python-tutorial/> Tutorial on operator overloading.
- <http://cafe.elharo.com/programming/operator-overloading-considered-harmful/> An article on why operagoer overloading may be harmful to good programming style.

23.8 Exercises

The aim of this exercise is to create a new numeric style class.

You should create a new user defined class called `Distance`. It will be very similar to `Quantity`.

You should be able to add two distances together, subtract one distance from another, divide a distance by an integer, multiply a distance by an integer etc.

You should therefore be able to support the following program:

```
d1 = Distance(6)
d2 = Distance(3)

print(d1 + d2)
print(d1 - d2)
print(d1 / 2)
print(d2 // 2)
print(d2 * 2)
```

Note that the division and multiplication operators work with a distance and an integer; you will therefore need to think about how to implement the special methods for these operators.

The output from this might be:

```
Distance[9]
Distance[3]
Distance[3.0]
Distance[1]
Distance[6]
```

Chapter 24

Python Properties



24.1 Introduction

Many object-oriented languages have the explicit concept of encapsulation; that is the ability to hide data within an object and only to provide specific gateways into that data. These gateways are methods defined to *get* or *set* the value of an attribute (often referred to as getters and setters). This allows more control over access to the data; for example, it is possible to check that only a positive integer above zero, but below 120, is used for a person's age, etc.

In many languages such as Java and C# attributes can be hidden from external access using specific keywords (such as `private`) that indicate the data should be made private to the object.

Python does not explicitly have the concept of encapsulation; instead it relies on two things; a standard convention used to indicate that an attribute should be considered private and a concept called a property which allows setters and getters to be defined for an attribute.

24.2 Python Attributes

All object attributes are publicly available in Python; that is, they are all visible to any code using the object.

For example, given the following definition of the class `Person` both `name` and `age` are part of the public interface of the class `Person`;

```
class Person:
    def __init__(self, name, age):
        self.name = name
        self.age = age
```

```

def __str__(self):
    return 'Person[' + str(self.name) + '] is ' +
           str(self.age)

```

Because name and age are part of the class's public interface it means that we can write:

```

person = Person('John', 54)
person.name = 42
person.age = -1
print(person)

```

Which is of course a bit bizarre as the person now has the name '42' and an age of -1, thus the output from this is:

```

Person[42] is -1

```

We can indicate that we want to treat age and name as being private to the object by prefixing the attribute names with an underbar ('_') as shown below:

```

class Person:
    def __init__(self, name, age):
        self._name = name
        self._age = age

    def __str__(self):
        return 'Person[' + str(self._name) + '] is ' + s
                str(self._age)

```

This tells Python programmers that we want to consider `_name` and `_age` as being *private*. However, it should be noted that this is only a convention; albeit a *very strongly* adhered to convention. There is still nothing here that would stop someone writing:

```

person = Person('John', 54)
person._age = -1
print(person)

```

However, the developer of the class `Person` is at liberty to change the internals of the class (such as `_age`) without notice and most would consider that anyone who had ignored the convention and now had a problem had only themselves to blame.

24.3 Setter and Getter Style Methods

This of course raises the question; how should we now get hold of a Person's name and age in an acceptable way?

The answer is that a developer should provide *getter* methods and *setter* methods that can be used to access the values.

We can update the `Person` class with some getter methods and a single setter method:

```
class Person:
    def __init__(self, name, age):
        self._name = name
        self._age = age

    def get_age(self):
        return self._age

    def set_age(self, new_age):
        if isinstance(new_age, int) & new_age > 0 & new_age < 120:
            self._age = new_age

    def get_name(self):
        return self._name

    def __str__(self):
        return 'Person[' + str(self._name) + '] is ' + str(self._age)
```

The two getter methods have the format of `get_` followed by the name of the attribute they are getting. Thus, we have `get_age` and `get_name`. Typically, all that getters do is to return the attribute being used (as is the case here).

The single setter method is a little different; it validates the data that has been provided to check that it is appropriate (i.e., that it is an Integer using `isinstance(new_age, int)` and that it is a value over zero but under 120). Only if the data passes these checks is it used as the new value of the person's age, for example if we try to set a persons age to `-1`:

```
person = Person('John', 54)
person.set_age(-1)
print(person)
```

Then this is ignored, and the person's age remains as it was, thus the output of this is:

```
Person[John] is 54
```

It should be noted that this might be considered *silent failure*; that is we tried to set the `age` and it failed but no one knows this. In many cases rather than fail silently we would prefer to notify someone of the error by throwing some form of Error object; this will be discussed in the next chapter on Error and Exception handling.

You might well ask at this point where is the setter for the `_name` attribute? The answer is that we want to make the `_name` attribute a *read only* attribute and therefore we have not provided a setter style method. This is a common idiom followed in Python—you can have read-write attributes and read-only attributes depending on whether they have getter and setter methods or not. It is also possible to a *write-only* attribute, but this is very rare and only has a few use cases.

24.4 Public Interface to Properties

Although we now have a more formal interface to the attributes held by an instance of the class `Person`; it is rather ungainly:

```
person = Person('John', 54)
print(person)
print(person.get_age())
print(person.get_name())
```

We end up having to write more code and although there is an argument that it makes the code more obvious (i.e., `person.get_age()` can be read as get the age of the person object); it is somewhat verbose and you have to remember to include the parentheses `('')`.

To get around this a concept known as Properties was introduced in Python 2.2. In the original syntax for this it was possible to add an addition line of code to the class that told Python that you wanted to provide a new property and that specific methods were to be used to set and get the values of this property.

The syntax for defining a property in this way is:

```
<property_name> = property(fget=None, fset=None, fdel=None,
doc=None)
```

Where `fget` indicates the getter function, `fset` the setter function `fdel` the function to be used to delete a value and `doc` provides documentation on the property.

We can modify our `Person` class so that `age` is now a property (note a common convention is that if the attribute is named `_age`, the methods are named `get_age` and `set_age` and the property will be called `age`):

```
class Person:
    def __init__(self, name, age):
        self._name = name
        self._age = age

    def get_age(self):
        return self._age

    def set_age(self, new_age):
        if isinstance(new_age, int) & new_age > 0 & new_age < 120:
            self._age = new_age

age = property(get_age, set_age, doc="An age property")

    def get_name(self):
        return self._name

name = property(get_name, doc="A name property")

    def __str__(self):
        return 'Person[' + str(self._name) + ']' is ' +
            str(self._age)
```

We can now write:

```
person = Person('John', 54)
print(person)
print(person.age)
print(person.name)
person.age = 21
print(person)
```

Notice how we can now write `person.age` and `person.age = 21`; in both these cases we are accessing the *property* `age` that results in the method `get_age()` and `set_age()` being executed respectively. Thus, the setter is still protecting the update to the underlying `_age` attribute that is actually used to hold the actual value.

Also note that if a method is not provided for one of the `fget`, `fset`, `fdel` methods then this is not an error; it merely indicates that the property does not support that type of accessor. Thus, the `name` property is a *read-only* property as it does not define a setter method.

A delete method can be used to release the memory associated with an attribute; in the case of an `int` it is not required but it may be required for a more complex, user-defined type.

We could therefore write:

```
def del_name(self):
    del self._name
name = property(get_name, fdel=del_name, doc="A name property")
```

Note that we are using a keyword reference for the delete method as we have skipped the setter and cannot therefore rely on positional arguments.

24.5 More Concise Property Definitions

The example shown in the previous section works but it is still quite verbose itself; while this is on the class writers' side it still seems somewhat heavyweight.

To overcome this a more concise option has been available since Python 2.4. This approach uses what are known as decorators. Decorators represent meta data (that is information about your code that the Python interpreter can use to work out what you want it to do with certain things).

Python 2.6 introduced three new decorators `@property`, `@<propertyname>.setter` and `@<propertyname>.deleter`. These decorators are added to the start of a method definition to indicate that the method should be used to provide access to a property (and define that property), define a setter for the property or a deleter for the property.

We will now update our `Person` class to use the decorators:

```
class Person:
    def __init__(self, name, age):
        self._name = name
        self._age = age

    @property
    def age(self):
        """ The docstring for the age property """
        print('In age method')
        return self._age

    @age.setter
    def age(self, value):
        print('In set_age method')
        if isinstance(value, int) & value > 0 & value < 120:
            self._age = value

    @property
    def name(self):
        print('In name')
        return self._name

    @name.deleter
    def name(self):
        del self._name

    def __str__(self):
        return 'Person[' + str(self._name) + '] is ' +
            str(self._age)
```

Notice three important things about this example:

- The name of the methods is no longer a `set_age` and `get_age`; instead both methods are now just `age` and the decorator distinguishes their role. Also notice that we no longer have a separate statement that declares the property—it is now implicit in the use of the `@property` decorator and the name of the associated method.
- The `@property` decorator is used to define the name of the property (in this case `age`) and to define further decorators which will be named after the property with a `.setter` or `.deleter` element e.g. `@age.setter`.
- The documentation string is now defined in the method associated with the `@property` decorator (providing this documentation string is usually considered good practice).

However, we do not need to change the program that used the class `Person`, as the interface to the class remained the same.

24.6 Online Resources

Some online resources on properties are:

- https://www.python-course.eu/python3_properties.php a discussion on Python properties versus getters and setters.
- <https://www.journaldev.com/14893/python-property-decorator> A short introduction to the @property decorator.

24.7 Exercises

In this exercise you will add *properties* to an existing class.

Return to the `Account` class that you created several chapters ago; convert the `balance` into a read only property using decorators, then verify that the following program functions correctly:

```
acc1 = CurrentAccount('123', 'John', 10.05, 100.0)
acc2 = DepositAccount('345', 'John', 23.55, 0.5)
acc3 = acc3 = InvestmentAccount('567', 'Phoebe', 12.45, 'high
risk')

print(acc1)
print(acc2)
print(acc3)

acc1.deposit(23.45)
acc1.withdraw(12.33)
print('balance:', acc1.balance)

print('Number of Account instances created:'
, Account.instance_count)

print('balance:', acc1.balance)
acc1.withdraw(300.00)
print('balance:', acc1.balance)
```

The output from this might be:

```
Creating new Account
Creating new Account
Creating new Account
Account[123] - John, current account = 10.05overdraft limit: -100.0
Account[345] - John, savings account = 23.55interest rate: 0.5
Account[567] - Phoebe, investment account = 12.45
balance: 21.17
Number of Account instances created: 3
balance: 21.17
Withdrawal would exceed your overdraft limit
balance: 21.17
```

Chapter 25

Abstract Base Classes



25.1 Introduction

This chapter presents Abstract Base Classes (also known as ABCs) which were originally introduced in Python 2.6.

An Abstract Base Class is a class that you cannot instantiate and that is expected to be extended by one or more subclassed. These subclasses will then fill in any the gaps left the base class.

Abstract Base Classes are very useful for creating class hierarchies with a high level of reuse from the root class in the hierarchy.

25.2 Abstract Classes as a Concept

An abstract class is a class from which you cannot create an object. It is typically missing one or more elements required to create a fully functioning object.

In contrast a non-abstract (or concrete) class leaves nothing undefined and can be used to create a working object.

You may therefore wonder what use an abstract class is?

The answer is that you can group together elements which are to be shared among a number of classes, without providing a complete implementation. In addition, you can force subclasses to provide specific methods ensuring that implementers of a subclass at least supply appropriately named methods. You should therefore use abstract classes when:

- you wish to specify data or behaviour common to a set of classes, but insufficient for a single instance,
- you wish to force subclasses to provide specific behaviour.

In many cases, the two situations go together. Typically, the aspects of the class to be defined as abstract are specific to each class, while what has been implemented is common to all classes.

25.3 Abstract Base Classes in Python

Abstract Base Classes (or ABCs as they are sometimes referred to) cannot be instantiated themselves but can be extended by subclasses. These subclasses can be concrete classes or can themselves be Abstract Base Classes (that extend the concept defined in the root Abstract Base Class).

Abstract Base Classes can be used to define generic (potentially abstract) behaviour that can be mixed into other Python classes and act as an abstract root of a class hierarchy. They can also be used to provide a more formal way of specifying behaviour that must be provided by a concrete class.

Abstract Base Classes can have:

- Zero or more abstract methods or properties (but are not required to)
- Zero or more concrete methods and properties.
- Both *private* and *protected* attributes (following the single underscore and double underscore conventions).

ABCs can also be used to specify a specific interface or formal protocol. If an ABC defines any abstract methods or abstract properties, then the subclasses must provide implementations for all such abstract elements.

There are many built-in ABCs including (but not limited to):

- Data structures (`collection` module),
- Numbers module,
- Streams (IO module).

In fact, ABCs are widely used internally within Python itself, and many developers use ABCs without ever knowing they exist or understanding how to define them.

Indeed, ABCs are not widely used by developers building systems with Python although this is in part because they are most appropriate to those building libraries, particularly those that are expected to be extended by developers themselves.

25.3.1 *Subclassing an ABC*

Typically, an Abstract Base Class will need to be imported from the module in which it is defined; of course, if the ABC is defined in the current module then this will not be necessary.

As an example, the `collections.MutableSequence` class is an ABC; this is an ABC for a sequence of elements that can be modified (mutable) and iterated

over. We can use this as the base class for our own type of collection which we will call a Bag, for example:

```
from collections import MutableSequence

class Bag(MutableSequence):
    pass
```

In this example we are importing the `MutableSequence` from the `collections` module. We then define the class `Bag` as extending the `MutableSequence` Abstract Base Class. For the moment we are using the special Python keyword `pass` as a place holder for the body of the class.

However, this means that the `Bag` class is actually also an abstract class as it does not implement any of the abstract methods in the `MutableSequence` ABC.

Python, however, does not validate this at *import* time; instead it validates it at *runtime* when an instance of the type is to be created.

In this case, the `Bag` class does not implement the abstract methods in `MutableSequence` and thus if a program attempts to create an instance of `Bag`, then the following error would be raised:

```
Traceback (most recent call last):
  File "/pythonintro/abstract/Bag.py", line 10, in <module>
    main()
    File "/pythonintro/abstract/Bag.py", line 7, in main
      bag = Bag()
TypeError: Can't instantiate abstract class Bag with abstract
methods __delitem__, __getitem__, __len__, __setitem__, insert
```

As can be seen this is a rather more *formal* requirement than we saw with protocols; if you don't implement all the methods defined as abstract in the parent class then you can't create an instance of the class you are defining (because it is also abstract).

We can define a method for each of the abstract classes in the `Bag` class and then we will be able to create an instance of the class, for example:

```
from collections import MutableSequence

class Bag(MutableSequence):

    def __getitem__(self, index):
        pass

    def __delitem__(self, index):
        pass

    def __len__(self):
        pass

    def __setitem__(self, index, value):
        pass

    def insert(self, index, value):
        pass
```

This version of `Bag` meets all the requirements imposed on it by the ABC `MutableSequence`; that is, it implements each of the special methods listed and the `insert` method. The class `Bag` can now be considered to be a concrete class.

However, in this case the methods themselves don't do anything (they again use the Python keyword `pass` which acts as a placeholder for the code to implement each method). However, we can now write:

```
bag = Bag()
```

And the application will not generate an error message.

At this point we could now progress to implementing each method such that it provides an appropriate implementation of the `Bag`.

25.3.2 *Defining an Abstract Base Class*

An Abstract Base Class can be defined by specifying that the class has a metaclass; typically, `ABCMeta`. The `ABCMeta` metaclass is provided by the `abc` module.

The metaclass is specified using the `metaclass` attribute of the parent class list. This will create a class that can be used as an ABC.

This is exactly how the ABCs in the `_collections_abc.py` file are implemented.

The following code snippet illustrates this idea. The `ABCMeta` class is imported from the `abc` module. It is then used with the class `Shape` via the `metaclass` attribute of the class inheritance list:

```
from abc import ABCMeta

class Shape(metaclass=ABCMeta):

    def __init__(self, id):
        self.id = id
```

Note that at this point although `Shape` is an ABC, it does not define any abstract elements and so can actually be instantiated just like any other concrete class. However, we will next define some abstract methods for the `Shape` ABC.

To define an abstract method, we need to also import the `abstractmethod` decorator from the `abc` module (if we want to define an abstract property then we need to add `@property` to an appropriate abstract method). Import the `abstractmethod` decorator is illustrated below:

```
from abc import ABCMeta, abstractmethod
```

We can now expand the definition of our Shape class:

```
from abc import ABCMeta, abstractmethod

class Shape(metaclass=ABCMeta):
    def __init__(self, id):
        self._id = id

    @abstractmethod
    def display(self): pass

    @property
    @abstractmethod
    def id(self): pass
```

The class Shape is now an Abstract Base Class and requires that any subclass must provide an implementation of the method `display()` and the property `id` (otherwise the subclass will automatically become abstract).

The class Circle is a concrete subclass of the Shape ABC; it thus provides a `__init__()` initialization method, a `display` method and an `id` property (the `__init__()` method is used to allow the `_id` attribute in the base class to be initialized).

```
class Circle(Shape):
    def __init__(self, id):
        super().__init__(id)

    def display(self):
        print('Circle: ', self._id)

    @property
    def id(self):
        """ the id property """
        return self._id
```

We can now use the class Circle in an application:

```
c = Circle("circle1")
print(c.id)
c.display()
```

We can instantiate the Circle class as it is concrete and we know we can call the method `display()` and access the property `id` on instances of Circle. The output from the above code is thus:

```
circle1
Circle: circle1
```

Alternatively you can just use the super class ABC instead of the long hand `metaclass = ABCMeta`, for example:

```
from abc import ABC

class Shape(ABC):

    def __init__(self, id):
        self._id = id

    @abstractmethod
    def display(self): pass

    @property
    @abstractmethod
    def id(self): pass
```

This is a more concise form and is typically used in modern Python.

25.4 Defining an Interface

Many languages such as Java and C# have the concept of an *interface* definition; this is a contract between the implementors of an interface and the user of the implementation guaranteeing that certain facilities will be provided.

Python does *not* explicitly have the concept of an interface contract (note here interface refers to the interface between a class and the code that utilizes that class).

However, it does have Abstract Base Classes.

Any Abstract Base Class that only has abstract methods or properties in it can be treated as a contract that must be implemented (it can of course also have concrete methods, properties and attributes; it is up to the developer). However, as we know that Python will guarantee that any instances can only be created from concrete classes, we can treat an ABC as behaving like a contract between a class and those using that class.

This is an approach that is adopted by numerous frameworks and libraries within Python.

25.5 Virtual Subclasses

In most object-oriented programming languages, for one class to be *treated* as a subclass of another class it is necessary for the subclass to *extend* the parent class. However, Python has a more relaxed approach to typing, as is illustrated by the idea of *Duck Typing*.

In some situations however, it is useful to be able to confirm that one type is a subclass of another or that an instance is an instance of a specific type (which may come from the object's class hierarchy) at runtime.

Indeed, in Python, it is not necessary to be an actual subclass of a parent class to be considered a subclass—instead *Virtual* subclasses allow one class to be treated as a subclass of another even though there is no direct inheritance relationship between them. The key here is that the *virtual* subclass must match the required *interface* presented by the virtual parent class.

This is done by registering one class as a Virtual Subclass of an Abstract Base Class. That is, the virtual parent class must be an ABC and then the subclass can be registered (at runtime) as a *virtual* subclass of the ABC. This is done using a method called `register()`.

Once a class is registered as a subclass of an ABC the `issubclass()` and `isinstance()` methods will return `True` for that class with respect to the virtual parent class.

For example, given the following two currently independent classes:

```

from abc import ABC

class Person(ABC):
    def __init__(self, name, age):
        self.name = name
        self.age = age

    def birthday(self):
        print('Happy Birthday')

class Employee(object):
    def __init__(self, name, age, id):
        self.name = name
        self.age = age
        self.id = id

    def birthday(self):
        print('Its your birthday')

```

If we now check to see if `Employee` is a subclass of `Person` we will get the value `False` returned. We will of course also get `False` if we check to see an instance of `Employee` is actually an instance of the class `Person`:

```

print(issubclass(Employee, Person))
e = Employee('Megan', 21, 'MS123')
print(isinstance(e, Person))

```

This will generate the output:

```

False
False

```

However, if we now register the class `Employee` as a virtual subclass of the class `Person`, the two test methods will return `True`:

```

Person.register(Employee)
print(issubclass(Employee, Person))
e = Employee('Megan', 21, 'MS123')
print(isinstance(e, Person))

```

Which now generates the output:

```
True
True
```

This provides a very useful level of flexibility that can be exploited when using existing libraries and frameworks.

25.6 Mixins

A *mixin* is a class that represents some (typically concrete) functionality that has the potentially to be useful in multiple situations but on its own is not something that would be instantiated.

However, a *mixin* can be mixed into other classes and can extend the data and behavior of that type and can access data and methods provided by those classes.

Mixins are a common category of Abstract Base Classes; although they are implicit in their use (and naming) rather than being a concrete construct within the Python language itself.

For example, let us define a `PrinterMixing` class that provides a utility method to be used with other classes. It is not something that we want developers to instantiate itself so we will make it an ABC, but it does not define any abstract methods or properties (so there is nothing for a subclass to have to implement).

```
from abc import ABC

class PrinterMixin(ABC):
    def print_me(self):
        print(self)
```

We can now use this with a class `Employee` that extends the class `Person` and mixes in the `PrinterMixin` class:

```
class Person(object):
    def __init__(self, name):
        self.name = name

class Employee(Person, PrinterMixin):
    def __init__(self, name, age, id):
        super().__init__(name)
        self.age = age
        self.id = id

    def __str__(self):
        return 'Employee(' + self.id + ')' + self.name
+ '[' + str(self.age) + ']'
```

This now means that when we instantiate the class `Employee` we can call the `print_me()` method on the `Employee` object:

```
e = Employee('Megan', 21, 'MS123')
e.print_me()
```

which will print out.

```
Employee(MS123)Megan[21]
```

One point to note about the `PrinterMixin` is that it is completely independent of the class it is mixed into. However, mixins can also impose some *constraints* on the classes they will be mixed into. For example, the `IDPrinterMixin` shown below assumes that the class it will be mixed into has an attribute or property called `id`.

```
class IDPrinterMixin(ABC):
    def print_id(self):
        print(self.id)
```

This means that it cannot be mixed successfully into the class `Person`—if it was then when the `print_id()` method was called an error would be generated.

However, the class `Employee` does have an `id` attribute and thus the `IDPrinterMixin` can be mixed into the `Employee` class:

```
class Employee(Person, PrinterMixin, IDPrinterMixin):

    def __init__(self, name, age, id):
        super().__init__(name)
        self.age = age
        self.id = id

    def __str__(self):
        return 'Employee(' + self.id + ')' + self.name
+ '[' + str(self.age) + ']'
```

Which means that we can now call write:

```
e = Employee('Megan', 21, 'MS123')
e.print_me()
e.print_id()
```

Which will generate:

```
Employee(MS123)Megan[21]
MS123
```

25.7 Online Resources

Some online references for Abstract Base Classes include:

- <https://docs.python.org/3/library/abc.html> The standard Library documentation on Abstract Base Classes.
- <https://pymotw.com/3/abc/index.html> The Python Module of the Week page for Abstract Base Classes.
- <https://www.python.org/dev/peps/pep-3119/> Python PEP 3119 that introduced Abstract Base Classes.

Chapter 26

Error and Exception Handling



26.1 Introduction

This chapter considers exception and error handling and how it is implemented in Python. You are introduced to the object model of exception handling, to the throwing and catching of exceptions, and how to define new exceptions and exception-specific constructs. It also considers exception groups and exception notes.

26.2 Errors and Exceptions

When something goes wrong in a computer program someone needs to know about it. One way of informing other parts of a program (and potentially those running a program) is by generating an error object and propagating that through the code until either something *handles* the error and sorts thing out or the point at which the program is entered is found.

If the error propagates out of the program, then the user who ran the program needs to know that something has gone wrong. They are notified of a problem via a short report on the error that occurred and a stack trace of where that error can be found.

You may have seen these yourself already when writing your own programs. For example, the following screen dump illustrates a programming *error* where someone has tried to concatenate a string and a number together using the ‘+’ operator. This *error* has propagated out of the program and a stack trace of the code that was called is presented (in this case the function `print` used the `__str__()` method of the class `Person`). Note the line numbers are included which helps with *debugging* the problem.

```

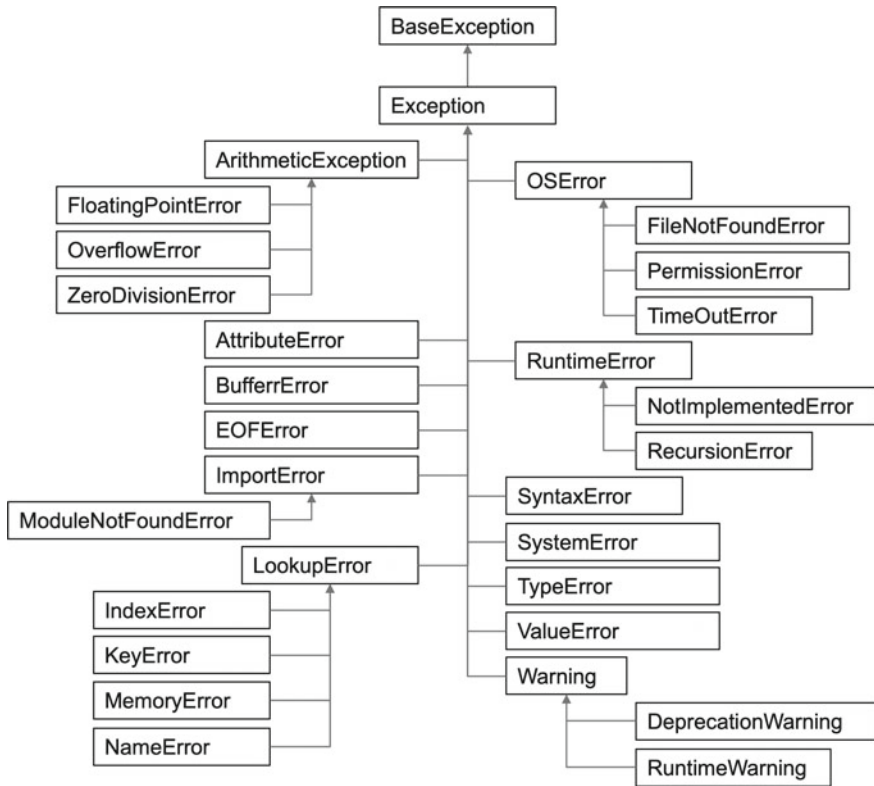
Run: properties
/Users/Shared/workspaces/pycharm/venv/bin/python /Users/Shared/workspaces/pycharm/pythonintro/properties/properties.py
Traceback (most recent call last):
  File "/Users/Shared/workspaces/pycharm/pythonintro/properties/properties.py", line 32, in <module>
    print(person)
  File "/Users/Shared/workspaces/pycharm/pythonintro/properties/properties.py", line 28, in __str__
    return 'Person[' + str(self._name) +'] is ' + self._age
TypeError: Can't convert 'int' object to str implicitly
Process finished with exit code 1

```

In Python the terms error and exception are used inter-changeably; although from a style point of view exceptions might be used to represent issues with operations such as arithmetic exceptions and errors might be associated with functional issues such as a file not being found.

26.3 What is an Exception?

In Python, everything is a type of object, including integers, strings, Booleans and indeed exceptions and errors. In Python the Exception/Error types are defined in a class hierarchy with the root of this hierarchy being the `BaseException` type. All built-in errors and exceptions eventually extend from the `BaseException` type. It has a subclass `Exception` which is the root of all user defined exceptions (as well as many built-in exceptions). In turn `ArithmeticException` is the base class for all built-in exceptions associated with arithmetic errors.



The above diagram illustrates the class hierarchy for some of the common types of errors and exceptions.

When an exception occurs, this is known as *raising* an exception and when it is passed to code to handle this is known as *throwing* an exception. These are terms that will hopefully become more obvious as this chapter progresses.

26.4 What is Exception Handling?

An exception moves the flow of control from one place to another. In most situations, this is because a problem occurs which cannot be handled locally but that can be handled in another part of the system.

The problem is usually some sort of error (such as dividing by zero), although it can be any problem (for example, identifying that the postcode specified with an address does not match). The purpose of an exception, therefore, is to handle an error condition when it happens at runtime.

It is worth considering why you should wish to handle an exception; after all the system does not allow an error to go unnoticed. For example, if we try to divide by

zero, then the system generates an error for you. This may mean that the user has entered an incorrect value, and we do not want users to be presented with a dialog suggesting that they enter the system debugger. We can therefore use exceptions to force the user to correct the mistake and rerun the calculation.

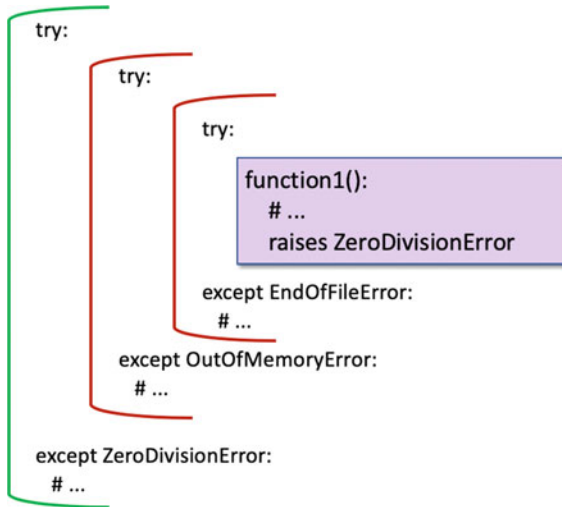
The following table illustrates terminology typically used with exception/error handling in Python.

Exception	An error which is generated at runtime
Raising an exception	Generating a new exception
Throwing an exception	Triggering a generated exception
Handling an exception	Processing code that deals with the error
Handler	The code that deals with the error (referred to as the catch block)
Signal	A particular type of exception (such as <i>out of bounds</i> or <i>divide by zero</i>)

Different types of error produce different types of exception. For example, if the error is caused by dividing an integer by zero, then the exception is a *arithmetic* exception. The type of exception is identified by objects and can be caught and processed by exception handlers. Each handler can deal with exceptions associated with its class of error or exception (and its subclasses).

An exception is instantiated when it is raised. The system searches back up the execution stack (the set of functions or methods that have been invoked in reverse order) until it finds a handler which can deal with the exception. The associated handler then processes the exception. This may involve performing some remedial action or terminating the current execution in a controlled manner. In some cases, it may be possible to restart executing the code.

As a handler can only deal with an exception of a specified class (or subclass), an exception may pass through a number of handler blocks before it finds one that can process it.



The above figure illustrates a situation in which a divide by zero exception called `ZeroDivisionError` is raised. This exception is passed up the execution stack where it encounters an exception handler defined for an end of file exception. This handler cannot handle the `ZeroDivisionError` and so it is passed further up the execution stack. It then encounters a handler for an out of memory exception. Again, it cannot deal with a `ZeroDivisionError`, and the exception is passed further up the execution stack until it finds a handler defined for the `ZeroDivisionError`. This handler then processes the exception.

26.5 Handling an Exception

You can catch an exception by implementing the `try-except` construct. This construct is broken into four parts:

- `try` block. The `try` block indicates the code which is to be monitored for the exceptions listed in the `except` expressions.
- `except` clause. You can use an optional `except` clause to indicate what to do when certain classes of exception/error occur (e.g., resolve the problem or generate a warning message). There can be any number of `except` clauses in sequence checking for different types of error/exceptions.
- `else` clause. This is an optional clause which will be run if and only if no exception was thrown in the `try` block. It is useful for code that must be executed if the `try` clause does not raise an exception.
- `finally` clause. The optional `finally` clause runs after the `try` block exits (whether or not this is due to an exception being raised). You can use it to clean up any resources, close files, etc.

This language construct may at first seem confusing, however, once you have worked with it for a while you will find it less daunting.

As an example, consider the following function which divides a number by zero; this will raise the `ZeroDivisionError` when it is run for any number:

```
def runcalc(x):
    x/0
```

If we now call this function, we will get the error traceback in the standard output:

```
runcalc(6)
```

This is shown below:

```
/Library/Frameworks/Python.framework/Versions/3.7/bin/python3.7 /Users/Shared/workspaces/pycharm/python
Traceback (most recent call last):
  File "/Users/Shared/workspaces/pycharm/pythonintro/exceptions/exceptions.py", line 67, in <module>
    main()
  File "/Users/Shared/workspaces/pycharm/pythonintro/exceptions/exceptions.py", line 39, in main
    runcalc(0)
  File "/Users/Shared/workspaces/pycharm/pythonintro/exceptions/exceptions.py", line 2, in runcalc
    x / 0
ZeroDivisionError: division by zero
```

However, we can handle this by wrapping the call to `runcalc` within a `try` statement and providing an `except` clause. The syntax for a `try` statement with an `except` clause is

```
try:
    <code to monitor>
except <type of exception to monitor for>:
    <code to call if exception is found>
```

A concrete example of this is given below for a `try` statement that will be used to monitor a call to `runcalc`:

```
try:
    runcalc(6)
except ZeroDivisionError:
    print('oops')
```

which now results in the string `'oops'` being printed out. This is because when `runcalc` is called the `'/'` operator throws a the `ZeroDivisionError` which is passed back to the calling code which has an `except` clause specifying this type of exception. This *catches* the exception and runs the associated code block which in this case prints out the string `'oops'`.

In fact, we don't have to be as precise as this; the `except` clause can be given a class of exception to look for and it will match any exception that is of that type or is an instance of a subclass of the exception. We therefore can also write

```
try:
    runcalc(6)
except Exception:
```

```
print('oops')
```

The `Exception` class is a grandparent of the `ZeroDivisionError`, thus any `ZeroDivisionError` object is also a type of `Exception`, and thus the `except` block matches the exception passed. This means that you can write one `except` clause and that clause can handle a whole range of exceptions.

However, if you don't want to have a common block of code handling your exceptions, you can define different behaviours for different types of exception. This is done by having a series of `except` clauses; each monitoring a different type of exception:

```
try:
    runcalc(6)
except ZeroDivisionError:
    print('oops')
except IndexError:
    print('arrgh')
except FileNotFoundError:
    print('huh!')
except Exception:
    print('Duh!')
```

In this case the first `except` monitors for a `ZeroDivisionError` but the other `excepts` monitor for other types of exception. Note that the `except Exception` is the last `except` clause in the list as `ZeroDivisionError`, `IndexError` and `FileNotFoundError` are all eventual subclasses of `Exception`, and thus this clause would catch any of these types of exception. As only one `except` clause is allowed to run; if this `except` handler came first the other `except` handlers would never be run.

26.5.1 Accessing the Exception Object

It is possible to gain access to the exception object being caught by the `except` clause using the `as` keyword. This follows the exception type being monitored and can be used to bind the exception object to a variable, for example:

```
try:
    runcalc(6)
except ZeroDivisionError as exp:
    print(exp)
    print('oops')
```

Which produces

```
division by zero
oops
```

If there are multiple `except` clauses each `except` clause can decide whether to bind the exception object to a variable or not (and each variable can have a different name):

```

try:
    runcalc(6)
except ZeroDivisionError as exp:
    print(exp)
    print('oops')
except IndexError as e:
    print(e)
    print('arrgh')
except FileNotFoundError:
    print('huh')
except Exception as exception:
    print(exception)
    print('Duh!')

```

In the above example three of the four `except` clauses bind the exception to a variable (each with a different name—although they could all have the same name) but one, the `IndexError` `except` clause does not bind the exception to a variable.

26.5.2 *Jumping to Exception Handlers*

One of the interesting features of exception handling in Python is that when an error or an exception is raised it is immediately *thrown* to the exception handlers (the `except` clauses). Any statements that follow the point at which the exception is raised are not run. This means that a function may be terminated early, and further statements in the calling code may not be run.

As an example, consider the following code. This code defines a function `my_function()` that prints out a string and performs a division operation which will cause a `ZeroDivisionError` to be raised if the `y` value is zero, and then it has a further `print` statement. This function is called from within a `try` statement. Notice that there is a `print` statement each side of the call to `my_function()`. There is also a handler for the `ZeroDivisionError`.

```

def my_function(x, y):
    print('my_function in')
    result = x / y
    print('my_function out')
    return result

print('Starting')

```

```

try:
    print('Before my_function')
    my_function(6, 2)
    print('After my_function')
except ZeroDivisionError as exp:
    print('oops')

print('Done')

```

When we run this the output is

```

Starting
Before my_function
my_function in
my_function out
After my_function
Done

```

Which is what would probably be expected; we have run every statement with the exception of the `except` clause as the `ZeroDivisionError` was not raised.

If we now change the call to `my_function()` to pass in 6 and 0 we will raise the `ZeroDivisionError`.

```

print('Starting')

try:
    print('Before my_function')
    my_function(6, 0)
    print('After my_function')
except ZeroDivisionError as exp:
    print('oops')

print('Done')

```

Now the output is

```

Starting
Before my_function
my_function in
oops
Done

```

The difference is that the second `print` statement in `my_function()` has not been run; instead after `print 'my_function in'` and then raising the error we have jumped straight to the `except` clause and run the `print` statement in the associated block of code.

This is partly why the term *throwing* is used with respect to error and exception handling; because the error or exception is raised in one place and thrown to the point where it is handled, or it is thrown out of the application if no `except` clause is found to handle the error/exception.

26.5.3 Catch Any Exception

It is also possible to specify an `except` clause that can be used to catch any type of error or exception, for example:

```
try:
    my_function(6, 0)
except IndexError as e:
    print(e)
except:
    print('Something went wrong')
```

This must be the *last* `except` clause as it omits the exception type and thus acts as a wildcard. It can be used to ensure that you get notified that an error did occur—although you do not know what type of error it actually was; therefore, use this feature with caution.

26.5.4 The Else Clause

The `try` statement also has an optional `else` clause. If this is present, then it must come after all `except` clauses. The `else` clause is executed if and only if no exceptions were raised. If any exception was raised the `else` clause will not be run. An example of the `else` clause is shown below:

```
try:
    my_function(6, 2)
except ZeroDivisionError as e:
    print(e)
else:
    print('Everything worked OK')
```

In this case the output is

```
my_function in
my_function out
Everything worked OK
```

As you can see the `print` statement in the `else` clause has been executed, however, if we change the `my_function()` call to pass in a zero as the second parameter (which will cause the function to raise a `ZeroDivisionError`), then the output is

```
my_function in
division by zero
```

As you can see the `else` clause was not run but the `except` handler was executed.

26.5.5 *The Finally Clause*

An optional `finally` clause can also be provided with the `try` statement. This clause is the last clause in the statement and must come after any `except` classes as well as the `else` clause.

It is used for code that you want to run whether an exception occurred or not. For example, in the following code snippet:

```
try:
    my_function(6, 2)
except ZeroDivisionError as e:
    print(e)
else:
    print('Everything worked OK')
finally:
    print('Always runs')
```

The `try` block will run, if no error is raised then the `else` clause will be executed and last of all the `finally` code will run, we will therefore have as output:

```
my_function in
my_function out
Everything worked OK
Always runs
```

If, however, we pass in 6 and 0 to `my_function()`:

```
try:
    my_function(6, 0)
except ZeroDivisionError as e:
    print(e)
else:
    print('Everything worked OK')
finally:
    print('Always runs')
```

We will now raise an exception in `my_function()` which means that the `try` block will execute, then the `ZeroDivisionError` will be raised, it will be handled by the `except` clause, and then the `finally` clause will run. The output is now:

```
my_function in
division by zero
Always runs
```

As you can see in both cases the `finally` clause is executed.

The `finally` clause can be very useful for general housekeeping type activities such as shutting down or closing any resources that your code might be using, even if an error has occurred.

26.6 Raising an Exception

An error or exception is raised using the keyword `raise`. The syntax of this is

```
raise <Exception/Error type to raise>()
```

For example:

```
def function_bang():
    print('function_bang in')
    raise ValueError('Bang!')
    print('function_bang')
```

In the above function the second statement in the function body will create a new instance of the `ValueError` class and then raise it so that it is thrown allowing it to be caught by any exception handlers that have been defined.

We can handle this exception by writing a `try` block with an `except` clause for the `ValueError` class. For example:

```
try:
    function_bang()
except ValueError as ve:
    print(ve)
```

This generates the output.

```
function_bang in
Bang!
```

Note that if you just want to raise an exception without providing any constructor arguments, then you can just provide the name of the exception class to the `raise` keyword:

```
raise ValueError # short hand for raise ValueError()
```

You can also *reraise* an error or an exception; this can be useful if you merely want to note that an error has occurred and then rethrow it so that it can be handled further up in your application:

```
try:
    function_bang()
except ValueError:
    print('oops')
    raise
```

This will reraise the `ValueError` caught by the `except` clause. Note here we did not even bind it to a variable; however, we could have done this if required.

```
try:
    function_bang()
except ValueError as ve:
    print(ve)
    raise
```


26.7 Defining an Custom Exception

You can define your own errors and exceptions, which can give you more control over what happens in particular circumstances. To define an exception, you create a subclass of the exception class or one of its subclasses.

For example, to define a `InvalidAgeException`, we can extend the `Exception` class and generate an appropriate message:

```
class InvalidAgeException(Exception):
    """ Valid Ages must be between 0 and 120 """
```

This class can be used to explicitly represent an issue when an age is set on a `Person` which is not within the acceptable age range.

