O'REILLY®

Head First GO

A Brain-Friendly Guide



Learn to write simple, maintainable code



Avoid embarrassing type errors



Bend your mind around more than 40 Go exercises



A Learner's Guide to Go Programming



Focus on the features that will make you most productive

Run functions concurrently with goroutines



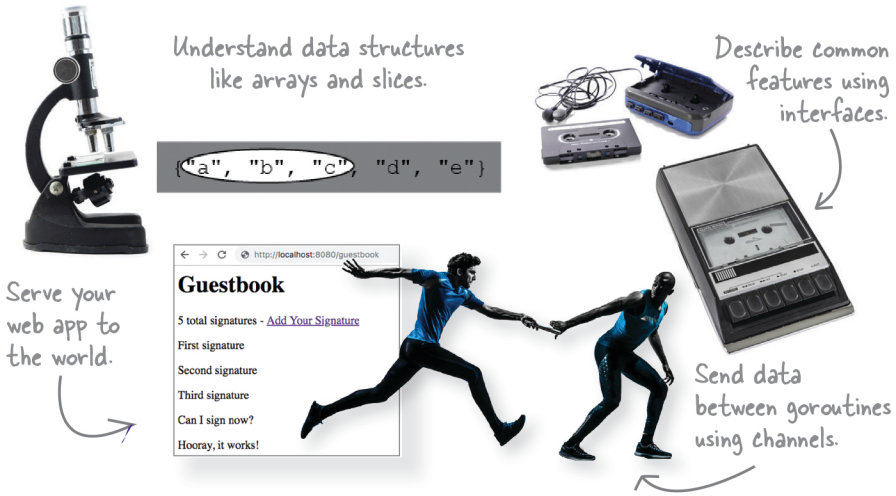
Jay McGavren

O'REILLY®

Go

What will you learn from this book?

Go makes it easy to build software that's simple, reliable, and efficient. And this book makes it easy for programmers like you to get started. Go is designed for high-performance networking and multiprocessing, but—like Python and JavaScript—the language is easy to read and use. With this practical hands-on guide, you'll learn how to write Go code using clear examples that demonstrate the language in action. Best of all, you'll understand the conventions and techniques that employers want entry-level Go developers to know.



Why does this book look so different?

Based on the latest research in cognitive science and learning theory, *Head First Go* uses a visually rich format to engage your mind rather than a text-heavy approach that puts you to sleep. Why waste your time struggling with new concepts? This multisensory learning experience is designed for the way your brain really works.

"*Head First Go* strikes the right tone for those of us who just want to get on and do things rather than wring our hands over endless syntax and technical decisions. Given just a little of your time, you'll learn something new and useful here, even if your work predominantly focuses on other languages."

—*Peter Cooper*
Editor, Golang Weekly

US $59.99 CAN $79.99
ISBN: 978-1-491-96955-7





Twitter: @oreillymedia
facebook.com/oreilly

oreilly.com

Head First
Go



Jay McGavren

Beijing • Boston • Farnham • Sebastopol • Tokyo

O'REILLY®

Head First Go

by Jay McGavren

Copyright © 2019 Jay McGavren. All rights reserved.

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly Media books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (*http://oreilly.com*). For more information, contact our corporate/institutional sales department: (800) 998-9938 or *corporate@oreilly.com*.

Series Creators:	Kathy Sierra, Bert Bates
Editor:	Jeff Bleiel
Cover Designer:	Randy Comer
Production Editor:	Kristen Brown
Production Services:	Rachel Monaghan
Indexer:	Lucie Haskins
Head First logo:	Eric Freeman

Printing History:

April 2019: First Edition.

The O'Reilly logo is a registered trademark of O'Reilly Media, Inc. The *Head First* series designations, *Head First Go*, and related trade dress are trademarks of O'Reilly Media, Inc.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and O'Reilly Media, Inc., was aware of a trademark claim, the designations have been printed in caps or initial caps.

While every precaution has been taken in the preparation of this book, the publisher and the author assume no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.

Code for this book was developed using 100% recycled electrons.

ISBN: 978-1-491-96955-7
[LSI] [2019-07-26]

To my eternally patient Christine.

Author of Head First Go



Jay McGavren

Jay McGavren is the author of *Head First Ruby* and *Head First Go*, both published by O'Reilly. He also teaches software development at Treehouse.

His home in the Phoenix suburbs houses himself, his lovely wife, and an alarmingly variable number of kids and dogs.

You can visit Jay's personal website at *http://jay.mcgavren.com*.

Table of Contents (Summary)

	Intro	xxv
1	Let's Get Going: *Syntax Basics*	1
2	Which Code Runs Next?: *Conditionals and Loops*	31
3	Call Me: *Functions*	79
4	Bundles of Code: *Packages*	113
5	On the List: *Arrays*	149
6	Appending Issue: *Slices*	175
7	Labeling Data: *Maps*	205
8	Building Storage: *Structs*	231
9	You're My Type: *Defined Type*	265
10	Keep It to Yourself: *Encapsulation and Embedding*	289
11	What Can You Do?: *Interfaces*	321
12	Back on Your Feet: *Recovering from Failure*	349
13	Sharing Work: *Goroutines and Channels*	379
14	Code Quality Assurance: *Automated Testing*	401
15	Responding to Requests: *Web Apps*	425
16	A Pattern to Follow: *HTML Templates*	445
A	Understanding os.OpenFile: *Opening Files*	481
B	Six Things We Didn't Cover: *Leftovers*	495

Table of Contents (the real thing)

Intro

Your brain on Go. Here *you* are trying to *learn* something, while here your *brain* is, doing you a favor by making sure the learning doesn't *stick*. Your brain's thinking, "Better leave room for more important things, like which wild animals to avoid and whether naked snowboarding is a bad idea." So how *do* you trick your brain into thinking that your life depends on knowing how to program in Go?

