

THIRD EDITION

PYTHON

PROGRAMMING:

AN INTRODUCTION TO COMPUTER SCIENCE

JOHN ZELLE




FRANKLIN, BEEDLE
[INDEPENDENT PUBLISHERS SINCE 1985]

PYTHON PROGRAMMING

AN INTRODUCTION TO COMPUTER SCIENCE

THIRD EDITION



John M. Zelle
Wartburg College

Publisher	Tom Sumner (tsumner@fbeedle.com)
Editor	Brenda Jones
Production Associate	Jaron Ayres
Cover Photography	Jim Leisy ©2012

Printed in the U.S.A.

Names of all products herein are used for identification purposes only and are trademarks and/or registered trademarks of their respective owners. Franklin, Beedle & Associates Inc. makes no claim of ownership or corporate association with the products or companies that own them.

©2017 Franklin, Beedle & Associates Incorporated. No part of this book may be reproduced, stored in a retrieval system, transmitted, or transcribed, in any form or by any means—electronic, mechanical, telepathic, photocopying, recording, or otherwise—without prior written permission of the publisher. Requests for permission should be addressed as follows:

Rights and Permissions
Franklin, Beedle & Associates Incorporated
2154 NE Broadway, Suite 100
Portland, Oregon 97232

Library of Congress Cataloging-in-Publication data

Names: Zelle, John M., author.
Title: Python programming : an introduction to computer science / John M. Zelle, Wartburg College.
Description: Third edition. | Portland, Oregon : Franklin, Beedle & Associates Inc., [2016] | Includes bibliographical references and index.
Identifiers: LCCN 2016024338 | ISBN 9781590282755
Subjects: LCSH: Python (Computer program language)
Classification: LCC QA76.73.P98 Z98 2016 | DDC 005.13/3--dc23
LC record available at <https://lcn.loc.gov/2016024338>

Contents

Foreword, by Guido van Rossum	ix
Preface	x

Chapter 1 Computers and Programs **1**

1.1 The Universal Machine.....	1
1.2 Program Power	3
1.3 What Is Computer Science?	3
1.4 Hardware Basics.....	5
1.5 Programming Languages.....	6
1.6 The Magic of Python.....	9
1.7 Inside a Python Program.....	15
1.8 Chaos and Computers	18
1.9 Chapter Summary	20
1.10 Exercises.....	21

Chapter 2 Writing Simple Programs **27**

2.1 The Software Development Process.....	27
2.2 Example Program: Temperature Converter	28
2.3 Elements of Programs	31
2.3.1 Names.....	31
2.3.2 Expressions.....	32
2.4 Output Statements	34
2.5 Assignment Statements.....	36
2.5.1 Simple Assignment	37
2.5.2 Assigning Input.....	39
2.5.3 Simultaneous Assignment.....	41
2.6 Definite Loops	43

2.7	Example Program: Future Value.....	47
2.8	Chapter Summary	50
2.9	Exercises.....	51
Chapter 3 Computing with Numbers		57
3.1	Numeric Data Types	57
3.2	Type Conversions and Rounding	62
3.3	Using the Math Library	65
3.4	Accumulating Results: Factorials	68
3.5	Limitations of Computer Arithmetic	71
3.6	Chapter Summary	75
3.7	Exercises.....	76
Chapter 4 Objects and Graphics		83
4.1	Overview.....	83
4.2	The Object of Objects.....	84
4.3	Simple Graphics Programming.....	85
4.4	Using Graphical Objects	91
4.5	Graphing Future Value	96
4.6	Choosing Coordinates.....	103
4.7	Interactive Graphics	107
	4.7.1 Getting Mouse Clicks.....	107
	4.7.2 Handling Textual Input	109
4.8	Graphics Module Reference	112
	4.8.1 GraphWin Objects	113
	4.8.2 Graphics Objects.....	115
	4.8.3 Entry Objects	119
	4.8.4 Displaying Images	120
	4.8.5 Generating Colors	121
	4.8.6 Controlling Display Updates (Advanced)	121
4.9	Chapter Summary	122
4.10	Exercises.....	123
Chapter 5 Sequences: Strings, Lists, and Files		129
5.1	The String Data Type.....	129
5.2	Simple String Processing.....	133
5.3	Lists as Sequences.....	136
5.4	String Representation and Message Encoding	139
	5.4.1 String Representation.....	139
	5.4.2 Programming an Encoder	141
5.5	String Methods.....	142
	5.5.1 Programming a Decoder	142
	5.5.2 More String Methods.....	146
5.6	Lists Have Methods, Too	147
5.7	From Encoding to Encryption.....	150

5.8	Input/Output as String Manipulation	151
5.8.1	Example Application: Date Conversion	151
5.8.2	String Formatting	154
5.8.3	Better Change Counter	157
5.9	File Processing	158
5.9.1	Multi-line Strings	158
5.9.2	File Processing	159
5.9.3	Example Program: Batch Usernames	163
5.9.4	File Dialogs (Optional)	164
5.10	Chapter Summary	167
5.11	Exercises	168