We can use this with the class `Person` that we defined earlier in the book; this version of the `Person` class defined age as a property and attempted to validate that an appropriate age was being set:

```
class Person:
    def __init__(self, name, age):
        self._name = name
        self._age = age

    @property
    def age(self):
        """ The docstring for the age property """
        print('In age method')
        return self._age

    @age.setter
    def age(self, value):
        print('In set_age method(', value, ')')
        if isinstance(value, int) & (value > 0 & value <
120):
            self._age = value
        else:
            raise InvalidAgeException(value)

    @property
    def name(self):
        print('In name')
        return self._name

    @name.deleter
    def name(self):
        del self._name

    def __str__(self):
        return 'Person[' + str(self._name) + ']'
is ' + self._age
```

Note that the age setter method now throws an `InvalidAgeException`, so if we write

```

try:
    p = Person('Adam', 21)
    p.age = -1
except InvalidAgeException:
    print('In here')

```

We can capture the fact that an invalid age has been specified.

However, in the exception handler we don't know what the invalid age was. We can of course provide this information by including it in the data held by the `InvalidAgeException`.

If we now modify the class definition such that we provide an initializer to allow parameters to be passed into the new instance of the `InvalidAgeException`:

```

class InvalidAgeException(Exception):
    """ Valid Ages must be between 0 and 120 """

    def __init__(self, value):
        self.value = value

    def __str__(self):
        return 'InvalidAgeException(' + str(self.value)
+ ') '

```

We have also defined a suitable `__str__()` method to convert the exception into a string for printing purposes.

We do of course need to update the setter to provide the value that has caused the problem:

```

@age.setter
def age(self, value):
    print('In set_age method(', value, ')')
    if isinstance(value,int) & (value > 0 & value < 120):
        self._age = value
    else:
        raise InvalidAgeException(value)

```

We can now write

```

try:
    p = Person('Adam', 21)
    p.age = -1
except InvalidAgeException as e:
    print(e)

```

Now if the exception is raised a message will be printed out giving the actual value that caused the problem:

```
In set_age method( -1 )
InvalidAgeException(-1)
```

26.8 Chaining Exceptions

One final feature that can be useful when creating your own exceptions is to chain them to a generic underlying exception. This can be useful when a generic exception is raised, for example, by some library or by the Python system itself, and you want to convert it into a more meaningful application exception.

For example, let us say that we want to create an exception to represent a specific issue with the parameters passed to a function `divide`, but we don't want to use the generic `ZeroDivisionException`, instead we want to use our own `DivideByYWhenZeroException`. This new exception could be defined as

```
class DivideByYWhenZeroException(Exception):
    """ Sample Exception class """
```

And we can use it in a function `divide`:

```
def divide(x, y):
    try:
        result = x /y
    except Exception as e:
        raise DivideByYWhenZeroException from e
```

We have used the `raise` and `from` keywords when we are instantiating the `DivideByYWhenZeroException`. This chains our exception to the original exception that indicates the underlying problem.

We can now call the `divide` method as below:

```
def main():
    divide(6, 0)
```

This produces a Traceback as given below:

```
Traceback (most recent call last):
  File "/Users/Shared/workspaces/pycharm/pythonintro/
exceptions/exceptions.py", line 43, in divide
    result = x /y
ZeroDivisionError: division by zero
The above exception was the direct cause of the following
exception:
Traceback (most recent call last):
  File "/Users/Shared/workspaces/pycharm/pythonintro/
exceptions/exceptions.py", line 136, in <module>
    main()
  File "/Users/Shared/workspaces/pycharm/pythonintro/
exceptions/exceptions.py", line 79, in main
    divide(6, 0)
```



```

ValueError("not a positive value"),
FileNotFoundError("Unknown file
data.txt"),
AttributeError("date"))

raise eg
do_something ()

```

In this example, the `do_something()` function creates an `ExceptionGroup` with the message ‘multiple exceptions occurred’ and a tuple of four exceptions made up of two `ValueErrors`, a `FileNotFoundError` and an `AttributeError`.

Note that an `ExceptionGroup` is part of the `Exception` class hierarchy. It can thus be used within the exception handling mechanism of Python and this it is possible to raise the exception group using the `raise` keyword.

When the `do_something()` function is called it will raise the `ExceptionGroup` and as at the moment this is not handled the program will terminate with a traceback, for example:

```

+ Exception Group Traceback (most recent call last):
  | File "exception_groups_example.py", line 8, in <module>
  |   do_something()
something |   | File "exception_groups_example.py", line 6, in do_
  |   |   raise e
  |   | ExceptionGroup: multiple exceptions occurred (4 sub-
exceptions)
  +-+----- 1 -----
  | ValueError: incorrect value
  +----- 2 -----
  | ValueError: not a positive value
  +----- 3 -----
  | FileNotFoundError: Unknown file data.txt
  +----- 4 -----
  | AttributeError: date
  +-----

```

Note that this is an `ExceptionGroup Traceback` and thus it knows about the exception group and displays the message associated with the group and each of the exceptions listed in the tuple used to create the group in the order they were defined in.

26.10 Exception Group Handling

Exceptions groups can be handled using the standard `try except` mechanism used with traditional exceptions and errors. Thus we can write

```

try:
    do_something()
except ExceptionGroup as e:
    print(e.message)
    print(e.exceptions)

```

If we run this code the output is

```

multiple exceptions occurred
(ValueError('incorrect value'), ValueError('not a positive
value'), FileNotFoundError('Unknown file data.txt'),
AttributeError('date'))

```

And the program does not crash, note that the `exceptions` property contains a tuple of the exceptions.

However, if we now wish to do something specific depending upon whether there was a `ValueError` in the list of exceptions or not we would need to handle the list of exceptions ourselves and see if one of them was a `ValueError`, etc.

In Python 3.11 a new `except*` keyword was added. This keyword can match exceptions that are within an `ExceptionGroup`. We could therefore look to see what exceptions are within the `ExceptionGroup` and perform different actions depending upon the members of the group. For example:

```

try:
    do_something()
except* ValueError as exp_group:
    print('ValueErrors:')
    print(exp_group.exceptions)
except* FileNotFoundError as exp_group:
    print('FileNotFoundErrors:')
    for e in exp_group.exceptions:
        print(e)
except* AttributeError as exp_group:
    print('AttributeError')
    print(exp_group.exceptions)

```

In this example we are handling the `ValueErrors`, `FileNotFoundError` and the `AttributeError` separately. Note that the result of matching on a particular exception/error type is another `ExceptionGroup` containing one or more exceptions in it. Thus the first `except* ValueError` will generate an `ExceptionGroup` with two exceptions in it. In turn the `except* FileNotFoundError` will generate an `ExceptionGroup` with one exception in it and so on.

The output from this program is therefore

```

ValueErrors:
(ValueError('incorrect value'), ValueError('not a positive
value'))
FileNotFoundErrors:
Unknown file data.txt
AttributeError

```

```
(AttributeError('date'),)
```

Note that unlike the traditional try except block more than one except* block can run as we might want to handle each of the exceptions in the group separately but within the same try except* structure.

Also note that we have had to handle all of the exceptions in the exception group. If we had missed one out then at least one exception would not have been handled and will cause the program to terminate. For example, if we did not include the `AttributeError` except* clause then the output of the program would have been

```
ValueErrors:
  (ValueError('incorrect value'), ValueError('not a positive
value'))
FileNotFoundError:
Unknown file data.txt
+ Exception Group Traceback (most recent call last):
  |       File "exception_groups_example.py", line 11, in
<module>
  |           do_something()
  |
  |       File "exception_groups_example.py", line 7, in do_
something
  |           raise eg
  | ExceptionGroup: multiple exceptions occurred (1 sub-
exception)
  +-+----- 1 -----
  | AttributeError: date
  +-----
```

As you can see from this both the `ValueErrors` and the one `FileNotFoundError` have been handled but the `AttributeError` has not, and thus the program terminates and generates traceback for the unhandled exception.

26.11 Exception Notes

Exception notes are another new feature in Python 3.11. Exception notes are described in PEP678 (see <https://peps.python.org/pep-0678/>).

This new feature allows a programmer to add a note to an exception when it is being handled before *rethrowing* it. This avoids the need to create application specific exceptions to allow additional information to be provided.

The note can then be viewed in any error trace back or examined in client code further up the stack. This can be very useful to anyone trying to debug a problem.

To add a note to an exception the `add_note()` method is used for example:

```

def do_something():
    try:
        # code raising a ValueError
        raise ValueError('Value Incorrect')
    except ValueError as err:
        err.add_note('Invalid age')
        raise err

do_something()

```

In this example the initial `ValueError` is handled within the function `do_something()` and a note is added to the error 'Invalid Age'.

When we run this code the output generated is

```

Traceback (most recent call last):
  File "exception_notes_example.py", line 11, in <module>
    do_something()
  File "exception_notes_example.py", line 8, in do_something
    raise err
  File "exception_notes_example.py", line 5, in do_something
    raise ValueError('Value Incorrect')
ValueError: Value Incorrect
Invalid age

```

As you can see at the end of this traceback the note added 'invalid age' is printed after the initial 'Value Incorrect' message.

We can of course handle this error in the same way as any other error or exception, for example:

```

try:
    do_something()
except ValueError as err:
    print(err)
    print(err.__notes__)

```

Running this code generates

```

Value Incorrect
['Invalid age']

```

Note that when you print the `ValueError` instance using `print(err)` you only see the original message ('Value Incorrect'). To access the notes you must use the `err.__notes__` special attribute. This returns a list of any notes that have been added to the exception.

26.12 Generate a Traceback Programmatically

You may have noted that when you handle an exception in your own code, and for example, print out the exception all you get is the message associated with the exception, for example:


```

try:
    print(6 / 0)
except Exception as exp:
    print('oops')
    print(exp)

```

The following code when run generates this output:

```

oops
division by zero

```

Whereas when you allow an exception to escape from your program (i.e., you don't handle it) you get a so called *traceback* which provides information on where the exception happened, etc. This can be very useful when trying to debug a problem in your code.

The `traceback` module provides the `print_exc()` function that can be used to print a traceback associated with the current exception. To use this module all you need to do is to import it (as it comes as part of the Python environment).

For example we can modify our earlier program by importing the `traceback` module and calling the `print_exc()` function within the exception handling code block:

```

import traceback

try:
    print(6 / 0)
except Exception as exp:
    print('oops')
    print(exp)
    traceback.print_exc()

```

Now when we run this program, we will get a traceback printed out when the exception is handled:

```

oops
division by zero
Traceback (most recent call last):
  File "traceback-example.py", line 6, in <module>
    print(6 / 0)
    ~~~^
ZeroDivisionError: division by zero

```

26.13 Online Resources

For more information on Python's errors and exceptions sets:

- <https://docs.python.org/3/library/exceptions.html> The Standard Library documentation for built-in exceptions

- <https://docs.python.org/3/tutorial/errors.html> The Standard Python documentation tutorial on errors and exceptions.
- https://www.tutorialspoint.com/python/python_exceptions.htm An alternative tutorial on Python exception handling.

26.14 Exercises

This exercise involves adding error handling support to the `CurrentAccount` class.

In the `CurrentAccount` class it should not be possible to withdraw or deposit a negative amount.

Define an exception/error class called `AmountError`. The `AmountError` should take the account involved and an error message as parameters.

Next update the `deposit()` and `withdraw()` methods on the `Account` and `CurrentAccount` class to raise an `AmountError` if the amount supplied is negative.

You should be able to test this using:

```
try:
    acc1.deposit(-1)
except AmountError as e:
    print(e)
```

This should result in the exception 'e' being printed out, for example:

```
AmountError (Cannot deposit negative amounts) on Account[123]
- John, current account = 21.17overdraft limit: -100.0
```

Next modify the class such that if an attempt is made to withdraw money which will take the balance below the overdraft limit threshold an error is raised.

The error should be a `BalanceError` that you define yourself. The `BalanceError` exception should hold information on the account that generated the error.

Test your code by creating instances of `CurrentAccount` and taking the balance below the overdraft limit.

Write code that will use `try` and `except` blocks to catch the exception you have defined.

You should be able to add the following to your test application:

```
try:
    print('balance:', acc1.balance)
    acc1.withdraw(300.00)
    print('balance:', acc1.balance)
except BalanceError as e:
    print('Handling Exception')
    print(e)
```

Chapter 27

Python Modules and Packages



27.1 Introduction

Modules and packages are two constructs used in Python to organize larger programs. This chapter introduces modules in Python, how they are accessed, how they are defined, how Python finds modules, etc. It also explores Python packages and sub-packages.

27.2 Modules

A module allows you to group together related functions, classes and code in general. You can think of a module as being like a library of code (although in fact many libraries are themselves composed of several modules as for example a library may have optional or extensions to the core functionality).

It is useful to organize your code into modules when the code either becomes large or when you want to reuse some elements of the code base in multiple projects.

Breaking up a large body of code from a single file helps with simplifying code, maintenance and comprehensibility of code, testing, reuse and scoping code. These are explored below:

- *Simplicity*—Focusing on a subset of an overall problem helps us to develop solutions that work for the subset and can be combined together to solve the overall problem. This means that individual modules can be simpler than the overall solution.
- *Maintenance*—Modules typically make it easier to define logical boundaries between one body of code and another. This means it is easier to see what comprises a module and to verify that the module works appropriately even when modified. It also helps to distinguish one body of code from another so makes it easier to work out where changes should go.


- *Testing*—As one module can be made independent of another module there are less dependencies and cross overs. This means that a module can be tested in isolation and even before other modules, and the overall application, have been written.
- *Reusability*—Defining a function or class on one module means that it is easier to reuse that function or class in another module, as the boundaries between the one module and another are clear.
- *Scoping*—Modules typically also define a namespace, that is a scope within which each function or class is unique. Think of a namespace a bit like a surname; within a classroom there may be several people with the first name ‘John’, but we can distinguish each person by using their full name, for example ‘John Hunt’, ‘John Jones’, ‘John Smith’ and ‘John Brown’; each surname in this example provides a namespace which ensures each John is unique (and can be referenced uniquely).

27.3 Python Modules

In Python a module equates to a file containing Python code. A module can contain.

- Functions
- Classes
- Variables
- Executable code
- Attributes associated with the module such as its name

The name of a module is the name of the file that it is defined in (minus the suffix ‘.py’). For example, the following diagram illustrates a function and a class defined within a file called `utils.py`:



Thus, the `printer()` function and the class `Shape` are defined in the `utils` module. They can be referenced via the name of the `utils` module.

As an example, let us look at a definition for our `utils` module defined in the file `util.py`:

```

"""This is a test module"""
print('Hello I am the utils module')

def printer(some_object):
    print('printer')
    print(some_object)
    print('done')

class Shape:
    def __init__(self, id):
        self._id = id

    def __str__(self):
        return 'Shape - ' + self._id

    @property
    def id(self):
        """ The docstring for the id property """
        print('In id method')
        return self._id

    @id.setter
    def id(self, value):
        print('In set_age method')
        self._id = id

default_shape = Shape('square')

```

The module has a comment which is at the start of the file—this is useful documentation for anyone working with the module. In some cases, the comment at the start of the module can provide extensive documentation on the module such as what it provides, how to use the features in the module and examples that can be used as a reference.

The module also has some executable code (the print statement just after the comment) that will be run when the module is loaded / initialized by Python. This will happen when the module is first referenced in an application.

A variable `default_shape` is also initialized when the module is loaded and can also be referenced outside the module in a similar manner to the module's function and class. These variables can be useful in setting up defaults or predefined data that can be used by developers working with the module.

27.4 Importing Python Modules

27.4.1 Importing a Module


A user-defined module is *not* automatically accessible to another file or script; it is necessary to import the module. Importing a module makes the functions, classes and variables defined in the module visible to the file they are imported into.

For example, to import all the contents of the `utils` module into a file called `my_app.py` we can use:

```
import utils
```

Note that we do not give the file name (i.e. `utils.py`) instead we give the module name to import (which does not include the `.py`).

The result is as shown below:



Once the definitions within the `utils` module are visible inside `my_app.py` we can use them as if they had been defined in the current file. Note that due to scoping the function, class and variable defined within the `utils` module will be prefixed by the name of the module; i.e. `utils.printer` and `utils.Shape` etc.:

```
import import utils

utils.printer(utils.default_shape)
shape = utils.Shape('circle')
utils.printer(shape)
```

When we run the `my_app.py` file we will produce the following output:

```
Hello I am the utils module
printer
Shape - square
done
printer
Shape - circle
```

Notice that the first line of the output is the output from the `print` statement in the `utils` module (i.e., the line `print('Hello I am the utils module')`) this illustrates the ability in Python to define behaviour that will be run

when a module is loaded; typically this might execute some housekeeping code or set up behaviour required by the module. It should be noted that the module is only initialized the first time that it is loaded and thus the executable statements are only run once.

If we forgot to include the `import` statement at the start of the file, then Python would not know to use the `utils` module. We would thus generate an error indicating that the `utils` element is not known: `NameError: name 'utils' is not defined`.

There can be any number of modules imported into a file; these can be both user defined modules and system provided modules. The modules can be imported via separate `import` statements or by supplying a comma separated list of modules:

```
import utils
import support
import module1, module2, module3
```

A common convention is to place all `import` statements at the start of a file; however, this is only a convention and `import` statements can be placed anywhere in a file prior to where the modules features are required. It is a common enough convention, however, that tools such as PyCharm will indicate a style issue if you do not put the imports at the start of the file.

27.4.2 *Importing from a Module*

One issue with the previous example is that we have had to keep referring to the facilities provided by the `utils` module as `utils.< thing of interest >`; which while making it very clear that these features come from the `utils` module, is a little tedious. It is the equivalent of referring to a person via their full name every time we speak to them. However, if they are the only other person in the room then we don't usually need to be so formal; we could just use their first name.

A variant of the `import` statement allows us to import *everything* from a particular module and remove the need to prefix the modules functions or classes with the module name, for example:

```
from <module name> import *
```

Which can be read as `from < module name > import everything` in that module and make it directly available.

```
from utils import *

printer(default_shape)
shape = Shape('circle')
printer(shape)
```

As this shows we can reference `default_shape`, the `printer()` function and the `Shape` class directly. Note the `*` here is often referred to as a wildcard; meaning that it represents everything in the module.

The problem with this form of `import` is that it can result in name clashes as it brings into scope all the elements defined in the `utils` module. However, we may only actually be interested in the `Shape` class; in which case we can choose to only bring that feature in, for example:

```
from utils import Shape
s = Shape('rectangle')
print(s)
```

Now only the `Shape` class has been imported into the file and made directly available.

27.4.3 Importing Aliases

You can even give an alias for an element being imported from a module using the `import` statement. For example, you can alias the whole package:

```
import <module_name> as <alternative_module_name>
```

for example:

```
import utils as utilities
utilities.printer(utilities.default_shape)
```

In this example, instead of referring to the module we are importing as `utils`, we have given it an alias and called it `utilities`.

We can also alias individual elements of a module, for example, a function can be given an alias, as can a variable or a class. The syntax for this is.

```
from <module_name> import <element> as <alias>
```

For example:

```
from utils import printer as myfunc
```

In this case the `printer` function in the `utils` module is being aliased to `myfunc` and will thus be known as `myfunc` in the current file.

We can even combine multiple `from` imports together with some of the elements being imported having aliases:

```
from utils import Shape, printer as myfunc
s = Shape('oval')
myfunc(s)
```


27.4.4 *Hiding Some Elements of a Module*

By default, any element in a module whose name starts with an underbar ('_') is hidden when a wildcard `import` of the contents of a module is performed.

In this way certain named items can be hidden unless they are explicitly imported. Thus, if our `utils` module now included a function:

```
"""This is a test module"""
print('Hello I am the utils module')

# as before

def _special_function():
    print('Special function')
```

And then tried to import the whole module using the wildcard `import` and access `_special_function`:

```
from utils import *
_special_function()
```

We will get an error:

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

However, if we explicitly import the function then we can still reference it:

```
from utils import _special_function
_special_function()
```

Now the code works:

```
Hello I am the utils module
Special function
```

This can be used to hide features that are either not intended to be used externally from a module (developers then use them at their own peril) or to make advanced features only available to those who really want them.

27.4.5 *Importing Within a Function*

In some cases, it may be useful to limit the scope of an `import` to a function; thus, avoiding any unnecessary use of, or name clashes with, special features.

To do this you merely add an `import` into the body of a function, for example:

```
def my_func():
    from util import Shape
    s = Shape('line')
```

In this case the `Shape` class is only accessible within the body of `my_func()`.

27.5 Module Properties

Every module has a set of properties that can be used to find what features it provides, what its name is, what (if any) its documentation string is etc.

These properties are considered special as they all start, and end, with a double underbar ('__'). These are:

- `__name__` the name of the module
- `__doc__` the documentation for the module
- `__file__` the file in which the module was defined.

You can also obtain a list of the contents of a module once it has been imported using the `dir(<module-name >)` function. For example:

```
import utils
print(utils.__name__)
print(utils.__doc__)
print(utils.__file__)
print(dir(utils))
```

Which produces:

```
Hello I am the utils module
utils
This is a test module
utils.py
['Shape', '__builtins__', '__cached__', '__doc__', '__file__', '__loader__', '__name__', '__package__', '__spec__', '__special_function__', 'default_shape', 'printer']
```

Note that the executable `print` statement is still executed as this is run when the module is loaded into Python's current runtime; even though all we do it then access some of the module's properties.

Mostly module properties are used by tools to help developers; but they can be a useful reference when you are first encountering a new module.

27.6 Standard Modules

Python comes with many built-in modules as well as many more available from third parties.

Of particular use is the `sys` module, which contains a number of data items and functions that relate to the execution platform on which a program is running.

Some of the `sys` module's features are shown below, including `sys.path()`, which lists the directories that are searched to resolve a module when an `import`

statement is used. This is a writable value, allowing a program to add directories (and hence modules) before attempting an import.

```
import sys

print('sys.version: ', sys.version)
print('sys.maxsize: ', sys.maxsize)
print('sys.path: ', sys.path)
print('sys.platform: ', sys.platform)
```

which produces the following output on a Mac:

```
sys.version: 3.11.3 (main, Apr 7 2023, 19:25:52) [Clang 14.0.0
(clang-1400.0.29.202)]
sys.maxsize: 9223372036854775807
sys.path: ['/beginnerspython3_2nd/chapter27_modules_
and_packages', '/beginnerspython3_2nd', '/usr/local/Cellar/
python@3.11/3.11.3/Frameworks/Python.framework/Versions/
3.11/lib/python311.zip', '/usr/local/Cellar/python@3.11/
3.11.3/Frameworks/Python.framework/Versions/3.11/lib/
python3.11', '/usr/local/Cellar/python@3.11/3.11.3/Frameworks/
Python.framework/Versions/3.11/lib/python3.11/lib-dynload', '/
Users/Shared/workspaces/pycharm/beginnerspython3_2nd/venv/lib/
python3.11/site-packages']
sys.platform: darwin
```

27.7 Python Module Search Path

One point that we have glossed over so far is how does Python find these modules? The answer is that it uses a special environment variable called the PYTHONPATH. This is a variable that can be set up prior to running Python that tells it where to look for to find any named modules.

It is hinted at in the previous section where the sys module's path variable is printed out. On the machine that this code was run on the output of this variable was:

```
sys.path: ['/beginnerspython3_2nd/chapter27_modules_and_
packages', '/beginnerspython3_2nd', '/usr/local/Cellar/
python@3.11/3.11.3/Frameworks/Python.framework/Versions/
3.11/lib/python311.zip', '/usr/local/Cellar/python@3.11/
3.11.3/Frameworks/Python.framework/Versions/3.11/lib/
python3.11', '/usr/local/Cellar/python@3.11/3.11.3/Frameworks/
Python.framework/Versions/3.11/lib/python3.11/lib-dynload', '/
Users/Shared/workspaces/pycharm/beginnerspython3_2nd/venv/lib/
python3.11/site-packages']
```

This is actually a list of the locations that Python would look in to find a module; these include the PyCharm project holding the modules related to the examples used in this book (which is the current directory), then a top level PyCharm location and a series of Python locations (where locations equal directories on the host machine).

Python will search through each of these locations in turn to find the named module; it will use the first module it finds. The actual search algorithm is:

- The current directory.
- If the module isn't found, Python then searches each directory in the shell variable `PYTHONPATH`.
- If all else fails, Python checks the default path. On Unix / Linux this default path is normally `/usr/local/lib/python/`.

One point to note about this search order is that it is possible to hide a system provided module by creating a user defined one and giving it the same name as the system module; this is because your own module will be found first and will hide the system provided one.

In this case the `PYTHONPATH` will be used to find built in Python modules under the Python installation, while our own user defined modules will be found within the PyCharm project workspace directories.

It is also possible to override the default `PYTHONPATH` variable. It is what is known as an environment variable and so it can be set as a Unix / Linux or Windows operating system environment variable which can then be picked up by your Python environment. The syntax used to set the `PYTHONPATH` depends on whether you are using Windows or Unix/Linux as it is set at the operating system level:

Here is a typical `PYTHONPATH` from a Windows system:

```
set PYTHONPATH = c:\python20\lib;
```

And here is a typical `PYTHONPATH` from a Unix /Linux system:

```
set PYTHONPATH = /usr/local/lib/python
```

27.8 Modules as Scripts

Any Python file is not only a module but also a Python *script*. This means that it can be executed directly if required. For example, the following is the contents of a file called `module1.py`:

```
"""This is a test module"""
print('Hello I am module 1')

def f1():
    print('f1[1]')

def f2():
    print('f2[1]')
```

```
x = 1 + 2
print('x is', x)
f1()
f2()
```

When this file is run directly, or loaded into the Python REPL, then the free-standing code will be executed and the output generated will be:

```
Hello I am module 1
x is 3
f1[1]
f2[1]
```

This looks fine until you try to use `module1` with your own code; the free-standing code will still execute this if we now write:

```
import module1

module1.f1()
```

where you might expect only to see the result of running the `f1()` function from `module1` you actually get:

```
Hello I am module 1
x is 3
f1[1]
f2[1]
f1[1]
```

The first 4 lines are run when `module1` is loaded as they are free standing executable statements within the module.

We can of course remove the free-standing code; but what if we sometimes want to run `module1` as a script and sometimes use it as a module imported into other modules?

In Python we can distinguish between when a file is loaded as a module and when it is being run as a standalone script. This is because Python sets the module property `__name__` to the name of the module when it is being loaded as a module; but if a file is being run as a standalone script (or the entry point of an application) then the `__name__` is set to the string `__main__`. This is partly because historically `main()` has been the entry point for applications in numerous other languages such as C, C++, Java and C#.

We can now determine whether a module is being loaded or run as a script / main application by checking the `__name__` property of the module (which is directly accessible from within a module). If the module is the `__main__` entry point run some code; if it is not then do something else.

For example,

```

"""This is a test module"""
print('Hello I am module 1')

def f1():
    print('f1[1]')

def f2():
    print('f2[1]')

if __name__ == '__main__':
    x = 1 + 2
    print('x is', x)
    f1()
    f2()

```

Now the code within the `if` statement will only be executed when this module is loaded as the starting point for an application / script. Note that the `print` statement at the top of the module will still be executed in both scenarios; this can be useful as it allows set up or initialization behaviour to still be executed where appropriate.

A common pattern, or idiom is to place the code to be run when a file is being loaded directly (rather than as a module) in a function called `main()` and to call that function from within the `if` statement. This helps to clarify which behaviour is intended to run when, thus our module's final version is:

```

"""This is a test module"""
print('Hello I am module 1')

def f1():
    print('f1[1]')

def f2():
    print('f2[1]')

def main():
    x = 1 + 2
    print('x is', x)
    f1()
    f2()

if __name__ == '__main__':
    main()

```

This version is what would now be called *idiomatic* Python or *Pythonic* in style.

27.9 Python Packages


27.9.1 Package Organization

Python allows developers to organize modules together into packages, in a hierarchical structure based on directories.

A package is defined as.

- A *directory* containing one or more Python source files and
- An *optional* source file named `__init__.py`. This file may also contain code that is executed when a module is imported from the package.

For example, the following picture illustrates a package `utils` containing two modules `classes` and `functions`.



In this case the `__init__.py` file contains package level initialization code:

```
print('utils package')
```

The contents of the `__init__.py` file will be run once, the first time either module within the package is referenced.

The `functions` module then contains several function definitions; while the `classes` module contains several class definitions.

We refer to elements of the package relative to the package name—as shown below:

```
from utils.functions import *
f1()
from utils.classes import *
p = Processor()
```

Here we are importing both the `functions` module and the `classes` module from the `utils` package. The function `f1()` is defined in the `functions` module while the `Processor` class is defined within the `classes` module.

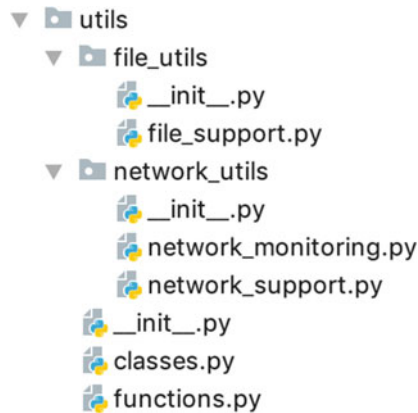
You can use all the *from* and *import* styles we have already seen, for example, you can import a function from a module in a package and give it an alias:

```
from util.functions import f1 as myfunc
myfunc()
```

It is possible to import all the modules from a package by simply importing the package name. If you wish to provide some control over what is imported from a package when this happens you can define a variable `all` in the `__init__.py` file that will indicate what will be imported in this situation.

27.9.2 Sub Packages

Packages can contain sub packages to any depth you require. For example, the following diagram illustrates this idea:



In the above diagram `utils` is the root package, this contains two sub packages `file_utils` and `network_utils`. The `file_utils` package has an initialization file and a `file_support` module. The `network_utils` package also has a package initialization file and two modules `network_monitoring` and `network_support`. Of course, the root package also has its own initialization file and its own modules `classes` and `functions`.

To import the sub package, we do the same as before but each package in the *path* to the module is separate by a dot, for example:

```
import utils.file_utils.file_support
```

or

```
from utils.file_utils.file_support import file_logger
```


27.10 Online Resources

See the Python Standard Library documentation for:

- <https://docs.python.org/3/tutorial/modules.html#standard-modules> the standard modules.
- <https://docs.python.org/3/tutorial/modules.html#packages> packages.
- <https://docs.python.org/3/tutorial/stdlib.html> brief tour of the standard library part 1.
- <https://docs.python.org/3/tutorial/stdlib2.html> brief tour of the standard library part 2.
- <https://pymotw.com/3/> The Python Module of the Week site listing very many modules and an extremely useful reference.

27.11 Exercise

The aim of this exercise is to create a module for the classes you have been developing.

You should move your `Account`, `CurrentAccount`, `DepositAccount` and `BalanceError` classes into a separate module (file) called `accounts`. Save this file into a new Python package called `fintech`.

Separate out the test application from this module so that you can import the classes from the package.

Your test application will now look like:

```
import fintech.accounts as accounts

acc1 = accounts.CurrentAccount('123', 'John', 10.05, 100.0)
acc2 = accounts.DepositAccount('345', 'John', 23.55, 0.5)
acc3 = accounts.InvestmentAccount('567', 'Phoebe', 12.45, 'high
risk')

print(acc1)
print(acc2)
print(acc3)

acc1.deposit(23.45)
acc1.withdraw(12.33)
print('balance:', acc1.balance)

print('Number of Account instances created:',
accounts.Account.instance_count)

try:
    print('balance:', acc1.balance)
    acc1.withdraw(300.00)
    print('balance:', acc1.balance)
except accounts.BalanceError as e:
    print('Handling Exception')
    print(e)
```

You could of course also use `from accounts import *` to avoid prefixing the accounts related classes with `accounts`.

Chapter 28

Protocols, Polymorphism and Descriptors



28.1 Introduction

In this chapter we will explore the idea of an implicit contract between an object and the code that uses that object. As part of this discussion we will explore what is meant by Duck Typing. Following this we will introduce the Python concept called a *protocol*. We will explore its role within Python programming and look at two commonly occurring protocols; the Context Manager Protocol and the Descriptor Protocol.

28.2 Implicit Contracts

Some programming languages (most notable Java and C#) have the idea of an explicit contract between a class and the user of that class; this contract provides a guarantee of the methods that will be provided and the types that will be used for parameters and return values from these methods. In these languages it helps to guarantee that a method is only called with the appropriate type of values and only in appropriate situations. Slightly confusingly these contracts are referred to as *interfaces* in Java and C# but are intended to describe the application programming *interface* presented by the class.

Python is a much more flexible and free flowing language than either Java or C# and so has no explicit concept of an *interface*. This can however make things more complex at times; for example, consider the very simple `Calculator` class given below:

```
class Calculator:
    def add(self, x, y):
        return x + y
```

What are the valid values that can be passed into the `add` method and used for the parameters `x` and `y`?

Initially it might seem that numeric values such as 1, 2 and 3.4, 5.77, etc., would be the only things that can be used with the add method:

```
calc = Calculator()

print('calc.add(3, 4):', calc.add(3, 4))
print('calc.add(3, 4.5):', calc.add(3, 4.5))
print('calc.add(4.5, 6.2):', calc.add(4.5, 6.2))
print('calc.add(2.3, 7):', calc.add(2.3, 7))
print('calc.add(-1, 4):', calc.add(-1, 4))
```

This generates the following output:

```
calc.add(3, 4): 7
calc.add(3, 4.5): 7.5
calc.add(4.5, 6.2): 10.7
calc.add(2.3, 7): 9.3
calc.add(-1, 4): 3
```

However, this actually represents a contract that the values passed in to a calculator method will support the *plus* operator. In an earlier chapter we explored a class `Quantity` which implemented this operator (among others) and this we can also use quantity objects with the calculator's `add()` method:

```
q1 = Quantity(5)
q2 = Quantity(10)
print(calc.add(q1, q2))
```

Which prints out

```
Quantity[15]
```

This implied contract says that the `Calculator.calc()` method will work with anything that supports the numeric add operator (or to put it another); anything that is numeric like.

This is also known as *Duck Typing*; this is described in the next section.

28.3 Duck Typing

This rather strange term comes from an old saying:

if it walks like a duck, swims like a duck and quacks like a duck then it's a Duck!

In Python *Duck Typing* (also known as shape typing or structural typing) implies that if a object can perform the required set of operations that it's a suitable thing to use for what you want. For example, if your type can be used with the add, multiply, divide and subtract operators than it can be treated as a Numeric type (even if it isn't).

This is a very powerful feature of Python and allows code originally written to work with a specific set of types, to also be used with a completely new set of types; as long as they meet the *implicit* contract defined within the code.

It is also interesting to note, that a particular set of methods may have a *super set* of requirements on a particular type, but you only need to implement as much as is required for the functionality you will actually use.

For example, let us modify the `Calculator` class a bit and add some more methods to it:

```
class Calculator:
    """ Simple Calculator class """
    def add(self, x, y):
        return x + y

    def subtract(self, x, y):
        return x - y

    def multiply(self, x, y):
        return x * y

    def divide(self, x, y):
        return x/y
```

At first sight this may indicate that anything being used with the `Calculator` must implement all four operators `+`, `-`, `/` and `*`. However, this is only true if you need to execute all four of the methods defined in the class.

For example, consider the type `Distance`:

```
class Distance:
    def __init__(self, d):
        self.value = d
    def __add__(self, other):
        return Distance(self.value + other.value)
    def __sub__(self, other):
        return Distance(self.value - other.value)
    def __str__(self):
        return 'Distance[' + str(self.value) + ']'
```

This defines a class that implements only the `__add__()` and `__sub__()` methods and will thus only support the `+` and `-` operators.

Can instances of `Distance` be used with the `Calculator` class? The answer is that they can but only with the `add` and `subtract` methods (as they only meet part of the implied contract between the `Calculator` class and any types used with that class).

We can thus write

```
d1 = Distance(6)
d2 = Distance(3)
print(calc.add(d1, d2))
print(calc.subtract(d1, d2))
```

And obtain the output:

```
Distance[9]
Distance[3]
```

However, if we try to use the `multiply()` or `divide()` methods, we will get an error, for example:

```
print(calc.divide(d1, d2))
```

Then we will get an error stack trace:

```
Traceback (most recent call last):
  File "Calculator.py", line 46, in <module>
    print(calc.divide(d1, d2))
  File "Calculator.py", line 15, in divide
    return x / y
TypeError: unsupported operand type(s) for /: 'Distance' and
'Distance'
```

You may well have already encountered these while exploring Python. Basically, it is telling you that the operator `/` is not supported when used with the `Distance` type.

28.4 Polymorphism

Polymorphism is the ability to send the same message (request to run a method) to different objects, each of which appear to perform the same function. However, the way in which the message is handled depends on the object's class.

Polymorphism is a strange sounding word, derived from Greek, for a relatively simple concept. It is essentially the ability to request that the same operation be performed by a wide range of different types of things. How the request is processed which depends on the thing that receives the request. The programmer need not worry about how the request is handled, only that it is. This is illustrated below.

```
def night_out(p):
    p.eat()
    p.drink()
    p.sleep()
```

In this example, the parameter passed into the `night_out()` function expects to be given something that will respond to the methods `eat()`, `drink()` and `sleep()`. Any object that meets this requirement can be used with the function.

We can define multiple classes that meet this informal contract, for example we can define a class hierarchy that provides these methods, or completely separate classes that implement the methods. In the case of the class hierarchy the methods may or may not override those from the parent class.

Effectively, this means that you can ask many different things to perform the same action. For example, you might ask a range of objects to provide a printable string describing themselves. In fact in Python this is exactly what happens. For example, if you ask an instance of a `Manager` class, a compiler object or a database object to return such a string, you use the same method name call (`__str__()`, in Python).

The name polymorphism is unfortunate and often leads to confusion. It makes the whole process sound rather grander than it actually is.

Note this is one of the most significant and flexible features of Python; it does not tie a variable to a specific type; instead via Duck Typing as long as the object provided meets the implied contract, then we are good.

The following classes all meet the contract implied by the `night_out()` function:

```
class Person:
    def eat(self): print('Person - Eat')
    def drink(self): print('Person - Drink')
    def sleep(self): print('Person - Sleep')

class Employee(Person):
    def eat(self): print('Employee - Eat')
    def drink(self): print('Employee - Drink')
    def sleep(self): print('Employee - Sleep')

class SalesPerson(Employee):
    def eat(self): print('SalesPerson - Eat')
    def drink(self): print('SalesPerson - Drink')

class Dog:
    def eat(self): print('Dog - Eat')
    def drink(self): print('Dog - Drink')
    def sleep(self): print('Dog - Sleep')
```

This means that instances of all of these classes can be used with the `night_out()` function.

Note that the `SalesPerson` class meets the implied contract partly via inheritance (the `sleep()` method is inherited from `Employee`).

28.5 Protocols

As mentioned above, Python does not have any formal mechanism for stating what is required between the supplier of some functionality and the user or consumer of that functionality. Instead a far less formal approach termed *Duck Typing* is adopted instead.

This, however, raises the question; how do you know what is required? How do you know that you must provide the numeric operators for an object to be used with the `Calculator` class?

The answer is that a concept known as a *protocol* is used.

A protocol is an informal description of the programmer interface provided by something in Python (for example a class but it could also be a module or a set of stand-alone functions).

It is defined solely via documentation (and thus the `Calculator` class should have a class documentation string defining its protocol).

Based on the information provided by the protocol if a function or method requires an object to provide a specific operation (or method) then if it all works great; if not an error will be thrown, and the type is not compatible.

It is one of the key elements in Python which allows the Object-Oriented concept of *Polymorphism* to operate.

28.6 The Length Protocol

There are many types that can be used with the `len()` function to determine their length, for example strings, tuples, lists, dictionaries can all be used with the `len()` function. How is it that this works? This is where a protocol comes in; there is the length (or `__len__`) protocol. This protocol is used by the `len()` function. It assumes that whatever has been given to it will support the `__len__()` protocol, in other words, that whatever it is given has a method on it called `__len__()`. This means we can write

```
name = 'John' # string
data = [1, 2, 3, 4] # list
info = ('a', 'b', 'c') # tuple
cities = cities = {'Wales': 'Cardiff', 'England': 'London'} #
Dictionary

print(f'len(name): {len(name)}')
print(f'len(data): {len(data)}')
print(f'len(info): {len(info)}')
print(f'len(cities): {len(cities)}')
```

The output from this code block is

```
len(name): 4
len(data): 4
len(info): 3
len(cities): 2
```

However, if we create a simple class `Bag` and try and find out its length:

```
class Bag:
    def __init__(self, owner, items):
        self.owner = owner
        self.items = items

bag1 = Bag('John', [])
```



```
print(f'len(bag1): {len(bag1)}')
```

Then when we run this code we will be too that there is a `TypeError` and that the object does not have (support) the `len()` function:

```
Traceback (most recent call last):
  File "len_protocol_example.py", line 18, in <module>
    print(f'len(bag1): {len(bag1)}')
                                ^^^^^^^^^^^
TypeError: object of type 'Bag' has no len()
```

To overcome this and make the custom `Bag` type (class) work with the `len()` function it is necessary to implement the `__len__()` function. The revised class `Bag` is given below with a `__len__()` function that uses the length of the items list as the length of the `Bag`:

```
class Bag:
    def __init__(self, owner, items):
        self.owner = owner
        self.items = items

    def __len__(self):
        return len(self.items)
```

Now when we rerun the code the output of the print statement is

```
len(bag1): 0
```

And no error is raised.