Who is this book for?	xxvi
We know what you're thinking	xxvii
We know what your brain is thinking	xxvii
Metacognition: thinking about thinking	xxix
Here's what WE did	xxx
Read me	xxxii
Acknowledgments	xxxiii

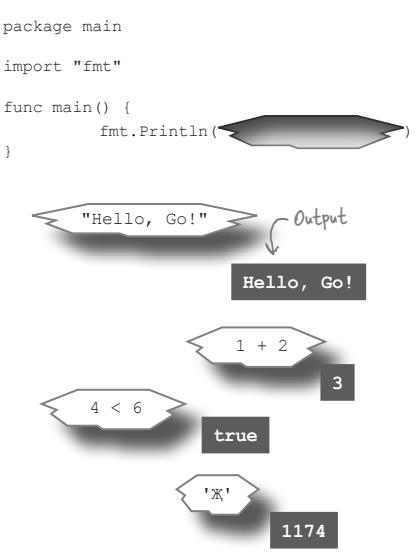
let's get going

1 Syntax Basics

Are you ready to turbo-charge your software? Do you want a **simple** programming language that **compiles fast**? That **runs fast**? That makes it **easy to distribute** your work to users? Then **you're ready for Go**!

Go is a programming language that focuses on **simplicity** and **speed**. It's simpler than other languages, so it's quicker to learn. And it lets you harness the power of today's multicore computer processors, so your programs run faster. This chapter will show you all the Go features that will make **your life as a developer easier**, and make your **users happier**.

Ready, set, Go!	2
The Go Playground	3
What does it all mean?	4
What if something goes wrong?	5
Calling functions	7
The Println function	7
Using functions from other packages	8
Function return values	9
A Go program template	11
Strings	11
Runes	12
Booleans	12
Numbers	13
Math operations and comparisons	13
Types	14
Declaring variables	16
Zero values	17
Short variable declarations	19
Naming rules	21
Conversions	22
Installing Go on your computer	25
Compiling Go code	26
Go tools	27
Try out code quickly with "go run"	27
Your Go Toolbox	28



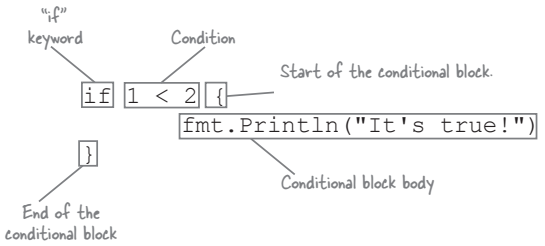
which code runs next?

2 Conditionals and Loops

Every program has parts that apply only in certain situations.

"This code should run *if* there's an error. Otherwise, that other code should run." Almost every program contains code that should be run only when a certain *condition* is true. So almost every programming language provides **conditional statements** that let you determine whether to run segments of code. Go is no exception.

You may also need some parts of your code to run *repeatedly*. Like most languages, Go provides **loops** that run sections of code more than once. We'll learn to use both conditionals and loops in this chapter!



Calling methods	32
Making the grade	34
Multiple return values from a function or method	36
Option 1: Ignore the error return value with the blank identifier	37
Option 2: Handle the error	38
Conditionals	39
Logging a fatal error, conditionally	42
Avoid shadowing names	44
Converting strings to numbers	46
Blocks	49
Blocks and variable scope	50
We've finished the grading program!	52
Only one variable in a short variable declaration has to be new	54
Let's build a game	55
Package names vs. import paths	56
Generating a random number	57
Getting an integer from the keyboard	59
Comparing the guess to the target	60
Loops	61
Init and post statements are optional	63
Using a loop in our guessing game	66
Breaking out of our guessing loop	69
Revealing the target	70
Congratulations, your game is complete!	72
Your Go Toolbox	74

call me

3 Functions

You've been missing out. You've been calling functions like a pro. But the only functions you could call were the ones Go defined for you. Now, it's your turn. We're going to show you how to create your own functions. We'll learn how to declare functions with and without parameters. We'll declare functions that return a single value, and we'll learn how to return multiple values so that we can indicate when there's been an error. And we'll learn about **pointers**, which allow us to make more memory-efficient function calls.



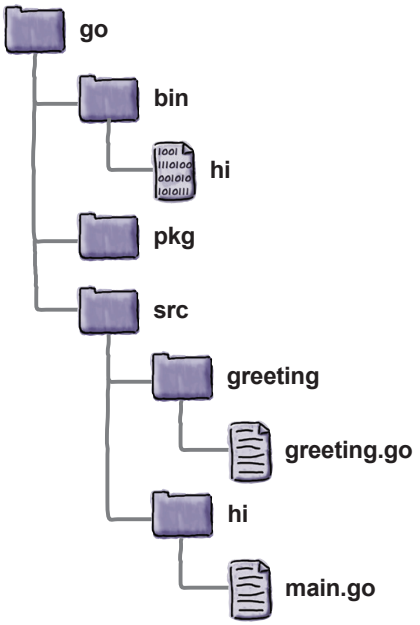
Some repetitive code	80
Formatting output with Printf and Sprintf	81
Formatting verbs	82
Formatting value widths	83
Formatting fractional number widths	84
Using Printf in our paint calculator	85
Declaring functions	86
Declaring function parameters	87
Using functions in our paint calculator	88
Functions and variable scope	90
Function return values	91
Using a return value in our paint calculator	93
The paintNeeded function needs error handling	95
Error values	96
Declaring multiple return values	97
Using multiple return values with our paintNeeded function	98
Always handle errors!	99
Function parameters receive copies of the arguments	102
Pointers	103
Pointer types	104
Getting or changing the value at a pointer	105
Using pointers with functions	107
Fixing our "double" function using pointers	108
Your Go Toolbox	110

bundles of code

4 Packages

It's time to get organized. So far, we've been throwing all our code together in a single file. As our programs grow bigger and more complex, that's going to quickly become a mess.

In this chapter, we'll show you how to create your own **packages** to help keep related code together in one place. But packages are good for more than just organization. Packages are an easy way to *share code between your programs*. And they're an easy way to *share code with other developers*.