Chapter 6 Defining Functions 175

6.1	The Function of Functions	175
6.2	Functions, Informally	177
6.3	Future Value with a Function	181
6.4	Functions and Parameters: The Exciting Details	183
6.5	Functions That Return Values	187
6.6	Functions that Modify Parameters	193
6.7	Functions and Program Structure	199
6.8	Chapter Summary	202
6.9	Exercises	203

Chapter 7 Decision Structures 209

7.1	Simple Decisions	209
7.1.1	Example: Temperature Warnings	210
7.1.2	Forming Simple Conditions	212
7.1.3	Example: Conditional Program Execution	214
7.2	Two-Way Decisions	216
7.3	Multi-Way Decisions	220
7.4	Exception Handling	223
7.5	Study in Design: Max of Three	227
7.5.1	Strategy 1: Compare Each to All	228
7.5.2	Strategy 2: Decision Tree	230
7.5.3	Strategy 3: Sequential Processing	231
7.5.4	Strategy 4: Use Python	234
7.5.5	Some Lessons	234
7.6	Chapter Summary	235
7.7	Exercises	236

Chapter 8 Loop Structures and Booleans 243

8.1	For Loops: A Quick Review	243
8.2	Indefinite Loops	245
8.3	Common Loop Patterns	247
8.3.1	Interactive Loops	247
8.3.2	Sentinel Loops	249

8.3.3	File Loops.....	252
8.3.4	Nested Loops.....	254
8.4	Computing with Booleans	256
8.4.1	Boolean Operators.....	256
8.4.2	Boolean Algebra	260
8.5	Other Common Structures	262
8.5.1	Post-test Loop.....	262
8.5.2	Loop and a Half	264
8.5.3	Boolean Expressions as Decisions	266
8.6	Example: A Simple Event Loop	269
8.7	Chapter Summary	275
8.8	Exercises.....	277
 Chapter 9 Simulation and Design		283
9.1	Simulating Racquetball.....	283
9.1.1	A Simulation Problem.....	284
9.1.2	Analysis and Specification	284
9.2	Pseudo-random Numbers	286
9.3	Top-Down Design	288
9.3.1	Top-Level Design	289
9.3.2	Separation of Concerns	291
9.3.3	Second-Level Design	291
9.3.4	Designing simNGames.....	293
9.3.5	Third-Level Design.....	295
9.3.6	Finishing Up	298
9.3.7	Summary of the Design Process	300
9.4	Bottom-Up Implementation.....	301
9.4.1	Unit Testing	301
9.4.2	Simulation Results.....	303
9.5	Other Design Techniques	304
9.5.1	Prototyping and Spiral Development	304
9.5.2	The Art of Design.....	306
9.6	Chapter Summary	306
9.7	Exercises.....	307
 Chapter 10 Defining Classes		313
10.1	Quick Review of Objects	313
10.2	Example Program: Cannonball	314
10.2.1	Program Specification.....	314
10.2.2	Designing the Program	315
10.2.3	Modularizing the Program	319
10.3	Defining New Classes	321
10.3.1	Example: Multi-sided Dice	321
10.3.2	Example: The Projectile Class	325
10.4	Data Processing with Class	327
10.5	Objects and Encapsulation	331

10.5.1	Encapsulating Useful Abstractions.....	331
10.5.2	Putting Classes in Modules	333
10.5.3	Module Documentation	333
10.5.4	Working with Multiple Modules	335
10.6	Widgets.....	337
10.6.1	Example Program: Dice Roller	337
10.6.2	Building Buttons.....	338
10.6.3	Building Dice.....	342
10.6.4	The Main Program	345
10.7	Animated Cannonball.....	346
10.7.1	Drawing the Animation Window.....	347
10.7.2	Creating a ShotTracker	348
10.7.3	Creating an Input Dialog.....	350
10.7.4	The Main Event Loop.....	353
10.8	Chapter Summary	355
10.9	Exercises.....	356

Chapter 11 Data Collections 363

11.1	Example Problem: Simple Statistics.....	363
11.2	Applying Lists.....	365
11.2.1	Lists and Arrays.....	366
11.2.2	List Operations.....	367
11.2.3	Statistics with Lists.....	370
11.3	Lists of Records	375
11.4	Designing with Lists and Classes	379
11.5	Case Study: Python Calculator	385
11.5.1	A Calculator as an Object.....	385
11.5.2	Constructing the Interface.....	385
11.5.3	Processing Buttons	388
11.6	Case Study: Better Cannonball Animation	392
11.6.1	Creating a Launcher.....	393
11.6.2	Tracking Multiple Shots.....	396
11.7	Non-sequential Collections.....	401
11.7.1	Dictionary Basics	401
11.7.2	Dictionary Operations	402
11.7.3	Example Program: Word Frequency	404
11.8	Chapter Summary	409
11.9	Exercises.....	410