28.7 The Sequence Protocol

There are numerous commonly occurring protocols that can be found in Python. For example, there is a protocol for defining sequences, such as a container that can be accessed an item at a time.

This protocol requires that any type that will be held in the container must provide the `__len__()` and `__getitem__()` methods.

Thus any class that implements these two methods meets the requirements of the protocol.

However, because protocols are informal and unenforced in Python it is not actually necessary to implement all the methods in a protocol. For example, if it is known that a class will only be used with iteration then it may only be necessary to implement the `__getitem__()` method.

28.8 The Context Manager Protocol

Another concrete example of a protocol is that of the *Context Manager Protocol*. This protocol was introduced in Python 2.5 so is very well established by now.

It is associated with the `with as` statement. This statement is typically used with classes which will need to allocate and release so called *resources*.

These resources could be files, or database connections, etc. In each of these cases a connection needs to be made (for example to a file or a database) before the associated object can be used.

However, the connection should then be closed and released before we finish using the object. This is because dangling connection to things such as files and databases can hang around and cause problems later on (typically only a limited number of concurrent connections are allowed to a file or a database at one time and if they are not closed properly a program can run out of available connections).

The `with as` statement ensures that any setup steps are performed before an object is available for use and that any shut down behaviour is invoked when it is finished with.

The syntax for the use of the `with as` statement is

```
with <managed object> as <localname>:
    Code to use managed object via <localname>
```

For example:

```
with ContextManagedClass() as cmc:
    print('In with block', cmc)
    print('Existing')
```

Note that in this case the object referenced by `cmc` is only in scope within the lines indented after the `with as` statement; after this the `cmc` variable is no longer accessible.

How does this work? In fact what the `with as` statement does is to call a special method when the `with as` statement is entered (just after the `:` above); this method is the `__enter__()` method. It then also calls another special method just as the `with as` statement is exited (just after the last indented statement). This second method is the `__exit__()` method.

- The `__enter__()` method is expected to do any setup/resource allocation/making connections, etc. It is expected to return an object that will be used within the block of statements that form that `with as` statement. It is common to return `self` although it is not a requirement to do so (this flexibility allows the `< managed object >` to act as a factory for other objects if required).
- The `__exit__()` method is called on the `< managed object >` and is passed information about any exceptions that might have been generated during the body of the `with as` statement. Note that the `__exit__()` method is called whether an exception has been thrown or not. The `__exit__()` method returns a `bool`, if it returns `True` then the exception is swallowed (that is it is suppressed and

not passed onto the calling code). If it returns `False` then the exception is also passed back to whatever code called the `'with as'` statement.

A class that can be used with the `'with as'` statement (that meets the *Context Manager Protocol* requirements) is given below:

```
class ContextManagedClass(object):
    def __init__(self):
        print('__init__')

    def __enter__(self):
        print('__enter__')
        return self

    # Args exception type, exception value and traceback
    def __exit__(self, *args):
        print('__exit__:', args)
        return True

    def __str__(self):
        return 'ContextManagedClass object'
```

The above class implements the Context Manager Protocol in that it defines both the `__enter__()` method and the `__exit__()` method.

We can now use this class with the `with as` statement:

```
print('Starting')

with ContextManagedClass() as cmc:
    print('In with block', cmc)
    print('Exiting')

print('Done')
```

The output from this is

```
Starting
__init__
__enter__
In with block ContextManagedClass object
Exiting
__exit__: (None, None, None)
Done
```

From this you can see that the `__enter__()` method is called before the code in the block and `__exit__()` is called after the code in the block.

28.9 The Descriptor Protocol

Another protocol that you may come across is the *descriptor* protocol that was introduced back in Python 2.2.

Descriptors can be used to create what are known as *managed attributes*. A managed attribute is an object attribute that is managed (or protected) from direct access by external code via the *descriptor*. The *descriptor* can then take whatever action is appropriate such as validating the data, checking the format, logging the action, updating a related attribute, etc.

The descriptor protocol defines four methods (as usual they are considered special methods and thus start with a double underbar ‘__’):

- `__get__(self, instance, owner)` This method is called when the value of an attribute is accessed. The `instance` is the instance being modified, and the `owner` is the class defining the object. This method should return the (computed) attribute value or raise an `AttributeError` exception.
- `__set__(self, instance, value)` This is called when the value of an attribute is being set. The parameter `value` is the new value being set.
- `__delete__(self, instance)` Called to delete the attribute.
- `__set_name__(self, owner, name)` Called at the time the owning class `owner` is created. The descriptor has been assigned to `name`. This method was added to the protocol in Python 3.6.

The following class `Logger` implements the *Descriptor* protocol. It can therefore be used with other classes to log the creation, access and update of whatever attribute it is defined for.

```
class Logger(object):
    """ Logger class implementing the descriptor protocol """
    def __init__(self, name):
        self.name = name

    def __get__(self, inst, owner):
        print('__get__:', inst, 'owner', owner,
              ', value', self.name, '=',
              str(inst.__dict__[self.name]))
        return inst.__dict__[self.name]

    def __set__(self, inst, value):
        print('__set__:', inst, '-', self.name, '=',
              value)
        inst.__dict__[self.name] = value

    def __delete__(self, instance):
        print('__delete__', instance)

    def __set_name__(self, owner, name):
        print('__set_name__', 'owner', owner, 'setting',
              name)
```

Each of the methods defined for the protocol prints out a message so that access can be monitored.

The `Logger` class is used with the following `Cursor` class.

```
class Cursor(object):
    # Set up the descriptors at the class level
    x = Logger('x')
    y = Logger('y')

    def __init__(self, x0, y0):
        # Initialize the attributes
        # Note use of __dict__ to avoid using self.x
notation
        # which would invoke the descriptor behaviour
        self.__dict__['x'] = x0
        self.__dict__['y'] = y0

    def move_by(self, dx, dy):
        print('move_by', dx, ',', dy)
        self.x = self.x + dx
        self.y = self.y + dy

    def __str__(self):
        return 'Point[' + str(self.__dict__['x']) +
            ', ' + str(self.__dict__['y']) + '']'
```

There are several points to note about this class definition including:

The *Descriptors* must be defined at the class level not at the object/instance level. The `x` and `y` attributes of `Cursor` object are defined as having `Logger` descriptors within the class (not within the `__init__()` method). If you try to define them using `self.x` and `self.y` the descriptors will not be registered.

The `Cursor` `__init__()` method uses the `__dict__` dictionary to initialize the instance/object attributes `x` and `y`. This is an alternative approach to accessing an objects' attributes; it is used internally by an object to hold the actual attribute values. It by passes the normal attribute look up mechanism invoked when you use the *dot* notation (such as `curser.x = 10`). This means that it will not be intercepted by the Descriptor. This has been done because the logger uses the `__str__()` method to print out the instance holding the attribute which uses the current values of `x` and `y`. When the value of `x` is initially set there will be no value for `y` and thus an error would be generated by the `__str__()`.

The `__str__()` method also uses the `__dict__` dictionary to access the attributes as it is not necessary to log this access. It would also become recursive if the `Logger` also used the method to print out the instance.

We can now use instances of the `Cursor` object without knowing that the descriptor will intercept access to the attributes `x` and `y`:

```

cursor = Cursor(15, 25)
print('-' * 25)

print('p1:', cursor)
cursor.x = 20
cursor.y = 35
print('p1 updated:', cursor)
print('p1.x:', cursor.x)
print('-' * 25)

cursor.move_by(1, 1)
print('-' * 25)

del cursor.x

```

The output from this illustrates how the descriptors have intercepted access to the attributes. Note that the `move_by()` method accesses both the getter and the setter descriptor methods as this method reads the current value of the attributes and then updates them.

```

__set_name__ owner <class '__main__.Cursor'> setting x
__set_name__ owner <class '__main__.Cursor'> setting y
-----
p1: Point[15, 25]
__set__: Point[15, 25] - x = 20
__set__: Point[20, 25] - y = 35
p1 updated: Point[20, 35]
__get__: Point[20, 35] owner <class '__main__.Cursor'>, value
x = 20
p1.x: 20
-----
move_by 1, 1
__get__: Point[20, 35] owner <class '__main__.Cursor'>, value
x = 20
__set__: Point[20, 35] - x = 21
__get__: Point[21, 35] owner <class '__main__.Cursor'>, value
y = 35
__set__: Point[21, 35] - y = 36
-----
__delete__ Point[21, 36]

```

28.10 Online Resources

The following online resources focusing on Python protocols are available:

- https://ref.readthedocs.io/en/latest/understanding_python/interfaces/existing_protocols.html Documentation on Python's default (native) protocols including the comparison, hash, attribute access and sequence protocols.
- <https://docs.python.org/3/library/stdtypes.html#context-manager-types> for context manager types.

- <https://pymotw.com/3/contextlib/index.html> The Python module of the week for context manager utilities.
- [https://en.wikipedia.org/wiki/Polymorphism_\(computer_science\)](https://en.wikipedia.org/wiki/Polymorphism_(computer_science)) Wikipedia page on Polymorphism.

28.11 Exercises

This exercise involves implementing the Context Manager Protocol.

Return to your Account related classes.

Modify the Account class such that it implements the *Context Manager* Protocol.

This means that you will need to implement the `__enter__()` and `__exit__()` methods.

Place `print` messages within the methods so that you can see when they are run.

The new methods you have defined will be inherited by each of the subclasses you have created; namely `CurrentAccount`, `DepositAccount` and `InvestmentAccount`.

Now test out your modified calculator using:

```
with accounts.CurrentAccount ('891', 'Adam', 5.0, 50.0) as
acc:
    acc.deposit(23.0)
    acc.withdraw(12.33)
    print(acc.balance)
```

Which should produce output similar to:

```
Creating new Account
__enter__
15.5
__exit__: (None, None, None)
```

Chapter 29

Decorators



29.1 Introduction

The idea behind Decorators comes from the *Gang of Four* Design Patterns book (so-called as there were four people involved in defining these design patterns). In this book numerous commonly occurring object oriented design patterns are presented. One of these design patterns is the Decorator design pattern.

The Decorator pattern addressed the situation where it is necessary to add additional behaviour to specific objects. One way to add such additional behaviour is to decorate the objects created with types that provide the extra functionality. These decorators wrap the original element but present exactly the same interface to the user of that element. Thus the Decorator Design pattern extends the behaviour of an object without using sub classing. This decoration of an object is transparent to the decorators' clients.

In Python Decorators are functions that take another function (or callable object) and return a third function representing the decorated behaviour.


This chapter introduces decorators, how they are defined, how they are used and presents built-in decorators.

29.2 What Are Decorators?

A Decorator is a piece of code, that is used to mark a callable object (such as a function, method, class or object) typically to enhance or modify its behaviour (potentially replacing it). It thus *decorates* the original behaviour.

Decorators are in fact callable objects themselves and as such behave more like macros that can be applied to callable objects that then return a new callable object (typically a new function).

The basic idea is illustrated in the following diagram:



This diagram illustrates a decorator wrapping a callable object (in this case a function). Note that the decorator presents exactly the same interface to the user of the decorator as the original function would present; that is, it takes the same parameters and either returns nothing (`None`) or something.

It should also be noted that the decorator is also at liberty to completely replace a callable object rather than to wrap it; it's a design decision made by the implementor of the decorator.

29.3 Defining a Decorator

To define a decorator you need to define a callable object such as a function that takes another function as a parameter and returns a new function.

An example of the definition of a very simple logger decorator function is given below.

```
def logger(func):
    def inner():
        print('calling ', func.__name__)
        func()
        print('called ', func.__name__)
    return inner
```

In this case the logger decorator wraps the original function within a new function, here called `inner`. When this function is executed, a statement is logged before and after the original function is executed.

Every function has an attribute `__name__` that provides the functions *name* and this is used in the `inner()` function above to printout the actual function about to be invoked.

Note that the function `inner()` is defined inside the `logger()` function (this is completely legal). A reference to the `inner()` function is then returned as the result of the `logger()` function. The `inner()` function is not executed above!

29.4 Using Decorators

To see what the effect of applying a decorator is; it is useful to explore the basic (explicit) approach to it use. This can be done by defining a function (we will call `target`) that prints out a simple message:

```
def target():
    print('In target function')
```

We can explicitly apply the `logger` decorator to this function by passing the reference to the `target` function (without invoking it), for example:

```
t1 = logger(target)
t1()
```

When we run this code, we actually execute the `inner()` function which was returned by the decorator. This function in turn prints out a message and then calls the function passed into the `logger`. Once this passed in function has executed, it prints another message. The effect of executing the `t1()` function is this to call the `inner()` function which calls the `target` function, thus printing out:

```
calling target
In target function
called target
```

This illustrates what happens when a decorator style function is *executed*.

Python provides some syntactic sugar that allows the definition of the function and the association with the decorator to be declared together using the '@' syntax, for example:

```
@logger
def target():
    print('In target function')

target()
```

This has the same effect as passing `target` into the `logger` function; but illustrates the role of the `logger` in a rather more Pythonic way. It is thus the more common use of Decorators.

The output of this function is the same as the previous version:

29.5 Functions with Parameters

Decorators can be applied to functions that take parameters; however the decorator functions must also take these parameters as well.

For example, if you have a function such as

```
@logger
```

```
def my_func(x, y):
    print(x, y)

my_func(4, 5)
```

Then the decorator used with this function must also take two parameters, for example:

```
def logger(func):
    def inner(x, y):
        print('calling ', func.__name__, 'with', x, 'and',
y)
        func(x, y)
        print('returned from ', func.__name__)
    return inner
```

29.6 Stacked Decorators

Decorators can be stacked; that is more than one decorator can be applied to the same callable object. When this occurs, each function is wrapped inside another function; this idea is illustrated by the following code:

```
# Define the decorator functions
def make_bold(fn):
    def makebold_wrapped():
        return "<b>" + fn() + "</b>"
    return makebold_wrapped

def make_italic(fn):
    def makeitalic_wrapped():
        return "<i>" + fn() + "</i>"
    return makeitalic_wrapped

# Apply decorators to function hello
@make_bold
@make_italic
def hello():
    return 'hello world'


# Call function hello
print(hello())
```

In this example, the function `hello()` is marked with two decorators, `@make_bold` and `@make_italic`.

This means that the function `hello()` is first passed into the `make_italic()` function and wrapped by the `makeitalic_wrapped` function. This function is then returned from the `make_italic` decorator.

The `makeitalic_wrapped` is then passed into the `make_bold()` function which then wraps it inside the `makebold_wrapped` function, which is returned by the `make_bold` decorator.

This means that the function invoked when `hello()` is called is the `makebold_wrapped` function which calls two further functions as shown below:



The end result is that the string returned by the function passed in is first wrapped by `<i>` and `</i>` (indicating italics) and then by `` and `` (indicating bold) in HTML.

Thus the output from `print(hello())` is:

```
<b><i>hello world</i></b>
```

29.7 Parameterized Decorators

Decorators can also take parameters; however, the syntax for such decorators is a little different; there is essentially an extra layer of indirection. The decorator function takes one or more parameters and returns a function that can use the parameter and takes the callable object that is being wrapped. For example:

```
def register(active=True):
    def wrap(func):
        def wrapper():
            print('Calling ', func.__name__, ' decorator param
', active)

            if active:
                func()
                print('Called ', func.__name__)
            else:
                print('Skipped ', func.__name__)
        return wrapper
    return wrap

@register()
def func1():
    print('func1')

@register(active=False)
```

```
def func2():
    print('func2')

func1()
print('-' * 10)
func2()
```

In this example, the wrapped function will only be called if the active parameter is `True`. This is the default and so for `func1()` it is not necessary to specify the parameter (although note that it is now necessary to provide the round brackets).

For `func2()` the `@register` decorator is defined with the active parameter set to `False`. This means that the wrapper function will not call the provided function.

Note that the usage of the decorator only differs in the need to include the round brackets even if no parameters are being specified; even though there are now two inner functions defined within the `register` decorator.

29.8 Method Decorators

29.8.1 *Methods Without Parameters*

It is also possible to decorate methods as well as functions (as they are also callable objects). However, it is important to remember that methods take the special parameter `self` as the first parameter which is used to reference the object that the method is being applied to. It is therefore necessary for the decorator to take this parameter into account, i.e., the *inner* wrapped function must take at least one parameter representing `self`:

```
def pretty_print(method):
    def method_wrapper(self):
        return "<p>{0}</p>".format(method(self))

    return method_wrapper
```

The `pretty_print` decorator defines an inner function that takes as its first (and in this case only) parameter the reference to the object (which by convention uses the parameter `self`). This is then passed into the actual method when it is called.

The `pretty_print` decorator can now be used with any method that only takes the `self` parameter, for example:

```
class Person:
    def __init__(self, name, surname, age):
        self.name = name
        self.surname = surname
        self.age = age

    def print_self(self):
```

```

        print('Person - ', self.name, ', ', self.age)

    @pretty_print
    def get_fullname(self):
        return self.name + " " + self.surname

```

In the above class the `get_fullname()` method is decorated with `pretty_print`. If we now call `get_fullname()` on an object, the resulting string will be wrapped in `<p>` and `</p>` (which is HTML markup for a paragraph):

```

print('Starting')
p = Person('John', 'Smith', 21)
p.print_self()
print(p.get_fullname())
print('Done')

```

This generates the output:

```

Starting
Person - John , 21
<p>John Smith</p>
Done

```

29.8.2 *Methods with Parameters*

As with functions, methods that take parameters in addition to `self` can also be decorated. In this case the decorator must take not only the `self` parameter but also any parameters passed to the method. For example:

```

def trace(method):
    def method_wrapper(self, x, y):
        print('Calling', method, 'with', x, y)
        method(self, x, y)
        print('Called', method, 'with', x, y)
    return method_wrapper

```

Now this `trace` decorator defines an inner function that takes the parameter `self` and two additional parameters. It can be used with any method that also takes two parameters such as the method `move_to()` below:

```

class Point:
    def __init__(self, x, y):
        self.x = x
        self.y = y

    @trace
    def move_to(self, x, y):
        self.x = x
        self.y = y

```

```
def __str__(self):
    return 'Point - ' + str(self.x) + ', ' + str(self.y)
```

When a Point object is created below, we can call the `move_to` method and see the result:

```
p = Point(1, 1)
print(p)
p.move_to(5, 5)
print(p)
```

The output from this is:

```
Point - 1,1
Calling <function Point.move_to at 0x110288b70> with 5 5
Called <function Point.move_to at 0x110288b70> with 5 5
Point - 5,5
```

29.9 Class Decorators

As well as being able to decorate functions and methods; it is possible to decorate classes.

A class can be decorated to add required functionality that may be external to that class.

As an example, a common class level operation is to want to indicate that a class should implement the singleton design pattern. The Singleton Design Pattern (again from the Gang of Four Design Patterns book) describes a type that can only have one object constructed for it. That is, unlike other objects it should not be possible to obtain more than one instance within the same program. Thus, the *Singleton* design pattern ensures that only one instance of a class is created. All objects that use an instance of that type use the same instance.

We can define a decorator that implements the singleton design pattern, for example:

```
def singleton(cls):
    print('In singleton for: ', cls)
    instance = None

    def get_instance():
        nonlocal instance
        if instance is None:
            instance = cls()
        return instance

    return get_instance
```

This decorator returns the `get_instance` function. This function checks to see if the variable `instance` is set to `None` or not; if it is set to `None` it instantiates the class passed into the decorator and stores this in the `instance` variable. It then returns the instance. If the instance is already set it merely returns the instance.

We can apply this decorator to whole classes such as `Service` and `Foo` below:

```
@singleton
class Service(object):
    def print_it(self):
        print(self)

@singleton
class Foo(object):
    pass
```

We can now use the classes `Service` and `Foo` as normal; however only one instance of `Service` and one instance of `Foo` will ever be created in the same program:

```
print('Starting')
s1 = Service()
print(s1)
s2 = Service()
print(s2)
f1 = Foo()
print(f1)
f2 = Foo()
print(f2)
print('Done')
```

In the above `main` function, it looks as if we have created two new `Service` objects and two `Foo` objects; however, the `@singleton` decorator will restrict the number of instances created to one and will reuse that instance whenever a request is made to instantiate the given class. Thus when we run this example, we can see that the hexadecimal number representing the location of the object in memory is the same for the two `Service` objects and the same for the two `Foo` objects:

```
In singleton for: <class '__main__.Service'>
In singleton for: <class '__main__.Foo'>
Starting
<__main__.Service object at 0x10ac3f780>
<__main__.Service object at 0x10ac3f780>
<__main__.Foo object at 0x10ac3f7b8>
<__main__.Foo object at 0x10ac3f7b8>
Done
```


29.10 When is a Decorator Executed?

An important feature of decorators is that they are executed right after the decorator function is defined. This is usually at *import time* (i.e., when a module is loaded by Python).

```
def logger(func):
    print('In Logger')
    def inner():
        print('In inner calling ', func.__name__)
        func()
        print('In inner called ', func.__name__)
    print('Finishing Logger')
    return inner

@logger
def print_it():
    print('Print It')

print('Start')
print_it()
print('Done')
```

For example, the decorator `logger` shown above prints out ‘In Logger’ and ‘Finished Logger’ when it is executed. If the output is examined, it can be seen that this output occurs before the main function prints ‘Start’.

```
In Logger
Finishing Logger
Start
In inner calling print_it
Print It
In inner called print_it
Done
```

Note that the decorated function and the wrapped function only execute when they are explicitly invoked.

This highlights the difference between what Pythonistas call *import time* and *runtime*.

29.11 Built-In Decorators

There are numerous built-in decorators in Python 3; some of which we have already seen such as `@classmethod`, `@staticmethod` and `@property`. We also saw some decorators when talking about abstract methods and properties. There are also decorators associated with unit testing and asynchronous operations.

29.12 FuncTools Wrap

One issue with decorated functions may become apparent when debugging or trying to trace what is happening. The problem is that by default the attributes associated with the function being called are actually those of the inner function returned by the decorator function. That is the `name`, `doc` and `module` of the function are those of the function returned by the decorator. The `name` and documentation of the original, decorated function, have been lost.

For example, returning to the original logger decorator we have:

```
def logger(func):
    def inner():
        print('calling ', func.__name__)
        func()
        print('called ', func.__name__)
    return inner
@logger
def get_text(name):
    """returns some text"""
    return "Hello "+name

print('name:', get_text.__name__)
print('doc: ', get_text.__doc__)
print('module; ', get_text.__module__)
```

When we run this code we get:

```
name: inner
doc: None
module; __main__
```

It appears that the `get_text` function is called `inner` and has no *docstring* associated with it. However, if we look at the function it should be called `get_text()` and have a *docstring* of 'returns some text'!

Python (since version 2.5) has included the `functools` module which contains the `functools.wraps` decorator which can be used to overcome this problem.

`Wraps` is a decorator for updating the attributes of the wrapping function(`inner`) to those of the original function (in this case `get_text()`). This is as simple as decorating the 'inner' function with `@wraps(func)`.

```
from functools import wraps

def logger(func):
    @wraps(func)
    def inner():
        print('calling ', func.__name__)
        func()
        print('called ', func.__name__)
    return inner
```

The end result is that in the above example the name and doc are now updated to the name of the wrapped function and the documentation associated with that function.

If we now rerun the earlier example, we get:

```
name: get_text
doc: returns some text
module; __main__
```

29.13 Online Resources

For further information on decorators see:

- https://www.python-course.eu/python3_decorators.php a short introduction to Python decorators
- <https://www.python.org/dev/peps/pep-0318/> PEP 318 considering decorators for functions and methods.
- <https://www.python.org/dev/peps/pep-3129/> PEP 3129 introducing class decorators.
- <https://docs.python.org/3/library/functools.html> Python Standard Library documentation for `functools`.
- <https://pymotw.com/3/functools/index.html> Python module of the Week for `functools`.
- <https://github.com/lord63/awesome-python-decorator> A useful page that lists many Python decorators as well as third-party contributions.
- <https://wiki.python.org/moin/PythonDecoratorLibrary> which provides a central repository for decorator examples.

29.14 Book Reference

For more information on the Decorator and Singleton design patterns see the “Patterns” book by the Gang of Four (E. Gamma, R. Helm, R. Johnson and J. Vlissades, *Design Patterns: Elements of Reusable Object-Oriented Software*, Addison-Wesley, 1995).

29.15 Exercises

The aim of this exercise is to develop your own decorator.

You will write a `timer` decorator to be used with methods in a class that take the first parameter `self`, followed by one other parameter.

The decorator should log how long a method takes to execute.

To do this you can use the `time` module and import the `default_timer`.

```
from time import default_timer
```

You can then obtain a `default_timer` object for the start and end of a function call and use these values to generate the time taken, for example:

```
start = default_timer()
func(self, value)
end = default_timer()
print('returned from ', func, 'it took', end - start, 'seconds')
```

You can then apply the decorator to the `deposit()` and `withdraw()` methods defined in the `Account` class. For example,

```
@timer
def deposit(self, amount):
    self._balance += amount

@timer
def withdraw(self, amount):
    self._balance -= amount
```

These methods will be inherited by the `DepositAccount` and `InvestmentAccount` classes. In the `CurrentAccount` class the `withdraw` method is over written so you will also need to decorate that method with `@timer` as well.

Now when you run your sample application you should get timing information printed out for the `deposit` and `withdraw` methods:

```
calling deposit on Account[123] - John, current
account = 10.05overdraft limit: -100.0 with 23.45
returned from deposit it took 8.009999999999268e-07
seconds
calling withdraw on Account[123] - John, current
account = 33.5overdraft limit: -100.0 with 12.33
returned from withdraw it took 1.141999999999116e-06
seconds
```

Chapter 30

Iterables and Iterators



30.1 Introduction

There are two protocols that you are very likely to use or will possibly need to implement at some point or other; these are the *iterable* protocol and the *iterator* protocols.

These two closely related protocols are very widely used and supported by a large number of types.

One of the reasons that iterators and iterables are important is that they can be used with `for` statements in Python; this makes it very easy to integrate an iterable into code which needs to process a sequence of values in turn.

30.2 Iteration

30.2.1 Iterables

The *iterable* protocol is used by types where it is possible to process their contents one at a time in turn. An iterable is something that will supply an iterator that can be used to perform this processing. As such it is not the iterator itself; but the provider of the iterator.

There are many iterable types in Python including lists, sets, dictionaries, tuples, etc. These are all iterable *containers* that will supply an iterator.

To be an *iterable* type; it is necessary to implement the `__iter__()` method (which is the only method in the iterable protocol). This method must supply a reference to the iterator object. This reference could be to the data type itself or it could be to another type that implements the iterator protocol.

30.2.2 Iterators

An *iterator* is an object that will return a sequence of values. Iterators may be finite in length or infinite (although many container-oriented iterators provide a fixed set of values).

The iterator protocol specifies the `__next__()` method. This method is expected to return the next item in the sequence to return or to raise the `StopIteration` exception. This is used to indicate that the iterator has finished supplying values.

30.2.3 The Iteration Related Methods

To summarize then we have

- `__iter__()` from the iterable protocol which is used to return the iterator object,
- `__next__()` from the iterator protocol which is used to obtain the next value in a sequence of values.

Any data type can be both an iterable and an iterator; but that is not required. An iterable could return a different object that will be used to implement the iterator or it can return itself as the iterator—it's the designers choice.

30.2.4 The Iterable Evens Class

To illustrate the ideas behind iterables and iterators we will implement a simple class; this class will be an `Evens` class that is used to supply a set of even values from 0 to some limit. This illustrates that it is not only data containers that can be iterable/iterators.

It also illustrates a type that is both an iterable and an iterator.

```
class Evens(object):

    def __init__(self, limit):
        self.limit = limit
        self.val = 0

    # Makes this class iterable
    def __iter__(self):
        return self

    # Makes this class an iterator
    def __next__(self):
        if self.val > self.limit:
            raise StopIteration
```

```

else:
    return_val = self.val
    self.val += 2
return return_val

```

There are a few things to note about this class.

- The `__iter__()` method returns `self`; this is a very common pattern and assumes that the class also implements the iterator protocol
- The `__next__()` method either returns the next value in the sequence or it raises the `StopIteration` exception to indicate that there are no more values available.

30.2.5 Using the Evens Class with a For Loop

Now that we have implemented both the *iterable* and *iterator* protocols for the class `Evens` we can use it with a `for` statement:

```

print('Start')
for i in Evens(6):
    print(i, end=', ')

print('Done')

```

Which generates the output:

```

Start
0, 2, 4, 6, Done

```

This makes it look as if the `Evens` type is a built-in type as it can be used with an existing Python structure; however, the `for` loop merely expects to be given an *iterable*; as such `Evens` is compatible with the `for` loop.

30.3 The Itertools Module

The `itertools` module provides a number of useful functions that return iterators constructed in various ways. It can be used to provide an iterator over a selection of values from a data type that is iterable; it can be used to combine iterables together, etc.

30.4 Online Resources

See the following for further information.

- <https://docs.python.org/3/library/stdtypes.html#iterator-types> iterator-types for iterator types.

Chapter 31

Generators and Coroutines



31.1 Introduction

This chapter introduces generators and coroutines. A generator is a data generator, and a coroutine is a data consumer. However, many developers get these two constructors confused as they both employ the `yield` keyword; although the purpose of `yield` in a generator is very different to the purpose of `yield` within a coroutine. This chapter will discuss both Python language constructs.

31.2 Generators

In many cases it is not appropriate (or possible) to obtain all the data to be processed up front (for performance reasons, for memory reasons, etc.). Instead lazily creating the data to be iterated over based on some underlying dataset may be more appropriate.

Generators are a *special* function that can be used to *generate* a sequence of values to be iterated over on demand (that is when the values are needed) rather than produced up front.

The only thing that makes a generator a *generator function* is the use of the `yield` keyword (which was introduced in Python 2.3).

The `yield` keyword can only be used inside a function or a method. Upon its execution the function is suspended, and the value of the `yield` statement is returned as the current *cycle* value. If this is used with a `for` loop, then the loop runs once for this value. Execution of the generator function is then resumed after the loop has cycled once, and the next cycle value is obtained.

The *generator* function will keep supplying values until it returns (which means that an infinite sequence of values can be generated).

31.2.1 *Defining a Generator Function*

A very simple example of a generator function is given below. This function is called the `gen_numbers()` function:

```
def gen_numbers():
    yield 1
    yield 2
    yield 3
```

This is a *generator* function as it has at least one `yield` statement (in fact it has three). Each time the `gen_numbers()` function is called within a `for` statement it will return one of the values associated with a `yield` statement; in this case the value 1, then the value 2 and finally the value 3 before it returns (terminates).

31.2.2 *Using a Generator Function in a for Loop*

We can use the `gen_numbers()` function with a `for` statement as shown below:

```
for i in gen_numbers():
    print(i)
```

Which produces 1, 2 and 3 as output.

It is common for the body of a generator to have some form of loop itself. This loop is typically used to generate the values that will be *yielded*. However, as is shown above that is not necessary, and here a `yield` statement is repeated three times.

Note that `gen_numbers()` is a function but it is a special function as it returns a generator *object*.

This is a generator function returns a generator object which wraps up the *generation* of the values required but this is hidden from the developer.

31.2.3 *When Do the Yield Statements Execute?*

It is interesting to consider what happens within the generator function; it is actually suspended each time a `yield` statement supplies a value and is only resumed when the next request for a value is received. This can be seen by adding some additional `print` statements to the `gen_numbers()` function:

```
def gen_numbers2():
    print('Start')
    yield 1
    print('Continue')
    yield 2
    print('Final')
```

```
    yield 3
    print('End')

for i in gen_numbers():
    print(i)
```

When we run this code snippet, we get

```
Start
1
Continue
2
Final
3
End
```

Thus the generator executes the `yield` statements on an as needed basis and not all at once.

31.2.4 An Even Number Generator

We could have used a generator to produce a set of even numbers up to a specific limit, as we did earlier with the `Evens` class, but without the need to create a class (and implement the two special methods `__iter__()` and `__next__()`). For example:

```
def evens_up_to(limit):
    value = 0
    while value <= limit:
        yield value
        value += 2

for i in evens_up_to(6):
    print(i, end=', ')
```

This produces

```
0, 2, 4, 6,
```

This illustrates the potential benefit of a generator over an iterator; the `evens_up_to()` function is a lot simpler and concise than the `Evens` class.

31.2.5 Nesting Generator Functions

You can even nest generator functions as each call to the generator function is encapsulated in its own generator object which captures all the state information needed by that generator invocation. For example:

```

for i in evens_up_to(4):
    print('i:', i)

for j in evens_up_to(6):
    print('j:', j, end=', ')

print("")

```

Which generates

```

i: 0
j: 0, j: 2, j: 4, j: 6,
i: 2
j: 0, j: 2, j: 4, j: 6,
i: 4
j: 0, j: 2, j: 4, j: 6,

```

As you can see from this the loop variable `i` is bound to the values produced by the first call to `evens_up_to()` (which produces a sequence up to 4) while the `j` loop variable is bound to the values produced by the second call to `evens_up_to()` (which produces a sequence of values up to 6).

31.2.6 Using Generators Outside a for Loop

You do not need a `for` loop to work with a generator function; the generator object actually returned by the generator function supports the `next()` function. This function takes a generator object (returned from the generator function) and returns the next value in sequence.

For example:

```

evens = evens_up_to(4)
print(next(evens), end=', ')
print(next(evens), end=', ')
print(next(evens))

```

This produces

```

0, 2, 4

```

Subsequent calls to `next(evens)` return no value; if required the generator can throw an error/exception.

31.3 Coroutines

Coroutines were introduced in Python 2.5 but are still widely misunderstood.

Much documentation introduces coroutines by saying that they are similar to generators, however, there is a fundamental difference between generators and coroutines:

1. Generators are data producers,
2. Coroutines are data *consumers*.

That is coroutines *consume* data produced by something else; whereas a generator produces a sequence of values that something else can process.

The `send()` function is used to send values to a coroutine. These data items are made available within the coroutine; which will wait for values to be supplied to it. When a value is supplied then some behaviour can be triggered. Thus, when a coroutine consumes a value it triggers some behaviour to be processed.

Part of the confusion between generators and coroutines is that the `yield` keyword is reused within a coroutine; it is used within a coroutine to cause the coroutine to wait until a value has been sent. It will then supply this value to the coroutine.

It is also necessary to *prime* a coroutine using with `next()` or `send(None)` functions. This advances the coroutine to the call to `yield` where it will then wait until a value is sent to it.

A coroutine may continue forever unless `close()` is sent to it. It is possible to pick up on the coroutine being closed by catching the `GeneratorExit` exception; you can then trigger some shut down behaviour if required.

An example of a coroutine is given by the `grep()` function below:

```
def grep(pattern):
    print('Looking for', pattern)
    try:
        while True:
            line = (yield)
            if pattern in line:
                print(line)
    except GeneratorExit:
        print('Exiting the Co-routine')
```

This coroutine will wait for input data; when data is sent to the coroutine, then that data will be assigned to the `line` variable. It will then check to see if the pattern used to initialize the coroutine function is present in the `line`; if it is it will print the `line`; it will then loop and wait for the next data item to be sent to the coroutine. If while it is waiting the coroutine is closed, then it will catch the `GeneratorExit` exception and print out a suitable message.

The `grep()` coroutine is used below, notice that the coroutine function returns a coroutine object that can be used to submit data:

```
print('Starting')
# Initialise the coroutine
g = grep('Python')

# prime the coroutine
next(g)

# Send data to the coroutine
g.send('Java is cool')
g.send('C++ is cool')
g.send('Python is cool')

# now close the coroutine
g.close()
print('Done')
```

The output from this is

```
Starting
Looking for Python
Python is cool
Exiting the Co-routine
Done
```

31.4 Exercises

These exercises focus on the creation of a generator.

Write a *prime number* generator; you can use the prime number program you wrote earlier in the book but convert it into a generator. The generator should take a *limit* to give the maximum size of the loop you use to generate the prime numbers. You could call this `prime_number_generator()`.

You should be able to run the following code:

```
number = input('Please input the number:')
if number.isnumeric():
    num = int(number)
    if num <= 2:
        print('Number must be greater than 2')
    else:
        for prime in prime_number_generator(num):
            print(prime, end=', ')
else:
    print('Must be a positive integer')
```

If the user enters 27 then the output would be

```
Please input the number:27
```

2, 3, 5, 7, 11, 13, 17, 19, 23,

Now create the `infinite_prime_number_generator()`; this generator does not have a limit and will keep generating prime numbers until it is no longer used.

You should be able to use this prime number generator as follows:

```
prime = infinite_prime_number_generator()
print(next(prime))
print(next(prime))
print(next(prime))
print(next(prime))
print(next(prime))
```

Chapter 32

Collections, Tuples and Lists



32.1 Introduction

Earlier in this book we looked at some Python built-in types such as `string`, `int` and `float` as well as `bools`. These are not the only built-in types in Python; another group of built-in types are collectively known as collection types. This is because they represent a collection of other types (such as a collection of strings or integers).

A collection is a single object representing a group of objects (such as a list or dictionary). That is, they are a collection of other objects. Collections may also be referred to as *containers* (as they contain other objects). These collection classes are often used as the basis for more complex or application specific data structures and data types.

These collection types support various types of data structures (such as lists and maps) and ways to process elements within those structures. This chapter introduces the Python collection types.

32.2 Python Collection Types

There are four classes in Python that provide container like behaviour; that is, data types for holding collections of other objects, these are

- **Tuples** A Tuple represents a collection of objects that are ordered and immutable (cannot be modified). Tuples allow duplicate members and are indexed.
- **Lists** Lists hold a collection of objects that are ordered and mutable (changeable), and they are indexed and allow duplicate members.
- **Sets** Sets are a collection that is unordered and unindexed. They are mutable (changeable) but do not allow duplicate values to be held.

- **Dictionary** A dictionary is an unordered collection that is indexed by a *key* which references a *value*. The value is returned when the *key* is provided. No duplicate keys are allowed. Duplicate values are allowed. Dictionaries are mutable containers.

Remember that everything in Python is actually a type of object, integers are instances/objects of the type `int`, strings are instances of the type `string`, etc. Thus, container types such as `Set` can hold collections of any type of thing in Python.

32.3 Characterising Collections

Collections, or containers as they are sometimes known, allow multiple values to be held together, but as indicated in the previous section they have various characteristics which can be used to help identify the most appropriate collection to use at any one time. These characteristics are

- *Ordered* or *unordered* that is there a maintained order to how the members of the collection are held.
- *Mutable* or *immutable* which means whether the collection can be modified once out has been created or not.
- Does the collection allow duplicate values to be held or not.
- Are the values held in the collection associated with a numerical index or by a key.
- Can the size of the collection grow once created or not.

By using these characteristics you can choose an appropriate container type/collection to use.

Looking at the four collection types introduced in the previous section we can consider them against these characteristics.

Tuples		Sets	
	Ordered		Unordered
	Immutable		Mutable
	Allow duplicates		Do not allow duplicates
	Indexed		Not indexed
	Fixed size		Growable
Lists		Dictionaries	
	Ordered		Ordered (from 3.7)
	Mutable		Mutable
	Allow duplicates		Allow duplicate values
	Indexed		Keyed
	Growable		Growable

32.4 Tuples

Tuples, along with Lists, are probably one of Python's most used container types. They will be present in almost any non-trivial Python program.

Tuples are an immutable ordered collection of objects; that is, each element in a tuple has a specific position (its index) and that position does not change over time. Indeed, it is not possible to add or remove elements from the tuple once it has been created.


The Python type that represents a tuple is `tuple`.

32.4.1 Creating Tuples

Tuples are defined using parentheses (i.e., round brackets `'()'`) around the elements that make up the tuple, for example:

```
tup1 = (1, 3, 5, 7)
```

This defines a new `tuple` which is referenced by the variable `tup1`. The `tuple` contains exactly 4 elements (in this case integers) with the first element in the tuple (the integer 1) having the index 0 and the last element in the tuple (the integer 7) having the index 3. This is illustrated below:



32.4.2 The `tuple()` Constructor Function

The `tuple()` function can also be used to create a new tuple from an iterable. An iterable is something that implements the iterable protocol (see the last chapter).

This means that a new tuple can be created from a `set`, a `list`, a `dict` (as these are all *iterable* types) or any type that implements the iterable protocol.

The syntax of the `tuple()` function is

```
tuple(iterable)
```

For example:

```
list1 = [1, 2, 3]
t1 = tuple(list1)
print(t1)
```

which generates the output:

```
(1, 2, 3)
```

Note that in the above the square brackets are used to represent a *list* of things (the list container type is described in more detail later in this chapter).