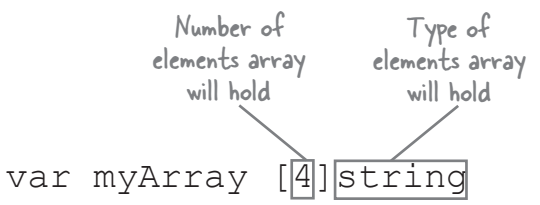
Different programs, same function	114
Sharing code between programs using packages	116
The Go workspace directory holds package code	117
Creating a new package	118
Importing our package into a program	119
Packages use the same file layout	120
Package naming conventions	123
Package qualifiers	123
Moving our shared code to a package	124
Constants	126
Nested package directories and import paths	128
Installing program executables with "go install"	130
Changing workspaces with the GOPATH environment variable	131
Setting GOPATH	132
Publishing packages	133
Downloading and installing packages with "go get"	137
Reading package documentation with "go doc"	139
Documenting your packages with doc comments	141
Viewing documentation in a web browser	143
Serving HTML documentation to yourself with "godoc"	144
The "godoc" server includes YOUR packages!	145
Your Go Toolbox	146

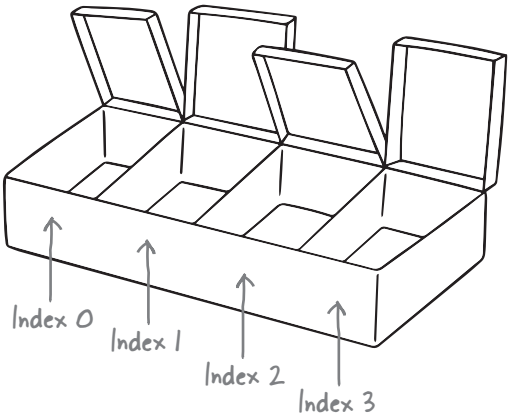
on the list

5 Arrays

A whole lot of programs deal with lists of things. Lists of addresses. Lists of phone numbers. Lists of products. Go has *two* built-in ways of storing lists. This chapter will introduce the first: **arrays**. You'll learn about how to create arrays, how to fill them with data, and how to get that data back out again. Then you'll learn about processing all the elements in array, first the *hard* way with `for` loops, and then the *easy* way with `for...range` loops.

Arrays hold collections of values	150
Zero values in arrays	152
Array literals	153
Functions in the "fmt" package know how to handle arrays	154
Accessing array elements within a loop	155
Checking array length with the "len" function	156
Looping over arrays safely with "for...range"	157
Using the blank identifier with "for...range" loops	158
Getting the sum of the numbers in an array	159
Getting the average of the numbers in an array	161
Reading a text file	163
Reading a text file into an array	166
Updating our "average" program to read a text file	168
Our program can only process three values!	170
Your Go Toolbox	172



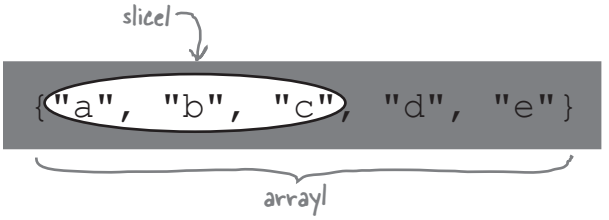


appending issue

6 Slices

We've learned we can't add more elements to an array. That's a real problem for our program, because we don't know in advance how many pieces of data our file contains. But that's where Go **slices** come in. Slices are a collection type that can grow to hold additional items—just the thing to fix our current program! We'll also see how slices give users an easier way to provide data to *all* your programs, and how they can help you write functions that are more convenient to call.





Slices	176
Slice literals	177
The slice operator	180
Underlying arrays	182
Change the underlying array, change the slice	183
Add onto a slice with the "append" function	184
Slices and zero values	186
Reading additional file lines using slices and "append"	187
Trying our improved program	189
Returning a nil slice in the event of an error	190
Command-line arguments	191
Getting command-line arguments from the os.Args slice	192
The slice operator can be used on other slices	193
Updating our program to use command-line arguments	194
Variadic functions	195
Using variadic functions	197
Using a variadic function to calculate averages	198
Passing slices to variadic functions	199
Slices have saved the day!	201
Your Go Toolbox	202

labeling data

7 Maps

Throwing things in piles is fine, until you need to find something again. You've already seen how to create lists of values using *arrays* and *slices*. You've seen how to apply the same operation to *every value* in an array or slice. But what if you need to work with a *particular* value? To find it, you'll have to start at the beginning of the array or slice, and *look through Every. Single. Value.*

What if there were a kind of collection where every value had a label on it? You could quickly find just the value you needed! In this chapter, we'll look at **maps**, which do just that.



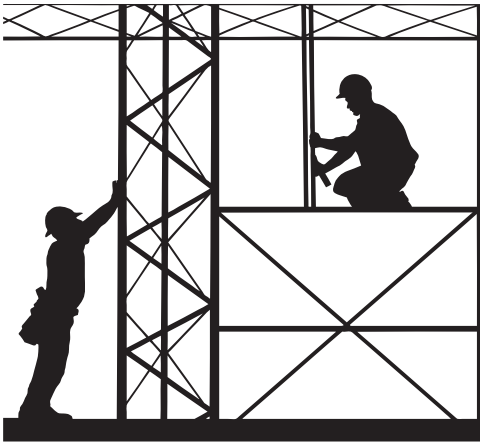
Counting votes	206
Reading names from a file	207
Counting names the hard way, with slices	209
Maps	212
Map literals	214
Zero values within maps	215
The zero value for a map variable is nil	215
How to tell zero values apart from assigned values	216
Removing key/value pairs with the "delete" function	218
Updating our vote counting program to use maps	219
Using for ... range loops with maps	221
The for ... range loop handles maps in random order!	223
Updating our vote counting program with a for ... range loop	224
The vote counting program is complete!	225
Your Go Toolbox	227

building storage

8 Structs

Sometimes you need to store more than one type of data.