Chapter 12 Object-Oriented Design 419

12.1	The Process of OOD.....	419
12.2	Case Study: Racquetball Simulation	422
12.2.1	Candidate Objects and Methods	422
12.2.2	Implementing SimStats	424
12.2.3	Implementing RBallGame.....	426
12.2.4	Implementing Player	429

	12.2.5	The Complete Program.....	430
12.3		Case Study: Dice Poker.....	433
	12.3.1	Program Specification.....	433
	12.3.2	Identifying Candidate Objects	434
	12.3.3	Implementing the Model	436
	12.3.4	A Text-Based UI.....	440
	12.3.5	Developing a GUI.....	443
12.4		OO Concepts	451
	12.4.1	Encapsulation	452
	12.4.2	Polymorphism.....	453
	12.4.3	Inheritance.....	453
12.5		Chapter Summary	455
12.6		Exercises.....	456

Chapter 13 Algorithm Design and Recursion 459

13.1		Searching.....	460
	13.1.1	A Simple Searching Problem	460
	13.1.2	Strategy 1: Linear Search	461
	13.1.3	Strategy 2: Binary Search	462
	13.1.4	Comparing Algorithms	463
13.2		Recursive Problem Solving	465
	13.2.1	Recursive Definitions.....	466
	13.2.2	Recursive Functions	468
	13.2.3	Example: String Reversal	469
	13.2.4	Example: Anagrams	471
	13.2.5	Example: Fast Exponentiation	472
	13.2.6	Example: Binary Search	473
	13.2.7	Recursion vs. Iteration	474
13.3		Sorting Algorithms.....	477
	13.3.1	Naive Sorting: Selection Sort.....	477
	13.3.2	Divide and Conquer: Merge Sort	479
	13.3.3	Comparing Sorts.....	481
13.4		Hard Problems.....	484
	13.4.1	Tower of Hanoi.....	484
	13.4.2	The Halting Problem	489
	13.4.3	Conclusion.....	492
13.5		Chapter Summary	493
13.6		Exercises.....	494

Appendix A Python Quick Reference 503

Appendix C Glossary 513

Index 525

Foreword

When the publisher first sent me a draft of this book, I was immediately excited. Disguised as a Python textbook, it is really an introduction to the fine art of programming, using Python merely as the preferred medium for beginners. This is how I have always imagined Python would be most useful in education: not as the only language, but as a first language, just as in art one might start learning to draw using a pencil rather than trying to paint in oil right away.

The author mentions in his preface that Python is near-ideal as a first programming language, without being a “toy language.” As the creator of Python I don’t want to take full credit for this: Python was derived from ABC, a language designed to teach programming in the early 1980s by Lambert Meertens, Leo Geurts, and others at CWI (National Research Institute for Mathematics and Computer Science) in Amsterdam. If I added anything to their work, it was making Python into a non-toy language, with a broad user base and an extensive collection of standard and third-party application modules.

I have no formal teaching experience, so I may not be qualified to judge its educational effectiveness. Still, as a programmer with nearly 30 years experience, reading through the chapters I am continuously delighted by the book’s clear explanations of difficult concepts. I also like the many good exercises and questions which both test understanding and encourage thinking about deeper issues.

Reader of this book, congratulations! You will be well rewarded for studying Python. I promise you’ll have fun along the way, and I hope you won’t forget your first language once you have become a proficient software developer.

—*Guido van Rossum*

Preface

This book is designed to be used as a primary textbook in a college-level first course in computing. It takes a fairly traditional approach, emphasizing problem solving, design, and programming as the core skills of computer science. However, these ideas are illustrated using a non-traditional language, namely Python. In my teaching experience, I have found that many students have difficulty mastering the basic concepts of computer science and programming. Part of this difficulty can be blamed on the complexity of the languages and tools that are most often used in introductory courses. Consequently, this textbook was written with a single overarching goal: to introduce fundamental computer science concepts as simply as possible without being simplistic. Using Python is central to this goal.

Traditional systems languages such as C++, Ada, and Java evolved to solve problems in large-scale programming, where the primary emphasis is on structure and discipline. They were not designed to make writing small- or medium-scale programs easy. The recent rise in popularity of scripting (sometimes called “agile”) languages, such as Python, suggests an alternative approach. Python is very flexible and makes experimentation easy. Solutions to simple problems are simply and elegantly expressed. Python provides a great laboratory for the neophyte programmer.