32.4.3 Accessing Elements of a Tuple

The elements of a `tuple` can be accessed using an *index* in square brackets. The index returns the object at that position, for example:

```
print('tuple[0]:\t', tup1[0])
print('tuple[1]:\t', tup1[1])
print('tuple[2]:\t', tup1[2])
print('tuple[3]:\t', tup1[3])
```

which generates the output

```
tup1[0]: 1
tup1[1]: 3
tup1[2]: 5
tup1[3]: 7
```

32.4.4 Creating New Tuples from Existing Tuples

It is also possible to return what is known as a slice from a `tuple`. This is a new `tuple` which is comprised of a subset of the original `tuple`. This is done by providing the start and end indexes for the slice, separated by a colon, within the index square brackets. For example:

```
print('tuple[1:3]:\t', tup1[1:3])
```

Which returns a new `tuple` of two elements containing the elements from index 1 up to (but not including) element 3. Note that the original `tuple` is not affected in any way (remember its immutable so cannot be modified). The output of the above is thus:

```
tup1[1:3]: (3, 5)
```

There are in fact numerous variations on the use of the slicing indices. For example, if the first index is omitted it indicates that the slice should start from the beginning of the `tuple`, while omitting the last index indicates it should go to the end of the `tuple`.

```
print('tup1[:3]:\t', tup1[:3])
print('tup1[1:]:\t', tup1[1:])
```

which generates:

```
tup1[:3]: (1, 3, 5)
tup1[1:]: (3, 5, 7)
```

You can reverse a `tuple` using the `::-1` notation (again this returns a new `tuple` and has no effect on the original `tuple`):

```
print('tup1[::-1]:\t', tup1[::-1])
```

This thus produces

```
tup1[::-1]: (7, 5, 3, 1)
```

32.4.5 *Tuples Can Hold Different Types*

Tuples can also contain a mixture of different types; that is, they are not restricted to holding elements all of the same type. You can therefore write a `tuple` such as

```
tup2 = (1, 'John', Person('Phoebe', 21), True, -23.45)
print(tup2)
```

Which produces the output:

```
(1, 'John', <__main__.Person object at 0x105785080>, True, -23.45)
```

32.4.6 *Iterating Over Tuples*

You can iterate over the contents of a `tuple` (that is, process each element in the `tuple` in turn). This is done using the `for` loop that we have already seen; however, it is the `tuple` that is used for the value to which the loop variable will be applied:

```
tup3 = ('apple', 'pear', 'orange', 'plum', 'apple')
for x in tup3:
    print(x)
```

This prints out each of the elements in the `tuple` `tup3` in turn:

```
apple
pear
orange
plum
apple
```

Note that again the order of the elements in the `tuple` is persevered.

32.4.7 *Tuple Related Functions*

You can also find out the length of a `tuple`

```
print('len(tup3):\t', len(tup3))
```

You can count how many times a specified value appears in a `tuple` (remember `tuples` allow duplicates);

```
print(tup3.count('apple')) # returns 2
```

You can also find out the (first) index of a value in a `tuple`:

```
print(tup3.index('pear')) # returns 1
```

Note that both `index()` and `count()` are methods defined on the class `tuple`, whereas `len()` is a function that the `tuple` is passed into. This is because `len()` is a generic function and can be used with other types as well such as strings.

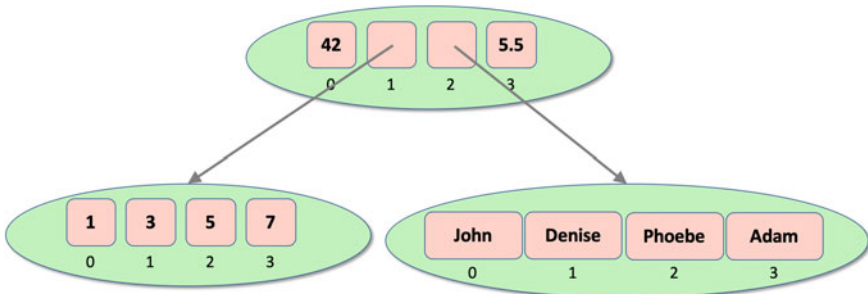
32.4.8 *Checking if an Element Exists*

You can check to see if a specific element exists in a `tuple` using the `in` operator, for example:

```
if 'orange' in tup3:
    print('orange is in the tuple')
```

32.4.9 Nested Tuples

Tuples can be nested within tuples; that is a tuple can contain, as one of its elements, another tuple. For example, the following diagram illustrates the nesting of a tree of tuples:



In code we could define this structure as

```
tuple1 = (1, 3, 5, 7)
tuple2 = ('John', 'Denise', 'Phoebe', 'Adam')
tuple3 = (42, tuple1, tuple2, 5.5)
print(tuple3)
```

The output from this is

```
(42, (1, 3, 5, 7), ('John', 'Denise', 'Phoebe', 'Adam'), 5.5)
```

Note the *nesting* of round brackets in the printout illustrating where one tuple is contained within another.

This feature of tuples (and other containers) allows for arbitrarily complex data structures to be constructed as required by an application.

In fact, a tuple can have nested within it not just other tuples but any type of container, and thus it can contain lists, sets, dictionaries, etc. This provides for a huge level of flexibility when constructing data structures for use in Python programs.

32.4.10 Things You Can't Do with Tuples

It is not possible to add or remove elements from a tuple; they are *immutable*. It should be particularly noted that none of the functions or methods presented above actually change the original tuple they are applied to; even those that return a subset of the original tuple actually return a new instance of the class `tuple` and have no effect on the original tuple.

32.5 Lists

Lists are mutable ordered containers of other objects. They support all the features of the `tuple` but as they are mutable it is also possible to add elements to a `list`, remove elements and modify elements. The elements in the list maintain their order (until modified).

32.5.1 Creating Lists

Lists are created using square brackets positioned around the elements that make up the list. For example:

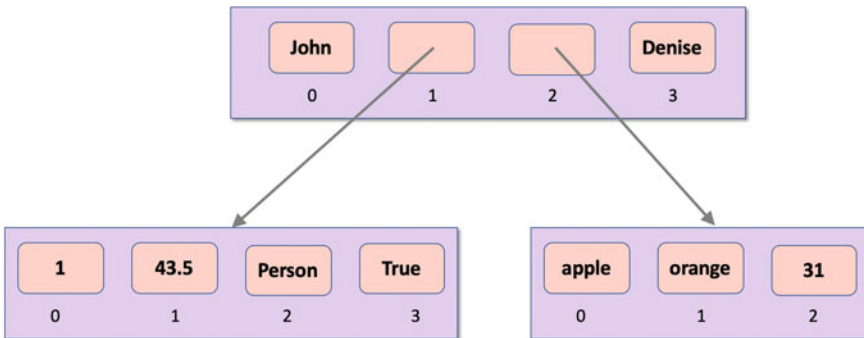
```
list1 = ['John', 'Paul', 'George', 'Ringo']
```

In this case we have created a list of four elements with the first element being indexed from zero, we thus have



As with tuples we can have nested lists and lists containing different types of elements.

We can this create the following structure of nested Lists:



In code this can be defined as


```
l1 = [1, 43.5, Person('Phoebe', 21), True]
l2 = ['apple', 'orange', 31]
root_list = ['John', l1, l2, 'Denise']
print(root_list)
```

When the root_list is printed we get

```
['John', [1, 43.5, <Person object at 0x1042ba4a8>, True],
 ['apple', 'orange', 31], 'Denise']
```

Note the square brackets inside the outer square brackets indicating nested lists.

We can of course also nest tuples in lists and lists in tuples. For example, the following structure shows tuples (the ovals) hold references to lists (the rectangles) and vice versa:



In code this would look like:

```
t1 = (1, 'John', 34.5)
l1 = ['Smith', 'Jones']
l2 = [t1, l1]
t2 = (l2, 'apple')
print(t2)
```

which produces

```
((1, 'John', 34.5), ['Smith', 'Jones']), 'apple')
```


32.5.2 List Constructor Function

The `list()` function can be used to construct a list from an iterable; this means that it can construct a list from a tuple, a dictionary or a set. It can also construct a list from anything that implements the iterable protocol.

The signature of the `list()` function is

```
list(iterable)
```

For example:

```
vowelTuple = ('a', 'e', 'i', 'o', 'u')
print(list(vowelTuple))
```

produces

```
['a', 'e', 'i', 'o', 'u']
```

32.5.3 Accessing Elements from a List

You can access elements from a list using an index (within square brackets). The index returns the object at that position, for example:

```
list1 = ['John', 'Paul', 'George', 'Ringo']
print(list1[1])
```

This will print out the element at index 1 which is Paul (lists are indexed from zero so the first element is the zeroth element).

If you use a negative index such as `-1` then the index is reversed so an index of `-1` starts from the end of the list (`-1` returns the last element, `-2` the second to last etc.).

It is also possible to extract a slice (or sublist) from a list. This is done by providing a starting and end index to within the square brackets separated by a colon. For example `[1:4]` indicates a slice starting at the *oneth* element and extending up to (but not including) the fourth element. If either of the indexes is missed for a slice then that indicates the start or end of the list respective.

The following illustrates some of these ideas:

```
list1 = ['John', 'Paul', 'George', 'Ringo']
print('list1[1]:', list1[1])
print('list1[-1]:', list1[-1])
print('list1[1:3]:', list1[1:3])
print('list[:3]:', list1[:3])
print('list[1:]:', list1[1:])
```

Which produces

```
list1[1]: Paul
list1[-1]: Ringo
list1[1:3]: ['Paul', 'George']
list[:3]: ['John', 'Paul', 'George']
list[1:]: ['Paul', 'George', 'Ringo']
```

32.5.4 Adding to a List

You can add an item to a list using the `append()` method of the `list` class (this changes the actual list; it does not create a copy of the list). The syntax of this method is

```
<alist>.append(<object>)
```

As an example, consider the following list of strings, to which we append a fifth string:

```
list1 = ['John', 'Paul', 'George', 'Ringo']
list1.append('Pete')
print(list1)
```

this will generate the output:

```
['John', 'Paul', 'George', 'Ringo', 'Pete']
```

You can also add all the items in a list to another list. There are several options here, we can use the `extend()` method which will add the items passed to it to the end of the list or we can use the `+=` operator which does the same thing:

```
list1 = ['John', 'Paul', 'George', 'Ringo', 'Pete']
print(list1)
list1.extend(['Albert', 'Bob'])
print(list1)
list1 += ['Ginger', 'Sporty']
print(list1)
```

The output from this code snippet is

```
['John', 'Paul', 'George', 'Ringo', 'Pete']
['John', 'Paul', 'George', 'Ringo', 'Pete', 'Albert', 'Bob']
['John', 'Paul', 'George', 'Ringo', 'Pete', 'Albert', 'Bob',
'Ginger', 'Sporty']
```

Which approach you prefer to use is up to you.

Note that strictly speaking both `extend()` and `+=` take what is known as an *iterable*.

32.5.5 *Inserting into a List*

You can also insert elements into an existing list. This is done using the `insert()` method of the `list` class. The syntax of this method is

```
<list>.insert(<index>, <object>)
```

The `insert()` method takes an index indicating where to insert the element and an object to be inserted.

For example, we can insert the string ‘Paloma’ in between the *Zeroth* and *oneth* item in the following list of names:

```
a_list = ['Adele', 'Madonna', 'Cher']
print(a_list)
a_list.insert(1, 'Paloma')
print(a_list)
```

The result is

```
['Adele', 'Madonna', 'Cher']
['Adele', 'Paloma', 'Madonna', 'Cher']
```

In other words, we have inserted the string ‘Paloma’ into the index position 1 pushing ‘Madonna’ and ‘Cher’ up one in the index within the list.

32.5.6 *List Concatenation*

It is possible to concatenate two lists together using the concatenation operator ‘+’:

```
list1 = [3, 2, 1]
list2 = [6, 5, 4]
list3 = list1 + list2
print(list3)
```

generates

```
[3, 2, 1, 6, 5, 4]
```

32.5.7 *Removing from a List*

We can remove an element from a list using the `remove()` method. The syntax for this method is

```
<list>.remove(<object>)
```

This will remove the object from the list; if the object is not in the list then an error will be generated by Python.

```
another_list = ['Gary', 'Mark', 'Robbie', 'Jason', 'Howard']
print(another_list)
another_list.remove('Robbie')
print(another_list)
```

The output from this is

```
['Gary', 'Mark', 'Robbie', 'Jason', 'Howard']
['Gary', 'Mark', 'Jason', 'Howard']
```

32.5.8 *The Pop() Method*

The syntax of the `pop()` method is

```
a.pop(index=-1)
```

It removes an element from the list; however, it differs from the `remove()` method in two ways:

- It takes an index which is the index of the item to remove from the list rather than the object itself.
- The method returns the item that was removed as its result.

An example of using the `pop()` method is given below:

```
list6 = ['Once', 'Upon', 'a', 'Time']
print(list6)
print(list6.pop(2))
print(list6)
```

Which generates

```
['Once', 'Upon', 'a', 'Time']
a
['Once', 'Upon', 'Time']
```

An overload of this method is just.

```
<list>.pop()
```

Which removes the last item in the list. For example:

```
list6 = ['Once', 'Upon', 'a', 'Time']
print(list6)
print(list6.pop())
print(list6)
```

with the output:

```
['Once', 'Upon', 'a', 'Time']
Time
['Once', 'Upon', 'a']
```

32.5.9 *Deleting from a List*

It is also possible to use the `del` keyword to delete elements from a list.

The `del` keyword can be used to delete a single element or a slice from a list.

To delete an individual element from a list use `del` and access the element via its index:

```
my_list = ['A', 'B', 'C', 'D', 'E']
print(my_list)
del my_list[2]
print(my_list)
```

which outputs

```
['A', 'B', 'C', 'D', 'E']
['A', 'B', 'D', 'E']
```

To delete a slice from within a list use the `del` keyword and the slice returned from the list.

```
my_list = ['A', 'B', 'C', 'D', 'E']
print(my_list)
del my_list[1:3]
print(my_list)
```

which deletes the slice from index 1 up to (but not including) index 3:

```
['A', 'B', 'C', 'D', 'E']
['A', 'D', 'E']
```

32.5.10 *List Methods*

Python has a set of built-in methods that you can use on lists.

Method	Description
<code>append()</code>	Adds an element at the end of the list
<code>clear()</code>	Removes all the elements from the list
<code>copy()</code>	Returns a copy of the list
<code>count()</code>	Returns the number of elements with the specified value
<code>extend()</code>	Add the elements of a list (or any iterable), to the end of the current list
<code>index()</code>	Returns the index of the first element with the specified value
<code>insert()</code>	Adds an element at the specified position
<code>pop()</code>	Removes the element at the specified position
<code>remove()</code>	Removes the item with the specified value
<code>reverse()</code>	Reverses the order of the list
<code>sort()</code>	Sorts the list

32.6 Online Resources

See the Python Standard Library documentation for:

- <https://docs.python.org/3/tutorial/datastructures.html> Python Tutorial on data structures.
- <https://docs.python.org/3/library/stdtypes.html#sequence-types-list-tuple-range> for lists and tuples.
- <https://docs.python.org/3/tutorial/datastructures.html#tuples-and-sequences> the online tuples tutorial.
- <https://docs.python.org/3/tutorial/datastructures.html#lists> the online list tutorial.

32.7 Exercises

The aim of this exercise is to work with a collection/container such as a list.

To do this we will return to your account related classes.

You should modify your `Account` class such that it is able to keep a history of transactions.

A *transaction* is a record of a deposit or withdrawal along with an amount.

Note that the initial amount in an account can be treated as an initial deposit.

The history could be implemented as a *list* containing an ordered sequence to transactions. A transaction itself could be defined by a class with an action (deposit or withdrawal) and an amount.

Each time a withdrawal or a deposit is made a new transaction record which should be added to a transaction history list.

Now provide support for iterating through the transaction history of the account such that each deposit or withdrawal can be reviewed. You can do this by implementing the iterable protocol—refer to the last chapter if you need to check how to do this. Note that it is the transaction history that we want to be able to iterate through—so you can use the history list as the basis of your iterable.

You should be able to run this following code at the end of your accounts application:

```
for transaction in acct:  
    print(transaction)
```

Depending upon the exact set of transactions you have performed (deposits and withdrawals) you should get a list of those transactions being printed out:

```
Transaction[deposit: 10.05]  
Transaction[deposit: 23.45]  
Transaction[withdraw: 12.33]
```

Chapter 33

Sets



33.1 Introduction

Python provides many different container or collection types. This chapter focuses on the set collection. This collection is used to hold unique *immutable* collections of objects. It is a mutable collection and so the values it holds can change over time but does not allow duplicate values. It also supports basic set like operations such as intersection, union and difference.

33.2 Creating a Set

A set is defined using curly brackets (e.g., '{ }').

For example,

```
basket = {'apple', 'orange', 'apple', 'pear',  
         'orange', 'banana'}  
print(basket)
```

When run this code will show that *apple* is only added once to the set:

```
{'banana', 'orange', 'pear', 'apple'}
```

Note that because a set is unordered it is not possible to refer to elements of the set using an index.

33.3 The Set() Constructor Function

As with tuples and lists Python provides a predefined function that can convert any iterable type into a `set`. The function signature is

```
set(iterable)
```

Given an iterable object, this function returns a new `set` based on the values obtained from the iterable. This means that a `set` can be easily created from a list, tuple or dictionary as well as any other data type that implements the iterable protocol. For example:

```
set1 = set((1, 2, 3))
print(set1)
```

which prints out

```
{1, 2, 3}
```

33.4 Accessing Elements in a Set

Unlike lists it is not possible to access elements from a `set` via an index; this is because they are unordered containers, and thus there are no indexes available. However, they are iterable containers.

Elements of a `set` can be iterated over using the `for` statement:

```
for item in basket:
    print(item)
```

This applies the `print` statement to each item in the list in turn.

33.5 Working with Sets

33.5.1 Checking for Presence of an Element

You can check for the presence of an element in a `set` using the `in` keyword, for example:

```
print('apple' in basket)
```

This will print `true` if `'apple'` is a member of the `set` `basket`.

33.5.2 *Adding Items to a Set*

It is possible to add items to a set using the `add()` method:

```
basket = {'apple', 'orange', 'banana'}
basket.add('apricot')
print(basket)
```

This generates

```
{'orange', 'apple', 'banana', 'apricot', 'pear'}
```

If you want to add more than one item to a set you can use the `update()` method:

```
basket = {'apple', 'orange', 'banana'}
basket.update(['apricot', 'mango', 'grapefruit'])
print(basket)
```

generating

```
{'orange', 'apple', 'mango', 'banana', 'apricot', 'grapefruit'}
```

The argument to `update` can be a set, a list, a tuple or a dictionary. The method automatically converts the parameter into a set if it is not a set already and then adds the value to the original set.

33.5.3 *Changing Items in a Set*

It is not possible to change the items already in a set.

33.5.4 *Obtaining the Length of a Set*

As with other collection/container classes; you can find out the length of a set using the `len()` function.

```
basket = {'apple', 'orange', 'apple', 'pear', 'orange', 'banana'}
print(len(basket)) # generates 4
```

33.5.5 *Obtaining the Max and Min Values in a Set*

You can also obtain the maximum or minimum values in a set using the `max()` and `min()` functions:

```
print(max(a_set))
print(min(a_set))
```

33.5.6 *Removing an Item*

To remove an item from a set, use the `remove()` or `discard()` functions. The `remove` function removes a single item from a set but generates an error if that item was not in the set. The `discard` function also removes a single item from a set but does not throw an error if it was not initially present in the set.

```
basket = {'apple', 'orange', 'apple', 'pear', 'orange', 'banana'}
print(basket)
basket.remove('apple')
basket.discard('apricot')
print(basket)
```

This generates

```
{'pear', 'banana', 'orange', 'apple'}
{'pear', 'banana', 'orange'}
```

There is also a method `pop()` that can be used to remove an item (and return that item as a result of running the method); however it removes the last item in the set (although as a set is unordered you will not know which item will be).

The method `clear()` is used to remove all elements from a set:

```
basket = {'apple', 'orange', 'banana'}
basket.clear()
print(basket)
```

which prints out

```
set()
```

Which is used to represent an empty set.

33.6 Nesting Lists Within Sets

It is not possible to hold a list within a set, this is because lists are mutable objects and sets can only hold immutable things. Thus it is necessary to convert a list into an immutable version whenever they are added to a set. This can be done very easily

using the `tuple()` function. This function can be used to convert any iterable into a tuple. In effect a tuple is an immutable list. Thus if we need to store a list within a set all we have to do is convert it into a tuple, for example:

```
# Convert lists into a tuple
list1 = [55, 60, 43, 55, 65]
results_set = {tuple(list1), 'John', 'CS110'}
print(f'results_set: {results_set}')
```

The output from this code snippet is

```
results_set: {(55, 60, 43, 55, 65), 'CS110', 'John'}
```


33.7 Set Operations

The `set` container also supports *set like* operations such as (`|`), intersection (`&`), difference (`-`) and symmetric difference (`^`). These are based on simple set theory.

Given the two sets:

```
s1 = {'apple', 'orange', 'banana'}
s2 = {'grapefruit', 'lime', 'banana'}
```

For example, the Union of two sets represents the combination of all the values in the two sets:




```
print('Union:', s1 | s2)
```

This would print out

```
Union: {'apple', 'lime', 'banana', 'grapefruit', 'orange'}
```

The intersection of two sets represents the common values between two sets:



```
print('Intersection:', s1 & s2)
```

This generates

```
Intersection: {'banana'}
```

The difference between two sets is the set of values in the first set that are not in the second set:




```
print('Difference:', s1 - s2)
```

which produces the output

```
Difference: {'apple', 'orange'}
```

The symmetric difference represents all the unique values in the two sets (that is it is the inverse of the intersection):



```
print('Symmetric Difference:', s1 ^ s2)
```

The output from this final operation is

```
Symmetric Difference: {'orange', 'apple', 'lime', 'grapefruit'}
```

In addition to the operators there are also method versions:

- `s1.union(s2)` is the equivalent of `s1 | s2`
- `s1.intersection(s2)` is the equivalent of `s1 & s2`
- `s1.difference(s2)` is the equivalent of `s1 - s2`
- `s1.symmetric_difference(s2)` is the equivalent of `s1 ^ s2`

33.8 Set Methods

Python has a set of built-in methods that you can use on sets.

Method	Description
<code>add()</code>	Adds an element to the set
<code>clear()</code>	Removes all the elements from the set
<code>copy()</code>	Returns a copy of the set
<code>difference()</code>	Returns a set containing the difference between two or more sets
<code>difference_update()</code>	Removes the items in this set that are also included in another, specified set
<code>discard()</code>	Remove the specified item
<code>intersection()</code>	Returns a set, that is the intersection of two other sets
<code>intersection_update()</code>	Removes the items in this set that are not present in other, specified set(s)
<code>isdisjoint()</code>	Returns whether two sets have a intersection or not
<code>issubset()</code>	Returns whether another set contains this set or not
<code>issuperset()</code>	Returns whether this set contains another set or not
<code>pop()</code>	Removes an element from the set
<code>remove()</code>	Removes the specified element
<code>symmetric_difference()</code>	Returns a set with the symmetric differences of two sets
<code>symmetric_difference_update()</code>	Inserts the symmetric differences from this set and another
<code>union()</code>	Return a set containing the union of sets
<code>update()</code>	Update the set with the union of this set and others

33.9 Online Resources

Online resources on sets are listed below:

- <https://docs.python.org/3/tutorial/datastructures.html> Python tutorial on data structures.

- <https://docs.python.org/3/tutorial/datastructures.html#sets> the online set tutorial.
- https://www.python-course.eu/python3_sets_frozensets.php A tutorial on sets and frozen sets.
- <https://docs.python.org/3/library/stdtypes.html#set-types-set-frozenset> for sets.

33.10 Exercises

The aim of this exercise is to use a set.

Create two sets of students, one for those who took an exam and one for those that submitted a project. You can use simple strings to represent the students, for example:

```
# set up sets
exam = {'Andrew', 'Kirsty', 'Beth', 'Emily', 'Sue'}
project = {'Kirsty', 'Emily', 'Ian', 'Stuart'}

# Output the basic sets
print('exam:', exam)
print('project:', project)
```

Using these sets answer the following questions:

- Which students took both the exam and submitted a project?
- Which students only took the exam?
- Which students only submitted the project?
- List all students who took either (or both) of the exam and the project.
- List all students who took either (but *not* both) of the exam and the project.

Chapter 34

Dictionaries

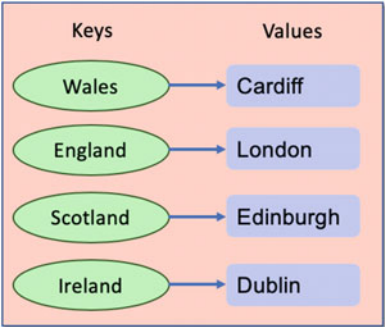


34.1 Introduction

A dictionary (technically called a `dict` in Python) is a set of associations between a key and a value that is unordered, changeable (mutable) and indexed. This collection type allows for keyed access to values. That is, a value is stored in a dictionary using a key. This key is then used to retrieve that value at a later date. Dictionaries are very widely used in Python programs.

34.2 What is a Dictionary

Pictorially we might view a dictionary as shown below for a set of countries and their capital cities. Note that in a dictionary the keys must be unique but the values do not need to be unique.



34.3 Creating a Dictionary

A dict is created using curly brackets (‘{ }’) where each entry in the dictionary is a key:value pair:

```
cities = {'Wales': 'Cardiff',
         'England': 'London',
         'Scotland': 'Edinburgh',
         'Northern Ireland': 'Belfast',
         'Ireland': 'Dublin'}
print(cities)
```

This creates a dictionary referenced by the variable `cities` which holds a set of key: value pairs for the capital cities of the UK and Ireland. When this code is run we see:

```
{'Wales': 'Cardiff', 'England': 'London', 'Scotland':
 'Edinburgh', 'Northern Ireland': 'Belfast', 'Ireland': 'Dublin'}
```

34.3.1 The dict() Constructor Function

The `dict()` function can be used to create a new dictionary object from an iterable or a sequence of key:value pairs. The syntax of this function is

```
dict(**kwarg)
dict(mapping, **kwarg)
dict(iterable, **kwarg)
```

This is an overloaded function with three version that can take different types of arguments:

- The first option takes a sequence of key: value pairs.
- The second takes a mapping and a (optionally) a sequence of key: value pairs.
- The third version takes an iterable of key-value pairs and an optional sequence of key: value pairs.

Some examples are given below for reference:

```
# note keys are not strings
dict1 = dict(uk='London', ireland='Dublin', france='Paris')
print('dict1:', dict1)
# key value pairs are tuples
dict2 = dict([('uk', 'London'), ('ireland', 'Dublin'), ('france',
'Paris')])
print('dict2:', dict2)
# key value pairs are lists
dict3 = dict([('uk', 'London'], ['ireland', 'Dublin'], ['france',
'Paris']))
```

```
print('dict3:', dict3)
```

The output printed by these examples is

```
dict1: {'uk': 'London', 'ireland': 'Dublin', 'france': 'Paris'}
dict2: {'uk': 'London', 'ireland': 'Dublin', 'france': 'Paris'}
dict3: {'uk': 'London', 'ireland': 'Dublin', 'france': 'Paris'}
```

34.4 Working with Dictionaries

34.4.1 Accessing Items via Keys

You can access the values held in a `dict` using their associated key. This is specified using either the square bracket (`[]`) notation (where the key is within the brackets) or the `get()` method:

```
print('cities[Wales]:', cities['Wales'])
print('cities.get(Ireland):', cities.get('Ireland'))
```

The output of this is

```
cities[Wales]: Cardiff
cities.get(Ireland): Dublin
```

34.4.2 Adding a New Entry

A new entry can be added to a dictionary by providing the key in square brackets and the new value to be assigned to that key:

```
cities['France'] = 'Paris'
```

34.4.3 Changing a Keys Value

The value associated with a key can be changed by reassigning a new value using the square bracket notation, for example:

```
cities['Wales'] = 'Swansea'
```

```
print(cities)
```

which would now show ‘Swansea’ as the capital of wales:

```
{'Wales': 'Swansea', 'England': 'London', 'Scotland':
'Edinburgh', 'Northern Ireland': 'Belfast', 'Ireland': 'Dublin'}
```

34.4.4 Removing an Entry

An entry into the dictionary can be moved using one of the methods `pop()` or `popitem()` method or the `del` keyword.

- The `pop(<key>)` method removes the *entry* with the specified *key*. This method returns the value of the key being deleted. If the key is not present then a default value (if it has been set using `setdefault()`) will be returned. If no default value has been set an error will be generated.
- The `popitem()` method removes the last inserted item in the dictionary (although prior to Python 3.7 a random item in the dictionary was deleted instead!). The *key:value* pair being deleted is returned from the method.
- The `del` keyword removes the entry with the specified key from the dictionary. This keyword just deletes the item; it does not return the associated value. It is potentially more efficient than `pop(<key>)`.

Examples of each of these are given below:

```
cities = {'Wales': 'Cardiff',
          'England': 'London',
          'Scotland': 'Edinburgh',
          'Northern Ireland': 'Belfast',
          'Ireland': 'Dublin'}
print(cities)
cities.popitem() # Deletes 'Ireland' entry
print(cities)
cities.pop('Northern Ireland')
print(cities)
del cities['Scotland']
print(cities)
```

The output from this code snippet is thus:

```
{'Wales': 'Cardiff', 'England': 'London', 'Scotland':
'Edinburgh', 'Northern Ireland': 'Belfast', 'Ireland': 'Dublin'}
{'Wales': 'Cardiff', 'England': 'London', 'Scotland':
'Edinburgh', 'Northern Ireland': 'Belfast'}
{'Wales': 'Cardiff', 'England': 'London', 'Scotland':
'Edinburgh'}
{'Wales': 'Cardiff', 'England': 'London'}
```

In addition the `clear()` method empties the dictionary of all entries:

```
cities = {'Wales': 'Cardiff',
```

```

        'England': 'London',
        'Scotland': 'Edinburgh',
        'Northern Ireland': 'Belfast',
        'Ireland': 'Dublin' }
print(cities)
cities.clear()
print(cities)

```

Which generates the following output:

```

{'Wales': 'Cardiff', 'England': 'London', 'Scotland':
'Edinburgh', 'Northern Ireland': 'Belfast', 'Ireland': 'Dublin'}
{}

```

Note that the empty dictionary is represented by the `{}` above which as the empty set was represented as `set()`.

34.4.5 Iterating Over Keys

You can loop through a dictionary using the `for` loop statement. The `for` loop processes each of the *keys* in the dictionary in turn. This can be used to access each of the values associated with the keys, for example:

```

for country in cities:
    print(country, end=', ')
    print(cities[country])

```

Which generates the output:

```

Wales, Cardiff
England, London
Scotland, Edinburgh
Northern Ireland, Belfast
Ireland, Dublin

```

If you want to iterate over all the values directly, you can do so using the `values()` method. This returns a collection of all the values, which of course you can then iterate over:

```

for e in d.values():
    print(e)

```

34.4.6 Values, Keys and Items

There are three methods that allow you to obtain a view onto the contents of a dictionary, these are `values()`, `keys()` and `items()`.

- The `values()` method returns a view onto the dictionary's values.

- The `keys()` method returns a view onto a dictionary's keys.
- The `items()` method returns a view onto the dictionary's items ((key, value) pairs).

A view provides a dynamic window onto the dictionary's entries, which means that when the dictionary changes, the view reflects these changes.

The following code uses the `cities` dictionaries with these three methods:

```
print(cities.values())
print(cities.keys())
print(cities.items())
```

The output makes it clear that these are all related to a dictionary by indicating that the type is `dict_values`, `dict_keys` or `dict_items`, etc:

```
dict_values(['Cardiff', 'London', 'Edinburgh', 'Belfast',
'Dublin'])
dict_keys(['Wales', 'England', 'Scotland', 'Northern Ireland',
'Ireland'])
dict_items([('Wales', 'Cardiff'), ('England', 'London'), ('Scotland',
'Edinburgh'), ('Northern Ireland', 'Belfast'), ('Ireland',
'Dublin')])
```

34.4.7 *Checking Key Membership*

You can check to see if a key is a member of a dictionary using the `in` syntax (and that it is not in a dictionary using the `not in` option), for example:

```
print('Wales' in cities)
print('France' not in cities)
```

Which both print out `True` for the `cities` dictionary.

34.4.8 *Obtaining the Length of a Dictionary*

Once again, as with other collection classes; you can find out the length of a dictionary (in terms of its key:value pairs) using the `len()` function.

```
cities = {'Wales': 'Cardiff',
         'England': 'London',
         'Scotland': 'Edinburgh',
         'Northern Ireland': 'Belfast',
         'Ireland': 'Dublin'}
print(len(cities)) # prints 5
```

34.4.9 Nesting Dictionaries

The *key* and *value* in a dictionary must be an object; however, everything in Python is an object and thus anything can be used as a key or a value.

One common pattern is where the *value* in a dictionary is itself a container such as a list, tuple, set or even another dictionary.

The following example uses tuples to represent the months that make up the seasons:

```
seasons = { 'Spring': ('Mar', 'Apr', 'May'),
           'Summer': ('June', 'July', 'August'),
           'Autumn': ('September', 'October', 'November'),
           'Winter': ('December', 'January', 'February') }
print(seasons['Spring'])
print(seasons['Spring'][1])
```

The output is

```
('Mar', 'Apr', 'May')
Apr
```

Each season has a `Tuple` for the value element of the entry. When this `Tuple` is returned using the key it can be treated just like any other `Tuple`.

Note in this case we could easily have used a `List` or indeed a `Set` instead of a `tuple`.

34.4.10 Checking to See if a Key is Present

To check to see if a key is present within a dictionary use the `in` keyword, for example with an `if` statement as shown below:

```
if key in dict:
    print('Key is present, value =', dict[key])
else:
    print('Key not present')
```

34.5 Dictionary Operators Added in 3.9

Back in Python 3.9 two additional dictionary operators were added to Python, these were the merge `|` and the update `|=` operators. The merge operator combines two dictionaries together to form a new third dictionary.

In turn the update operator modifies the left hand dictionary by adding the contents of the right hand dictionary.

Examples of both the merge and the update operator are given below:

```
offices = {'John': 1, 'Denise': 21}
labs = {'Cyber': 5, 'Games': 33}
print('offices', offices)
print('labs', labs)

# Merge two dictionaries to create a new one
all_rooms = offices | labs
print('all_rooms', all_rooms)

# Update one dictionary with contents of another
offices |= labs
print('offices', offices)
```

The output from this program is

```
offices {'John': 1, 'Denise': 21}
labs {'Cyber': 5, 'Games': 33}
all_rooms {'John': 1, 'Denise': 21, 'Cyber': 5, 'Games': 33}
offices {'John': 1, 'Denise': 21, 'Cyber': 5, 'Games': 33}
```

34.6 A Note on Dictionary Key Objects

A class whose objects are to be used as the key within a dictionary should consider implementing two special methods, these are `__hash__()` and `__eq__()`. The hash method is used to generate a hash number that can be used by the dictionary container, and the equals method (`__eq__()`) is used to test if two objects are equal. For example:

```
print('key.__hash__():', key.__hash__())
print("key.__eq__('England'):", key.__eq__('England'))
```

The output from these two lines for an example run is

```
key.__hash__(): 8507681174485233653
key.__eq__('England'): True
```

Python has two rules associated with these methods:

- If two objects are equal, then their hashes should be equal.
- In order for an object to be hashable, it must be immutable.

It also has two properties associated with the hashcodes of an object that should be adhered to:

- If two objects have the same hash, then they are likely to be the same object.

- The hash of an object should be cheap to compute.

Why do you need to care about these methods?

For built-in type you do not need to worry; however for user defined classes/types then if these types are to be used as keys within a dictionary then you should consider implementing these methods.

This is because a dictionary uses

- The *hashing* method to manage how values are organized and
- The *equals* method to check to see if a key is already present in the dictionary.

As an aside if you want to make a class something that cannot be used as a key in a dictionary, that is it is not hashable, then you can define this by setting the `__hash__()` method to `None`.

```
class NotHashableThing(object):
    __hash__ = None
```

34.7 Dictionary Methods

Python has a set of built-in methods that you can use on dictionaries.

Method	Description
clear()	Removes all the elements from the dictionary
copy()	Returns a copy of the dictionary
fromkeys()	Returns a dictionary with the specified keys and values
get()	Returns the value of the specified key
items()	Returns a list containing the tuple for each key-value pair
keys()	Returns a list containing the dictionary's keys
pop()	Removes the element with the specified key
popitem()	Removes the last inserted key-value pair
setdefault()	Returns the value of the specified key. If the key does not exist: insert the key, with the specified value
update()	Updates the dictionary with the specified key-value pairs
values()	Returns a list of all the values in the dictionary

34.8 Online Resources

Online resources on dictionaries are listed below:

- <https://docs.python.org/3/tutorial/datastructures.html> Python tutorial on data structures.
- https://www.python-course.eu/python3_dictionaries.php A tutorial on dictionaries in Python.
- <https://docs.python.org/3/library/stdtypes.html#mapping-types-dict> for dictionaries.
- <https://docs.python.org/3/tutorial/datastructures.html#dictionaries> the online dictionary tutorial.
- https://en.wikipedia.org/wiki/Hash_table For more information on defining hash functions and how they are used in containers such as dictionary.

34.9 Exercises

The aim of this exercise is to use a dictionary as a simple form of data cache.

Calculating the factorial for a very large number can take some time. For example calculating the factorial of 150,000 can take several seconds. We can verify this using a timer decorator similar to that we created back in chapter on decorators.

The following program runs several factorial calculations on large numbers and prints out the time taken for each:

```

from timeit import default_timer

def timer(func):
    def inner(value):
        print('calling', func.__name__, 'with', value)
        start = default_timer()
        func(value)
        end = default_timer()
        print('returned from', func.__name__, 'it took', int(end -
start), 'seconds')

    return inner

@timer
def factorial(num):
    if num == 0:
        return 1
    else:
        Factorial_value = 1
        for i in range(1, num + 1):
            Factorial_value = factorial_value * i
        return factorial_value

```

```
print(factorial(150000))
print(factorial(80000))
print(factorial(120000))
print(factorial(150000))
print(factorial(120000))
print(factorial(80000))
```

An example of the output generated by this program is given below:

```
calling factorial with 150000
returned from factorial it took 5 seconds
None
calling factorial with 80000
returned from factorial it took 1 seconds
None
calling factorial with 120000
returned from factorial it took 3 seconds
None
calling factorial with 150000
returned from factorial it took 5 seconds
None
calling factorial with 120000
returned from factorial it took 3 seconds
None
calling factorial with 80000
returned from factorial it took 1 seconds
None
```

As can be seen from this, in this particular run, calculating the factorial of 150,000 took 5 s, while the factorial of 80,000 took just over 1 1/4 s, etc.

In this particular case we have decided to rerun these calculations so that we have actually calculated the factorial of 150,000, 80,000 and 120,000 at least twice.

The idea of a *cache* is that it can be used to save the previous calculations and reuse those if appropriate rather than have to perform the same calculation multiple times. The use of a cache can greatly improve the performance of systems in which these repeat calculations occur.

There are many commercial caching libraries available for a wide variety of languages including Python. However, at their core they are all somewhat dictionary like; that is, there is a *key* which is usually some combination of the operation invoked and the parameter values used. In turn the *value* element is the result of the calculation.

These caches usually also have eviction policies so that they do not become overly large; these eviction policies can usually be specified so that they match the way in which the cache is used. One common eviction policy is the Least Recently Used (or LRU) policy. When using this policy once the size of the cache reaches a predetermined limit the Least Recently Used value is evicted, etc.

For this exercise you should implement a simple caching mechanism using a dictionary (but without an eviction policy).

The cache should use the parameter passed into the `factorial()` function as the key and return the stored value if one is present.

The logic for this is usually:

1. Look in the cache to see if the key is present
2. If it is return the value
3. If not perform the calculation
4. Store the calculated result for future use
5. Return the value

Note as the `factorial()` function is exactly that a function; you will need to think about using a global variable to hold the cache.

Once the cache is used with the `factorial()` function, then each subsequent invocation of the function using the previous value should return almost immediately. This is shown in the sample output below where subsequent method calls return in less than a second.

```
calling factorial with 150000
returned from factorial it took 5 seconds
None
calling factorial with 80000
returned from factorial it took 1 seconds
None
calling factorial with 120000
returned from factorial it took 3 seconds
None
calling factorial with 150000
returned from factorial it took 0 seconds
None
calling factorial with 120000
returned from factorial it took 0 seconds
None
calling factorial with 80000
returned from factorial it took 0 seconds
None
```

Chapter 35

Frozensets



35.1 Introduction

A frozenset is an immutable version of a set. That is, a frozenset is a set that cannot be modified once created.

Frozensets are useful as a set can only hold an immutable object, thus it is not possible to *nest* a set within a set (as sets are mutable things). To get around this it is possible to convert any set (and in fact any iterable) into a frozenset.

35.2 Frozenset Characteristics

A frozenset has all the characteristics of a set with the exception of being modifiable. Thus the characteristics of a frozenset are

- Immutable.
- Do not allow duplicates.
- They are not indexed.
- They are of fixed size.

A frozenset can be useful in any situation where you require an immutable type with a set like set of operations.

35.3 Creating a Frozenset

Frozensets are created using the `frozenset()` function. This function takes an (optional) iterable element and creates the new frozenset based on that iterable. The iterable could be a list, a set, a dict or a tuple, etc.