We learned about slices, which store a list of values. Then we learned about maps, which map a list of keys to a list of values. But both of these data structures can only hold values of *one* type. Sometimes, you need to group together values of *several* types. Think of billing receipts, where you have to mix item names (strings) with quantities (integers). Or student records, where you have to mix student names (strings) with grade point averages (floating-point numbers). You can't mix value types in slices or maps. But you *can* if you use another type called a **struct**. We'll learn all about structs in this chapter!



Slices and maps hold values of ONE type	232
Structs are built out of values of MANY types	233
Access struct fields using the dot operator	234
Storing subscriber data in a struct	235
Defined types and structs	236
Using a defined type for magazine subscribers	238
Using defined types with functions	239
Modifying a struct using a function	242
Accessing struct fields through a pointer	244
Pass large structs using pointers	246
Moving our struct type to a different package	248
A defined type's name must be capitalized to be exported	249
Struct field names must be capitalized to be exported	250
Struct literals	251
Creating an Employee struct type	253
Creating an Address struct type	254
Adding a struct as a field on another type	255
Setting up a struct within another struct	255
Anonymous struct fields	258
Embedding structs	259
Our defined types are complete!	260
Your Go Toolbox	261

you're my type

9 Defined Types

There's more to learn about defined types. In the previous chapter, we showed you how to define a type with a struct underlying type. What we *didn't* show you was that you can use *any* type as an underlying type.

And do you remember methods—the special kind of function that's associated with values of a particular type? We've been calling methods on various values throughout the book, but we haven't shown you how to define your *own* methods. In this chapter, we're going to fix all of that. Let's get started!

Type errors in real life	266
Defined types with underlying basic types	267
Defined types and operators	269
Converting between types using functions	271
Fixing our function name conflict using methods	274
Defining methods	275
The receiver parameter is (pretty much) just another parameter	276
A method is (pretty much) just like a function	277
Pointer receiver parameters	279
Converting Liters and Milliliters to Gallons using methods	283
Converting Gallons to Liters and Milliliters using methods	284
Your Go Toolbox	285



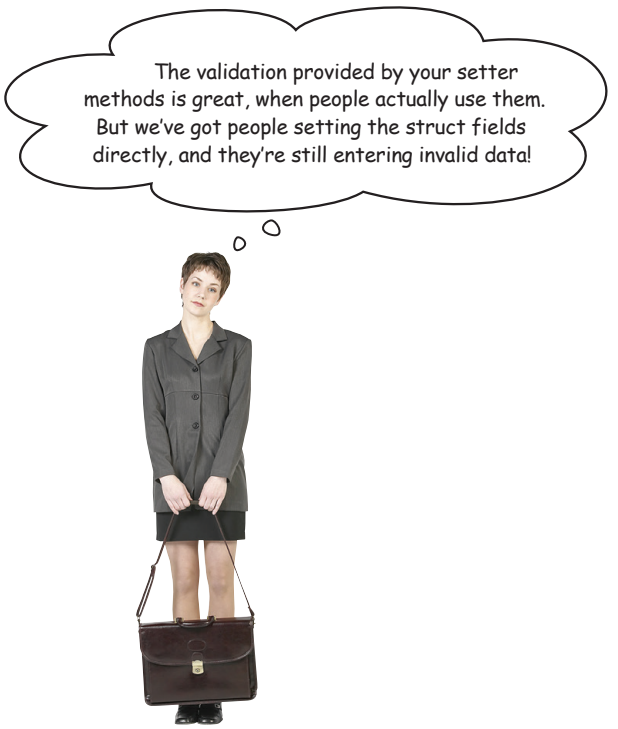
keep it to yourself

10 Encapsulation and Embedding

Mistakes happen. Sometimes, your program will receive invalid data from user input, a file you're reading in, or elsewhere. In this chapter, you'll learn about **encapsulation**: a way to protect your struct type's fields from that invalid data. That way, you'll know your field data is safe to work with!

We'll also show you how to **embed** other types within your struct type. If your struct type needs methods that already exist on another type, you don't have to copy and paste the method code. You can embed the other type within your struct type, and then use the embedded type's methods just as if they were defined on your own type!

Creating a Date struct type 290
People are setting the Date struct field to invalid values! 291
Setter methods 292
Setter methods need pointer receivers 293
Adding the remaining setter methods 294
Adding validation to the setter methods 296
The fields can still be set to invalid values! 298
Moving the Date type to another package 299
Making Date fields unexported 301
Accessing unexported fields through exported methods 302
Getter methods 304
Encapsulation 305
Embedding the Date type in an Event type 308
Unexported fields don't get promoted 309
Exported methods get promoted just like fields 310
Encapsulating the Event Title field 312
Promoted methods live alongside the outer type's methods 313
Our calendar package is complete! 314
Your Go Toolbox 316



what can you do?

11 Interfaces

Sometimes you don't care about the particular type of a value. You don't care about what it *is*. You just need to know that it will be able to *do* certain things. That you'll be able to call *certain methods* on it. You don't care whether you have a `Pen` or a `Pencil`, you just need something with a `Draw` method. You don't care whether you have a `Car` or a `Boat`, you just need something with a `Steer` method.

That's what Go **interfaces** accomplish. They let you define variables and function parameters that will hold *any* type, as long as that type defines certain methods.

Two different types that have the same methods 322
A method parameter that can only accept one type 323
Interfaces 325
Defining a type that satisfies an interface 326
Concrete types, interface types 327
Assign any type that satisfies the interface 328
You can only call methods defined as part of the interface 329
Fixing our playList function using an interface 331
Type assertions 334
Type assertion failures 336
Avoiding panics when type assertions fail 337
Testing TapePlayers and TapeRecorders using type assertions 338
The "error" interface 340
The Stringer interface 342
The empty interface 344
Your Go Toolbox 347



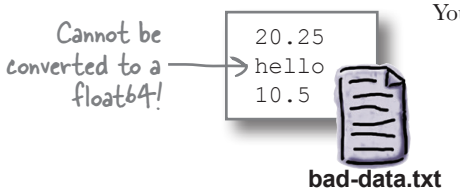
back on your feet

12 Recovering from Failure

Every program encounters errors. You should plan for them.