Python has a number of features that make it a near-perfect choice as a first programming language. The basic structures are simple, clean, and well designed, which allows students to focus on the primary skills of algorithmic thinking and program design without getting bogged down in arcane language details. Concepts learned in Python carry over directly to subsequent study of

systems languages such as C++ and Java. But Python is not a “toy language.” It is a real-world production language that is freely available for virtually every programming platform and comes standard with its own easy-to-use integrated programming environment. The best part is that Python makes learning to program fun again.

Although I use Python as the language, teaching Python is not the main point of this book. Rather, Python is used to illustrate fundamental principles of design and programming that apply in any language or computing environment. In some places I have purposely avoided certain Python features and idioms that are not generally found in other languages. There are many good books about Python on the market; this book is intended as an introduction to computing. Besides using Python, there are other features of this book designed to make it a gentler introduction to computer science. Some of these features include:

- **Extensive use of computer graphics.** Students love working on programs that include graphics. This book presents a simple-to-use graphics package (provided as a Python module) that allows students both to learn the principles of computer graphics and to practice object-oriented concepts without the complexity inherent in a full-blown graphics library and event-driven programming.
- **Interesting examples.** The book is packed with complete programming examples to solve real problems.
- **Readable prose.** The narrative style of the book introduces key computer science concepts in a natural way as an outgrowth of a developing discussion. I have tried to avoid random facts or tangentially related sidebars.
- **Flexible spiral coverage.** Since the goal of the book is to present concepts simply, each chapter is organized so that students are introduced to new ideas in a gradual way, giving them time to assimilate an increasing level of detail as they progress. Ideas that take more time to master are introduced in early chapters and reinforced in later chapters.
- **Just-in-time object coverage.** The proper place for the introduction of object-oriented techniques is an ongoing controversy in computer science education. This book is neither strictly “objects early” nor “objects late,” but gradually introduces object concepts after a brief initial grounding in the basics of imperative programming. Students learn multiple design

techniques, including top-down (functional decomposition), spiral (prototyping), and object-oriented methods. Additionally, the textbook material is flexible enough to accommodate other approaches.

- **Extensive end-of-chapter problems.** Exercises at the end of every chapter provide ample opportunity for students to reinforce their mastery of the chapter material and to practice new programming skills.

Changes in the Second and Third Editions

The first edition of the textbook has aged gracefully, and the approach it takes remains just as relevant now as when it was first published.

While fundamental principles do not change, the technology environment does. With the release of Python 3.0, updates to the original material became necessary. The second edition was basically the same as the original textbook, except that it was updated to use Python 3. Virtually every program example in the book had to be modified for the new Python. Additionally, to accommodate certain changes in Python (notably the removal of the string library), the material was reordered slightly to cover object terminology before discussing string processing. A beneficial side effect of this change was an even earlier introduction of computer graphics to pique student interest.

The third edition continues the tradition of updating the text to reflect new technologies while maintaining a time-tested approach to teaching introductory computer science. An important change to this edition is the removal of most uses of `eval` and the addition of a discussion of its dangers. In our increasingly connected world, it's never too early to begin considering computer security issues.

Several new graphics examples, developed throughout chapters 4–12, have been added to introduce new features of the graphics library that support animations, including simple video game development. This brings the text up to date with the types of final projects that are often assigned in modern introductory classes.

Smaller changes have been made throughout the text, including:

- Material on file dialogs has been added in Chapter 5.
- Chapter 6 has been expanded and reorganized to emphasize value-returning functions.

- Coverage has been streamlined and simplified to use IDLE (the standard “comes-with-Python” development environment) consistently. This makes the text more suitable for self-study as well as for use as a classroom textbook.
- Technology references have been updated.
- To further accommodate self-studiers, end-of-chapter solutions for this third edition are freely available online. Classroom instructors wishing to use alternative exercises can request those from the publisher. Self-studiers and instructors alike can visit <https://fbedle.com> for details.

Coverage Options

In keeping with the goal of simplicity, I have tried to limit the amount of material that would not be covered in a first course. Still, there is probably more material here than can be covered in a typical one-semester introduction. My classes cover virtually all of the material in the first 12 chapters in order, though not necessarily covering every section in depth. One or two topics from Chapter 13 (“Algorithm Design and Recursion”) are generally interspersed at appropriate places during the term.

Recognizing that different instructors prefer to approach topics in different ways, I have tried to keep the material relatively flexible. Chapters 1–4 (“Computers and Programs,” “Writing Simple Programs,” “Computing with Numbers,” “Objects and Graphics”) are essential introduction and should probably be covered in order. The initial portions of Chapter 5 (“Sequences: Strings, Lists, and Files”) on string processing are also fundamental, but the later topics such as string formatting and file processing can be delayed until needed later on. Chapters 6–8 (“Defining Functions,” “Decision Structures,” and “Loop Structures and Booleans”) are designed to stand independently and can be taken in virtually any order. Chapters 9–12 on design approaches are written to be taken in order, but the material in Chapter 11 (“Data Collections”) could easily be moved earlier, should the instructor want to cover lists (arrays) before various design techniques. Instructors wishing to emphasize object-oriented design need not spend much time on Chapter 9. Chapter 13 contains more advanced material that may be covered at the end or interspersed at various places throughout the course.