Some examples of these conversions are shown below:

```

tuple1 = (1, 2, 3, 4, 3, 2)
list1 = ['A', 'B', 'C', 'B', 'A']
dict1 = {'Wales': 'Cardiff',
         'England': 'London',
         'Scotland': 'Edinburgh'}
set1 = {'apple', 'orange', 'banana'}

fs1 = frozenset(tuple1)
print(f'tuple1: {tuple1}')
print(f'fs1: {fs1}')

print('-' * 10)

fs2 = frozenset(list1)
print(f'list1: {list1}')
print(f'fs2: {fs2}')

print('-' * 10)

fs3 = frozenset(dict1)
print(f'dict1: {dict1}')
print(f'fs3: {fs3}')

print('-' * 10)

fs4 = frozenset(set1)
print(f'set1: {set1}')
print(f'fs4: {fs4}')

```

When this code is run the output generated is

```

tuple1: (1, 2, 3, 4, 3, 2)
fs1: frozenset({1, 2, 3, 4})
-----
list1: ['A', 'B', 'C', 'B', 'A']
fs2: frozenset({'A', 'B', 'C'})
-----
dict1: {'Wales': 'Cardiff', 'England': 'London', 'Scotland':
'Edinburgh'}
fs3: frozenset({'Wales', 'Scotland', 'England'})
-----
set1: {'banana', 'apple', 'orange'}
fs4: frozenset({'banana', 'apple', 'orange'})

```

There are a few things to note about this:

- The first is that the `frozenset` is still a set, and thus duplicates are not allowed. This is illustrated for `tuple1` and `list1` where both of these containers possess duplicate values, but the `frozensets` generated from them do not.
- The second thing to notice is that when a `frozenset` is created from a `dict` it is the keys that are extracted not the values.
- The third point to note is that the `frozenset` is unordered (just as a set is unordered). Thus the order that the values are held in a tuple, list or dictionary may not be the same as the order in which they are held in the `frozenset`.

35.4 Frozenset Operations

The operations available on a frozenset are similar to the set with the obvious exception that you cannot modify the frozenset. However, you can copy a frozenset, apply set like operations such as union and compare frozensets, etc.

Some of these are illustrated below:

```
# frozenset operations
brands = frozenset(['ford', 'jaguar', 'bmw'])
manufacturers = brands.copy()
print(f'brands: {brands}')
print(f'manufacturers: {manufacturers}')

# Equality
print(f'brands == manufacturers: {brands == manufacturers}')
print(f'brands == fs3: {brands == fs3}')

print('-' * 10)

fruit = frozenset(['apple', 'pear', 'banana'])
veg = frozenset(['carrot', 'cabbage', 'peas'])
print(f'fruit.union(veg): {fruit.union(veg)}')

breakfast = frozenset(['apple', 'strawberry', 'banana'])
print(f'fruit.intersection(breakfast):
{fruit.intersection(breakfast)}')

print(f'fruit.difference(breakfast):
{fruit.difference(breakfast)}')

owner = frozenset(['jaguar', 'bmw'])
print(f'breakfast.issubset(fruit): {owner.issubset(brands)}')
```

The output from this code is

```
brands: frozenset({'bmw', 'ford', 'jaguar'})
manufacturers: frozenset({'bmw', 'ford', 'jaguar'})
brands == manufacturers: True
brands == fs3: False
-----
fruit.union(veg): frozenset({'peas', 'cabbage', 'apple', 'pear',
'carrot', 'banana'})
fruit.intersection(breakfast): frozenset({'apple', 'banana'})
fruit.difference(breakfast): frozenset({'pear'})
breakfast.issubset(fruit): True
```

In the above example we have used:

- `copy` returns a new instance of the frozenset containing the same members as the original.
- `union` returns the union of frozensets as a new frozenset.
- `intersection` returns those members in both frozensets as a frozenset.
- `difference` returns members that are not in either of the frozensets as a frozenset.
- `issubset` returns `True` or `False` depending upon whether the receiving frozensets' members are fully contained within the second frozenset.

In addition there is

- `issuperset` which indicates whether the receiving frozenset contains all the members of the second frozenset.
- `isdisjoint` returns `True` if the two frozensets have no members in common. Otherwise returns `False`.
- `symmetric_difference` Return the symmetric difference of two frozensets as a new frozenset.

35.5 Using Frozensets with Sets

The most common usage of a frozenset is when wanting to place a set within a set (that is to nest sets).

It is possible to hold any *immutable* object within a set. This means that a set can contain a reference to a tuple (as that is immutable). We can thus write:

```
s1 = { (1, 2, 3) }
print(s1)
```

This prints out

```
{(1, 2, 3)}
```

However, we cannot nest lists or other sets within a set as these are not immutable types. The following would both generate a runtime error in Python:

```
# Can't have the following
s2 = { {1, 2, 3} }
print(s2)

s3 = { [1, 2, 3] }
print(s3)
```

However we can use `frozensets` and nest these within sets. For example:

```
# Need to convert sets and lists into frozensets
s2 = { frozenset ({1, 2, 3}) }
print(s2)

s3 = { frozenset ([1, 2, 3]) }
print(s3)
```

This generates

```
{frozenset({1, 2, 3})}
{frozenset([1, 2, 3])}
```


Chapter 36

Collection Related Modules



36.1 Introduction

The chapter introduces a feature known as a list comprehension in Python. It then introduces the `collections` and `itertools` modules.

36.2 List Comprehension

This is a very powerful mechanism that can be used to generate new lists. The syntax of the list comprehension is

```
[ <expression> for item in iterable <if optional_condition > ]
```

The new list is formed of the results from the expression. Note that the whole `for` statement and expression is surrounded by the square brackets normally used to create a List.

When a list comprehension is executed it generates a new list by applying the expression to the items in another collection.

For example, given one list of integers, another (new) list can be created using the list comprehension format:

```
list1 = [1, 2, 3, 4, 5, 6]
print('list1:', list1)

list2 = [item + 1 for item in list1]
print('list2:', list2)
```

which produces the output:

```
list1: [1, 2, 3, 4, 5, 6]
list2: [2, 3, 4, 5, 6, 7]
```

Here the new list is generated by adding 1 to each element in the initial list `list1`. Essentially, we iterate (loop) over all the elements in the initial list and bind each element in the list to the `item` variable in turn. The result of the expression `item + 1` is then captured, in order, in the new list.

This feature is not limited to processing values in a list; any iterable collection can be used such as tuples or sets as the source of the values to process.

Another feature of the list comprehension is the ability to filter the values passed to the expression using the optional `if` condition.

For example, if we wish to only consider even numbers for the expression then we can use the optional `if` statement to filter out all odd numbers:

```
list3 = [item + 1 for item in list1 if item % 2 == 0]
print('list3:', list3)
```

The output from these two lines of code is

```
list3: [3, 5, 7]
```

Thus only the even numbers were passed to the expression `item + 1` resulting in only the values 3, 5 and 7 being generated for the new list.

A new feature introduced in Python 3.12 allows a variable that will be used in the list comprehension but will not be modified during the loop to be initialized using the assignment expressions (`:=`). For example:

```
[(b := 1) for a, b.prop in some_iter]
```

the assignment to `b` is now allowed. Note that assigning to variables stored to in the target part of comprehensions (like `a`) is still disallowed.

36.3 The Collections Module

The collections module extends that basic features of the collection-oriented data types within Python with high performance container data types. It provides many useful containers such as

Name	Purpose
namedtuple()	Factory function for creating tuple subclasses with named fields
deque	List-like container with fast appends and pops on either end
ChainMap	Dict-like class for creating a single view of multiple mappings
Counter	Dict subclass for counting hashable objects
OrderedDict	Dict subclass that remembers the order entries were added
Defaultdict	Dict subclass that calls a factory function to supply missing values

(continued)

(continued)

Name	Purpose
UserDict	Wrapper around dictionary objects for easier dict subclassing
UserList	Wrapper around list objects for easier list subclassing
UserString	Wrapper around string objects for easier string subclassing

As this is not one of the default modules that are automatically loaded for you by Python; you will need to `import` the required collection.

In this section we will look at the `Counter`, `Defaultdict` and `namedtuple` containers as examples.

36.3.1 The Counter Container

We will start off by looking at the `Counter` type. It is designed to efficiently hold multiple copies of the same element. It is efficient because it only holds one copy of each element but keeps a count of the number of times that element has been added to the collection:

```
import collections

fruit = collections.Counter(['apple', 'orange', 'pear', 'apple',
                             'orange', 'apple'])
print(fruit)
print(fruit['orange'])
```

The output of this is

```
Counter({'apple': 3, 'orange': 2, 'pear': 1})
2
```

Which makes the counting behaviour associated with the `Counter` class quite clear.

There are many uses of such a class, for example, it can be used to find out the most frequently used word in an essay; all you have to do is add each word in an essay to the `Counter` and then retrieve the word with the highest count. This can be done using the `Counter` class's `most_common()` method. This method takes a parameter `n` that indicates how many of the *most common* elements should be returned. If `n` is omitted (or set to `None`) then the method returns an ordered list of elements. Thus to obtain the most common fruit from the above `Counter` collection we can use:

```
print('fruit.most_common(1):', fruit.most_common(1))
```

Which generates:

```
fruit.most_common(1): [('apple', 3)]
```

You can also perform some mathematical operations with multiple Counter objects. For example, you can add and subtract Counter objects. You can also obtain a combination of Counters that combines the maximum values from two Counter objects. You can also generate an intersection of two Counters. These are all illustrated below:

```
fruit1 = collections.Counter(['apple', 'orange', 'pear',
                              'orange'])
fruit2 = collections.Counter(['banana', 'apple', 'apple'])

print('fruit1:', fruit1)
print('fruit2:', fruit2)

print('fruit1 + fruit2:', fruit1 + fruit2)
print('fruit1 - fruit2:', fruit1 - fruit2)
# Union (max(fruit1[n], fruit2[n]))
print('fruit1 | fruit2:', fruit1 | fruit2)
# Intersection (min(fruit1[n], fruit2[n]))
print('fruit1 & fruit2:', fruit1 & fruit2)
```

Which produces

```
fruit1: Counter({'orange': 2, 'apple': 1, 'pear': 1})
fruit2: Counter({'apple': 2, 'banana': 1})
fruit1 + fruit2: Counter({'apple': 3, 'orange': 2, 'pear': 1,
                          'banana': 1})
fruit1 - fruit2: Counter({'orange': 2, 'pear': 1})
fruit1 | fruit2: Counter({'apple': 2, 'orange': 2, 'pear': 1,
                          'banana': 1})
fruit1 & fruit2: Counter({'apple': 1})
```

Once a Counter object has been created you can test it to see if an item is present using the `in` keyword, for example:

```
print('apple' in fruit)
```

You can also add items to a Counter object by accessing the value using the item as the key, for example:

```
fruit['apple'] = 1 # initializes the number of apples
fruit['apple'] = + 1 # Adds one to the number of apples
fruit['apple'] = - 1 # Subtracts 1 from the number of apples
```

36.3.2 The Defaultdict Container

The standard `dict` container in Python will throw an exception if you try to access a key that is not currently in the set of dictionary entries if you use the `[]` access operator. In turn it will return `None` if you use the `get()` method for a key that is currently not present.

The first of these could be dangerous and the second may not be a suitable default value (although this second issue can be changed in `dict`).

This is where the `collections.DefaultDict` comes in. This container is a subclass of the default dictionary type and provides a similar key-value dictionary behaviour. However, `DefaultDict` will *never* raise a `KeyError` if a function is provided to generate the default value. When creating a `DefaultDict` you should specify a function that will be used to generate the default value so that you can customize its behaviour.

Two examples are shown below, one using a lambda function (a very common approach) and one using a named function:

```
days = defaultdict(lambda: -1)
days['Monday'] = 1
days['Tuesday'] = 2

print(days['Monday'])
print(days['Tuesday'])
print(days['Wednesday'])

def default_value():
    return 'Not Known'

# Defining the dict
cities = defaultdict(default_value)
cities['England'] = 'London'
cities['Wales'] = 'Cardiff'

print(cities['England'])
print(cities['Wales'])
print(cities['Ireland'])
```

In both cases there is a request to retrieve a value for a missing key—‘Wednesday’ in the first example and ‘Ireland’ in the second example. However, neither will raise a `KeyError` as both have a function that will be used to generate the missing value. In the first example the lambda function will return `-1` and in the second example the `default_value()` named function will return the string ‘Not Known’.

The output from this code snippet is

```
1
2
-1
London
Cardiff
Not Known
```

36.3.3 The Namedtuple Container

A `namedtuple` is a type of tuple (and thus immutable) that has a set of names for each member in the tuple (although the use of positional integer indexes is still available). As such it has some aspects of a dictionary in that each member is a key: value pair.

To create a `namedtuple` you provide the name of the tuple type being created and then an iterable of a list of names for that tuple, for example:

```
from collections import namedtuple

# Declaring namedtuple()
Student = namedtuple('Student', ['name', 'age', 'degree'])

# Adding values
undergrad_student = Student('John', '21', 'Computer Science')
postgrad = Student(name='Denise', age=25, degree='PhD')

print(undergrad_student)
print(postgrad)

# Access using index
print(f'Name via index (student[0]): {undergrad_student[0]}')

# Access using name
print(f'Name via named member (student.name): {undergrad_
student.name}')

# Can field out the field names using _fields
print(undergrad_student._fields)
```

This code example creates a new `namedtuple` called 'Student' and stores it into a variable called `Student`. Note the capitalization, `Student` here represents a new `namedtuple` type and thus follows the capitalization convention for user defined types.

The `Student` `namedtuple` had the members `name`, `age` and `degree`.

When we create a named student instance we use the `Student` variable (which contains the newly created named tuple) and provide a value for the `name`, `age` and `degree` members. The first instance created for the undergraduate uses positional parameter passing to initialize the member values. The second example uses *named* parameter passing to do the same thing.

When the `namedtuple` is printed out the default representation will illustrate the name of the type and each of the key-value pairs.

Following this we illustrate how the values in the `namedtuple` can be accessed using both *positional* indexing and *named* member access. This `[0]` and `.name` both access the *name* element of the `namedtuple` `Student`.

To finish the example, the code illustrates that it is possible to access a tuple of names in the `namedtuple`—this can be useful for iteration or display to a user.

The output from this program is

```
Student(name='John', age='21', degree='Computer Science')
Student(name='Denise', age=25, degree='PhD')
Name via index (student[0]): John
Name via named member (student.name): John
('name', 'age', 'degree')
```

It is also possible to iterate over the values in the named tuple using a standard for loop, for example:

```
for value in undergrad_student:
    print(value)
```

This produces the expected output of:

```
John
21
Computer Science
```

Note if you try to change a value in a named tuple after you have created it, for example as shown below:

```
postgrad.name = "Pete"
print(postgrad)
```

You will generate an `AttributeError` as named tuples are immutable:

```
Traceback (most recent call last):
  File "namedtuple_example.py", line 25, in <module>
    postgrad.name = "Pete"
    ^^^^^^^^^^^^^^^^^
AttributeError: can't set attribute
```

36.4 The Itertools Module

The `itertools` module is another module that it is worth being familiar with. This module provides a number of useful functions that return iterators constructed in various ways. As there are many different options available it is worth looking at the documentation for a complete list of available functions.

To give you a flavor of some of the facilities available look at the following listing:

```
import itertools

# Connect two iterators together
r1 = list(itertools.chain([1, 2, 3], [2, 3, 4]))
print(r1)

# Create iterator with element repeated specified number of
# times (possibly infinite)
```

```

r2 = list(itertools.repeat('hello', 5))
print(r2)

# Create iterator with elements from first iterator starting
# where predicate function fails
values = [1, 3, 5, 7, 9, 3, 1]
r3 = list(itertools.dropwhile(lambda x: x < 5, values))
print(r3)

# Create iterator with elements from supplied iterator between
# the two indexes (use 'None' for second index to go to end)
r4 = list(itertools.islice(values, 3, 6))
print(r4)

```

The output from this code is

```

[1, 2, 3, 2, 3, 4]
['hello', 'hello', 'hello', 'hello', 'hello']
[5, 7, 9, 3, 1]
[7, 9, 3]

```

36.5 Online Resources

Online resources on the `itertools` library are listed below:

- <https://docs.python.org/3/library/itertools.html> The standard library documentation on the `itertools` module.
- <https://pymotw.com/3/itertools/index.html> The Python module of the week page for `itertools`.

36.6 Exercises

- The aim of this exercise is to create a concordance program in Python using the `collections.Counter` class.
- For the purposes of this exercise a concordance is an alphabetical list of the words present in a text or texts with a count of the number of times that the word occurs.
- Your concordance program should include the following steps:
 - Ask the user to input a sentence.
 - Split the sentence into individual words (you can use the `split()` method of the string class for this).
 - Use the `Counter` class to generate a list of all the words in the sentence and the number of times they occur.
 - Produce an alphabetic list of the words with their counts. You can obtain a sorted list of the keys using the `sorted()` function. You can then use a `collection.OrderedDict` to generate a dictionary that maintains the order in which keys were added.

- Print out the alphabetically ordered list.

An example of how the program might operate is given below:

```
Please enter text to be analyzed: cat sat mat hat cat hat cat  
Unordered Counter  
Counter({'cat': 3, 'hat': 2, 'sat': 1, 'mat': 1})  
Ordered Word Count  
OrderedDict([('cat', 3), ('hat', 2), ('mat', 1), ('sat', 1)])
```

Chapter 37

ADTs, Queues and Stacks



37.1 Introduction

There are a number of common data structures that are used within computer programs that you might expect to see within Python’s list of collection or container classes; these include queues and stacks. However, in the basic collection classes these are missing. However, we can either create our own implementations or use one of the extension libraries that provide such collections. In this chapter we will explore implementing our own versions.

37.2 Abstract Data Types

The queue and stack are concrete examples of what are known as Abstract Data Types (or ADTs).

An Abstract Data Type (or ADT) is a model for a particular type of data, where a data type is defined by its behaviour (or semantics) from the point of view of the *user* of that data type. This behaviour is typically defined in terms of possible values, possible operations on the data of this type and behaviour of the operations provided by the data type.

An ADT is used to define a common concept that can be implemented by one or more concrete data structures. These implementations may use different internal representations of the data or different algorithms to provide the behaviour; by semantically they meet the descriptions provided by the ADT.

For example, a list ADT may be defined that defines the operations and behaviour that a list like data structure must provide. Concrete implementations may meet the semantics of a list using an underlying array of elements, or by linking elements together with pointers or using some form of hash table (all of which are different internal representations that could be used to implement a list).

37.3 Data Structures

We will look at how the Python collection types can be used as both a queue and a stack but first we need to define both these ADTs:

- **Queue** is an ADT that defines how a collection of entities are managed and maintained. Queues have what is known as a First-In-First-Out (or FIFO) behaviour that is the first entity added to a queue is the first thing removed from the queue. Within the queue the order in which the entities were added is maintained.
- **Stack** is another ADT but this time it has a Last-In-First-Out (or LIFO) behaviour. That is the most recently added entity to the stack will be the next entity to be removed. Within the stack the order that the entities were added in is maintained.


37.4 Queues

Queues are very widely used within Computer Science and in Software Engineering. They allow data to be held for processing purposes where the guarantee is that the earlier elements added will be processed before later ones.

There are numerous variations on the basic queue operations but in essence all queues provide

- Queue creation.
- Add an element to the back of the queue (known as enqueueing).
- Remove an element from the front of the queue (known as dequeueing).
- Find out the length of the queue.
- Check to see if the queue is empty.
- Queues can be of fixed size or variable (growable) in size.

The basic behaviour of a queue is illustrated by



Here there are five elements in the queue, one element has already been removed from the front and another is being added at the back. Note that when one element is removed from the front of the queue all other element moved forward one position. Thus, the element that was the second to the front of the queue becomes the front of the queue when the first element is dequeued.

Many queues also allow features such as

- *Peek* at the element at the front of the queue (that is see what the element is but do not remove it from the queue).
- Provide *priorities* so that elements with a higher priority are not added to the back of the queue but to a point in the middle of the queue related to their priority.

37.4.1 Python List as a Queue

The Python list container can be used as a queue using the existing operations such as `append()` and `pop()`, for example:

```
queue = [] # Create an empty queue
queue.append('task1')
print('initial queue:', queue)
queue.append('task2')
queue.append('task3')
print('queue after additions:', queue)
element1 = queue.pop(0)
print('element retrieved from queue:', element1)
print('queue after removal', queue)
```

The output of which is

```
initial queue: ['task1']
queue after additions: ['task1', 'task2', 'task3']
element retrieved from queue: task1
queue after removal ['task2', 'task3']
```

Note that each *task* was added to the end of the queue, but the first task obtained from the queue was task 1.

37.4.2 Defining a Queue Class

In the last section we used the `List` class as a way of providing a queue; this approach does work but it is not obvious that we are using the list as a queue (with the exception of the name of the variable that we are holding the `List` in). For example we have used `pop(0)` to dequeue an element from the queue and we have used `append()` to enqueue an element. In addition there is nothing to stop a programmer forgetting to use `pop(0)` and instead using `pop()` which is an easy mistake to make and which will remove the most recently added item from the queue.

It would be better to create a new data type and ensure that this data provides the queue like behaviour and hide the list inside this data type.

We can do this by defining our own `Queue` class in Python.

```
class Queue:
    def __init__(self):
        self._list = [] # initial internal data

    def enqueue(self, element):
        self._list.append(element)

    def dequeue(self):
        return self._list.pop(0)

    def __len__(self):
        """ Supports the len protocol """
        return len(self._list)

    def is_empty(self):
        return self.__len__() == 0

    def peek(self):
        return self._list[0]

    def __str__(self):
        return 'Queue: ' + str(self._list)
```

This `Queue` class internally holds a list. Note we are using the convention that the internal list instance variable name is preceded by an underbar (`'_'`), thereby indicating that no one should access it directly.

We have also defined methods for dequeuing and enqueueing elements to the queue. To complete the definition, we have also defined methods for checking the current length of the queue, whether the queue is empty or not, allowing the element at the front of the queue to be peeked at and of course proving a string version of the queue for printing.

Note that the `is_empty()` method uses the `__len__()` method when determining if the queue is empty; this is an example of an important idea; only define something once. As we want to use the length of the queue to help determine if the queue is empty, we reuse the `__len__()` method rather than the code implementing the length method; thus, if the internal representation changes we will not affect the `is_empty()` method.

The following short program illustrates how the `Queue` class can be used:

```
queue = Queue()
print('queue.is_empty():', queue.is_empty())
queue.enqueue('task1')
print('len(queue):', len(queue))
queue.enqueue('task2')
queue.enqueue('task3')
print('queue:', queue)
print('queue.peak():', queue.peak())
print('queue.dequeue():', queue.dequeue())
print('queue:', queue)
```

The output from this is:

```
initial queue: ['task1']
queue after additions: ['task1', 'task2', 'task3']
element retrieved from queue: task1
queue after removal ['task2', 'task3']
queue.is_empty(): True
len(queue): 1
queue: Queue: ['task1', 'task2', 'task3']
queue.peek(): task1
queue.dequeue(): task1
queue: Queue: ['task2', 'task3']
```

This provides a far more explicit and semantically more meaningful implementation of a `Queue` than the use of the raw `List` data structure.

Of course, Python understands this and provides a queue container class in the `collections` module called `deque`. This implementation is optimized to be more efficient than the basic `List` which is not very efficient when it comes to popping elements from the front of the list.

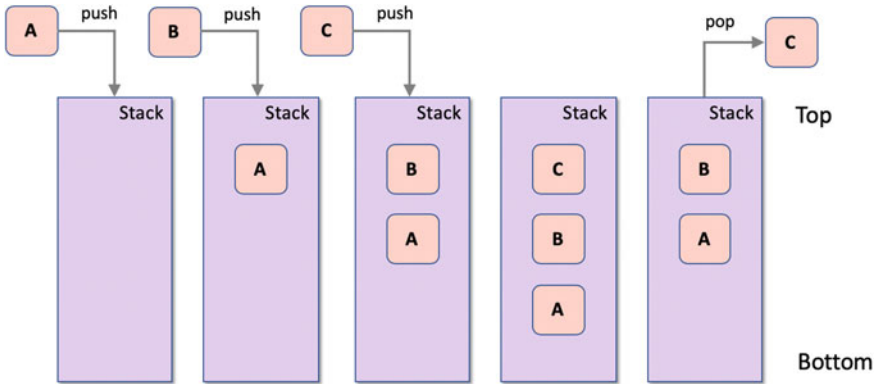
37.5 Stacks

Stacks are another very widely used ADT within computer science and in software applications. They are often used for evaluating mathematical expressions, parsing syntax, for managing intermediate results, etc.

The basic facilities provided by a `Stack` include

- Stack creation.
- Add an element to the top of the stack (known as pushing onto the stack).
- Remove an element from the top of the stack (known as popping from the stack).
- Find out the length of the stack.
- Check to see if the stack is empty.
- Stacks can be of fixed size or a variable (growable) stack.

The basic behaviour of a stack is illustrated by



This diagram illustrates the behaviour of a stack. Each time a new element is pushed onto the stack, it forces any existing elements further down the stack. Thus the most recent element is at the *top* of the stack, and the oldest element is at the *bottom* of the stack. To get to older elements you must first pop newer elements off the top, etc.

Many stacks also allow features such as

- *Top* which is often an operation that allows you to peek at the element at the top of the stack (that is see what the element is but do not remove it from the queue).

37.5.1 Python List as a Stack

A List may initially appear particularly well suited to being used as a Stack as the basic `append()` and `pop()` methods can be used to emulate the stack behaviour. Whatever was most recently appended to the list is the element that will be next returned from the `pop()` method, for example:

```
stack = [] # create an empty stack
stack.append('task1')
stack.append('task2')
stack.append('task3')
print('stack:', stack)
top_element = stack.pop()
print('top_element:', top_element)
print('stack:', stack)
```

Which produces the output:

```
stack: ['task1', 'task2', 'task3']
top_element: task3
stack: ['task1', 'task2']
```

This certainly works although when we print out the stack it does not make it clear that 'task3' is at the front of the stack.

In addition, as when using the `List` as a queue ADT it is still possible to apply any of the other methods defined on a `List` to this `stack` and thus we can still write `stack.pop(0)` which would remove the very first element added to the stack.

We could therefore implement a `Stack` class to wrap the list and provide suitable stack like behaviour as we did for the `Queue` class.

37.6 Online Resources

For more information on ADTs, queues and stacks see:

- https://en.wikipedia.org/wiki/Abstract_data_type Wikipedia page on Abstract Data Types (ADTs).
- [https://en.wikipedia.org/wiki/Queue_\(abstract_data_type\)](https://en.wikipedia.org/wiki/Queue_(abstract_data_type)) Wikipedia page on the queue data structure.
- [https://en.wikipedia.org/wiki/Stack_\(abstract_data_type\)](https://en.wikipedia.org/wiki/Stack_(abstract_data_type)) Wikipedia page on stacks.
- https://en.wikibooks.org/wiki/Data_Structures/Stacks_and_Queues Wikibooks tutorial on stack and queue data structures.

37.7 Exercises

Implement your own `Stack` class following the pattern used for the `Queue` class.

The `Stack` class should provide

- A `push(element)` method used to add an element to the `Stack`.
- A `pop()` method to retrieve the top element of the `Stack` (this method removes that element from the stack).
- A `top()` method that allows you to peek at the top element on the stack (it does not remove the element from the stack).
- A `__len__()` method to return the size of the stack. This method also meets the `len` protocol requirements.
- An `is_empty()` method that checks to see if the stack is empty.
- A `__str__()` method used to convert the stack into a string.

Once completed you should be able to run the following test application:

```
stack = Stack()
stack.push('T1')
stack.push('T2')
stack.push('T3')
print('stack:', stack)
print('stack.is_empty():', stack.is_empty())
print('stack.length():', stack.length())
print('stack.top():', stack.top())
print('stack.pop():', stack.pop())
```



```
print('stack:', stack)
```

An example of the type of output this might produce is

```
stack: ['task1', 'task2', 'task3']
top_element: task3
stack: ['task1', 'task2']
stack: Stack: ['T1', 'T2', 'T3']
stack.is_empty(): False
stack.length(): 3
stack.top(): T3
stack.pop(): T3
stack: Stack: ['T1', 'T2']
```

Chapter 38

Map, Filter and Reduce



38.1 Introduction

Python provides three functions that are widely used to implement Functional Programming style solutions in combinations with collection container types.

These functions are what are known as higher-order functions that take both a collection and a function that will be applied in various ways to that collection.

This chapter introduces three functions `filter()`, `map()` and `reduce()`.

38.2 Filter

The `filter()` function is a higher-order function that takes a function to be used to filter out elements from a collection. The result of the `filter()` function is a new iterable containing only those elements selected by the test function.

That is, the function passed into `filter()` is used to test all the elements in the collection that is also passed into `filter`. Those where the test filter returns `True` are included in the list of values returned. The result returned is a new iterable consisting of all elements of this list that satisfy the given test function. Note that the order of the elements is preserved.

The syntax of the `filter()` function is

```
filter(function, iterable)
```

Note that the second argument to the filter function is anything that implements the iterable protocol which includes all lists, tuples, sets and dictionaries or and many other types, etc.

The function passed in as the first argument is the test function; it can be a lambda (a function defined in line) or the name of an existing function. The result returned will be an iterable that can be used to create an appropriate collection.

Here are some examples of using `filter` with a simple list of integers:

```
data = [1, 3, 5, 2, 7, 4, 10]
print('data:', data)

# Filter for even numbers using a lambda function
d1 = list(filter(lambda i: i % 2 == 0, data))
print('d1:', d1)

def is_even(i):
    return i % 2 == 0

# Filter for even numbers using a named function
d2 = list(filter(is_even, data))
print('d2:', d2)
```

The output from this is

```
Data: [1, 3, 5, 2, 7, 4, 10]
d1: [2, 4, 10]
d2: [2, 4, 10]
```

One difference between the two examples is that it is more obvious what the role is of the test function in the second example as it is explicitly named (i.e., `is_even()`), that is, the function is testing the integer to see whether it is even or not. The in-line lambda function does exactly the same, but it is necessary to understand the test function itself to work out what it is doing.

It is also worth pointing out that defining a named function such as `is_even()` may actually pollute the namespace of the module as there is now a function that others might decide to use even though the original designer of this code never expected anyone else to use the `is_even()` function. This is why lambda functions are often used with `filter()` (and indeed `map()` and `reduce()`).

Of course, you are not just limited to fundamental built-in types such as integers, or real numbers or indeed strings; any type can be used. For example, if we have a class `Person` such as

```
class Person:

    def __init__(self, name, age):
        self.name = name
        self.age = age

    def __str__(self):
        return 'Person(' + self.name +
            ', ' + str(self.age) + ')'
```

Then we can create a list of instances of the class `Person` and then filter out all those over 21:

```
data = [Person('Alun', 54), Person('Niki', 21), Person('Megan',
19)]
for p in data:
    print(p, end=', ')

print('\n-----')

# Use a lambda to filter out People over 21
d3 = list(filter(lambda p: p.age <= 21, data))
for p in d3:
    print(p, end = ', ')
```

The output from this is

```
Person(Alun, 54), Person(Niki, 21), Person(Megan, 19),
-----
Person(Niki, 21), Person(Megan, 19),
```

38.3 Map

Map is another higher-order function available in Python. Map applies the supplied function to all items in the iterable(s) passed to it. It returns a new iterable of the results generated by the applied function.

It is the functional equivalent of a `for` loop applied to an iterable where the results of each iteration round the `for` loop are gathered up.

The `map` function is very widely used within the Functional Programming world and it is certainly worth becoming familiar with it.

The function signature of `map` is

```
map(function, iterable, ...)
```

Note that the second argument to the `filter` function is anything that implements the iterable protocol.

The function passed into the `map` function is applied to each item in the iterable passed as the second argument. The result returned from the function is then gathered up into the iterable object returned from `map`.

The following example applies a function that adds one to a number, to a list of integers:

```
data = [1, 3, 5, 2, 7, 4, 10]
print('data:', data)

# Apply the lambda function to each element in the list
# using the map function
d1 = list(map(lambda i: i + 1, data))
print('d1', d1)

def add_one(i):
    return i + 1

# Apply the add_one function to each element in the
# list using the map function
d2 = list(map(add_one, data))
print('d2:', d2)
```

The output of the above example is

```
data: [1, 3, 5, 2, 7, 4, 10]
d1 [2, 4, 6, 3, 8, 5, 11]
d2: [2, 4, 6, 3, 8, 5, 11]
```

As with the `filter()` function, the function to be applied can either be defined in line as a lambda or it can be named function as in `add_one()`. Either can be used, the advantage of the `add_one` named function is that it makes the intent of the function explicit; however, it does pollute the namespace of functions defined.

Note that more than one iterable can be passed to the `map` function. If multiple iterables are passed to `map`, then the function passed in must take as many parameters as there are iterables. This feature is useful if you want to merge data held in two or more collections into a single collection.

For example, let us assume that we want to add the numbers in one list to the numbers in another list, we can write a function that takes two parameters and returns the result of adding these two numbers together:

```
data1 = [1, 3, 5, 7]
data2 = [2, 4, 6, 8]

result = list(map(lambda x, y: x + y, data1, data2))
print(result)
```

The output printed by this is

```
[3, 7, 11, 15]
```

As with the `filter` function, it is not only built-in types such as numbers that can be processed by the function supplied to `map`; we can also use built-in types such as the class `Person`. For example, if we wanted to collect all the ages for a list of `Person` we could write:

```
data = [Person('John', 54), Person('Phoebe', 21), Person('Adam',
19)]
ages = list(map(lambda p: p.age, data))
print(ages)
```

Which creates a list of the ages of the three people:

```
[54, 21, 19]
```

38.4 Reduce

The `reduce()` function is the last higher-order function that can be used with collections of data that we will look at.

The `reduce()` function applies a function to an iterable and combines the result returned for each element together into a single result.

This function was part of the core Python 2 language but was not included into the core of Python 3. This is partly because Guido van Rossum believed (probably correctly) that the applicability of `reduce` is quite limited, but where it is useful it is very useful. Although it has to be said that some developers try and shoe horn `reduce()` into situations that just make the implementation very hard to understand—remember always aim to keep it simple.

To use `reduce` in Python 3 you need to import it from the `functools` module. One point that is sometimes misunderstood with `reduce()` is that the function passed into `reduce` takes two parameters, which are the previous result and the next value in the sequence; it then returns the result of applying some operation to these parameters.

The signature of the `functools.reduce` function is

```
functools.reduce(function, iterable[, initializer])
```

Note that optionally you can provide an initializer that is used to provide an initial value for the result.

One obvious use of `reduce` is to sum all the values in a list:

```
from functools import reduce

data = [1, 3, 5, 2, 7, 4, 10]
result = reduce(lambda total, value: total + value, data)
print(result)
```

The result printed out for this is 32.

Although it might appear that `reduce` is only useful for numbers such as integers; it can be used with other types as well. For example, let us assume that we want to calculate the average age for a list of people, we could use `reduce` to add together all the ages and then divide by the length of the data list we are processing:

```
data = [Person('John', 54), Person('Phoebe', 21),
        Person('Adam', 19)]

total_age = reduce(lambda running_total, person: running_total +
                  person.age, data, 0)

average_age = total_age // len(data)
print('Average age:', average_age)
```

In this code example, we have a data list of three people. We then use the `reduce` function to apply a lambda to the data list. The lambda takes a `running_total` and adds a person's age to that total. The value zero is used to initialize this running total. When the lambda is applied to the data list, we will add 54, 21 and 19 together. We then divide the final result returned by the length of the data list (3) using the `/` operator which will use floor division to return a whole integer (rather than a real number such as 3.11). Finally, we print out the average age:

```
Average age: 31
```

38.5 Online Resources

More information on *map*, *filter* and *reduce* can be found using the following online resources:

- http://book.pythontips.com/en/latest/map_filter.html Summary of map, filter and reduce.
- https://www.w3schools.com/python/ref_func_map.asp The W3C schools `map()` function tutorial.
- https://www.w3schools.com/python/ref_func_filter.asp The W3 schools `filter()` function tutorial.
- <https://pymotw.com/3/funcitools/index.html> The Python module of the week page including `reduce()`.
- <https://docs.python.org/3/library/funcitools.html> The Python Standard Library documentation for functors including `reduce()`.

38.6 Exercises

This exercise aims to allow you to use `map` and `filter` with your `Stack` class.

Take the `Stack` that you developed in the last chapter and make it iterable. This can be done by implementing the `__iter__()` method from the iterable protocol. As a list is held internally to the `Stack` this could be implemented by returning an iterable wrapper around the list, for example:

```
def __iter__(self):
    return iter(self._list)
```

Now define a function that will check to see if a string passed to it starts with 'Job'; if it does return `True` if not return `False`. Call this function `is_job()`.

Also define a function that will prepend the string 'item': to the string passed in and will then return this as the result of the function. Call this function `add_item()`.

You should now be able to use the `filter` and `map` functions with the `Stack` class as shown below:

```
stack = Stack()
stack.push('Task1')
stack.push('Task2')
stack.push('Job1')
stack.push('Task3')
stack.push('Job2')
print('stack contents:', stack)

# Apply functions to stack contents using map and filter
new_list = list(map(add_item, stack))
print('new_list:', new_list)

filtered_list = list(filter(is_job, stack))
print('filtered_list: ', filtered_list)
```


Chapter 39

Sorting and Higher-Order Functions



39.1 Introduction

There are two built-in ways of sorting within Python: using the `sorted()` function and the `sort()` method. They differ in one significant way, the `sort()` method modifies the receiving container, whereas the `sorted()` function creates a new sorted container leaving the original unchanged. In addition, both of these functions have several *optional* parameters which allow these facilities to be customized to provide two very powerful features.

In this chapter we will look at both the `sorted()` function and the `sort()` method and then consider how they can be customized using `lambda` (or indeed named) functions.

39.2 The `sorted()` Function

The `sorted()` function takes an *iterable* object and returns a sorted list. The values supplied by the iterable are sorted based on their natural ordering; thus for integers and floats this is in increasing (by default) numerical order. In turn for strings this is in increasing alphabetical order.

The basic syntax for the sorted function is:

```
sorted_list = sorted(iterable)
```

Some examples are given below.

The first example takes a tuple of integers and sorts them returning a list of the integers in increasing value:

```
data = (5, 2, 8, 1, 9, 3)
print(data)

sorted_data = sorted(data)
print(sorted_data)
```

The effect of running this code is given below:

```
(5, 2, 8, 1, 9, 3)
[1, 2, 3, 5, 8, 9]
```

Note that although the iterable provided to the `sorted()` function is a tuple the result is a list. This is because the `sorted()` function always returns a list (by default in ascending order).

Remember tuples are immutable containers and thus cannot be changed; thus the `sort()` method is not appropriate here.

We could repeat the process with any iterable, for example, a list:

```
data_list = [7, 5, 9, 2, 1, 3]
print(data_list)

sorted_data_list = sorted(data_list)
print(sorted_data_list)
```

The output from this is now:

```
[7, 5, 9, 2, 1, 3]
[1, 2, 3, 5, 7, 9]
```

Again the sorted list is in ascending order. Of course we can also sort floating point numbers:

```
data_list_2 = [2.3, 5.6, 1.2, 1.1, 5.5]
print(data_list_2)
sorted_data_list_2 = sorted(data_list_2)
print(sorted_data_list_2)
```

This produces the following output:

```
[2.3, 5.6, 1.2, 1.1, 5.5]
[1.1, 1.2, 2.3, 5.5, 5.6]
```

We can also sort strings as they have a natural ordering:

```
string_list = ["John", "Denise", "Phoebe", "Adam"]
print(string_list)

sorted_string_list = sorted(string_list)
print(sorted_string_list)
```

This produces this output:

```
['John', 'Denise', 'Phoebe', 'Adam']
['Adam', 'Denise', 'John', 'Phoebe']
```

As you can see the sorted list is in ascending alphabetic order.

All of the above examples have used lists containing values of the same type. What happens if you mix the types, for example, using integers and floating point numbers:

```
mixed_list = [5, 1.2, 7.34, 4]
print(mixed_list)
sorted_mixed_list = sorted(mixed_list)
print(sorted_mixed_list)
```

This generates the following output:

```
[5, 1.2, 7.34, 4]
[1.2, 4, 5, 7.34]
```

Actually, it is fine to mix any types as long as they understand the '<' operator between those types, and thus the Boolean value True can also be used with integers and floating point numbers as it has a logical integer value of 1:

```
mixed_list = [5, 1.2, 7.34, True, 4]
print(mixed_list)
sorted_mixed_list = sorted(mixed_list)
print(sorted_mixed_list)
```

This produces the output:

```
[5, 1.2, 7.34, True, 4]
[True, 1.2, 4, 5, 7.34]
```

However, it is not possible to mix numbers and strings together as it is not possible to compare a string and an integer or a floating point number using the '<' less than operator, for example, if you try this:

```
mixed_list = [5, 1.2, 7.34, True, 'John', 4]
print(mixed_list)
sorted_mixed_list = sorted(mixed_list)
print(sorted_mixed_list)
```

You will get a runtime error indicating that the '<' less than operator is not supported between a string and an int:

```
Traceback (most recent call last):
  File      "/Users/Shared/workspaces/pycharm/cern-python-
programming-samples/08-further-functions/sorting_examples.py",
  line 26, in <module>
    sorted_mixed_list = sorted(mixed_list)
    ~~~~~^~~~~~
TypeError: '<' not supported between instances of 'str' and 'int'
```

39.3 Sorting in Reverse Order

The examples in the previous section all show how to sort numbers and strings into ascending order, what about if you want to sort them into descending order? This is very simply done with the `sorted()` function as it takes an optional keyword parameter `reverse`. If this is set to `True` then the resulting sorted lists will be sorted into descending order.

Thus the syntax for the `sorted()` function can now be extended to be:

```
descending_sorted_list = sorted(iterable, reverse=<bool>)
```

Notice we are using *named* parameter passing for the `reverse` parameter, that is, because there is a third parameter between these two which we will introduce later in this chapter.