Sometimes handling an error can be as simple as reporting it and exiting the program. But other errors may require additional action. You may need to close opened files or network connections, or otherwise clean up, so your program doesn't leave a mess behind. In this chapter, we'll show you how to **defer** cleanup actions so they happen even when there's an error. We'll also show you how to make your program **panic** in those (rare) situations where it's appropriate, and how to **recover** afterward.

Reading numbers from a file, revisited	350
Any errors will prevent the file from being closed!	352
Deferring function calls	353
Recovering from errors using deferred function calls	354
Ensuring files get closed using deferred function calls	355
Listing the files in a directory	358
Listing the files in subdirectories (will be trickier)	359
Recursive function calls	360
Recursively listing directory contents	362
Error handling in a recursive function	364
Starting a panic	365
Stack traces	366
Deferred calls completed before crash	366
Using "panic" with scanDirectory	367
When to panic	368
The "recover" function	370
The panic value is returned from recover	371
Recovering from panics in scanDirectory	373
Reinstating a panic	374
Your Go Toolbox	376



sharing work

13 Goroutines and Channels

Working on one thing at a time isn't always the fastest way to finish a task. Some big problems can be broken into smaller tasks. **Goroutines** let your program work on several different tasks at once. Your goroutines can coordinate their work using **channels**, which let them send data to each other *and* synchronize so that one goroutine doesn't get ahead of another. Goroutines let you take full advantage of computers with multiple processors, so that your programs run as fast as possible!

Retrieving web pages	380
Multitasking	382
Concurrency using goroutines	383
Using goroutines	384
Using goroutines with our responseSize function	386
We don't directly control when goroutines run	388
Go statements can't be used with return values	389
Sending and receiving values with channels	391
Synchronizing goroutines with channels	392
Observing goroutine synchronization	393
Fixing our web page size program with channels	396
Updating our channel to carry a struct	398
Your Go Toolbox	399



code quality assurance

14 Automated Testing

Are you sure your software is working right now? Really sure? Before you sent that new version to your users, you presumably tried out the new features to ensure they all worked. But did you try the *old* features to ensure you didn't break any of them? *All* the old features? If that question makes you worry, your program needs **automated testing**. Automated tests ensure your program's components work correctly, even after you change your code. Go's `testing` package and `go test` tool make it easy to write automated tests, using the skills that you've already learned!

Automated tests find your bugs before someone else does	402
A function we should have had automated tests for	403
We've introduced a bug!	405
Writing tests	406
Running tests with the "go test" command	407
Testing our actual return values	408
More detailed test failure messages with the "Errorf" method	410
Test "helper" functions	411
Getting the tests to pass	412
Test-driven development	413
Another bug to fix	414
Running specific sets of tests	417
Table-driven tests	418
Fixing panicking code using a test	420
Your Go Toolbox	422



Pass.

For `[]slice{"apple", "orange", "pear"}`, `JoinWithCommas` should return `"apple, orange, and pear"`.

Fail!

For `[]slice{"apple", "orange"}`, `JoinWithCommas` should return `"apple and orange"`.

15

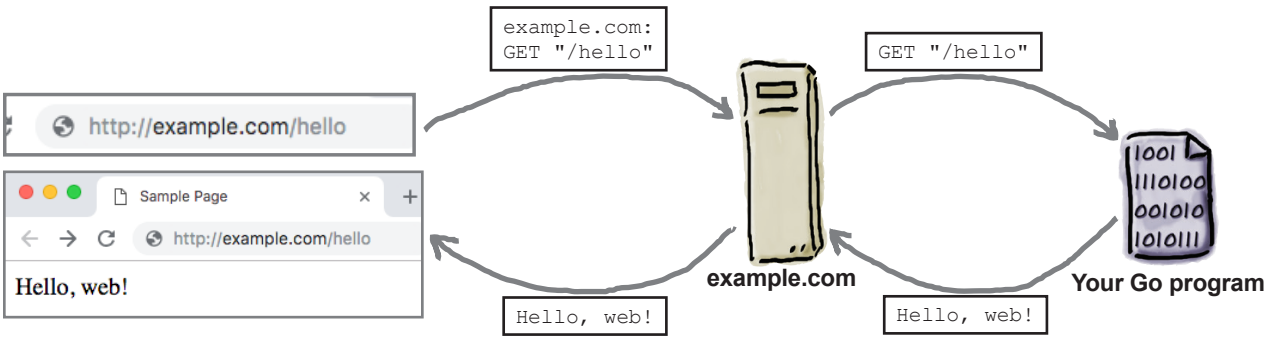
responding to requests

Web Apps

This is the 21st century. Users want web apps. Go's got you covered there, too! The Go standard library includes packages to help you host your own web applications and make them accessible from any web browser. So we're going to spend the final two chapters of the book showing you how to build web apps.

The first thing your web app needs is the ability to respond when a browser sends it a request. In this chapter, we'll learn to use the `net/http` package to do just that.

Writing web apps in Go	426
Browsers, requests, servers, and responses	427
A simple web app	428
Your computer is talking to itself	429
Our simple web app, explained	430
Resource paths	432
Responding differently for different resource paths	433
First-class functions	435
Passing functions to other functions	436
Functions as types	436
What's next	440
Your Go Toolbox	441



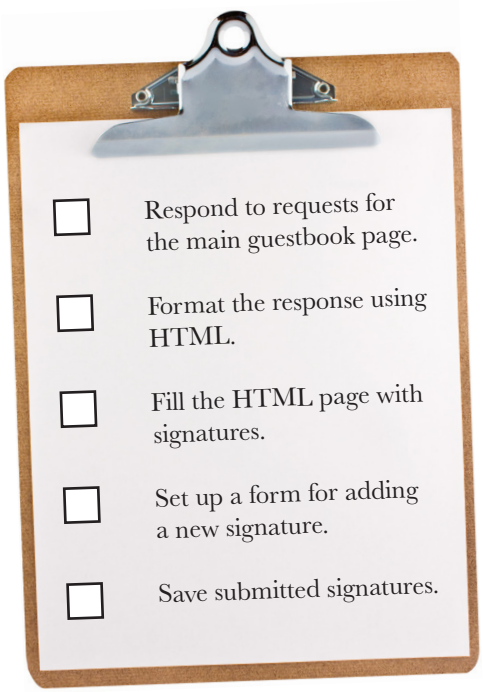
a pattern to follow

16 HTML Templates

Your web app needs to respond with HTML, not plain text.