Acknowledgments

My approach to CS1 has been influenced over the years by many fine textbooks that I have read and used for classes. Much that I have learned from those books has undoubtedly found its way into these pages. There are a few specific authors whose approaches have been so important that I feel they deserve special mention. A.K. Dewdney has always had a knack for finding simple examples that illustrate complex issues; I have borrowed a few of those and given them new legs in Python. I also owe a debt to wonderful textbooks from both Owen Astrachan and Cay Horstmann. The graphics library I introduce in Chapter 4 was directly inspired by my experience teaching with a similar library designed by Horstmann. I also learned much about teaching computer science from Nell Dale, for whom I was fortunate enough to serve as a TA when I was a graduate student at the University of Texas.

Many people have contributed either directly or indirectly to the production of this book. I have also received much help and encouragement from my colleagues (and former colleagues) at Wartburg College: Lynn Olson for his unflagging support at the very beginning; Josef Breutzmann, who supplied many project ideas; and Terry Letsche, who prepared PowerPoint slides for the first and third editions.

I want to thank the following individuals who read or commented on the manuscript for the first edition: Rus May, Morehead State University; Carolyn Miller, North Carolina State University; Guido Van Rossum, Google; Jim Sager, California State University, Chico; Christine Shannon, Centre College; Paul Tymann, Rochester Institute of Technology; Suzanne Westbrook, University of Arizona. I am grateful to Dave Reed at Capital University, who used early versions of the first edition, offered numerous insightful suggestions, and worked with Jeffrey Cohen at University of Chicago to supply alternate end-of-chapter exercises for this edition. Ernie Ackermann test drove the second edition at Mary Washington College. The third edition was test driven in classes by Theresa Migler at California Polytechnic State University in San Luis Obispo and my colleague Terry Letsche; and David Bantz provided feedback on a draft. Thanks to all for their valuable observations and suggestions.

I also want to acknowledge the fine folks at Franklin, Beedle, and Associates, especially Tom Sumner, Brenda Jones, and Jaron Ayres, who turned my pet project into a real textbook. This edition is dedicated to the memory of Jim Leisy, the founder of Franklin, Beedle and Associates, who passed away unex-

pectedly as the third edition was getting off the ground. Jim was an amazing man of unusually wide-ranging interests. It was his vision, guidance, relentless enthusiasm, and a fair bit of determined prodding, that ultimately molded me into a textbook author and made this book a success.

A special thanks also goes out to all my students, who have taught me so much about teaching, and to Wartburg College for giving me sabbatical support to work on the book. Last, but most importantly, I acknowledge my wife, Elizabeth Bingham, who has served as editor, advisor, and morale booster while putting up with me during my writing spells.

—JMZ

Chapter 1

Computers and Programs

Objectives

- To understand the respective roles of hardware and software in computing systems.
- To learn what computer scientists study and the techniques that they use.
- To understand the basic design of a modern computer.
- To understand the form and function of computer programming languages.
- To begin using the Python programming language.
- To learn about chaotic models and their implications for computing.

1.1 The Universal Machine

Almost everyone has used a computer at one time or another. Perhaps you have played computer games or used a computer to write a paper, shop online, listen to music, or connect with friends via social media. Computers are used to predict the weather, design airplanes, make movies, run businesses, perform financial transactions, and control factories.

Have you ever stopped to wonder what exactly a computer is? How can one device perform so many different tasks? These basic questions are the starting point for learning about computers and computer programming.

A modern computer can be defined as “a machine that stores and manipulates information under the control of a changeable program.” There are two key elements to this definition. The first is that computers are devices for manipulating information. This means we can put information into a computer, and it can transform the information into new, useful forms, and then output or display the information for our interpretation.

Computers are not the only machines that manipulate information. When you use a simple calculator to add up a column of numbers, you are entering information (the numbers) and the calculator is processing the information to compute a running sum which is then displayed. Another simple example is a gas pump. As you fill your tank, the pump uses certain inputs: the current price of gas per gallon and signals from a sensor that reads the rate of gas flowing into your car. The pump transforms this input into information about how much gas you took and how much money you owe.

We would not consider either the calculator or the gas pump as full-fledged computers, although modern versions of these devices may actually contain embedded computers. They are different from computers in that they are built to perform a single, specific task. This is where the second part of our definition comes into the picture: Computers operate under the control of a changeable program. What exactly does this mean?

A *computer program* is a detailed, step-by-step set of instructions telling a computer exactly what to do. If we change the program, then the computer performs a different sequence of actions, and hence, performs a different task. It is this flexibility that allows your PC to be at one moment a word processor, at the next moment a financial planner, and later on, an arcade game. The machine stays the same, but the program controlling the machine changes.