The examples from the previous section are rerun below but with the `reverse` parameter set to `True` (it is `False` by default).

```
data = (5, 2, 8, 1, 9, 3)
print(data)

sorted_data = sorted(data, reverse=True)
print(sorted_data)
print('--' * 25)

data_list = [7, 5, 9, 2, 1, 3]
print(data_list)

sorted_data_list = sorted(data_list, reverse=True)
print(sorted_data_list)
print('--' * 25)

data_list_2 = [2.3, 5.6, 1.2, 1.1, 5.5]
print(data_list_2)
sorted_data_list_2 = sorted(data_list_2, reverse=True)
print(sorted_data_list_2)
print('--' * 25)

string_list = ["John", "Denise", "Phoebe", "Adam"]
print(string_list)

sorted_string_list = sorted(string_list, reverse=True)
print(sorted_string_list)
print('--' * 25)
```

The output from running this code is shown below:

```
(5, 2, 8, 1, 9, 3)
[9, 8, 5, 3, 2, 1]
-----
[7, 5, 9, 2, 1, 3]
[9, 7, 5, 3, 2, 1]
-----
[2.3, 5.6, 1.2, 1.1, 5.5]
```

```
[5.6, 5.5, 2.3, 1.2, 1.1]
-----
['John', 'Denise', 'Phoebe', 'Adam']
['Phoebe', 'John', 'Denise', 'Adam']
-----
```

As you can see each of the sorted results lists is now in descending order.

This also works with mixed data types (again along the types support the '<' operator between themselves), for example:

```
mixed_list = [5, 1.2, 7.34, True, 4]
print(mixed_list)
sorted_mixed_list = sorted(mixed_list, reverse=True)
print(sorted_mixed_list)
```

Produces:

```
[5, 1.2, 7.34, True, 4]
[7.34, 5, 4, 1.2, True]
```

39.4 The sort() Method

As mentioned at the start of this chapter, there is an alternative way to sort (some) containers, that is, the `sort()` method. This method is invoked using the *dot notation* and is applied directly to the mutable container being used such as a `list`. It is not available on immutable containers such as `tuples` as it is not possible to modify them. However, on a mutable container such as a `list` it allows the original list to be sorted in situ. For example:

```
data_list = [5, 2, 8, 1, 9, 3]
print(data_list)
data_list.sort()
print(data_list)
```

Note that it is not necessary to store the result of invoking the `sort()` method into a new variable; this avoids additional variables being created. However, the original list in the variable `data_list` has now been sorted into ascending order using the natural ordering of the types held within the list. As these are *integers* they will be sorted into ascending numerical order, for example:

```
[5, 2, 8, 1, 9, 3]
[1, 2, 3, 5, 8, 9]
```

It is also possible to sort floating point numbers and strings in this way. For example a list of floating point numbers can be sorted into ascending order:

```
data_list_2 = [2.3, 5.6, 1.2, 1.1, 5.5]
print(data_list_2)
data_list_2.sort()
print(data_list_2)
```

which produces the following output:

```
[2.3, 5.6, 1.2, 1.1, 5.5]
[1.1, 1.2, 2.3, 5.5, 5.6]
```

Strings can also be sorted into ascending alphabetical order:

```
string_list = ['John', 'Denise', 'Phoebe', 'Adam']
print(string_list)
string_list.sort()
print(string_list)
print('-' * 25)
```

which when run generates:

```
['John', 'Denise', 'Phoebe', 'Adam']
['Adam', 'Denise', 'John', 'Phoebe']
```

We can also mix types again as long as they support the less than operator ('<'), for example:

```
mixed_list = [5, 1.2, 7.34, True, 4]
print(mixed_list)
mixed_list.sort()
print(mixed_list)
```

which produces:

```
[5, 1.2, 7.34, True, 4]
[True, 1.2, 4, 5, 7.34]
```

39.5 Reversing the sort() Method

As with the `sorted()` function it is possible to get the `sort()` method to sort the receiver into descending order using the `reverse` parameter, for example:

```
data_list = [5, 2, 8, 1, 9, 3]
print(data_list)
data_list.sort(reverse=True)
```

```
print(data_list)
```

Note that again we are using the *named* parameter passing style for the `reverse` parameter. This is again because there are other parameters before this one in the method which we are not using.

The output from this code is:

```
[5, 2, 8, 1, 9, 3]
[9, 8, 5, 3, 2, 1]
```

We could repeat this with each of the lists in the previous section, including the mixed list, for example:

```
data_list_2 = [2.3, 5.6, 1.2, 1.1, 5.5]
print(data_list_2)
data_list_2.sort(reverse=True)
print(data_list_2)
print('-' * 25)

string_list = ['John', 'Denise', 'Phoebe', 'Adam']
print(string_list)
string_list.sort(reverse=True)
print(string_list)
print('-' * 25)

mixed_list = [5, 1.2, 7.34, True, 4]
print(mixed_list)
mixed_list.sort(reverse=True)
print(mixed_list)
```

This produces:

```
[5, 1.2, 7.34, True, 4]
[7.34, 5, 4, 1.2, True]
```

39.6 Changing the Sort Key

So far all our examples have used the natural ordering of the values held in the iterable to be sorted. What if you want to order based on some other criteria, for example, based on the length of the string rather than the alphabetical order of strings? You can do this by specifying a function (either a named function or a lambda function) to be used to determine the value to be used when comparing elements in the list.

For example, the following code illustrates how a list of strings is sorted by default as well as how it can be sorted based on string length:

```
string_list = ['John', 'Denise', 'Phoebe', 'Adam']
print(f'string list: {string_list}')
sorted_string_list1 = sorted(string_list)
print(f'default ordering: {sorted_string_list1}')
```

```
# Sorted based on string length
sorted_string_list2 = sorted(string_list, key=lambda s: len(s))
print(f'sorting based on string length: {sorted_string_list2}')
```

The output from this code is:

```
string list: ['John', 'Denise', 'Phoebe', 'Adam']
default ordering: ['Adam', 'Denise', 'John', 'Phoebe']
sorting based on string length: ['John', 'Adam', 'Denise',
'Phoebe']
```

Note that the `key` parameter is provided with the function that takes one value (the value being considered) and returns a value to be used for the '`<`' (less than) comparison. The output of the sorted function thus changes from:

```
['Adam', 'Denise', 'John', 'Phoebe']
```

which is based on the alphabetic ordering of the strings to:

```
['John', 'Adam', 'Denise', 'Phoebe']
```

which is based on the length of the strings and thus 'John' is before 'Denise', etc.

This approach can also be used with the `sort()` method as it also has a `key` parameter as shown below:

```
string_list = ['John', 'Denise', 'Phoebe', 'Adam']
print(f'string list: {string_list}')
string_list.sort(key=lambda s: len(s))
print(f'sorting based on string length: {string_list}')
```

The output from this is:

```
string list: ['John', 'Denise', 'Phoebe', 'Adam']
sorting based on string length: ['John', 'Adam', 'Denise',
'Phoebe']
```

39.7 Sorting Containers of Containers

The approach of using the `key` parameter to sort elements within an iterable can be very useful where the iterable contains other containers. By default an iterable will be sorted based on the result of comparing two iterables using the '`<`' less than operator. For tuples for examples this will try to compare the tuples using their natural ordering which is probably not what you want. For example in the following code the inner tuples represent the marks achieved by a group of students, the instructor wants to order the inner tuples based on the grade achieved, and however just using `sorted()` or indeed `sort()` does not achieve the desired result:


```
data = [('John', 55), ('Denise', 90), ('Phoebe', 76)]
print(data)
sorted_data = sorted(data)
print(sorted_data)
```

The result of running this code is:

```
[('John', 55), ('Denise', 90), ('Phoebe', 76)]
[('Denise', 90), ('John', 55), ('Phoebe', 76)]
```

The ordered list is not sorted based on the grades. However, if we now use the `key` parameter we can provide a function that will extract the second value in the tuple (with index 1) which can then be used in the ordering comparison:

```
sorted_data2 = sorted(data, key=lambda i: i[1])
print(sorted_data2)
```

The result of these two lines is:

```
[('John', 55), ('Phoebe', 76), ('Denise', 90)]
```

We can of course reverse this as well using the `reverse` parameter:

```
sorted_data3 = sorted(data, key=lambda i: i[1], reverse=True)
print(sorted_data3)
```

With the resulting output:

```
[('Denise', 90), ('Phoebe', 76), ('John', 55)]
```

We could have used the `sort()` method in exactly the same way:

```
print('-' * 25)
data = [('John', 55), ('Denise', 90), ('Phoebe', 76)]
print(data)
data.sort(key=lambda i: i[1])
print(data)
data.sort(key=lambda i: i[1], reverse=True)
print(data)
```

The output from this is:

```
[('John', 55), ('Denise', 90), ('Phoebe', 76)]
[('John', 55), ('Phoebe', 76), ('Denise', 90)]
[('Denise', 90), ('Phoebe', 76), ('John', 55)]
```

39.8 Exercises

The aim of this exercise is to use the sorting function to handle temperature data.

The data you will be working with is a list of temperatures, for example:

```
temperatures = [12.4, 18.2, 15.4, 17.8, 13.1, 19.0, 12.3]
print(f'temperatures: {temperatures}')
```

You should now perform each of the following steps:

- Print out the result of sorting the temperatures into ascending order (without modifying the original list).
- Print out the result of sorting the temperatures into descending order (without modifying the original list).
- Print out the result of sorting the original list into ascending order.
- Print out the result of sorting the original list into descending order.

The output from this might look like:

```
temperatures: [12.4, 18.2, 15.4, 17.8, 13.1, 19.0, 12.3]
sorted temperatures: [12.3, 12.4, 13.1, 15.4, 17.8, 18.2, 19.0]
sorted temperatures descending order: [19.0, 18.2, 17.8, 15.4,
13.1, 12.4, 12.3]
sorted temperatures: [12.3, 12.4, 13.1, 15.4, 17.8, 18.2, 19.0]
sorted temperatures descending order: [19.0, 18.2, 17.8, 15.4,
13.1, 12.4, 12.3]
```

Next we will record the date and time of the temperature reading within a tuple, for example:

```
records = [('Monday', '12:00', 12.4),
           ('Tuesday', '12:00', 18.2),
           ('Wednesday', '13:00', 15.4),
           ('Thursday', '11:45', 17.8),
           ('Friday', '12:15', 13.1)]
print(f'records: {records}')
```

Now sort the records such that you can:

1. Print out the result of sorting the records into ascending order (without modifying the original list).
2. Print out the result of sorting the records into descending order (without modifying the original list).
3. Print out the result of sorting the records list into ascending order.
4. Print out the result of sorting the records list into descending order.

The output from this might look like:

```
records: [('Monday', '12:00', 12.4), ('Tuesday', '12:00', 18.2),
          ('Wednesday', '13:00', 15.4), ('Thursday', '11:45', 17.8),
          ('Friday', '12:15', 13.1)]
```

```
sorted records: [('Monday', '12:00', 12.4), ('Friday', '12:15',
13.1), ('Wednesday', '13:00', 15.4), ('Thursday', '11:45', 17.8),
('Tuesday', '12:00', 18.2)]
sorted records descending order: [('Tuesday', '12:00', 18.2),
('Thursday', '11:45', 17.8), ('Wednesday', '13:00', 15.4),
('Friday', '12:15', 13.1), ('Monday', '12:00', 12.4)]
sorted temperatures: [('Monday', '12:00', 12.4), ('Friday',
'12:15', 13.1), ('Wednesday', '13:00', 15.4), ('Thursday',
'11:45', 17.8), ('Tuesday', '12:00', 18.2)]
sorted temperatures descending order: [('Tuesday', '12:00',
18.2), ('Thursday', '11:45', 17.8), ('Wednesday', '13:00', 15.4),
('Friday', '12:15', 13.1), ('Monday', '12:00', 12.4)]
```

Chapter 40

Python Enumerated Values



40.1 Introduction

Introduced in Python 3.4 the `enum` module provides support for enumerated values or enums. This facility is a very useful one for creating sets of ordered values. In this chapter we discuss what an enumeration is, the two ways of creating an enumeration in Python, how enumerations are used and conclude by introducing `IntEnums` for integer-based enumeration.

40.2 What is an Enumeration?

An enumeration is a collection of entities that represent a complete, ordered set of all of the values in a set. For example, the days of the week could be represented by an enumeration as there is a finite set of days in the (working) week, and there is an ordering to the days of the week.

For example, Monday comes before Tuesday, and Wednesday comes after Tuesday, etc. In most languages it is then possible to create a variable that can hold a value from an enumeration and perform certain tests on it such as checking for *equality* or *in equality* as well as ordering style tests.

More formally, an enumeration is:

A set of symbolic names (members) bound to a given type (class) where each value has a unique value, where the values are ordered.

The values in an enumeration can be iterated over to return its members in definition order.

In Python enumerations can be created in two different ways, either using the class syntax or using the function call syntax. We will look at both approaches in this chapter.

40.3 Creating an Enumeration Using Class Syntax

A new enumerated type (class) can be created by subclassing the `enum.Enum` class. That is, the `enum` module provides a class specifically to be used with inheritance, allowing the programmer to create a concrete enumeration type.

For example, the following code creates a new enumerated type `DaysOfWeek` by subclassing the `Enum` type:

```
class DaysOfWeek (Enum) :
    MONDAY = 1
    TUESDAY = 2
    WEDNESDAY = 3
    THURSDAY = 4
    FRIDAY = 5
```

This creates a new type (or class) of things called `DayOfWeek`. Within this type there are five values defined representing each of the days of the week. These elements have an ordering (as indicated by the default values given to them). They are also constants and thus cannot be changed which is why they are defined using all uppercase (although the use of all uppercase letters is only a convention). The result is that you can refer to `DaysOfWeek` as an enumeration or more typically as an *Enum*.

Using the enumeration terminology we would say that `DaysOfWeek.Monday` and `DaysOfWeek.Tuesday` are enumeration (or enum) members. Each member has a *name* and a *value*, and the enum itself contains these enum *members*.

You can reference the individual members using the dot notation as shown below:

```
print (DaysOfWeek.MONDAY)
print (DaysOfWeek.TUESDAY)
print (DaysOfWeek.WEDNESDAY)

today = DaysOfWeek.THURSDAY
print (today)
```

In the above code we are printing out the values for Monday, Tuesday and Wednesday. We then create a variable `today` and store the value of Thursday into that variable and print it out. The output generated may not be what you expect:

```
DaysOfWeek.MONDAY
DaysOfWeek.TUESDAY
DaysOfWeek.WEDNESDAY
DaysOfWeek.THURSDAY
```

Notice that what is printed out is not the values 1, 2, 3 and 4 as you might have expected; instead the name (or label) associated with the enum member is printed out.

Also notice that an enum member can be assigned to a variable and used within your program just like any other value such as 43, True or 4.123.

Note if you did try and change the value associated with an *enum member* you would generate an `AttributeError`, for example:

```
DaysOfWeek.MONDAY = 6
```

Generates

```
Traceback (most recent call last):
  File "enum_examples.py", line 18, in <module>
    DaysOfWeek.MONDAY = 6
    ^^^^^^^^^^^^^^^^^^^^^
  File "enum.py", line 841, in __setattr__
    raise AttributeError('cannot reassign member %r' % (name, ))
AttributeError: cannot reassign member 'MONDAY'
```

40.4 Enum Values

It is also possible to create enums where the values used by the enum members are not 1, 2, 3 but any integer sequence you like. For example:

```
class Compass(Enum) :
    NORTH = 0
    EAST = 90
    SOUTH = 180
    WEST = 270

print(Compass.SOUTH)
```

In this example we have created a `Compass` enum with four members, and the values of these members are 0, 90, 180 and 270 (that is the points on a compass). When we run this code the output of course is just:

```
Compass.SOUTH
```

It is also possible to use string values for your Python enums. Here is another enum using strings:

```
class Sizes(Enum) :
    S = "small"
    M = "medium"
    L = "large"
    XL = "extra large"

print(Sizes.XL)
```

The output from this is:

```
Sizes.XL
```

Finally, it is also possible to create an enum of Boolean values, for example:

```
class SwitchPosition(Enum):
    ON = True
    OFF = False

switch = SwitchPosition.OFF
print(switch)
```

This produces:

```
SwitchPosition.OFF
```

Although it should be noted that this implies an ordering between ON and OFF!

It should also be noted that integer-based values can be in any order, and they do not need to progress 1, 2, 3 and can be whatever order is required; for example, the Compass enum could have been defined as:

```
class Compass(Enum):
    NORTH = 0
    SOUTH = 180
    EAST = 90
    WEST = 270
```

40.5 Using the Name and Value Attributes

Each enum member has an associated name and value property. These can be accessed and used for example to display options to a user or to use the value for some operation.

For example, we could access the name and value attributes for the Compass enum as shown below:

```
print(Compass.SOUTH.name)
print(Compass.SOUTH.value)
```

The output from this code is:

```
SOUTH
180
```

40.6 Creating an Enumeration Using the Function Call Syntax

It is not required to subclass the `Enum` class to create an enumerated type in Python; you can also use the functional interface. To do this you can use the `Enum` functional API from the `enum` module. This function takes the following parameters:

```
Enum (
    value,
    names,
    module=None,
    qualname=None,
    type=None,
    start=1
)
```

This returns a new enumerated type based on the parameters passed to it. The parameters have the following meaning:

- `value` (required field) represents the string name of the new enumeration type (class).
- `names` (required field) is a list of the member values that will be used for the enumeration.
- `module` (optional) takes the name of the module that defines the enumeration type.
- `qualname` (optional) holds the location of the module that defines the enumeration type.
- `type` (optional) holds the class to use as the parent for the enumerated type.
- `start` (optional) takes the starting value for the enumeration (by default it is 1).

The `names` argument can take any of the following:

- a string containing member names separated either with spaces or commas,
- an iterable (such as a list) of member names,
- an iterable (such as a list) of name-value pairs.

Most of the time, you'll just use the first two arguments (and possibly the last) to the `Enum` function when creating your enumerations. For example:

```
Weekend = Enum('Weekend', ['SATURDAY', 'SUNDAY'])
```

Alternatively if you wanted to start the enumeration from 10 rather than 1 you could write:

```
Weekend = Enum('Weekend',
               ['SATURDAY', 'SUNDAY'],
               start=10)
```


In both cases you can reference the new enum via the variable `Weekend` and access the enum members (`SATURDAY` and `SUNDAY`) as you will do normally, for example:

```
day = Weekend.SATURDAY
print(day)
```

The output from this is:

```
Weekend.SATURDAY
```

It is also possible to specify custom values for each of the members in the enumerated type. This is done using an iterable of name-value pairs (below we are using a list and a tuple for this). For example:

```
HTTPStatusCode = Enum(
    value="HTTPStatusCode",
    names=[
        ("OK", 200),
        ("CREATED", 201),
        ("BAD_REQUEST", 400),
        ("NOT_FOUND", 404),
        ("SERVER_ERROR", 500),
    ],
)

print(HTTPStatusCode.OK)
```

The output from this is:

```
HTTPStatusCode.OK
```

40.7 Creating Enumerations from Automatic Values

The `enum` module also provides an `auto()` function that can be used to automatically assign values to enumeration members. It does so by incrementing the previous value used. This function can be used in combination with user specified values to fill in the gaps, for example:

```
class Day(Enum):
    MONDAY = 5
    TUESDAY = auto()
    WEDNESDAY = 9
    THURSDAY = auto()
    FRIDAY = auto()
    SATURDAY = auto()
```

```
SUNDAY = 15

print(list(Day))
```

In this example MONDAY, WEDNESDAY and SUNDAY have had explicit integer literals specified for their values. However, TUESDAY, THURSDAY, FRIDAY and SATURDAY will all be autogenerated based on an increment of the preceding value. Thus the values of the enumerated members will be 5, 6, 9, 10, 11, 12 and 15.

To illustrate this all the values in the enumeration are gathered up and stored into a list which is then printed out. The result is:

```
[<Day.MONDAY: 5>, <Day.TUESDAY: 6>, <Day.WEDNESDAY: 9>,
<Day.THURSDAY: 10>, <Day.FRIDAY: 11>, <Day.SATURDAY: 12>,
<Day.SUNDAY: 15>]
```

40.8 Working with Enum Aliases

It is also possible to create enumerations in which two or more members have the same constant value; the redundant enum member is known as an *alias*. This can be useful in situation such as where the different names have the semantic meaning, for example:

```
class OperatingSystem(Enum):
    UBUNTU = "linux"
    MACOS = "darwin"
    WINDOWS = "win"
    DEBIAN = "linux"

print(OperatingSystem.UBUNTU)
print(OperatingSystem.DEBIAN)
print(OperatingSystem.MACOS)
print(OperatingSystem.WINDOWS)
```

In this example both DEBIAN and UBUNTU have the value “linux”; this DEBIAN is an alias for UBUNTU.

The result of running this code is:

```
OperatingSystem.UBUNTU
OperatingSystem.UBUNTU
OperatingSystem.MACOS
OperatingSystem.WINDOWS
```

40.9 Enumerations Unique Values

The previous section illustrates that two values within an enumeration can have, by default, the same value. However, what if this is not what you want then this could happen by accident! To avoid this you can mark the enumeration class with an `@enum.unique` decorator. This ensures that each enum member will have a unique value. If not then a `ValueError` will be generated, for example:

The following code is legal in Python:

```
class Lights(Enum):
    RED = 1
    AMBER = 2
    GREEN = 3
    RED_AMBER = 1

print(Lights.RED)
```

However, if we now mark it with `@unique` from the `enum` module as shown below:

```
@unique
class Lights(Enum):
    RED = 1
    AMBER = 2
    GREEN = 3
    RED_AMBER = 1

print(Lights.RED)
```

Then when we now try and run this program we will get a `ValueError` being generated:

```
Traceback (most recent call last):
  File "enum_examples.py", line 97, in <module>
    @unique
    ^^^^^^^
  File "enum.py", line 1564, in unique
    raise ValueError('duplicate values found in %r: %s' %
ValueError: duplicate values found in <enum 'Lights'>: RED_AMBER ->
RED
```

Thus illustrating that this enum must have unique values for each of its members.

40.10 Iterating Through Enumeration Values

Python enums are *iterable* things; this means that you can iterate through all the enum members very easily.

For example, to iterate through all the members within the `DaysOfWeek` enum we can write a for loop as shown below:

```
for day in DaysOfWeek:
    print(day)
```

The output from this is that each enum member will be printed out in turn, for example:

```
DaysOfWeek.MONDAY
DaysOfWeek.TUESDAY
DaysOfWeek.WEDNESDAY
DaysOfWeek.THURSDAY
DaysOfWeek.FRIDAY
```

40.11 Comparing Enumerations

By default enumerated types support two types of comparison operators:

- Identity, using the `is` and `is not` operators.
- Equality, using the `==` and `!=` operators.

The identity comparison relies on the fact that each enum member is a unique instance of its enumeration class (in object-oriented design terminology it is a singleton). This means that very fast identity comparisons can be performed using the `is` and `is not` operators.

For example:

```
class OperatingSystem(Enum):
    UBUNTU = "linux"
    MACOS = "darwin"
    WINDOWS = "win"
    DEBIAN = "linux"

os1 = OperatingSystem.UBUNTU
os2 = OperatingSystem.MACOS
os3 = OperatingSystem.UBUNTU
print(os1 is os3)
print(os1 is os2)
print(os1 is not os2)
```

The above checks that each of the `OperatingSystem` enumerated values in the variables `os1`, `os2` and `os3` is with the exact same instance of the `OperatingSystem` type or different. This is very fast and efficient.

This works because every enum member has its own identity, which is different from the identity of the other members of the enum. However, this rule doesn't apply to member aliases, because they're just references to existing members and share the same identity. Thus:

```
os1 = OperatingSystem.UBUNTU
os4 = OperatingSystem.DEBIAN
print(os1 is os4)
```

Also generates:

```
True
```

We mentioned at the start of this section that enums also support `==` and `!=` this is because they delegate their operation to `is` and `is not`. Thus you can use either `is` or `==` and you can use `is not` or `!=` to compare enums.

40.12 Integer Enumerations with IntEnum

Integer-based enumerations are so common that there is a special type called `IntEnum` that is provided by the `enum` module. This class is useful if you need the members of your enumerations to behave like integer numbers with comparison operators. Using `IntEnum` allows your enum members to be compared using `<`, `>`, etc. which can be very useful.

For example:

```
from enum import IntEnum

class Size(IntEnum):
    XS = 1
    S = 2
    M = 3
    L = 4
    XL = 5
    XXL = 6

print(Size.XS < Size.L)
print(Size.XL > Size.S)
print(Size.XL >= Size.S)
```

The output from this is:

```
True
True
True
```

40.13 Exercises

Write a short program to create an enumerated type for medals at the Olympics, that is, Gold, Silver and Bronze. Gold should come before Silver and Silver before Bronze. You should then extend your code so that an athlete can be awarded one of these medals. The program should ask the user for each athletes name in turn and then award the first three athletes with a gold, silver or bronze medal. The program can terminate when no athlete name is entered (i.e., the user just presses enter).

The output form of this program might look like:

```
Starting
Please enter the athletes name: Natalia
Natalia obtained Medals.GOLD
Please enter the athletes name: Adam
Adam obtained Medals.SILVER
Please enter the athletes name: Denise
Denise obtained Medals.BRONZE
Please enter the athletes name: John
John obtained no medal
Please enter the athletes name:
Done
```

Chapter 41

Structural Pattern Matching



41.1 Introduction

A new feature introduced in Python 3.10 was *structural pattern matching*. In this chapter we will look at the use of the `match` keyword and how it can be used with literal values, variables and containers such as lists and tuples.

41.2 Basic Pattern Matching

In its most basic form the `match` statement acts very much like `match` and `case` statements from other languages such as Java or JavaScript. That is, it is used to compare a value (typically held in a variable) with a set of literal (or fixed) values. As such it is intended as a more concise form of a conditional statement. For example, let us say that we want to display an appropriate welcome message depending on the users' language. To do this we could write a traditional `if` statement such as:

```
def greeter(language):
    if language == 'English':
        print('Hello')
    elif language == 'Spanish':
        print('Hola')
    elif language == 'French':
        print('Bonjour')
    elif language == 'German':
        print('Hallo')

greeter('Spanish')
```

In this initial function `greeter()` an `if` statement is used to test what the value of the parameter `language` is; the appropriate welcome is then printed out. In this case the `greeter` function is called with the value 'Spanish' and thus the output is:

```
Hola
```

However, the `if` statement is quite long winded, and it's not immediately obvious that all of the conditional elements of the `if` statement relate to the value of the parameter `language`. To see this it is necessary to read all the code. In this case that is not too difficult as each of the behaviours associated with each condition is just one line, but in some other situations it may be far less clear.

The basic use of the `match` statement tries to make this more explicit. The equivalent of the above function rewritten using a `match` statement is given below:

```
def greeter2(language):
    match language:
        case 'English':
            print('Hello')
        case 'Spanish':
            print('Hola')
        case 'French':
            print('Bonjour')
        case 'German':
            print('Hallo')

greeter2('Spanish')
```

The behaviour of this `greeter2()` function is exactly the same as the original `greeter()` function, but the implementation now uses the `match` syntax. The syntax for the `match` statement is:

```
match <value>:
    case <test>:
        actions
    ...
```

Thus it is easy to see that the whole `match` statement relates to the current value of the `language` parameter (as this is what follows the `match` keyword).

As always with Python, note the indentation! In Python the layout of the `match` statement is critical, with all `case` conditions indented relative to the `match` keyword and the actions to be performed when a case condition is met, indented another level again.

In this situation we can see that if the value of `language` is 'English' we print 'Hello', whereas if it is 'Spanish' we print 'Hola' and so on. Note that at most only one case condition will be triggered, thus once one case condition has run the whole `match` statement terminates, additional case conditions are not evaluated.

Also note that there can be any number of statements within the body of the case condition, for example:

```
def greeter2a(language):
    match language:
        case 'English':
            print('Hello')
        case 'Spanish':
            print('Hola')
            print('Buenas')
        case 'French':
            print('Bonjour')
        case 'German':
            print('Hallo')

greeter2a('Spanish')
```

The output from this is:

```
Hola
Buenas
```

41.3 Default Case

One issue with the above use of the `match` statement is that if a value is passed into the `match` which is not one of those listed, then nothing happens, there is no failure but there is also no action. This may of course be what you want, but if you need to pick up on this (possibly as it is unexpected behaviour) then you can use a *wildcard*. A wildcard must be the last `case` condition in the `match` statement and is represented by an underbar `'_'`, for example:

```
def get_status_message(status):
    match status:
        case 400:
            return 'Bad request'
        case 404:
            return 'Not found'
        case 418:
            return "I'm a teapot"
        case _:
            return 'Something is wrong'

print(get_status_message(404))
print(get_status_message(401))
```

The output from this example is:

```
Not found
Something is wrong
```

In this case the first time `get_status_message()` is called the string ‘Not found’ is returned, whereas when it is called with the value 401, there are no explicit case conditions matching that value and thus the wildcard case condition (the last one) runs. This means that the string ‘Something is wrong’ is returned (and subsequently printed out).

41.4 Capture Default Case

In some situations it would be really useful to capture the unknown value in a match statement and log that (possibly for further investigation at a later date). To do this instead of using the wildcard from the previous section, you can use a *variable*. The variable will be populated with the value that would have been compared in the case condition and that is then available within the case condition block of code. For example:

```
def greeter4(language):
    match language:
        case 'English':
            print('Hello')
        case 'Spanish':
            print('Hola')
        case 'French':
            print('Bonjour')
        case 'German':
            print('Hallo')
        case unknown:
            print(f'Language {unknown} is not yet known to us')

greeter4('Welsh')
```

In this case if the language specified is not ‘English’, ‘Spanish’, ‘French’ or ‘German’, then the value used will be bound to the variable `unknown`, and then this can be logged in the print statement.

Thus the output from this code example is:

```
Language Welsh is not yet known to us
```

Note that the name of the variable used in the case condition can be any legal Python variable; here we have used ‘unknown’ just as an example.

41.5 Alternative Values in Case Condition

Another option is to use an *or* condition (represented by a `|`) to indicate that if a value equals one or more values then the associated action should be performed. For example:

```
def greeter5(language):
    match language:
        case 'English' | 'American':
            print('Hello')
        case 'Spanish':
            print('Hola')
        case 'French':
            print('Bonjour')
        case 'German':
            print('Hallo')
        case unknown:
            print(f'Language {unknown} is not yet known to us')

greeter5('English')
greeter5('American')
```

In the above example, if the language is either 'English' or 'American' then the string 'Hello' will be printed out. The output from running this code is:

```
Hello
Hello
```

Note that there can be any number of values so if we wanted to use 'English', 'American' and 'Australian' we could write:

```
case 'English' | 'American' | 'Australian':
```

41.6 Matching Sequences/Containers

A sequence is something that contains a set of values such as a container/collection. A tuple is an example of a sequence or container. The structural aspect of the `match` statement allows you to match against the contents of a tuple, but also to structurally extract values from a tuple.

The following example illustrates several options,

- The first `case` matches the exact values in a tuple.
- The second matches on the first element of a two value tuple and extracts the *second* value and stores it into a local variable `y`.
- The third case condition does the opposite, it extracts the *first* value in the tuple and matches against the `0` in the second position.

- The fourth case condition can extract both the x and y value.
- The final case condition is the wildcard.

Note, however, that only one of these case conditions will actually execute. The one that does execute will be the first case condition to return `True`. Thus if a tuple of two zeros is entered then it will be the first case condition that *fires*. However, if the tuple passed in contains a 5 and a zero then it will be the third case condition that is met and the associated action that is performed.

```
def locator(point):
    match point:
        case (0, 0):
            print('Corner')
        case (0, y):
            print(f'Left edge at x=0, y={y}')
        case (x, 0):
            print(f'Top edge at x={x}, y=0')
        case (x, y):
            print(f'Point in grid at x={x}, y={y}')
        case _:
            raise ValueError('Not a point tuple')
```

This function can be used by passing in a tuple as a parameter, for example:

```
cursor = (0, 0)
locator(cursor)
cursor = (0, 10)
locator(cursor)
cursor = (10, 0)
locator(cursor)
cursor = (10, 5)
locator(cursor)
```

The output from this code snippet is:

```
Corner
Left edge at x=0, y=10
Top edge at x=10, y=0
Point in grid at x=10, y=5
```

41.7 Structural Matching and Lists

It is also possible to use lists with the match statement (they are also sequences/containers). Again it is possible to extract one or more values from the list:

```
def banking_operation1(action):
    match action:
        case ['balance']:
            print('Getting the balance')
        case ['deposit', amount]:
            print(f'Depositing {amount}')
        case ['withdraw', amount]:
            print(f'Withdrawing {amount}')
```

In the above example their action passed into the function is expected to be a list. Each of the case conditions checks for a particular pattern and in two cases extracts a value:

- The first case condition checks for a list containing a single string 'balance'.
- The second case condition checks for a list containing two values where the first value is the string 'deposit' and the second value will be extracted and stored in the variable `amount`. This can then be used in the body of the case condition.
- The third case condition looks for a list containing the string 'withdraw' and a second value which will be extracted into the variable `amount` which will then be available in the body of the case condition.

This function can be used by passing in a list of values when the function is invoked, for example:

```
banking_operation1(['balance'])
banking_operation1(['deposit', 12.55])
banking_operation1(['withdraw', 9.99])
```

The output from this code snippet is:

```
Getting the balance
Depositing 12.55
Withdrawing 9.99
```

We can be a little bit more flexible than this by indicating that a case condition can match on a list containing one or *many* values. This is done by preceding the variable name to be used for the extraction with an *asterisk*. This indicates that this variable will be a container that will contain zero or more matched value, for example:

```
case ['deposit', *amounts]:
    print(f'Depositing multiple amounts - {amounts}')
```

This small code sample indicates that this case condition will match onto a list containing the string 'deposit' and zero or more following values.

Placing this within a modified `bank_operation` function we get:

```
def banking_operation2(action):
    match action:
        case ['balance']:
            print('Getting the balance')
        case ['deposit', amount]:
            print(f'Depositing {amount}')
        case ['withdraw', amount]:
            print(f'Withdrawing {amount}')
        case ['deposit', *amounts]:
            print(f'Depositing multiple amounts - {amounts}')
```

Notice it has been placed after the earlier deposit case condition as the new case condition will match onto a two value list as well as a three four or more value list, etc.

If we call this with a few different options we can see the behaviour of the case condition:

```
banking_operation2(['deposit', 2.51, 9.44])
banking_operation2(['deposit', 2.51, 9.44, 3.21, 5.76])
banking_operation2(['deposit'])
```

The output from this code sample is:

```
Depositing multiple amounts - [2.51, 9.44]
Depositing multiple amounts - [2.51, 9.44, 3.21, 5.76]
Depositing multiple amounts - []
```

Note how the last example contains only a string 'deposit' but still matches on this new case condition.

41.8 Matching Alternative Sequences/Containers

It is also possible to match onto two or more list patterns in the same case condition using the *or* style case operation '|'. For example, if we want to do the same thing whether a list contains the string 'help' or the string 'info' we can write:

```
case ['help'] | ['info']:
    print('Help information')
```

This is shown used in the `banking_operation` function below:

```
def banking_operation3(action):
    match action:
        case ['help'] | ['info']:
            print('Help information')
        case ['balance']:
            print('Getting the balance')
        case ['deposit', amount]:
```

```

    print(f'Depositing {amount}')
case ['withdraw', amount]:
    print(f'Withdrawing {amount}')
case ['deposit', *amounts]:
    print(f'Depositing multiple amounts - {amounts}')

```

When used as shown below, both invocations of the function will print out 'Help information':

```

banking_operation3(['help'])
banking_operation3(['info'])

```

The output from this code snippet is:

```

Help information
Help information

```

Note that in this case both lists had only one element in them, but the list could contain any number of elements as required.

41.9 Matching Multiple Options Within a Container

In some cases we want to use the same case condition with a container such as a list but indicate that it can match onto several values. For example, we may wish to use the same functionality in a case condition in an adventure game. When the user indicates that want to go somewhere we can indicate that the case condition should match on north, south, east or west, for example:

```

def direction(actions):
    match actions:
        case ['go', ("north" | "south" | "east" | "west")]:
            print('Going north or south or east or west')

```

If we now call this direction function as follows:

```

direction(['go', 'north'])

```

We will see the output:

```

Going north or south or east or west

```

Of course we do not yet know which direction they wanted to go in; however we can bind the actual value the user supplied to a variables using the `as` keyword. For example:

```
def direction2(actions):
    match actions:
        case ['go', ("north" | "south" | "east" | "west") as towards]:
            print(f'Going {towards}')
```

This now captures whatever value was supplied for the second element of the list (if it's one of the values listed) and stores it into the variable `towards`. This variable is now available within the body of the case condition and can be used to log the direction the user is going in.

For example, by calling the function as shown below:

```
direction2(['go', 'north'])
```

We will obtain the output:

```
Going north
```

41.10 Guards on a Case Condition

In some situations it is useful to be able to put an additional *guard* onto a case condition. For example, if we want to check to see that an extracted value is a member of a group of values that are not hardcoded into the `match` statement then we can use a guard.

The following example uses a case condition guard to check that the extracted `direction` variable contains a value in a list of `VALID_DIRECTIONS`:

```
VALID_DIRECTIONS = ["north", "south", "east", "west"]

def handle_command(command):
    match command:
        case ['go', direction] if direction in VALID_DIRECTIONS:
            print(f'Going {direction}')
        case _:
            print('Wrong direction')
```

When the `handle_command()` function is invoked using a list of two values, if the first value is 'go' then the case condition will be triggered, the second value will be extracted into the `direction` variable, and the value it holds will be checked against the values in the `VALID_DIRECTIONS` list.

If the value is a member of this list then the case condition will be activated and the case condition body will be executed. However, if the value extracted is not a member of the `VALID_DIRECTIONS` list then the case condition will not be activated and the program will fall through to the wildcard.

This is illustrated by the following two lines of code:

```
handle_command(['go', 'east'])
handle_command(['go', 'north-east'])
```

The output from this code is:

```
Going east
Wrong direction
```

41.11 Exercises

Create a program to select the appropriate message for a user for each of a set of HTTP error codes. The codes you should support include:

- 400 Bad Request,
- 401 Unauthorized,
- 403 Forbidden,
- 404 Not Found,
- 418 I'm a teapot,
- 429 Too Many Requests.

These are all legal error codes (see https://en.wikipedia.org/wiki/List_of_HTTP_status_codes).

When the user enters an error code, it should tell them the associated message or that it is current unknown. The program should keep asking the user for an error code, until the user enters nothing. For example:

```
Starting
Please enter the error code: 400
Bad Request
Please enter the error code: 401
Unauthorized
Please enter the error code: 403
Forbidden
Please enter the error code: 404
Not Found
Please enter the error code: 418
I'm a teapot
Please enter the error code: 429
Too Many Requests
Please enter the error code: 440
440 - is an unknown error code
Please enter the error code:
Done
```

Chapter 42

Python Virtual Environments



42.1 Introduction

The Python runtime environment was originally designed a long time ago, and as such one of the aspects that is still a little old fashioned is how it deals with additional *third-party* libraries and how they are installed. In this chapter we will look at what we mean by third-party libraries, what happens when such a library is installed into your Python runtime environment, the purpose and use of virtual environments and the tool `pip` and conclude by looking at how an IDE such as PyCharm actually makes working with these environments straight forward.

42.2 Third-Party Libraries

The Python runtime environment comes with many built-in libraries (aka modules) such as the `sys`, `os`, `random` and `math`. Information on these can be found online at the Python Module Index (see <https://docs.python.org/3/py-modindex.html>). Note that colloquially we often use the term library to mean a Python module. Why is this? In the main it's because a library could be made up of one or more modules (or indeed packages). The term library is thus a looser (non-technical) term to indicate we are talking about something that isn't directly available by default. That is we must use some form of `import` statement to access the code. Thus if I want to access the `os` module (library) then I need to import it into my program, for example:

```
import os
print(os.uname())
```

which will return information on the current operating system, for example:

```
posix.uname_result(sysname='Darwin', nodename='Johns-
```

```
iMac.local', release='22.4.0', version='Darwin Kernel Version  
22.4.0: Mon Mar 6 21:00:17 PST 2023; root:xnu-8796.101.5~3/RELEASE_  
X86_64', machine='x86_64')
```

These modules only require the programmer to `import` them into their code as they are provided as part of the Python runtime environment (that is they are available by default within your Python system). You do not need to do anything else to make them available to your program other than install Python itself and import them.

However, there are many, many, many other modules/libraries available to the Python programmer. These are referred to as third-party libraries as they are provided by people and organizations other than `python.org` (that is the Python Foundation).

Some of these libraries are extremely well established and provide the foundation for many of the types of Python application the make it such a popular language. In fact right back in chapter two we talked about several of these such as those used in the Data Analytics world including libraries such as Pandas (see <https://pandas.pydata.org>), NumPy (which represents Numeric Python see <https://numpy.org/>), SciPy (which stands for Scientific Python see <https://scipy.org/>) and Matplotlib (see <https://matplotlib.org/>) which is a graphing library.

This means that if you are working within the Data Analytics arena and someone tells you that they are using Python to do their analysis what they really mean is that they are probably using something like:

- Python,
- Plus Pandas and NumPy,
- and typically with SciPy,
- and some form of graphing tool such as Matplotlib.