Plain text is fine for emails and social media posts. But your pages need to be formatted. They need headings and paragraphs. They need forms where your users can submit data to your app. To do any of that, you need HTML code.

And eventually, you'll need to insert data into that HTML code. That's why Go offers the `html/template` package, a powerful way to include data in your app's HTML responses. Templates are key to building bigger, better web apps, and in this final chapter, we'll show you how to use them!



A guestbook app	446
Functions to handle a request and check errors	447
Setting up a project directory and trying the app	448
Making a signature list in HTML	449
Making our app respond with HTML	450
The "text/template" package	451
Using the io.Writer interface with a template's Execute method	452
ResponseWriters and os.Stdout both satisfy io.Writer	453
Inserting data into templates using actions	454
Making parts of a template optional with "if" actions	455
Repeating parts of a template with "range" actions	456
Inserting struct fields into a template with actions	457
Reading a slice of signatures in from a file	458
A struct to hold the signatures and signature count	460
Updating our template to include our signatures	461
Letting users add data with HTML forms	464
Form submission requests	466
Path and HTTP method for form submissions	467
Getting values of form fields from the request	468
Saving the form data	470
HTTP redirects	472
Our complete app code	474
Your Go Toolbox	477

understanding os.openfile

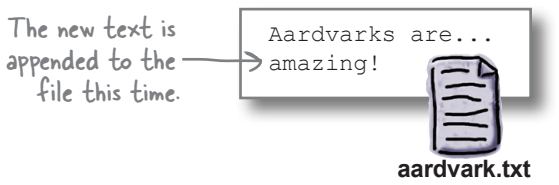
Appendix A: Opening Files

Some programs need to write data to files, not just read data.

Throughout the book, when we've wanted to work with files, you had to create them in your text editor for your programs to read. But some programs *generate* data, and when they do, they need to be able to *write* data to a file.

We used the `os.OpenFile` function to open a file for writing earlier in the book. But we didn't have space then to fully explore how it worked. In this appendix, we'll show you everything you need to know in order to use `os.OpenFile` effectively!

Understanding os.OpenFile	482
Passing flag constants to os.OpenFile	483
Binary notation	485
Bitwise operators	485
The bitwise AND operator	486
The bitwise OR operator	487
Using bitwise OR on the "os" package constants	488
Using bitwise OR to fix our os.OpenFile options	489
Unix-style file permissions	490
Representing permissions with the os.FileMode type	491
Octal notation	492
Converting octal values to FileMode values	493
Calls to os.OpenFile, explained	494

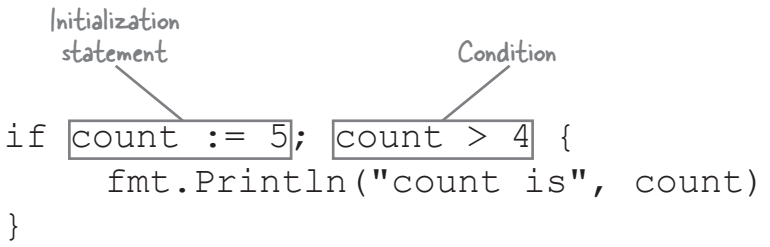


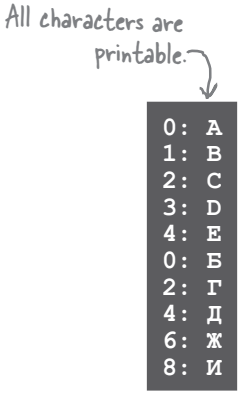
six things we didn't cover

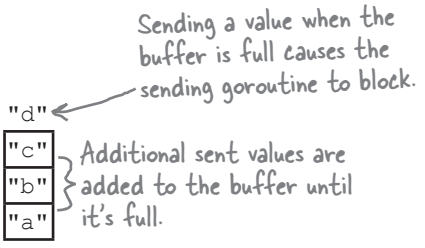
Appendix B: Leftovers

We've covered a lot of ground, and you're almost finished with this book. We'll miss you, but before we let you go, we wouldn't feel right about sending you out into the world without a *little* more preparation. We've saved six important topics for this appendix.

#1 Initialization statements for "if"	496
#2 The switch statement	498
#3 More basic types	499
#4 More about runes	499
#5 Buffered channels	503
#6 Further reading	506