Every computer is just a machine for *executing* (carrying out) programs. There are many different kinds of computers. You might be familiar with Macintoshes, PCs, laptops, tablets and smartphones, but there are literally thousands of other kinds of computers both real and theoretical. One of the remarkable discoveries of computer science is the realization that all of these different computers have the same power; with suitable programming, each computer can basically do all the things that any other computer can do. In this sense, the PC that you might have sitting on your desk is really a universal machine. It can do anything you want it to do, provided you can describe the task to be accomplished in sufficient detail. Now that’s a powerful machine!

1.2 Program Power

You have already learned an important lesson of computing: *Software* (programs) rules the *hardware* (the physical machine). It is the software that determines what any computer can do. Without software, computers would just be expensive paperweights. The process of creating software is called *programming*, and that is the main focus of this book.

Computer programming is a challenging activity. Good programming requires an ability to see the big picture while paying attention to minute detail. Not everyone has the talent to become a first-class programmer, just as not everyone has the skills to be a professional athlete. However, virtually anyone *can* learn how to program computers. With some patience and effort on your part, this book will help you to become a programmer.

There are lots of good reasons to learn programming. Programming is a fundamental part of computer science and is, therefore, important to anyone interested in becoming a computer professional. But others can also benefit from the experience. Computers have become a commonplace tool in our society. Understanding the strengths and limitations of this tool requires an understanding of programming. Non-programmers often feel they are slaves of their computers. Programmers, however, are truly in control. If you want to become a more intelligent user of computers, then this book is for you.

Programming can also be loads of fun. It is an intellectually engaging activity that allows people to express themselves through useful and sometimes remarkably beautiful creations. Believe it or not, many people actually write computer programs as a hobby. Programming also develops valuable problem-solving skills, especially the ability to analyze complex systems by reducing them to interactions of understandable subsystems.

As you probably know, programmers are in great demand. More than a few liberal arts majors have turned a couple of computer programming classes into a lucrative career option. Computers are so commonplace in the business world today that the ability to understand and program computers might just give you the edge over your competition regardless of your occupation. When inspiration strikes, you could be poised to write the next killer app.

1.3 What Is Computer Science?

You might be surprised to learn that computer science is not the study of computers. A famous computer scientist named Edsger Dijkstra once quipped that

computers are to computer science what telescopes are to astronomy. The computer is an important tool in computer science, but it is not itself the object of study. Since a computer can carry out any process that we can describe, the real question is “What processes can we describe?” To put it another way, the fundamental question of computer science is simply “What can be computed?” Computer scientists use numerous techniques of investigation to answer this question. The three main ones are *design*, *analysis*, and *experimentation*.

One way to demonstrate that a particular problem can be solved is to actually design a solution. That is, we develop a step-by-step process for achieving the desired result. Computer scientists call this an *algorithm*. That’s a fancy word that basically means “recipe.” The design of algorithms is one of the most important facets of computer science. In this book you will find techniques for designing and implementing algorithms.

One weakness of design is that it can only answer the question “What is computable?” in the positive. If I can devise an algorithm, then the problem is solvable. However, failing to find an algorithm does not mean that a problem is unsolvable. It may mean that I’m just not smart enough, or I haven’t hit upon the right idea yet. This is where analysis comes in.

Analysis is the process of examining algorithms and problems mathematically. Computer scientists have shown that some seemingly simple problems are not solvable by *any* algorithm. Other problems are *intractable*. The algorithms that solve these problems take too long or require too much memory to be of practical value. Analysis of algorithms is an important part of computer science; throughout this book we will touch on some of the fundamental principles. Chapter 13 has examples of unsolvable and intractable problems.

Some problems are too complex or ill-defined to lend themselves to analysis. In such cases, computer scientists rely on experimentation; they actually implement systems and then study the resulting behavior. Even when theoretical analysis is done, experimentation is often needed in order to verify and refine the analysis. For most problems, the bottom line is whether a working, reliable system can be built. Often we require empirical testing of the system to determine that this bottom line has been met. As you begin writing your own programs, you will get plenty of opportunities to observe your solutions in action.

I have defined computer science in terms of designing, analyzing, and evaluating algorithms, and this is certainly the core of the academic discipline. These days, however, computer scientists are involved in far-flung activities, all of which fall under the general umbrella of computing. Some examples




Figure 1.1: Functional view of a computer

include mobile computing, networking, human-computer interaction, artificial intelligence, computational science (using powerful computers to model scientific processes), databases and data mining, software engineering, web and multimedia design, music production, management information systems, and computer security. Wherever computing is done, the skills and knowledge of computer science are being applied.