Similarly if a developer is using Python to create a set of RESTful services (that is a remote service accessed over `https` and supplying some form of data on request), they are likely to be using Python plus a library such as `fastAPI` (<https://fastapi.tiangolo.com/>), `Flask` (<https://palletsprojects.com/p/flask/>) or `Django` (<https://www.djangoproject.com/>).

As such these third-party modules are key to some of the most successful application areas within which Python is used.

It is perhaps worth noting that anyone can create a third-party module and make it available to others to use, however not all modules are created equal. Almost all libraries are made available as open source, free of charge, software. As such it can be difficult to determine which libraries are appropriate to use and which are not.

In general the longer the module has been around, the larger the number of people using it, the more plentiful the documentation (both from the module provider but also in general on the web) and the greater the number of contributors and maintainers the better the module is likely to be.

For example, a module called something like, `johnsAPI-0.0.1alpha`, supported just by one person and downloaded only by a handful of people is probably not your best bet, whereas `numpy 1.24.3` which has extensive documentation on the `numpy.org` website as well as numerous tutorials and examples around the web, that has an extensive release history, is used by numerous other well-respected libraries, has 905 contributors (as of April 2023), has been maintained and updated over the last 17 years, with numerous regular releases over that time, and is a much better bet.

42.3 The Pip Tool

The Python environment comes with a tool that is designed to help you install third-party libraries into your runtime. This tool is called `pip` (which stands for ‘Pip Installs Python’). It had been a default part of the Python installation since Python 3.4 (prior to this it was itself considered a third-party tool). This tool is sometimes referred to as Python’s package manager, module installer or even the `pip` installer.

`Pip` can install any third-party library. To do this `pip` downloads modules by default from a central repository called `pypi.org` (which stands for the Python Package Index—hence `pypi`—but pronounced `pie pie.org`). This is useful as it means that any module provided in this central repository is available to all Python developers wherever they are in the world.

In fact `pip` can install modules from

- `pypi.org` central repository,
- version control system URLs (such as GitHub),
- local project directories,
- local or remote source archives.

In addition, `pip` can install from a “requirements file” which is a useful way to specify a whole environment to be installed.

The `pip` tool is also aware of dependencies between modules. That is if one module itself relies on another module being installed then it can install both modules. For example the `Pandas` module depends on (among others) the `NumPy` module, and thus when you install `Pandas`, then `NumPy` is installed for you (automatically).

The `pip` command comes with several options that can be used to install and uninstall modules, reinstall modules, list currently installed modules, generate a report on the modules installed and any dependencies for the modules you require.

42.3.1 Finding the Pip Version

You can easily check to see if `pip` is installed in your environment using the `pip --version` command, for example, at a Command Prompt or within a terminal window you can type:

```
pip --version
```

depending on your configuration and operating system you should see something similar to:

```
% pip --version
pip 23.0.1 from/usr/local/lib/python3.11/site-packages/pip
(python 3.11)
%
```

This indicates that in the installation on this Apple Mac, that the current version of `pip` is version 23.0.1.

42.3.2 Installing Pip

If you do not have `pip` installed then you can download it and install it yourself from the `pip` download page on the `pypi.org` site (see <https://pypi.org/project/pip/>).

Interestingly, the `pip` tool is treated as if it was a separate module by the Python environment, this means that you can also update the `pip` tool without updating the whole Python runtime. This may be useful if a bug fix is released for `pip` but it is not included in the current release of Python yet.

42.3.3 Updating Pip

To update `pip` you can use `pip` itself! That is `pip` can be used to install an update to `pip` (this may sound confusing but is actually perfectly fine as the `pip` tool will update the installation of `pip` and the next time you run `pip` the new version will automatically be picked up).

To update `pip` itself use:

```
pip install --upgrade pip
```

Note that this is the `pip` command being passed the request to install something, in this case the upgrade for `pip`. On Windows you might want to use a variation on this command as shown below:

```
python -m pip install --upgrade pip
```

This does exactly the same as the first version but tells Python to run the module (-m) pip for you. Indeed many people use this style on all operating systems as it is the recommended approach to running pip. However, in most cases just using the pip command on its own is sufficient on Linux and Macs.

An example of updating pip is shown below on an Apple Mac:

```
% pip install --upgrade pip
Requirement already satisfied: pip in /usr/local/lib/python3.11/
site-packages (23.0.1)
Collecting pip
  Downloading pip-23.1.2-py3-none-any.whl (2.1 MB)
-----2.1/2.1 MB 8.3 MB/s
eta 0:00:00
Installing collected packages: pip
  Attempting uninstall: pip
    Found existing installation: pip 23.0.1
    Uninstalling pip-23.0.1:
      Successfully uninstalled pip-23.0.1
  Successfully installed pip-23.1.2
%
```

From this you might see that the result of running this command in this environment is to update pip from 23.0.1 to 23.1.2. You can see the command downloading the 23.1.2 version and then uninstalling the old pip and installing the new version of pip. If we now check the version of pip installed we now see that it is version 23.1.2:

```
% pip --version
pip 23.1.2 from /usr/local/lib/python3.11/site-packages/pip
(python 3.11)
%
```

42.4 Default Installation

There is a warning that should be noted with pip; that is that by default it installs a module into the central Python installation on a machine. This is fine if you are only ever going to run one program requiring a specific set of modules. However, if you are going to work with different Python programs, which might require different modules, and in particular different version of a module, then this can be a problem.

To illustrate the issue it is important to understand that every module has a version number, such as NumPy 1.24.1; that is this is version 1.24.1 of NumPy, and it will differ in some way (however small) from version 1.24.0. The version number for Python packages typically uses something called *semantic versioning*. In semantic versioning a version number consists of three integers such that it is made up of `major.minor.patch`, that is:

- **Major** the main release number, 1.0.0, is usually the first release; all subsequent releases that are part of the release 1 will start 1.x.x. Typically release will be

2.x.x, and this indicates a major change in the module, etc. Note that there is no guarantee of backwards compatibility between one major release and the next. Backwards compatibility is the idea that code that previously worked should still work from one release to the next—this guarantee does not exist between one major version and the next.

- **Minor** such as 1.1.0 to 1.2.0 where the minor release typically includes larger new features or bug fixes but ensures backwards compatibility with other older major release versions (that is version 1.2.0 should be backwards compatible with version 1.1.0).
- **Patch** release this is typically used for bug fixes or other technical issues which should have no effect on the backwards compatibility of the code and do not change or extend the functionality of the library.

In some cases packages use a year-based number system such as 2020.1 which is the first release of 2020, although this is actually far less common in Python.

Now let us consider that on your computer you have two Python programs: one uses Pandas 1.5.0, and the other uses Pandas 2.0.0. Given what has been described above there is no guarantee that the code which was written to use Pandas 1.5.0 will work with Pandas 2.0.0. Therefore we want to run one program using one Pandas library and the other using the newer Pandas library. We could of course update the old code to work with the new version of Pandas but that could be time-consuming and non-trivial; and we may just want to run the code.

However, as `pip` installs these libraries centrally for all Python programs by default, we cannot have the two libraries installed at the same time!

Of course, this is not a new problem, and the solution comes in the form of *pip virtual environments*.

42.5 Pip Virtual Environments

A virtual environment is a configuration of Python modules that can be set up on a *per project* basis. That is each individual Python project (or program) can have its own virtual environment and that virtual environment can be activated as and when required. The virtual environment is then used to configure all the modules used, isolating one virtual environment's configuration from another.

Although it should be noted that there is no reason why several Python programs cannot share the same virtual environment, it is very common practice to have a virtual environment per program.

To use a virtual environment it is necessary to first create this virtual environment. This is done using the `python` command itself (rather than the `pip` command). For example:

```
python -m venv <name of virtual environment>
```

In the above the `python` command is being passed the `-m` option followed by `venv`. This indicates that you want Python to run the `venv` module. The name `venv` stands for *virtual environment*. This is followed by the name of the virtual environment you wish to create. A common pattern is to call this `venv`—this is then created in the current directory. This is fine, but it can be useful to give the virtual environment a more meaningful name such as `genesis_venv`—indicating that this is the virtual environment for the `genesis` project (whatever that might be).

Whatever name you choose to use, the tool will create a directory within the current directory called that name.

Let us first check to see what has been installed into the central Python runtime on a sample Mac computer:

```
% pip list
Package      Version
-----
pip          23.1.2
setuptools   67.6.1
wheel        0.40.0
```

This illustrates the typical basic Python installation. All that is contained within the Python system is the default built-in modules plus the `pip` module, the `setuptools` module and the `wheel` module. The `setuptools` module helps with package handling, and the `wheels` module supports wheel file which are essentially zip files with a special file name convention that tells installs what Python versions and platforms they support.

For example, let us create a virtual environment called `genesis_venv` under a `projects` directory:

```
% mkdir projects
% cd projects
projects % ls

projects % python -m venv genesis_venv
```

In the above terminal/command window session we have created the `projects` directory and then moved into that directory. Currently when we list the contents of the directory, it is empty. We then issue the command:

```
% python -m venv genesis_venv
```

If we now look at the contents of the directory we will see a subdirectory has been created called `genesis_venv`:

```
projects % ls
genesis_venv
```

At the moment this is not that interesting, but we can now *activate* that environment. This is done using the rather awkward incantation which accesses the `activate` script within the virtual environment directory's `bin` directory. On a Mac or Linux this would be done using the `source` command followed by the `<venv_name>/bin/activate` path, for example:

```
% source genesis_venv/bin/activate
```



```
(genesis_venv) projects %
```

Note that on this Mac you can see the virtual environment that has been activated as it is shown before the prompt in round brackets.

On a Windows machine this file will be a `bat` file and therefore you can activate it by directly running the `activate.bat` file, for example:

```
> genesis_venv\bin\activate.bat
```

If this is starting to all sound a bit much, don't worry the PyCharm editor will handle most of this for you and we will look at this at the end of the chapter; for now it is useful to see what happens *under the hood*.

We can also deactivate the virtual environment using the `deactivate` command which on Windows is `deactivate.bat`, for example, on Linux/Macs:

```
% deactivate
jeh@Johns-iMac projects %
```

Or on Windows

```
> deactivate.bat
```

42.6 Installing Modules

If we reactivate the virtual environment again we can now install modules into this environment. To do this we use the `pip install` command which has the following syntax:

```
pip install module-name
```

Let us first review the default installation using `pip list`:

```
% source genesis_venv/bin/activate
(genesi_venv) projects % pip list
Package      Version
-----
pip          23.1.2
setuptools   67.6.1
(genesi_venv) %
```

From this we can see that there are two modules installed into the virtual environment (`wheel` is not included in this case).

The following terminal session illustrates what happens when we do this for `Pandas`:

```
(genesis_venv) % pip install pandas
Collecting pandas
  Downloading pandas-2.0.1-cp311-cp311-macosx_10_9_x86_64.whl
  (11.6 MB)
-----11.6/11.6 MB 8.7
MB/s eta 0:00:00
```

```

Collecting python-dateutil>=2.8.2
  Using cached python_dateutil-2.8.2-py2.py3-none-any.whl (247 kB)
Collecting pytz>=2020.1
  Using cached pytz-2023.3-py2.py3-none-any.whl (502 kB)
Collecting tzdata>=2022.1
  Using cached tzdata-2023.3-py2.py3-none-any.whl (341 kB)
Collecting numpy>=1.21.0
  Downloading numpy-1.24.3-cp311-cp311-macosx_10_9_x86_64.whl
(19.8 MB)
-----19.8/19.8 MB 8.5
MB/s eta 0:00:00
Collecting six>=1.5
  Using cached six-1.16.0-py2.py3-none-any.whl (11 kB)
Installing collected packages: pytz, tzdata, six, numpy, python-
dateutil, pandas
Successfully installed numpy-1.24.3 pandas-2.0.1 python-dateutil-
2.8.2 pytz-2023.3 six-1.16.0 tzdata-2023.3

(genesis_venv) projects %

```

The command used here is `pip install` followed by the name of the module to install. In this case we have installed `pandas`. This will result in the latest version of `pandas` being installed. However we could have specified a specific version of `pandas` using the `==<version number>` option after the name of the module (no spaces), for example:

```
pip install pandas==1.5.2
```

This would install `pandas` version 1.5.2 rather than the latest version.

It is also worth noting that in addition to installing `Pandas` (version 2.0.1) several other modules have been installed, and these include `python-dateutil` (2.8.2), `pytz` (2020.1), `tzdata` (2022.1), `numpy` (1.21.0) and `six` (1.5).

If we now ask `pip` to list what has been installed we can now see these modules listed:

```

(genesis_venv) % pip list
Package           Version
-----
numpy             1.24.3
pandas            2.0.1
pip               23.1.2
python-dateutil  2.8.2
pytz              2023.3
setuptools        67.6.1
six               1.16.0
tzdata            2023.3
(genesis_venv) projects %

```

Interestingly if we deactivate the virtual environment, which would mean that we would then return to using the default installation setup and list the modules we will see that `pandas` and its associated dependencies are no longer listed, for example:

```
(genesis_venv) % deactivate
projects % pip list
Package      Version
-----
pip          23.1.2
setuptools  67.6.1
wheel       0.40.0
projects %
```

This is exactly what we should expect, we are no longer using the virtual environment, and thus the configuration of modules defined within that virtual environment is no longer available. We can of course reactivate the virtual environment, and we will get Pandas and its modules back again, etc.

This allows a developer to switch between different virtual environments for different programs/projects.

42.7 Listing and Freezing on Modules

There are three similar pip tools that can provide information on the modules installed in a virtual environment. The commands are:

- `pip list`—this generates a list of installed packages, and the list is ordered in ascending order based on the case sensitive module names.
- `pip inspect`—which generates a very detailed output file of all the modules installed.
- `pip freeze`—which generates output in *requirements* format that can be used to regenerate the configuration of the current virtual environment in another installation.

Each of these will be discussed below.

42.7.1 The `pip list` Command

As an example of using this command we will list the modules installed into the virtual environment we created above:

```
(genesis_venv) % pip list
Package      Version
-----
numpy        1.24.3
pandas       2.0.1
pip          23.1.2
python-dateutil 2.8.2
pytz         2023.3
setuptools   67.6.1
six          1.16.0
tzdata       2023.3
(genesi... %
```

This lists all the modules that have been installed and the version that was installed.

42.7.2 *The pip inspect Command*

If required, the module information can be generated as a JSON report file (JSON stands for JavaScript Object Notation) which is a widely used data format for information exchange and storage. This can be done using the `pip inspect` command, for example:

```
(genesis_venv) % pip inspect
{
  "version": "1",
  "pip_version": "23.1.2",
  "installed": [
    {
      "metadata": {
        "metadata_version": "2.1",
        "name": "pip",
        "version": "23.1.2",
        "summary": "The PyPA recommended tool for installing Python
packages.",
        "description": "pip - The Python Package Installer
...

```

The above illustrates the part of the output of the JSON data as it is quite extensive, and the above illustrates the point.

42.7.3 *The pip freeze Command*

The `pip freeze` command can be used to create a *requirements* format file. This file contains information on the installed modules that can be used by the `pip` command to configure another virtual environment to include the same modules (and version of those modules). This allows two users to share their virtual environment configurations.

For example:

```
(genesis_venv) % pip freeze
numpy==1.24.3
pandas==2.0.1
python-dateutil==2.8.2
pytz==2023.3
six==1.16.0
tzdata==2023.3
(genesi_venv) %
```

This prints to the output each of the modules installed. If we wish this information to be stored into a file we can use the redirection arrow, for example:

```
% pip freeze > johnsenv.txt
```

If we list the contents of this generated `johnsenv.txt` file we can see that it contains the same information as the output above:

```
% more johnsenv.txt
numpy==1.24.3
pandas==2.0.1
python-dateutil==2.8.2
pytz==2023.3
six==1.16.0
tzdata==2023.3
```

In the above example, the command `more` will list the contents of the file to the output.

We can now use this in another virtual environment to set that environment up to use the same set of modules, for example:

```
% python -m venv projenv
% source projenv/bin/activate
(projenv) % pip list
Package      Version
-----
pip          23.1.2
setuptools   67.6.1
(projenv) jeh@Johns-iMac projects % pip install -r johnsenv.txt
Collecting numpy==1.24.3 (from -r johnsenv.txt (line 1))
  Using cached numpy-1.24.3-cp311-cp311-macosx_10_9_x86_64.whl
(19.8 MB)
Collecting pandas==2.0.1 (from -r johnsenv.txt (line 2))
  Using cached pandas-2.0.1-cp311-cp311-macosx_10_9_x86_64.whl
(11.6 MB)
Collecting python-dateutil==2.8.2 (from -r johnsenv.txt (line 3))
  Using cached python_dateutil-2.8.2-py2.py3-none-any.whl (247 kB)
Collecting pytz==2023.3 (from -r johnsenv.txt (line 4))
  Using cached pytz-2023.3-py2.py3-none-any.whl (502 kB)
Collecting six==1.16.0 (from -r johnsenv.txt (line 5))
  Using cached six-1.16.0-py2.py3-none-any.whl (11 kB)
Collecting tzdata==2023.3 (from -r johnsenv.txt (line 6))
  Using cached tzdata-2023.3-py2.py3-none-any.whl (341 kB)
Installing collected packages: pytz, tzdata, six, numpy, python-
dateutil, pandas
Successfully installed numpy-1.24.3 pandas-2.0.1 python-dateutil-
2.8.2 pytz-2023.3 six-1.16.0 tzdata-2023.3
(projenv) % pip list
```

```

Package          Version
-----
numpy            1.24.3
pandas           2.0.1
pip              23.1.2
python-dateutil 2.8.2
pytz             2023.3
setuptools       67.6.1
six              1.16.0
tzdata           2023.3
(projvenv) %

```

The result is that `projvenv` and `genesis_venv` virtual environments have exactly the same configuration of the same modules.

42.8 Uninstalling Modules

In some cases it is very useful to be able to uninstall an unwanted package. For example, if we have installed Pandas by mistake into our virtual environment then we can remove it using the `pip uninstall` command:

```
pip uninstall <module name>
```

As an example, if we first list the packages installed in our virtual environment using `pip list` we will see that we have 8 modules installed:

```

(genesis_venv) % pip list
Package          Version
-----
numpy            1.24.3
pandas           2.0.1
pip              23.1.2
python-dateutil 2.8.2
pytz             2023.3
setuptools       67.6.1
six              1.16.0
tzdata           2023.3
(genesis_venv) jeh@Johns-iMac projects %

```

If we now uninstall `pandas`:

```

(genesis_venv) jeh@Johns-iMac projects % pip uninstall pandas
Found existing installation: pandas 2.0.1
Uninstalling pandas-2.0.1:
  Would remove:
    /Users/jeh/projects/genesis_venv/lib/python3.11/site-
packages/pandas-2.0.1.dist-info/*
    /Users/jeh/projects/genesis_venv/lib/python3.11/site-
packages/pandas/*

```

```
Proceed (Y/n)? y
Successfully uninstalled pandas-2.0.1
(genesis_venv) projects %
```

We can check to see which modules are still installed:

```
(genesis_venv) projects % pip list
Package              Version
-----
numpy                1.24.3
pip                  23.1.2
python-dateutil     2.8.2
pytz                  2023.3
setuptools           67.6.1
six                   1.16.0
tzdata                2023.3
(genesis_venv) projects %
```

As you will see there are now 7 modules listed! The pandas module is no longer installed as part of this virtual environment.

However, the other modules that were installed along with pandas when we ran `pip install pandas` have not been removed! What is going on here? This was a decision taken by the pip developers, although when you install Pandas all of its dependencies are installed, when you uninstall Pandas you must manually remove its dependencies.

Why is this? The main reason is that dependency management across modules is very hard to do. For example, let us say we installed `johns` module and that `johns` module also used `numpy` version 1.24.3, should `pip uninstall` remove `numpy` in this case? Probably not if we also have `pandas` installed! This scenario gets more complicated as we may have `module1` using `module2` which uses `module3` which then uses `numpy`, etc.

Therefore the decision was taken to require developers to manually handle un-installation of modules. In this case that means uninstalling `numpy`, `python-dateutil`, `putz`, `six` and `tzdata`.

42.9 Module Path Search Order

An important point to note is that Python will search for module using the following order:

- local directory,
- installed third-party modules,
- built-in modules.


It is therefore important never to name your own files using the name of a *built-in* module nor using the name of a third-party module. As such if you are writing some code that uses the `pandas` module do not call your own file `pandas.py` as that will hide the `pandas` you have installed into your virtual environment!

42.10 Using PyCharm with Virtual Environments

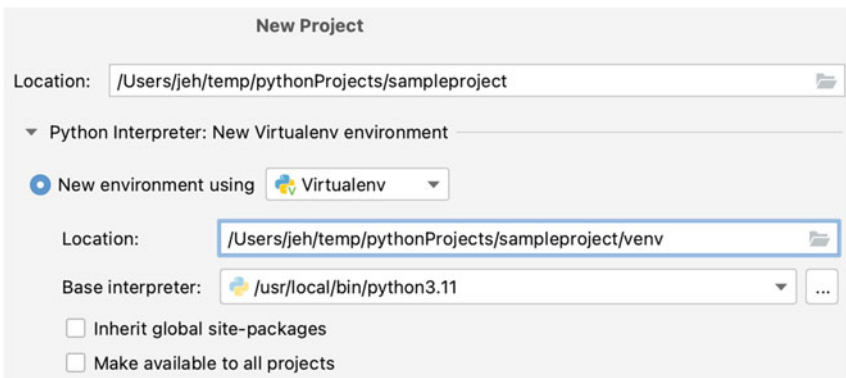
The above in this chapter might all feel a bit daunting as many people are not familiar with, or indeed comfortable using, a command line set of tools such as `pip`. Indeed we have suggested that you use the PyCharm editor to not only write your programs, but also to run them, yet here we have been describing how to use the `pip` command line tool.

The reason for this was to show what happens at the tooling level. However you will probably interact with all of these tools using an IDE such as PyCharm. In fact when you create any project in PyCharm you will have had the option to use a virtual environment!

To see this we will start a new PyCharm project. This is done by starting PyCharm up and selecting New Project from the ‘Welcome to PyCharm’ dialog, for example:

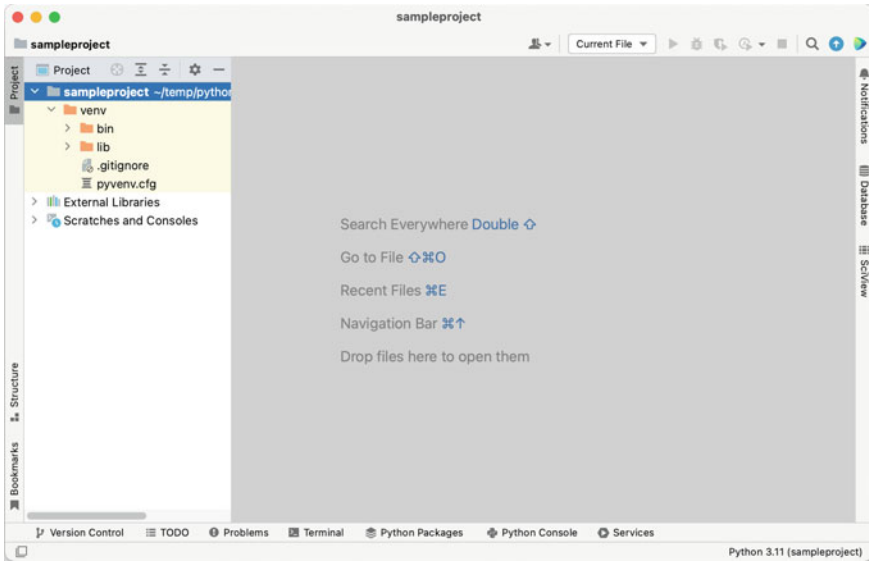


If you do this you will be presented with the ‘New Project’ dialog. This dialog has several fields on it, the main one being the location into which the project will be stored. However, below this is a drop down option for the ‘Python Interpreter: New Virtualenv environment’. If you open this drop down you will see that it provides a location for the virtual environment (by default using the name `venv` in the current directory) and the option to select which version of Python you wish to use, for example:



Note under the drop down next to ‘New environment using’ there are several other options but we are using *Virtualenv* which is the Python default.

If you now click on ‘Create’ you will create the new project workspace. This workspace will be created with a directory within it called `venv`, and this virtual environment will automatically be activated for you. For example:



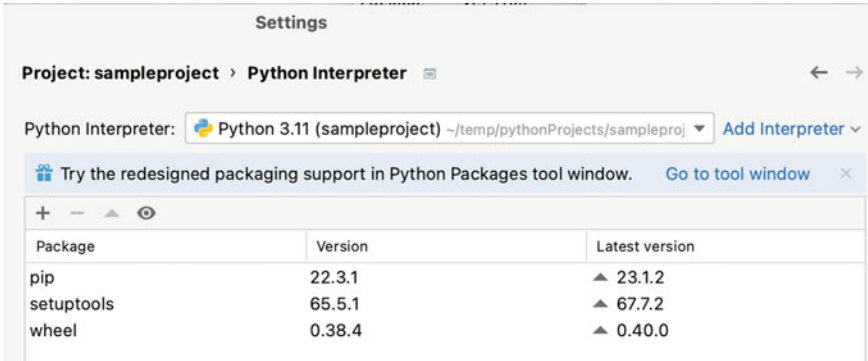
You can explore the `venv` directory structure, but do not under any circumstance edit the contents of this directory or add files to it. Your own code should be at the top level; be very careful not to accidentally add you own files to the `venv` directory!

Now if you want to see the modules installed into your virtual environment you can go to the ‘Settings’ menu option, this is under File on a Windows machine and under PyCharm menu on a Mac.

When you open the Settings dialog, select the ‘Project: < name of project >’ node in the left hand tree, and open the node for example:

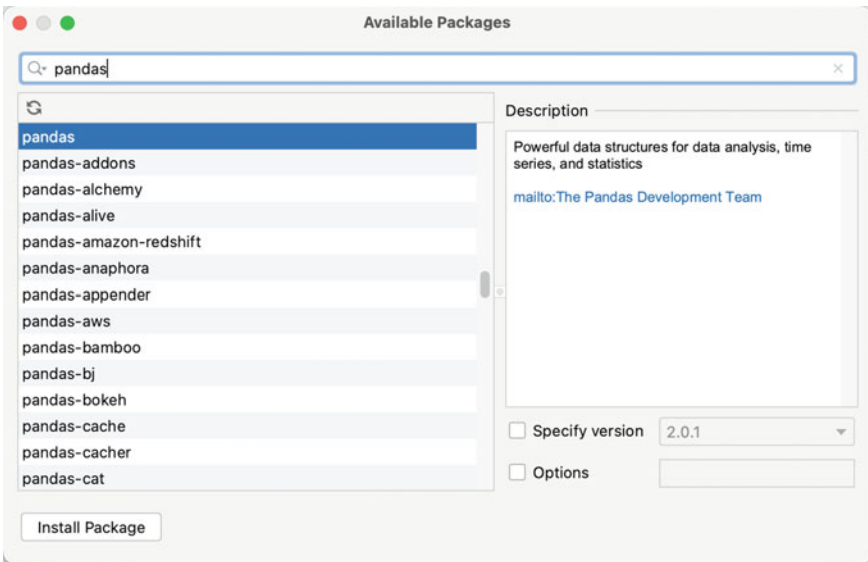


In here select the ‘Python Interpreter’ node. You should now see that on the right hand side of the editor the list of installed packages is presented along with the Python runtime you are using, for example:




This shows that we have three modules loaded including `pip`, `setuptools` and `wheel`. It also lists the version of the module and what the latest version is.

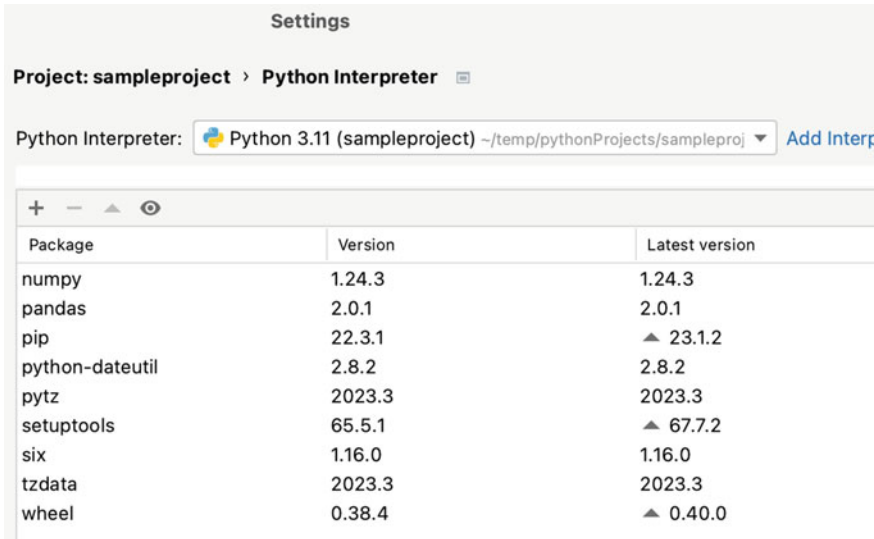
We can add additional modules to this virtual environment using the ‘+’ button. If you click on the ‘+’ button you will see the available packages dialog. On the top search bar, if you type in the name of the module you are interested in, then the modules with that name will be listed. For example, if you type in `pandas`, you will see:



You can select whichever module you are interested in. In this case we want to install the `pandas` module and that is the one selected, and you can therefore just click on the ‘Install Package’ button at the bottom of the dialog. Once the module/package is installed you will see a message indicating that the installation was successful:



If you close the dialog and return to the main Settings dialog you should now see the Pandas module has been installed along with all the modules it depends on, for example:



You can now close this dialog and can use the Pandas module in your own code, etc.

42.11 Online Resources

- <https://docs.python.org/3/py-modindex.html> the Python Module Index
- <https://fastapi.tiangolo.com/> FastAPI website
- <https://palletsprojects.com/p/flask/> Flask website
- <https://www.djangoproject.com/> Django website.

Chapter 43

Monkey Patching



43.1 Introduction

Monkey Patching is a term you might well come across when looking into Python further or when searching the web for Python related concepts. It relates to the ability in Python to extend the functionality associated with a class/type at runtime.

Although not directly related to Monkey Patching; how Python looks up attributes and how this process can be managed is a useful aspect to understand. In particular, how to handle unknown attributes can be very useful in managing situations in which Monkey Patching might be used to solve an initial attribute incompatibility.

This chapter explores both Monkey Patching and Python attribute lookup.

43.2 What is Monkey Patching?

Monkey Patching is the idea that it is possible to add behaviour to an existing object, at runtime, to meet some requirement that originally the type did not meet. This can happen for example as there is no fixed requirement for a class to implement all of protocol; in many cases a class may only implement as much of a protocol as is required to meet the current needs; if at a later stage, other elements of a protocol are required then they can be added.

Of course, if this is likely to be a common occurrence then the features can be added to the class for use by everyone; but if not, then those features can be added dynamically at runtime to an object itself. This avoids the public interface of the type becoming cluttered with rarely used features/functionality.

43.2.1 How Does Monkey Patching Work?

Python is a dynamic language which allows the definition of a type to change at runtime. As methods on objects are in essence just another attribute of a class, albeit one that can be executed, it is possible to add new functionality to a class by defining new attributes that will hold references to the new behaviour.

43.2.2 Monkey Patching Example

The following class, `Bag`, implements a initialization method `__init__()`, `__str__()` and the `__getitem__()` method used to support indexed access to a container (or collection) type.

```
class Bag():
    def __init__(self):
        self.data = ['a', 'b', 'c']

    def __getitem__(self, pos):
        return self.data[pos]

    def __str__(self):
        return 'Bag(' + str(self.data) + ')'

b = Bag()
print(b)
```

This creates a `Bag` object and prints out the contents of the `Bag`:

```
Bag(['a', 'b', 'c'])
```

However, if we try to run

```
print(len(b))
```

We will get a runtime error:

```
Traceback (most recent call last):
  File "Bag.py", line 12, in <module>
    print(len(b))
TypeError: object of type 'Bag' has no len()
```

This is because the `len()` function expects the object passed to it will implement the `__len__()` method that is be used to obtain its length. In this case the class `Bag` does not implement this method.

However, we can define a stand-alone function that behaves in the way we would need the `Bag` to calculate its length, for example:

```
def get_length(self):
```

```
return len(self.data)
```

This function takes a single parameter which we have called `self`. It then uses this parameter to reference an attribute called `data` which itself uses `len()` to return the length of the associated data items.

At present this function has no relationship to the `Bag` class other than the fact that it assumes that whatever is passed into it will have an attribute called `data`—which the `Bag` class does.

In fact, the `get_length()` function is compatible with any class that has an attribute `data` that can be used to determine its length.

We can now associate it with the `Bag` class; this can be done by assigning the function reference (in practice the function name) to an appropriate attribute on the `Bag` class. Since the `len` function expects a class to implement the `__len__()` method we can assign the `get_length()` function to the `__len__()` attribute. This effectively adds a *new* method to the `Bag` class with the signature `__len__(self)`:

```
# Monkey patching
Bag.__len__ = get_length
```

Now when we invoke

```
print(len(b))
```

We get the value 3 being printed out.

We have now Monkey Patched the class `Bag` so that the missing method becomes available.

43.2.3 *The Self Parameter*

One of the reasons that Monkey Patching works is because all methods receive the special first parameter (called `self` by convention) representing the object itself. This means that any function that treats the first parameter as being a reference to an object can potentially be used to define a method on a class.

If a function does not assume that the first parameter is a reference to an object (the one holding the method) then it cannot be used to add new functionality to a class.

43.2.4 Adding New Data to a Class

Monkey Patching is not just limited to functionality; it is also possible to add new data attributes to a class. For example, if we wanted each `Bag` to have a *name* then we could add a new attribute to the class to hold its name:

```
Bag.name = 'My Bag'  
print(b.name)
```

Which prints out the string `'My Bag'` which now acts as a default value of the name of any `Bag`. Once the attribute is added we can then change the name of this particular instance of a `Bag`:

For example, if we extend the above example:

```
Bag.name = 'My Bag'  
print(b.name)  
  
b.name = 'Johns Bag'  
print(b.name)  
  
b2 = Bag()  
print(b2.name)
```

We can now generate

```
My Bag  
Johns Bag  
My Bag
```

43.3 Online Resources

For further information on Monkey Patching the following resources may be of interest:

- https://en.wikipedia.org/wiki/Monkey_patch Wikipedia page on Monkey Patching.
- <http://net-informations.com/python/iq/patching.htm> A discussion on whether Monkey Patching should be considered good or bad practice.

Chapter 44

Attribute Lookup



44.1 Introduction

Compared to static compile languages such as Java and C# Python is a dynamic and flexible language. This is also true of how attributes within a class definition are handled. In this chapter we will look in more detail about how Python manages attributes and method lookups, how unknown attribute and method lookup are handled and how you can intercept such lookups yourself.

44.2 Attributes in Python

As shown in the last chapter, Python is very dynamic and it is easy to add attributes and methods to a class, but how does this work. It is worth briefly digressing and considering how Python manages attribute and method lookup for an object.

Python classes can have both *class* and *instance* oriented attributes, for example the following class `Student` has a *class* attribute `count` (that is associated with the class itself) and an *instance* or *object* attribute `name`. Thus each instance of the class `Student` will have its own `name` attribute.

```
class Student:
    count = 0

    def __init__(self, name):
        self.name = name
        Student.count += 1
```

Whenever an instance of the class `Student` is created the `Student.count` attribute will be incremented by 1.

To manage these attributes Python maintains internal dictionaries; one for *class* attributes and one for *object* attributes. These dictionaries are called `__dict__` and can be accessed either from the class `<class>.__dict__` (for class attributes) or from an instance of the class `<instance.>__dict__` (for object attributes). For example:

```
student = Student('John')

# Class attribute dictionary
print('Student.__dict__:', Student.__dict__)
# Instance / Object dictionary
print('student.__dict__:', student.__dict__)
```

Which produces the output shown below (note that the class dictionary holds more information than just the class attribute count):

```
Student.__dict__: {'__module__': '__main__', 'count': 1, '__init__': <function Student.__init__ at 0x10d515158>, '__dict__': <attribute '__dict__' of 'Student' objects>, '__weakref__': <attribute '__weakref__' of 'Student' objects>, '__doc__': None}
student.__dict__: {'name': 'John'}
```

To look up an attribute Python does the following for class attributes:

1. Search the class dictionary for an attribute
2. If the attribute is *not* found in step 1 then search the parent class(es) dictionaries.

For object attributes, Python first searches the instance dictionary and repeats the above steps, it thus performs these steps:

1. Search the object/instance dictionary
2. If the attribute was *not* found in step 1, then search the class dictionary for an attribute
3. If the attribute is *not* found in step 2, then search the parent class(es) dictionaries.

Thus given the following statements, different steps are taken each time:

```
student = Student('John')

print('Student.count:', Student.count) # class lookup
print('student.name:', student.name) # instance lookup
print('student.count:', student.count) # lookup finds class attribute
```

The output as expected is that either attempt to access the class attribute `count` will result in the value 1 whereas the object attribute `name` returns 'John'.

```
Student.count: 1
student.name: John
student.count: 1
```

As the dictionaries used to hold the class and object attributes are just that dictionaries, they provide another way to access the attributes of a class such as `Student`. That is you can write code that will access an attribute value using the appropriate `__dict__` rather than the more usual *dot* notation, for example the following are equivalent:

```
# class lookup
print('Student.count:', Student.count)
print("Student.__dict__['count']:", Student.__dict__['count'])

# Instance/Object Lookup
print('student.name:', student.name)
print("student.__dict__['name']:", student.__dict__['name'])
```

In both cases the end result is that same, either the class attribute `count` is accessed or the object/instance attribute `name` is accessed:

```
Student.count: 1
Student.__dict__['count']: 1
student.name: John
student.__dict__['name']: John
```

However, accessing attributes via the `__dict__` does not trigger a search process; it is instead a direct lookup in the associated dictionary container. Thus if you try to access a class variable via the objects `__dict__` then you will get an error. This is illustrated below where we attempt to access the `count` class variable via the *student* object:

```
# Attempt to look up class variable via object
print('student.name:', student.name)
print("student.__dict__['count']:", student.__dict__['count'])
```

This will generate a `KeyError` indicating that the object `__dict__` does not hold a key called `'count'`:

```
Traceback (most recent call last):
  File "Student.py", line 60, in <module>
    print("student.__dict__['count']:", student.__dict__
['count'])
KeyError: 'count'
```

44.3 Handling Unknown Attribute Access

Monkey Patching is of course very flexible and very useful when you know what you need to provide; however what happens when an attribute reference (or method invocation) occurs when it is not expected? By default an error is generated such as the `AttributeError` below:

```
student = Student('John')

res1 = student.dummy_attribute
print('p.dummy_attribute:', res1)
```

This generates an `AttributeError` for the `dummy_attribute`

```
Traceback (most recent call last):
  File "Student.py", line 51, in <module>
    res1 = student.dummy_attribute
AttributeError: 'Student' object has no attribute 'dummy_
attribute'
```

You can of course catch the `AttributeError` if you want; but the means wrapping your code in a `try` block. An alternative approach is to define a method called `__getattr__()`; this method will be called when an attribute is not found in the objects (and classes) `__dict__` dictionary. This method can then perform whatever action is appropriate such as logging a message or providing a default value, etc.

For example, if we modify the definition of the `Student` class to include a `__getattr__()` method such that a default value is returned:

```
class Student:
    count = 0

    def __init__(self, name):
        self.name = name
        Student.count += 1

    # Method called if attribute is unknown
    def __getattr__(self, attribute):
        print('__getattr__:', attribute)
        return 'default'
```

Now when we try and access the `dummy_attribute` on a student object we will get the string 'default' returned:

```
student = Student('John')

res1 = student.dummy_attribute
print('p.dummy_attribute:', res1)
```

Now generates

```
__getattr__: dummy_attribute
p.dummy_attribute: default
```

Note that the `__getattr__()` method is only called for unknown attributes as Python first looks in `__dict__` and thus if the attribute is found, no call is made to the `__getattr__()` method.

Also note that if an attribute is accessed directly from the `__dict__` (for example `student.__dict__['name']`) then the `__getattr__()` method is never invoked.

44.4 Handling Unknown Method Invocations

The `__getattr__()` method is also invoked if an unknown method is called. For example, if we call the method `dummy_method()` on a `Student` object then an error is raised (in fact this is again an `AttributeError`). However, if we define a `__getattr__()` method we can return a reference to a method to use as a default. For example, if we modify the `__getattr__()` method to return a method reference (i.e., the name of a method in the `Student` class):

```
class Student:
    count = 0

    def __init__(self, name):
        self.name = name
        Student.count += 1

    # Method called if attribute is unknown
    def __getattr__(self, attribute):
        print('__getattr__: ', attribute)
        return self.my_default

    def my_default(self):
        return 'default'
```

Now when a undefined method is invoked (such as `dummy-method()`) `__getattr__()` will be called. This method will return a reference to the method `my_default`. This will be run and the value returned as a side effect of the method invocation as indicated by the round baskets (the `()`) after the call to the original method:

```
student = Student('John')
res2 = student.dummy_method()
```

```
print(`student.dummy_method():`, res2)
```

Which produces the following output: rather than an error message

```
__getattr__: dummy_method
student.dummy_method(): default
```

44.5 Intercepting Attribute Lookup

It is also possible to always intercept attribute lookups using the *dot* nation (e.g., `student.name`) by implementing the `__getattr__()` method. This method will always be called instead of looking the attribute up in the objects dictionary. The `__getattr__()` method will only be called if an `AttributeError` is raised or the `__getattr__()` method explicitly calls the `__getattr__()` method.