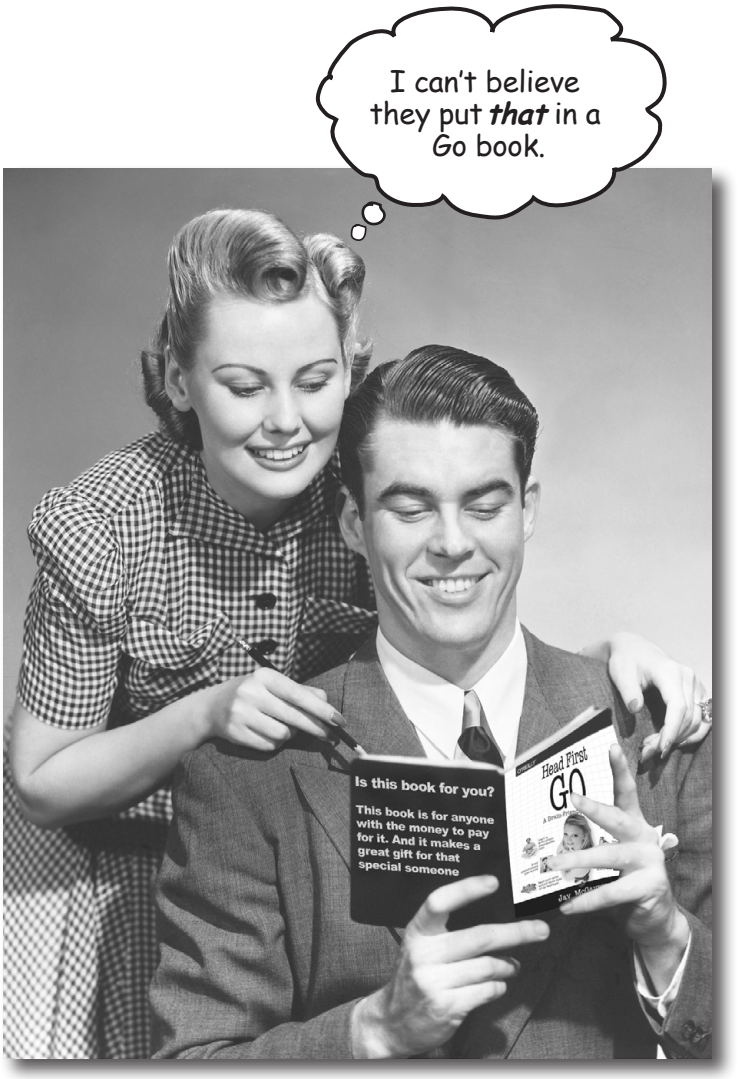






how to use this book

Intro



In this section, we answer the burning question:
"So why DID they put that in a book on Go?"

Who is this book for?

If you can answer "yes" to ***all*** of these:

1. Do you have access to a computer with a text editor?

2. Do you want to learn a programming language that makes development **fast** and **productive**?

3. Do you prefer **stimulating dinner-party conversation** to **dry, dull, academic lectures**?

this book is for you.

Who should probably back away from this book?

If you can answer "yes" to any ***one*** of these:

1. **Are you completely new to computers?**

   (You don't need to be advanced, but you should understand folders and files, how to open a terminal app, and how to use a simple text editor.)

2. Are you a ninja rockstar developer looking for a ***reference*** **book**?

3. Are you **afraid to try something new**? Would you rather have a root canal than mix stripes with plaid? Do you believe that a technical book can't be serious if it's full of bad puns?

this book is *not* for you.



[Note from Marketing: this book is for anyone with a valid credit card.]

We know what you're thinking

"How can *this* be a serious book on developing in Go?"

"What's with all the graphics?"

"Can I actually *learn* it this way?"

We know what your *brain* is thinking

Your brain craves novelty. It's always searching, scanning, *waiting* for something unusual. It was built that way, and it helps you stay alive.

So what does your brain do with all the routine, ordinary, normal things you encounter? Everything it *can* to stop them from interfering with the brain's *real* job—recording things that *matter*. It doesn't bother saving the boring things; they never make it past the "this is obviously not important" filter.

How does your brain *know* what's important? Suppose you're out for a day hike and a tiger jumps in front of you—what happens inside your head and body?

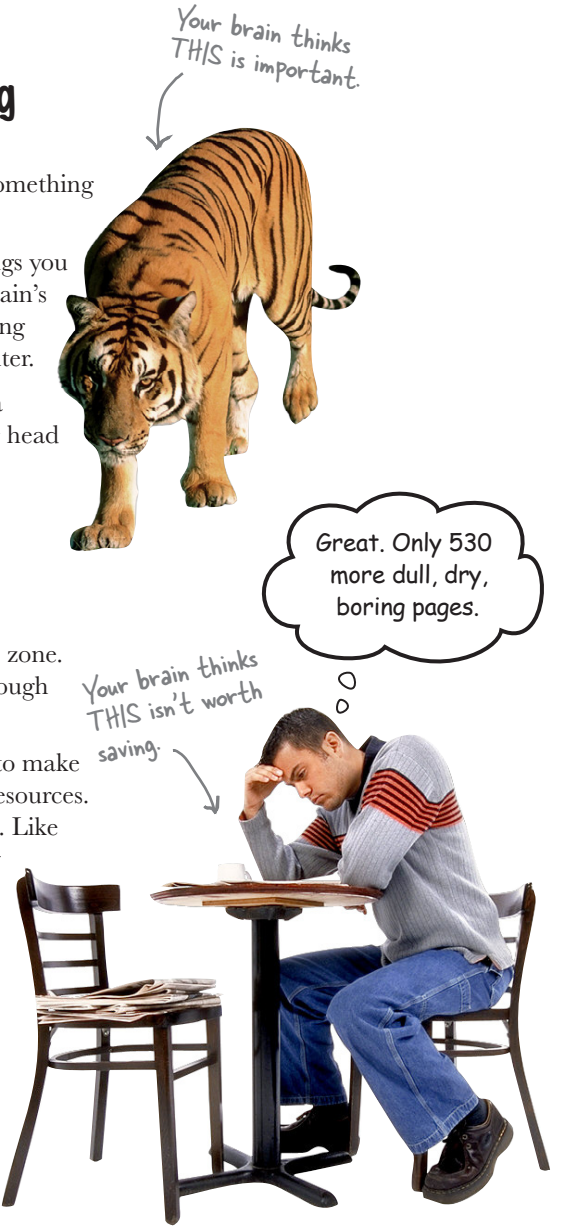
Neurons fire. Emotions crank up. *Chemicals surge.*

And that's how your brain knows…

This must be important! Don't forget it!

But imagine you're at home or in a library. It's a safe, warm, tiger-free zone. You're studying. Getting ready for an exam. Or trying to learn some tough technical topic your boss thinks will take a week, 10 days at the most.