1.4 Hardware Basics

You don't have to know all the details of how a computer works to be a successful programmer, but understanding the underlying principles will help you master the steps we go through to put our programs into action. It's a bit like driving a car. Knowing a little about internal combustion engines helps to explain why you have to do things like fill the gas tank, start the engine, step on the accelerator, and so on. You could learn to drive by just memorizing what to do, but a little more knowledge makes the whole process much more understandable. Let's take a moment to "look under the hood" of your computer.

Although different computers can vary significantly in specific details, at a higher level all modern digital computers are remarkably similar. Figure 1.1 shows a functional view of a computer. The *central processing unit* (CPU) is the "brain" of the machine. This is where all the basic operations of the computer are carried out. The CPU can perform simple arithmetic operations like adding two numbers and can also do logical operations like testing to see if two numbers are equal.

The memory stores programs and data. The CPU can directly access only information that is stored in *main memory* (called RAM for *Random Access Memory*). Main memory is fast, but it is also volatile. That is, when the power is turned off, the information in the memory is lost. Thus, there must also be some secondary memory that provides more permanent storage.

In a modern personal computer, the principal secondary memory is typically an internal hard disk drive (HDD) or a solid state drive (SSD). An HDD stores information as magnetic patterns on a spinning disk, while an SSD employs electronic circuits known as flash memory. Most computers also support removeable media for secondary memory such as USB memory “sticks” (also a form of flash memory) and DVDs (digital versatile discs), which store information as optical patterns that are read and written by a laser.

Humans interact with the computer through input and output devices. You are probably familiar with common devices such as a keyboard, mouse, and monitor (video screen). Information from input devices is processed by the CPU and may be shuffled off to the main or secondary memory. Similarly, when information needs to be displayed, the CPU sends it to one or more output devices.

So what happens when you fire up your favorite game or word processing program? First, the instructions that comprise the program are copied from the (more) permanent secondary memory into the main memory of the computer. Once the instructions are loaded, the CPU starts executing the program.

Technically the CPU follows a process called the *fetch-execute cycle*. The first instruction is retrieved from memory, decoded to figure out what it represents, and the appropriate action carried out. Then the next instruction is fetched, decoded, and executed. The cycle continues, instruction after instruction. This is really all the computer does from the time that you turn it on until you turn it off again: fetch, decode, execute. It doesn’t seem very exciting, does it? But the computer can execute this stream of simple instructions with blazing speed, zipping through billions of instructions each second. Put enough simple instructions together in just the right way, and the computer does amazing things.

1.5 Programming Languages

Remember that a program is just a sequence of instructions telling a computer what to do. Obviously, we need to provide those instructions in a language that a computer can understand. It would be nice if we could just tell a computer what to do using our native language, like they do in science fiction

movies. (“Computer, how long will it take to reach planet Alpha at maximum warp?”) Computer scientists have made great strides in this direction; you may be familiar with technologies such as Siri (Apple), Google Now (Android), and Cortana (Microsoft). But as anyone who has seriously used such systems can attest, designing a computer program to fully understand human language is still an unsolved problem.

Even if computers could understand us, human languages are not very well suited for describing complex algorithms. Natural language is fraught with ambiguity and imprecision. For example, if I say “I saw the man in the park with the telescope,” did I have the telescope, or did the man? And who was in the park? We understand each other most of the time only because all humans share a vast store of common knowledge and experience. Even then, miscommunication is commonplace.

Computer scientists have gotten around this problem by designing notations for expressing computations in an exact and unambiguous way. These special notations are called *programming languages*. Every structure in a programming language has a precise form (its *syntax*) and a precise meaning (its *semantics*). A programming language is something like a code for writing down the instructions that a computer will follow. In fact, programmers often refer to their programs as *computer code*, and the process of writing an algorithm in a programming language is called *coding*.

Python is one example of a programming language and is the language that we will use throughout this book.¹ You may have heard of some other commonly used languages, such as C++, Java, Javascript, Ruby, Perl, Scheme, or BASIC. Computer scientists have developed literally thousands of programming languages, and the languages themselves evolve over time yielding multiple, sometimes very different, versions. Although these languages differ in many details, they all share the property of having well-defined, unambiguous syntax and semantics.

All of the languages mentioned above are examples of *high-level* computer languages. Although they are precise, they are designed to be used and understood by humans. Strictly speaking, computer hardware can understand only a very low-level language known as *machine language*.

Suppose we want the computer to add two numbers. The instructions that the CPU actually carries out might be something like this:

¹This edition of the text was developed and tested using Python version 3.4. Python 3.5 is now available. If you have an earlier version of Python installed on your computer, you should upgrade to the latest stable 3.x version to try out the examples.

```

load the number from memory location 2001 into the CPU
load the number from memory location 2002 into the CPU
add the two numbers in the CPU
store the result into location 2003

```

This seems like a lot of work to add two numbers, doesn't it? Actually, it's even more complicated than this because the instructions and numbers are represented in *binary* notation (as sequences of 0s and 1s).