The `__getattr__()` method should therefore return an attribute value (which may be a default value) or raise an `AttributeError` if appropriate.

It is important to avoid implementing code that will recursively call itself (for example calling `self.name` within `__getattr__()` will result in a recursive call back to `__getattr__()`!). To avoid this the implementation of the method should either access the `__dict__` directory or call the base classes `__getattr__()` method.

An example is given below of a simple `__getattr__()` method that logs the call to a method and then passes the invocation onto the base class implementation:

```
class Student:
    count = 0

    def __init__(self, name):
        self.name = name
        Student.count += 1

    # Method called if attribute is unknown
    def __getattr__(self, attribute):
        print('__getattr__: ', attribute)
        return 'default'

    # Method will always be called when an attribute
    # is accessed, will only called __getattr__ if it
    # does so explicitly or if an AttributeError is raised
    def __getattr__(self, name):
        print('__getattr__()', name)
        return object.__getattr__(self, name)
```

```
def my_default(self):
    return 'default'
```

We can use this version of the class with the following code snippet

```
student = Student('Katie')

print('student.name:', student.name) # instance lookup

res1 = student.dummy_attribute # invoke missing attribute
print('student.dummy_attribute:', res1)
```

The output from this is now:

```
__getattr__() name
student.name: Katie
__getattr__() dummy_attribute
__getattr__: dummy_attribute
student.dummy_attribute: default
```

As you can see from this the `__getattr__()` method is called for both `student.name` and `student.dummy_attribute`. However the `__getattr__()` method is only called when the `dummy_attribute` is accessed.

Note that `__getattr__()` is only invoked for attribute access and not for method invocation (unlike `__getattr__()`).

44.6 Intercepting Setting an Attribute

It is also possible to intercept object/instance attribute assignment when the *dot* notation (e.g., `student.name = 'Bob'`) is being used. This can be done by implementing the `__setattr__()` method. This method is invoked instead of the assignment.

The `__setattr__()` method can perform any action required including storing the value supplied. However, to do this it should either insert the value directly into the objects dictionary (e.g., `student.__dict__['name'] = 'Bob'`) or preferably call the base class `__setattr__()` method, for example `object.__setattr__(self, name, value)` as shown below for the `Student` class:

```
class Student:
    count = 0

    def __init__(self, name):
        self.name = name
        Student.count += 1

    # Method will always be called when an attribute is set
    def __setattr__(self, name, value):
        print('__setattr__:', name, value)
```

```
object.__setattr__(self, name, value)
```

If we now define the follow program that uses the Student class:

```
student = Student('John')

student.name = 'Bob'
print('student.name:', student.name) # instance lookup
```

Running this could we generate the following output:

```
__setattr__: name John
__setattr__: name Bob
student.name: Bob
```

There are a few things to note about this output:

- Firstly the assignment of the Student's name within the `__init__()` method also invokes the `__setattr__()` method.
- Secondly, the assignment to the class variable `count` does not invoke the `__setattr__()` method.
- Thirdly the `student.name` assignment does of course invoke the `__setattr__()` method.

44.7 Exercises

This exercises focusses on attribute look up.

You should add a method to the `Account` class that can be used to handle how accounts should behave when an attempt is made to access an undefined attribute.

In this case you should log the attempt to access the attribute (which means print out a warning message) then return a default value of `-1`.

For example, if you had the following line to your application:

```
print('acc1.branch:', acc1.branch)
```

Then this should invoke the `__getattr__()` method for the undefined attribute `branch`. Print a warning message and then return the value `-1` which will be printed by the above statement.

The output of this statement should be something like:

```
__getattr__: unknown attribute accessed - branch
acc1.branch: -1
```

Chapter 45

The Python Bites Back



45.1 Fourteen Beginner Mistakes to Avoid in Python

Just like other programming languages it has its fair share of gotchas that can trip up a Python newbie. In this chapter we will look at 14 of the most common mistakes and how to solve them.

45.2 Mistake #1: Indentation

The very first thing most people notice about Python is that indentation is key to structuring your program. Unlike free form languages such as Java and C# the way in which you indent you program is meaningful and part of the program itself. For example, you indicate the statement or statements that are part of an if statement via indentation. One of the important things to note about indentation is that statements should be indented to the same level for the same language construct.

For example, the following code intends to indicate that the four print statements are all part of the for loop.

```
main.py
1 for i in range(0, 4):
2     print(i, end='')
3
4     print(i * i)
5     print(i ** i)
6     print(i * i * i)
7
```

Line 7, Column 1 Spaces: 4

However, the first print statement is indented to a different level than the following print statements (3 spaces vs. 4). This will be an error in Python and if we attempt to run this program it will result in a runtime error being generated:

```
% python main.py
File "main.py", line 4
    print(i * i)
    ^
IndentationError: unexpected indent
```

This indicates that there is an `IndentationError` at line 4—in other words the indentations do not match.

One issue that can catch out the unwary is that even if the layout looks correct in an editor you may still get a runtime error. For example, the following code looks as if it has been laid out correctly in Emacs:

The screenshot shows an Emacs editor window with a menu bar at the top containing 'File Edit Options Buffers Tools Python Help'. Below the menu bar, the following Python code is displayed:

```
for i in range(0, 4):
    print(i, end='')
    print(i * i)

    print(i ** i)
    print(i * i * i)
```

The code is visually aligned, but the indentation of the last two lines is achieved using a tab character, which is not visible in the image.

However, when we run this program, we get

```
python main.py
File "main.py", line 3
    print(i * i)
    ^
TabError: inconsistent use of tabs and spaces in indentation
```

This `TabError` error has been generated because a tab has been used for line 3 while spaces have been used for the other lines. This means that as far as Python is concerned the indentations are not the same.

By convention Python programmers should always use spaces and not tabs (although in practice within a single block of code all that is required is that you are consistent in your use of spaces or tabs). However, programmers should not mix the two.

To resolve this sort of issue you can replace the tab with the equivalent spaces. Many IDEs such as PyCharm do this for you either when you hit the tab key or when you ask the IDE to reformat the code. The above code is being written in EMACS which does not do any conversion but also does not display any visual difference, and thus this can be a difficult problem to resolve if it is the first time you encounter it.

45.3 Mistake #2: Case Sensitivity

Python is a very case sensitive programming language. In Python a variable with the name `My_Var` is a completely different variable to one called `my_var`. As Python dynamically declares variables this can cause confusion and lead to unexpected bugs in programs.

For example, in the following program two similarly named but completely different variables are declared:

```
my_var = None
print(my_var)
My_Var = input('Please input your name:')
print(my_var)
```

When the program is run and the user enters the number 5, the output generated is

```
None
Please input your name: John
None
```

This is probably not what the programmer expected. The problem is that the program declares *two* variables. One is declared on line 1 and initialized with the value `None`. The second is declared on line 3 and is initialized with the string entered by the user. Whereas the programmer probably planned to reuse the first variable from line 1 in line 3!

To solve this problem developers should take great care with how they name variables. In this case if the developer had followed the standard Python conventions as described by [PEP 8—Style Guide for Python Code](#) (note PEP actually stands for Python Enhancement Proposal), then they would have used lower case for all variable names and thus would have always referred to the variable as `my_var`.

45.4 Mistake #3: Default Initialization of Function/Method Parameters with Mutable Types

Python allows a parameter to a function or a method to have a default value. This makes the parameter optional and provides a great deal of flexibility in how a function may be invoked.

For example, the following function takes up to four parameters; the first is a required parameter, but the remaining parameters are all optional:

```
def greeter(name,
            title = 'Dr',
            prompt = 'Welcome',
            message = 'Live Long and Prosper'):
    print(prompt, title, name, '-', message)
```

When the `greeter()` function is invoked we can provide one or more arguments as shown below:

```
greeter('Eloise')
greeter('Lloyd', 'Mr')
greeter('Phoebe', 'Miss', 'Hi')
greeter('Adam', 'Mr', 'Hi', 'May the Force be with you')
```

As such default parameters are a valuable part of the language.

However, many programmers new to Python misunderstand when the initial value for a parameter is evaluated. In the above we used strings to illustrate the use of default parameters which is ideal as they are immutable types (that is once created they cannot be changed).

However, many other types in Python are mutable such as Lists. It is perfectly possible to modify a List after you have created it.

Such types, if used as default values for parameters, can lead to unexpected results. For example, consider the following function and class with a method:

```
def do_something(msg, data=[]):
    data.append(msg)
    print(data)

class Bag:
    def do_something(self, msg, data=[]):
        data.append(msg)
        print(data)
```

What happens if we run the following program that uses the above function and class:

```
do_something('John')
do_something('Paul')
do_something('Denise')

bag = Bag()
bag.do_something('Adam')
bag.do_something('Natalia')
bag.do_something('Phoebe')
```

The output from this program is

```
['John']
['John', 'Paul']
['John', 'Paul', 'Denise']
['Adam']
['Adam', 'Natalia']
['Adam', 'Natalia', 'Phoebe']
```

This may not be what you expected! The issue is that the default value is evaluated *only once*, at the point at which the function or method is defined. As the value specified for the default value of the data parameter is a list, then the *same* list object is reused each time the function is invoked with only the msg parameter being supplied.

A common work around for this problem is to use an *immutable* value for the default. Then within the function or method, the mutable object is assigned to the parameter if the default immutable value is being used. For example:

```
def do_something(msg, data=None):
    if data is None:
        data = []
    data.append(msg)
    print(data)

class Bag:
    def do_something(self, msg, data=None):
        if data is None:
            data = []
        data.append(msg)
        print(data)
```

If we now rerun the program that uses the above function and class we now get

```
['John']
['Paul']
['Bill']
['Denise']
['Adam']
['Jasmine']
['Theeban']
```

45.5 Mistake #4: Default Initialization of Function/Method Parameters Using a Function

This mistake is similar to the last one in that it relates to when the default value is evaluated. It occurs because, if a function is used to initialize a default parameter, it is also executed/evaluated at the point that the function is declared.

For example, the following code allows a programmer to log a message with a time stamp indicating when the message was printed. However, the programmer has made the `time_point` parameter optional by providing a default value for it, as supplied by the call to the `datetime.now()` function.

```
from datetime import datetime

def log_it(msg, time_point = datetime.now()):
    print(msg, time_point)

log_it('In here')
log_it('Went In here', datetime.now())
log_it('Now In here')
```

The output generated by this code is

```
In here 2021-06-29 14:28:25.100567
Went In here 2021-06-29 14:28:25.100793
Now In here 2021-06-29 14:28:25.100567
```

What you can see from this is that the first and third logged messages have the same time stamp, and the 2nd message has a different time stamp. This is because the `time_point` parameter's default value is evaluated when the function is defined, that is, when it is loaded by the Python runtime, and thus the time stamp generated for the `time_point` variable is the point in time when the function came into existence and is thus not related to when the function is called.

The fix for this is like the last mistake in that an *immutable* value should be used as the default value. If the default value is present then a new timestamp should be generated, for example:

```
def log_it(msg, time_point = None):
    if time_point is None:
        time_point = datetime.now()
    print(msg, time_point)
```

45.6 Mistake #5: Out by 1 Ranges

In Python ranges are used whenever a sequence of values is required. One typical use of a range is with a for loop. For example, to loop through a set of integers starting at zero we might write

```
for i in range(0, 5):
    print(i, end=', ')
print()
```

However, ranges are what are known as *open-closed*. That is the initial value specified for the range is *included* in the range, but the end value is *outside* or closed from the range, thus when we run this program the output is

```
0, 1, 2, 3, 4,
```

If you want to include the end value in the range a common *idiom* or pattern that Python programmers use is to explicitly specify the step value to use with the range and add that to the final value. For example:

```
step = 1
for i in range(0, 5 + step, step):
    print(i, end=', ')
print()
```

The range function can take several different optional values. In this case we are including the optional step value which is the third value passed in. We are now explicitly specifying the step value and adding the step value to the end value for the range. When we run this version of the program the output is

```
0, 1, 2, 3, 4, 5,
```

45.7 Mistake #5: Assignment Operator Versus Equality Operator

Python has two operators which may look very similar to a non-programmer; these are the = and the == operators. However, they have very different meanings; one means assignment and the other means equality. That is the = operator which is used to assign a value to a variable; while == is used to test for equality. However, it is possible in some situations to use a = when you meant to use a =. For example:

```
x = False
x == True

print(x)
```


The output from this program is

```
False
```

This is because the second statement above is a test to see if `x` has the value `True`. The result of this equality test is, however, thrown away as it is not assigned to anything and of course the value of `x` is not changed.

Most IDEs will warn you about this as shown below for the PyCharm tool:

```
x = False
x == True
```



This can help to rectify the error.

45.8 Mistake #6: Falsey and Truthy

Python is actually quite flexible when it comes to what actually is used to represent True and False, in fact the following rules apply:

- 0, '' (empty strings), None equate to False
- Non zero, non empty strings, any object equate to True

This flexibility can, however, lead to mistakes being made. For example, in the following code tests are performed against the temperature value input by the user:

```
temp = int(input('The temperature:'))
if temp and temp < 10:
    print(f'Its cold {temp}')
else:
    print(f'Its warm {temp}')
```

When this program is run with different inputs for example, 1, 0, and - 1 we get different outputs:

Run 1	Run 2	Run 3
The temperature: 1 Its cold 1	The temperature: 0 Its warm 0	The temperature: - 1 Its cold - 1

According to this when the temperature is 0 its warm!

The problem is that the conditional part of the if statement will always evaluate to false when the temperature that is input is 0. This is because 0 is treated as being equivalent to False.

This is probably not what the programmer intended. They were probably intending to guard against having a rouge value in temp. It would probably have been better to test for this directly.

In general, the best way to avoid this error is always to work with Boolean values directly using True or False. For example:

```

if temp is not None and temp < 10:
    print(f'It's cold {temp}')
else:
    print(f'It's warm {temp}')

```

45.9 Mistake #7: Hiding Built-In Functions and Classes

Python uses a specific search order to find the definitions of variables, functions or classes. This search order is known by the acronym *LEGB*, which stands for:

- (L)ocal within a function or class.
- (E)nclosed within an enclosing scope such as an enclosing function.
- (G)lobal defined at the global or outer most level.
- (B)uilt-in for built-in language elements.

This means that it is possible to hide a *built-in* function or class by giving something locally the same name.

For example, in the following code a programmer has decided to store some data in a local variable called `list`:

```

list = [1, 2, 3]
print(list)

```

If we run this code, then the output generated is

```
[1, 2, 3]
```

All seems fine until later in the code we try to convert a set to a List using the `list()` constructor function:

```

data = list((1, 2, 3))
print(data)

```

When we run this code, we get the following error generated by Python:

```

Traceback (most recent call last):
  File "main.py", line 61, in <module>
    data = list((1, 2, 3))
TypeError: 'list' object is not callable


```

This error indicates that `list()` is not a callable, that is it is not something that can be called (or executed). If we look up the Python documentation for the `list()` constructor function (see <https://docs.python.org/3/library/stdtypes.html#typeseq-list>) we can see that `list()` can be used to create a list from an iterable thing and sets are certainly iterable containers. So what is going on?

The problem is that the *name* of the variable `list` has hidden (or shadows) the *name* of the built-in function `list`. Thus when we attempt to call `list()` Python find the variable `list` containing the `List` data structure which it then tries to execute at which point the error is generated.

Most IDEs will warn you that you are shadowing/hiding a built-in name; you can use this to change the name of your variable, function or class. For example:

```
list = [1, 2, 3]
```



```
pr Shadows built-in name 'list'
    Rename the element ↵ ↵ More actions... ↵ ↵
da ten_things
pr list: list[int] = [1, 2, 3] variable
```

The above uses PyCharm, and the IDE is indicating that we are shadowing the built-in *name* `list` and that we can ‘rename the element’.

45.10 Mistake #8: Hiding Built-In Modules

The way in which Python finds modules to load is by searching a predefined series of locations. By default, this set of locations looks in the current directory before searching within the Python installation for built-in modules.

As Python modules are files, for example the module `sys` is defined in a file called `sys.py`, it is possible to create a file in the *current* directory with the same name as a system provided module. For example, `sys` is a built-in module however if we do have a file called `sys.py` in the current directory, if our code references the `sys` module:

```
import sys
```

Python will load the local file instead of the system module. When the programmer then tries to access the functions and classes defined within `sys` they would find that they were not available.

This can be very confusing the first time this happens. However, if the programmer hangs over the import IDEs will often indicate where the import is coming from:

```
import sys
sys.argv /Users/jeh/temp/pythonProjects/dsproject/sys.py
#
# import Size: 15 B
# Type: Python
# Modified: 29/06/2021, 16:36
# Created: 29/06/2021, 16:36
```

As we had expected `sys` to be a built-in module, the fact that it is being loaded from a local directory suggests this is not the correct module. If we renamed the `sys.py` file to something else and re-examine the import we can see that we are now presented with information on the module:

```
import sys
sys.argv Module sys
#
# Dynamic objects:
```

45.11 Mistake #9: Function Reference Versus Function Call

Python supports a lot of Functional Programming concepts. Key to the languages ability to do this is the fact that functions are callable objects. That is, they are objects in memory and can be referenced or executed (called).

However, this can lead to some confusion for new developers. For example, the following code is legal Python:

```
import random

messages = ['Hi', 'Hello', 'Welcome']

def get_mod():
    return messages[random.randint(0, len(messages) - 1)]

x = get_mod()
print(x)
y = get_mod
print(y)
```

The output when this code is run is

```
Welcome
<function get_mod at 0x1079c60d0>
```

The first line of the output may be the sort of thing you expected; that is one of the messages is printed out, in this case the 3rd one 'Welcome'. However, the second line of output may look very strange indeed.

What has happened is that for the declaration of the variable `x` we have assigned the result returned from calling (or executing) the function `get_mod()`. However, we have assigned the reference to the function `get_mod` to the variable `y`. In effect the variable `y` is now an *alias* for the function `get_mod`. Thus when we print out the value of `y` we get something indicating a reference to a function called `get_mod` at a particular location in memory.

From the point of view of the language this is not an error, that is, obtaining a reference to a function is a valid thing to do. However, it can be confusing for new programmers and leads to errors such as that seen above.

To correct that it is important to remember to include the `()` when invoking/calling a function, for example:

```
y = get_mod()
print(y)
```

45.12 Mistake #10: Modifying a List While Iterating Over It

There are several ways in which Python allows programmers to process each element in a list in turn. One of these is to use list index access within a for loop, such as

```
numbers = [1, 2, 3, 4]
for index in range(len(numbers)):
    print(index)
    if numbers[index] % 2 == 1:
        del numbers[index]
print(numbers)
```

This loop is being used to remove all odd numbers from the list held in the `numbers` variable.

When we run this program, however, we get an error generated:

```

0
1
2
Traceback (most recent call last):
  File "main.py", line 80, in <module>
    if numbers[index] % 2 == 1:
IndexError: list index out of range

```

The problem here is that we are changing the list that we are processing and thus the data in the list and their positions changes.

There are in fact several Pythonic ways to solve this problem including using a `filter` function, a list comprehension, iterating over a copy of the numbers list, etc. Some of these are shown below

```

numbers = [1, 2, 3, 4]
evens = list(filter(lambda i: i % 2 == 0, numbers))
print(evens)

evens2 = [i for i in numbers if i % 2 == 0]
print(evens2)

for item in numbers.copy():
    if item % 2 == 1:
        numbers.remove(item)
print(numbers)

```

The output from these examples is

```

[2, 4]
[2, 4]
[2, 4]

```

45.13 Mistake #11: Sets Can Only Contain Immutable Objects

It is possible to hold any *immutable* object within a set. This means that a `Set` can contain a reference to a `Tuple` (as that is *immutable*). We can declare a `Set` literal using the `{}` brackets. We can thus write

```

s1 = { (1, 2, 3) }
print(s1)

```

This prints out

```

{(1, 2, 3)}

```

However, we cannot nest lists or other sets within a set as these are not *immutable* types. The following would both generate a runtime error in Python:

```
# Can't have the following
s2 = { {1, 2, 3} }
print(s2)

s3 = { [1, 2, 3] }
print(s3)
```

To solve this problem, we can use `Frozensets` and `Tuples`. That is we can use `Frozensets` to nest one `Set` inside another. A `Frozenset` is exactly like a `Set` except that it is *immutable* and thus it can be nested within a `Set`. A `Set` can be converted into a `Frozenset` using the `frozenset()` constructor function. For example:

```
# Need to convert sets into frozensets
s2 = { frozenset({1, 2, 3}) }
print(s2)
```

For lists we can convert a list into a tuple to add them to a set, for example:

```
s3 = { tuple([1, 2, 3]) }
print(s3)
```

45.14 Mistake #12: Local Versus Global Variables

Global variables can be referenced from within a function. For example:

```
max = 100

def print_max():
    print(max)
print_max()
```

This prints out the value 100. However, if Python thinks you are declaring a *new local variable* that just happens to have the same name as a *global* variable it will create a local version which will **shadow** the global one. For example:

```
max = 100
def print_max():
    max = 95
    print(max)

print_max()
print(max)
```

The output from this is

```
95
100
```

However, if you wish to modify the global variable inside the function, for example as is done below:

```
max = 100
def print_max():
    max = max + 1
    print(max)

print_max()
```

Then when we try to run this code, we will get a runtime error:

```
Traceback (most recent call last):
  File "main.py", line 101, in <module>
    print_max()
  File "main.py", line 98, in print_max
    max = max + 1
UnboundLocalError: local variable 'max' referenced before assignment
```

This is because the variable `max` is local to the function, and there does not have a value assigned to when we try to access it to add 1 to the current `max`.

Why does it do this? To protect us from ourselves—Python is really saying ‘Do you really want to modify a global variable here?’ Instead, it is treating `max` as a local variable and as such it is being referenced before a value has been assigned to it.

To tell Python that we know what we are doing and that we want to reference the global variable at this point we need to use the keyword `global` with the name of the variable. For example:

```
max = 100

def print_max():
    global max
    max = max + 1
    print(max)

print_max()
print(max)
```

Now when we try to update the variable `max` inside the function `print_max()`, Python knows we mean the *global* version of the variable and uses that one. The result is that we now print out the value 101 and `max` is updated to 101 for everyone everywhere!

45.15 Mistake #13: Referencing Nonlocal Variable

It is possible to define functions inside other functions, and this can be very useful when we are working with collections of data and operations such as `map()` (which maps a function to all the elements of a collection in turn).

However, local variables are local to a specific function; even functions defined within another function cannot modify the outer functions local variable (as the inner function is a separate function). They can reference it, just as we could reference the global variable earlier; the issue is again modification.

The **global** keyword is no help here as the outer function's variables are not global, they are local to a function.

For example, if we define a nested function (inner) inside the parent outer function (outer) and want the inner function to modify the local field we have a problem:

```
def outer():
    title = 'original title'

    def inner():
        title = 'another title'
        print('inner:', title)

    inner()

    print('outer:', title)

outer()
```

In this example both `outer()` and `inner()` functions modify the `title` variable. However, they are not the same `title` variable and as long as this is what we need then that is fine; both functions have their own version of a `title` local variable.

This can be seen in the output where the `outer` function maintains its own value for `title`:

```
inner: another title
outer: original title
```

However, if what we want is for the `inner()` function to modify the `outer()` function's `title` variable then we have a problem.

This problem can be solved using the `nonlocal` keyword. This indicates that a variable is not global but is also not local to the current function and Python should look within the scope in which the function is defined to find a local variable with the same name:

If we now declare `title` as `nonlocal` in the `inner()` function, then it will use the `outer()` functions version of `title` (it will be shared between them) and thus when the `inner()` function changes the `title` it will change the it for both functions:

```
def outer():
    title = 'original title'

    def inner():
        nonlocal title
        title = 'another title'
        print('inner:', title)

    inner()

    print('outer:', title)

outer()
```

The result of running this is

```
inner: another title
outer: another title
```

45.16 Mistake #14: Copying a Mutable Type

Care needs to be taken when assigning a mutable type from one variable to another. For example, a `List` is a mutable type. When you assign a `List` from one variable to another you actually assign a reference to the `List` from one variable to another. This means that both variables will reference (point to) the same list. Thus, if the contents of the list are modified via one variable, then it will also be modified for the other variable. For example:


```
data1 = [1, 2, 3, 4]
data2 = data1
data2.append(5)
print(data1)
```

This produces the following output:

```
[1, 2, 3, 4, 5]
```

To avoid this, you can create a *copy* of the original list. This will create a new list and copy the values the original list contains over to the new list. For example:

```
data1 = [1, 2, 3, 4]
data2 = data1.copy()
data2.append(5)
print(data1)
```

Now the output from this is

```
[1, 2, 3, 4]
```

If the list contains elements that can themselves hold values you may want to use a `deepcopy()` for example `data1.deepcopy()`.

Chapter 46

TicTacToe Game



46.1 Introduction

In this chapter we will explore the creation of a simple TicTacToe (or Noughts and Crosses) game using an Object-Oriented approach. This example utilizes

- Python classes, methods and instance variables/attributes.
- Abstract Base Classes and an abstract method.
- Python properties.
- Python lists.
- A simple piece of game playing logic.
- While loops, for loops and if statements for flow of control behaviour.

The aim of the game is to make a line of 3 counters (either X or O) across a 3 by 3 grid. Each player takes a turn to place a counter. The first player to achieve a line of three (horizontal, vertically or diagonally) wins.

46.2 Classes in the Game

We will begin by identifying the key classes in the game. Note that there is not necessarily a right or wrong answer here; although one set of classes may be more obvious or easier to understand than another.


In our case we will start with what data we will need to represent for our TicTacToe game as recommended back in the 'Introduction to Object Orientation' chapter.

Our key data elements include

- the TicTacToe board itself,
- the players involved in the game (both computer and human),
- the state of the game, i.e., whose go it is and whether someone has won,

- the moves being made by the players, etc.
- the counters used which are traditionally O and X (hence the alternative name ‘Noughts and Crosses’).

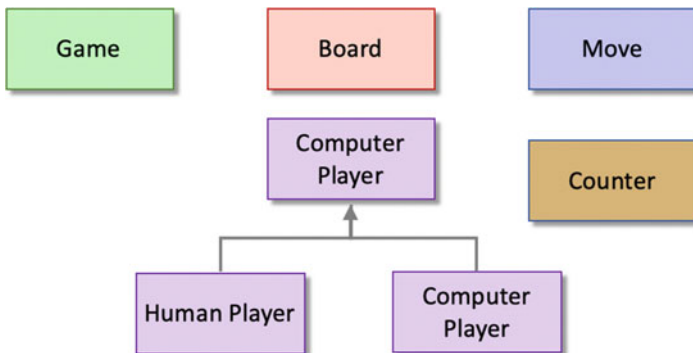
Based on an analysis of the data one possible set of classes is shown below:



In this diagram we have

- `Game` the class that will hold the board, players and the core game playing logic,
- `Board` this is a class that represents the current state of the TicTacToe board or grid within the game,
- `Human Player` this class represents the human player involved in the game,
- `Computer Player` this class represents the computer playing the game,
- `Move` this class represents a particular move made by a player,
- `Counter` which can be used to represent the counters to play with; this will be either X or Y.

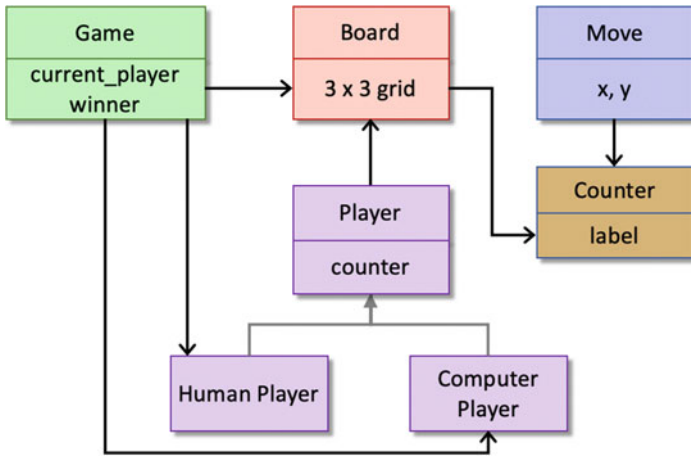
We can refine this a little further. For example, much of what constitutes a player will be common for both the human and the computer player. We can therefore introduce a new class `Player`, with both `Computer Player` and `Human Player` inheriting from this class, for example:



In terms of the data held by the classes we can say:

- Game has a board, a human and a computer player. It also has links to the current player and an attribute indicating whether a player has won.
- Board holds a 3 by 3 grid of cells. Each cell can be empty or contains a counter.
- Player Each player has a current counter and can see the board.
- Move represent a players selected move; it therefore holds the counter being played and the location to put the counter in.
- Counter holds a label indicating either X or O.

We can now update the class diagram with data and links between the classes:



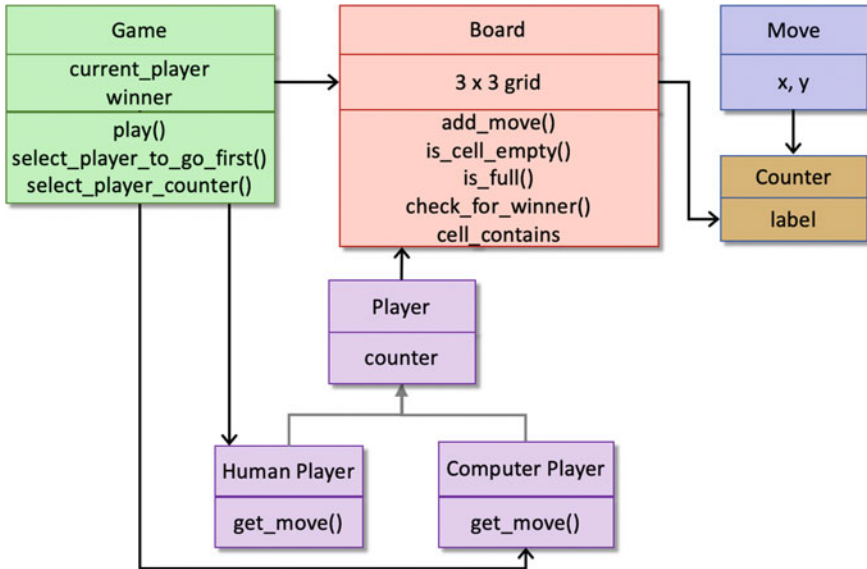
At this point it looks as though the HumanPlayer and ComputerPlayer classes are unnecessary as they do not hold any data of there own. However, in this case the behaviour for the HumanPlayer and the ComputerPlayer is quite different.

The HumanPlayer class will prompt the human user to select the next move. In contrast the ComputerPlayer class must implement an algorithm which will allow the next move to be generated within the program.

Other behavioural aspects of the classes are

- Game this must hold the overall logic of the game. It must also be able to select which player will go first. We will also allow the human player to select which counter they will play with (X or O). This logic will also be placed within the Game class.
- Board The Board class must allow a move to be made but it must also be able to verify that a move is legal (as cell is empty) and whether a game has been won or whether there is a draw. This latter logic could be located within the game instead; however, the Board holds the data necessary to determine a win or a draw, and thus we are locating the logic with the data.

We can now add the behavioural aspects of the classes to the diagram. Note we have followed the convention here for separating the data and behaviour into different areas within a class box:



We are now ready to look at the Python implementation of our class design.

46.3 Counter Class

The Counter class is given below; it is a data oriented class, sometimes referred to as a *value* type. This is because it holds *values* but does not include any behaviour.

```

class Counter:
    """ Represents a Counter used on the board """

    def __init__(self, string):
        self.label = string

    def __str__(self):
        return self.label

# Set up Counter Globals
X = Counter('X')
O = Counter('O')
  
```

We have also defined two constants X and Y to represent the X and O counters used in the game.

46.4 Move Class

The Move class is given below; it is another a data oriented class or *value* type.

```
class Move:
    """ Represents a move made by a player """

    def __init__(self, counter, x, y):
        self.x = x
        self.y = y
        self.counter = counter.
```

46.5 The Player Class

The root of the Player class hierarchy is presented below. This class is an *abstract* class in which the `get_move()` method is marked as being *abstract*. The class maintains a reference to the board and to a counter.

```
class Player(metaclass=ABCMeta):
    """ Abstract class representing a Player
        and their counter """

    def __init__(self, board):
        self.board = board
        self._counter = None

    @property
    def counter(self):
        """ Represents Players Counter - may be X or Y"""
        return self._counter

    @counter.setter
    def counter(self, value):
        self._counter = value

    @abstractmethod
    def get_move(self): pass

    def __str__(self):
        return self.__class__.__name__ + '[' + str(self.counter) +
',']
```

Note that counter is defined as a Python property.

The class Player is extended by the classes HumanPlayer and ComputerPlayer.

46.6 The HumanPlayer Class

This class extends the abstract `Player` class and defines the `get_move()` method. This method returns a `Move` object representing the counter to be placed and the location in the 3×3 grid in which to place the counter. Note that the `get_move()` method relies on a reference being maintained by the player to the board so that it can check that the selected location is empty.

To support the `get_move()` method a `_get_user_input()` method has been defined. This method could have been defined as a stand alone function as it is really independent of the `HumanPlayer`; however it has been defined within this class to keep the related behaviour together. It also follows the Python convention by starting the method name with an underbar (`_`) which indicates that the method is private and should not be accessed from outside of the class.

```
class HumanPlayer(Player):
    """ Represents a Human Player and their behaviour """

    def __init__(self, board):
        super().__init__(board)

    def _get_user_input(self, prompt):
        invalid_input = True
        while invalid_input:
            print(prompt)
            user_input = input()
            if not user_input.isdigit():
                print('Input must be a number')
            else:
                user_input_int = int(user_input)
                if user_input_int < 1 or user_input_int > 3:
                    print('input must be a number in the range 1 to 3')
                else:
                    invalid_input = False
        return user_input_int - 1

    def get_move(self):
        """ Allow the human player to enter their move """
        while True:
            row = self._get_user_input('Please input the row:')
            column = self._get_user_input('Please input the column:')

            if self.board.is_empty_cell(row, column):
                return Move(self.counter, row, column)
            else:
                print('That position is not free')
                print('Please try again')
```

46.7 The ComputerPlayer Class

This class provides an algorithmic implementation of the `get_move()` method. This algorithm tries to find the best empty grid location in which to place the counter. If it cannot find one of these locations free then it randomly finds an empty cell to fill. The `get_move()` method could be replaced with whatever game playing logic you want.

```
class ComputerPlayer(Player):
    """ Implements algorithms for playing game """

    def __init__(self, board):
        super().__init__(board)

    def randomly_select_cell(self):
        """ Use a simplistic random selection approach
        to find a cell to fill. """
        while True:
            # Randomly select the cell
            row = random.randint(0, 2)
            column = random.randint(0, 2)
            # Check to see if the cell is empty
            if self.board.is_empty_cell(row, column):
                return Move(self.counter, row, column)

    def get_move(self):
        """ Provide a very simple algorithm for selecting a move"""
        if self.board.is_empty_cell(1, 1):
            # Choose the center
            return Move(self.counter, 1, 1)
        elif self.board.is_empty_cell(0, 0):
            # Choose the top left
            return Move(self.counter, 0, 0)
        elif self.board.is_empty_cell(2, 2):
            # Choose the bottom right
            return Move(self.counter, 2, 2)
        elif self.board.is_empty_cell(0, 2):
            # Choose the top right
            return Move(self.counter, 0, 2)
        elif self.board.is_empty_cell(2, 0):
            # Choose the bottom left
            return Move(self.counter, 2, 0)
        else:
            return self.randomly_select_cell()
```


46.8 The Board Class

The Board class holds a 3 by 3 grid of cells in the form of a list of lists. It also defines the methods used to verify or make a move on the board. The `check_for_winner()` method determines if there is a winner given the current board positions.

```

class Board:
    """ The ticTacToe board """

    def __init__(self):
        # Set up the 3 by 3 grid of cells
        self.cells = [[' ', ' ', ' '], [' ', ' ', ' '], [' ', ' ', ' ']]
        self.separator = '\n' + ('-' * 11) + '\n'

    def __str__(self):
        row1 = ' ' + str(self.cells[0][0]) + ' | ' +
str(self.cells[0][1]) + ' | ' + str(self.cells[0][2])
        row2 = ' ' + str(self.cells[1][0]) + ' | ' +
str(self.cells[1][1]) + ' | ' + str(self.cells[1][2])
        row3 = ' ' + str(self.cells[2][0]) + ' | ' +
str(self.cells[2][1]) + ' | ' + str(self.cells[2][2])
        return row1 + self.separator + row2 + self.separator + row3

    def add_move(self, move):
        """ A a move to the board """
        row = self.cells[move.x]
        row[move.y] = move.counter

    def is_empty_cell(self, row, column):
        """ Check to see if a cell is empty or not """
        return self.cells[row][column] == ' '

    def cell_contains(self, counter, row, column):
        """ Check to see if a cell contains the provided counter
        """
        return self.cells[row][column] == counter

    def is_full(self):
        """ Check to see if the board is full or not """
        for row in range(0, 3):
            for column in range(0, 3):
                if self.is_empty_cell(row, column):
                    return False
        return True

    def check_for_winner(self, player):
        """ Check to see if a player has won or not """
        c = player.counter
        return (# across the top
            (self.cell_contains(c, 0, 0) and self.cell_contains(c,
0, 1) and self.cell_contains(c, 0, 2)) or
            # across the middle
            (self.cell_contains(c, 1, 0) and self.cell_contains(c,
1, 1) and self.cell_contains(c, 1, 2)) or

```

```

        # across the bottom
        (self.cell_contains(c, 2, 0) and self.cell_contains(c,
2, 1) and self.cell_contains(c, 2, 2)) or
        # down the left side
        (self.cell_contains(c, 0, 0) and self.cell_contains(c,
1, 0) and self.cell_contains(c, 2, 0)) or
        # down the middle
        (self.cell_contains(c, 0, 1) and self.cell_contains(c,
1, 1) and self.cell_contains(c, 2, 1)) or
        # down the right side
        (self.cell_contains(c, 0, 2) and self.cell_contains(c,
1, 2) and self.cell_contains(c, 2, 2)) or
        # diagonal
        (self.cell_contains(c, 0, 0) and self.cell_contains(c,
1, 1) and self.cell_contains(c, 2, 2)) or
        # other diagonal
        (self.cell_contains(c, 0, 2) and self.cell_contains(c,
1, 1) and self.cell_contains(c, 2, 0))

```

46.9 The Game Class

The Game class implements the main game playing loop. The `play()` method will loop until a winner is found. Each time round the loop one of the players takes a turn and makes a move. A check is then made to see if the game has been won.

```

class Game:
    """ Contains the Game Playing Logic """

    def __init__(self):
        self.board = Board()
        self.human = HumanPlayer(self.board)
        self.computer = ComputerPlayer(self.board)
        self.next_player = None
        self.winner = None

    def select_player_counter(self):
        """ Let the player select their counter """
        counter = ' '
        while not (counter == 'X' or counter == 'O'):
            print('Do you want to be X or O?')
            counter = input().upper()
            if counter != 'X' and counter != 'O':
                print('Input must be X or O')
        if counter == 'X':
            self.human.counter = X
            self.computer.counter = O
        else:
            self.human.counter = O
            self.computer.counter = X

```

```

def select_player_to_go_first(self):
    """ Randomly selects who will play first -
    the human or the computer. """
    if random.randint(0, 1) == 0:
        self.next_player = self.human
    else:
        self.next_player = self.computer

def play(self):
    """ Main game playing loop """
    print('Welcome to TicTacToe')
    self.select_player_counter()
    self.select_player_to_go_first()
    print(self.next_player, 'will play first first')
    while self.winner is None:
        # Human players move
        if self.next_player == self.human:
            print(self.board)
            print('Your move')
            move = self.human.get_move()
            self.board.add_move(move)
            if self.board.check_for_winner(self.human):
                self.winner = self.human
            else:
                self.next_player = self.computer
        # Computers move
        else:
            print('Computers move')
            move = self.computer.get_move()
            self.board.add_move(move)
            if self.board.check_for_winner(self.computer):
                self.winner = self.computer
            else:
                self.next_player = self.human
        # Check for a winner or a draw
        if self.winner is not None:
            print('The Winner is the' + str(self.winner))
        elif self.board.is_full():
            print('Game is a Tie')
            break
    print(self.board)

```

46.10 Running the Game

To run the game we need to instantiate the Game class and then call the play() method on the object obtained. For example:

```

def main():
    game = Game()
    game.play()
if __name__ == '__main__':
    main()

```

A sample output from running the game is given below in which the human users goes first.

```

Welcome to TicTacToe
Do you want to be X or O? X
ComputerPlayer[Y] will play first first
Computers move
| |
-----
| Y |
-----
| |
Your move
Please input the row: 1
Please input the column: 1
Computers move
X | |
-----
| Y |
-----
| | Y
Your move
Please input the row: 2
Please input the column: 1
Computers move
X | | Y
-----
X | Y |
-----
| | Y
Your move
Please input the row: 3
Please input the column: 1
The Winner is the HumanPlayer[X]
X | | Y
-----
X | Y |
-----
X | | Y

```