In a high-level language like Python, the addition of two numbers can be expressed more naturally: $c = a + b$. That's a lot easier for us to understand, but we need some way to translate the high-level language into the machine language that the computer can execute. There are two ways to do this: a high-level language can either be *compiled* or *interpreted*.

A *compiler* is a complex computer program that takes another program written in a high-level language and translates it into an equivalent program in the machine language of some computer. Figure 1.2 shows a block diagram of the compiling process. The high-level program is called *source code*, and the resulting *machine code* is a program that the computer can directly execute. The dashed line in the diagram represents the execution of the machine code (also known as “running the program”).




Figure 1.2: Compiling a high-level language

An *interpreter* is a program that simulates a computer that understands a high-level language. Rather than translating the source program into a machine language equivalent, the interpreter analyzes and executes the source code instruction by instruction as necessary. Figure 1.3 illustrates the process.

The difference between interpreting and compiling is that compiling is a one-shot translation; once a program is compiled, it may be run over and over again without further need for the compiler or the source code. In the interpreted

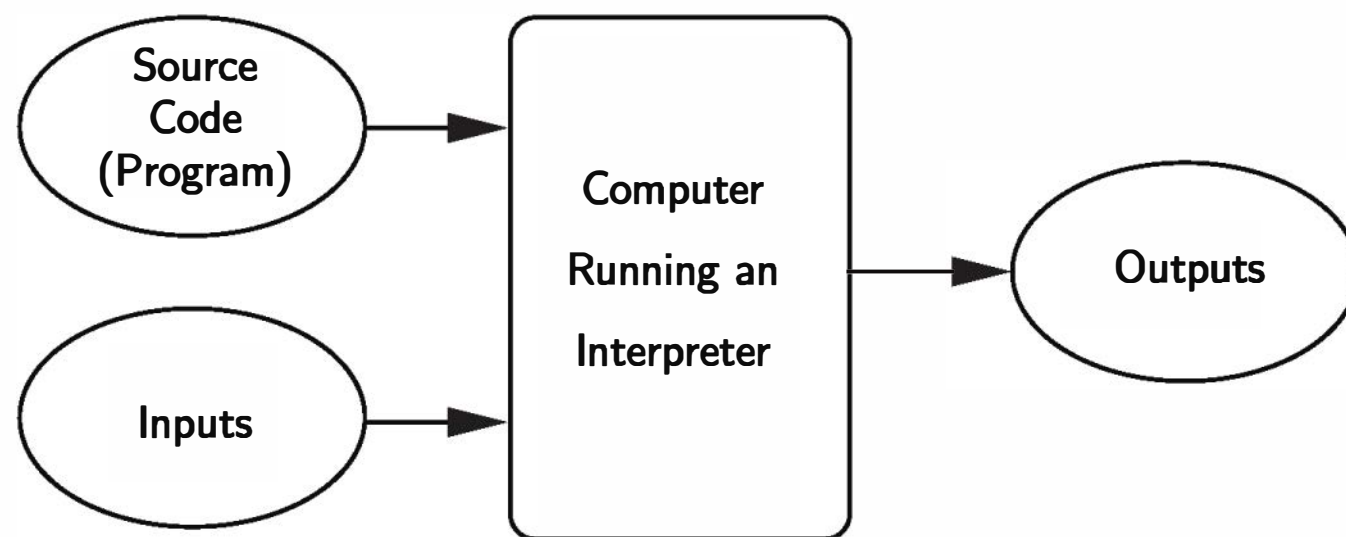


Figure 1.3: Interpreting a high-level language

case, the interpreter and the source are needed every time the program runs. Compiled programs tend to be faster, since the translation is done once and for all, but interpreted languages lend themselves to a more flexible programming environment as programs can be developed and run interactively.

The translation process highlights another advantage that high-level languages have over machine language: *portability*. The machine language of a computer is created by the designers of the particular CPU. Each kind of computer has its own machine language. A program for an Intel i7 Processor in your laptop won't run directly on an ARMv8 CPU in your smartphone. On the other hand, a program written in a high-level language can be run on many different kinds of computers as long as there is a suitable compiler or interpreter (which is just another program). As a result, I can run the exact same Python program on my laptop and my tablet; even though they have different CPUs, they both sport a Python interpreter.

1.6 The Magic of Python

Now that you have all the technical details, it's time to start having fun with Python. The ultimate goal is to make the computer do our bidding. To this end, we will write programs that control the computational processes inside the machine. You have already seen that there is no magic in this process, but in some ways programming *feels* like magic.

The computational processes inside the computer are like magical spirits that we can harness for our work. Unfortunately, those spirits only understand a very arcane language that we do not know. What we need is a friendly genie that can direct the spirits to fulfill our wishes. Our genie is a Python interpreter. We can give instructions to the Python interpreter, and it directs the underlying spirits