Just one problem. Your brain's trying to do you a big favor. It's trying to make sure that this *obviously* unimportant content doesn't clutter up scarce resources. Resources that are better spent storing the really *big* things. Like tigers. Like the danger of fire. Like how you should never have posted those party photos on your Facebook page. And there's no simple way to tell your brain, "Hey, brain, thank you very much, but no matter how dull this book is, no matter how little I'm registering on the emotional Richter scale right now, I really *do* want you to keep this stuff around."



We think of a "Head First" reader as a learner.

So what does it take to *learn* something? First, you have to *get* it, then make sure you don't *forget* it. It's not about pushing facts into your head. Based on the latest research in cognitive science, neurobiology, and educational psychology, *learning* takes a lot more than text on a page. We know what turns your brain on.

Some of the Head First learning principles:

Make it visual. Images are far more memorable than words alone, and make learning much more effective (up to 89% improvement in recall and transfer studies). They also make things more understandable. **Put the words within or near the graphics** they relate to, rather than on the bottom or on another page, and learners will be up to *twice* as likely to solve problems related to the content.

Use a conversational and personalized style. In recent studies, students performed up to 40% better on post-learning tests if the content spoke directly to the reader, using a first-person, conversational style rather than taking a formal tone. Tell stories instead of lecturing. Use casual language. Don't take yourself too seriously. Which would *you* pay more attention to: a stimulating dinner-party companion, or a lecture?

Get the learner to think more deeply. In other words, unless you actively flex your neurons, nothing much happens in your head. A reader has to be motivated, engaged, curious, and inspired to solve problems, draw conclusions, and generate new knowledge. And for that, you need challenges, exercises, and thought-provoking questions, and activities that involve both sides of the brain and multiple senses.

Get—and keep—the reader's attention. We've all had the "I really want to learn this, but I can't stay awake past page one" experience. Your brain pays attention to things that are out of the ordinary, interesting, strange, eye-catching, unexpected. Learning a new, tough, technical topic doesn't have to be boring. Your brain will learn much more quickly if it's not.

Touch their emotions. We now know that your ability to remember something is largely dependent on its emotional content. You remember what you care about. You remember when you *feel* something. No, we're not talking heart-wrenching stories about a boy and his dog. We're talking emotions like surprise, curiosity, fun, "what the…?", and the feeling of "I rule!" that comes when you solve a puzzle, learn something everybody else thinks is hard, or realize you know something that "I'm more technical than thou" Bob from Engineering *doesn't*.

Metacognition: thinking about thinking



If you really want to learn, and you want to learn more quickly and more deeply, pay attention to how you pay attention. Think about how you think. Learn how you learn.

Most of us did not take courses on metacognition or learning theory when we were growing up. We were *expected* to learn, but rarely *taught* to learn.

But we assume that if you're holding this book, you really want to learn how to write Go programs. And you probably don't want to spend a lot of time. If you want to use what you read in this book, you need to *remember* what you read. And for that, you've got to *understand* it. To get the most from this book, or *any* book or learning experience, take responsibility for your brain. Your brain on *this* content.

The trick is to get your brain to see the new material you're learning as Really Important. Crucial to your well-being. As important as a tiger. Otherwise, you're in for a constant battle, with your brain doing its best to keep the new content from sticking.

So just how *DO* you get your brain to treat programming like it's a hungry tiger?

There's the slow, tedious way, or the faster, more effective way. The slow way is about sheer repetition. You obviously know that you *are* able to learn and remember even the dullest of topics if you keep pounding the same thing into your brain. With enough repetition, your brain says, "This doesn't *feel* important to him, but he keeps looking at the same thing *over* and *over* and *over*, so I suppose it must be."

The faster way is to do ***anything that increases brain activity,*** especially different *types* of brain activity. The things on the previous page are a big part of the solution, and they're all things that have been proven to help your brain work in your favor. For example, studies show that putting words *within* the pictures they describe (as opposed to somewhere else in the page, like a caption or in the body text) causes your brain to try to make sense of how the words and picture relate, and this causes more neurons to fire. More neurons firing = more chances for your brain to *get* that this is something worth paying attention to, and possibly recording.

A conversational style helps because people tend to pay more attention when they perceive that they're in a conversation, since they're expected to follow along and hold up their end. The amazing thing is, your brain doesn't necessarily *care* that the "conversation" is between you and a book! On the other hand, if the writing style is formal and dry, your brain perceives it the same way you experience being lectured to while sitting in a roomful of passive attendees. No need to stay awake.

But pictures and conversational style are just the beginning…

Here's what WE did

We used ***pictures***, because your brain is tuned for visuals, not text. As far as your brain's concerned, a picture really *is* worth a thousand words. And when text and pictures work together, we embedded the text *in* the pictures because your brain works more effectively when the text is *within* the thing it refers to, as opposed to in a caption or buried in the body text somewhere.

We used ***redundancy***, saying the same thing in *different* ways and with different media types, and *multiple senses*, to increase the chance that the content gets coded into more than one area of your brain.

We used concepts and pictures in ***unexpected*** ways because your brain is tuned for novelty, and we used pictures and ideas with at least *some* ***emotional*** *content*, because your brain is tuned to pay attention to the biochemistry of emotions. That which causes you to *feel* something is more likely to be remembered, even if that feeling is nothing more than a little ***humor***, ***surprise***, or ***interest.***

We used a personalized, ***conversational style***, because your brain is tuned to pay more attention when it believes you're in a conversation than if it thinks you're passively listening to a presentation. Your brain does this even when you're *reading*.

We included ***activities***, because your brain is tuned to learn and remember more when you ***do*** things than when you *read* about things. And we made the exercises challenging-yet-doable, because that's what most people prefer.

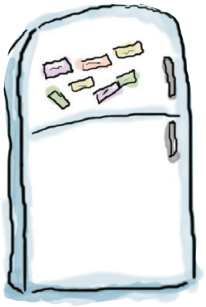
We used ***multiple learning styles***, because *you* might prefer step-by-step procedures, while someone else wants to understand the big picture first, and someone else just wants to see an example. But regardless of your own learning preference, *everyone* benefits from seeing the same content represented in multiple ways.

We include content for ***both sides of your brain***, because the more of your brain you engage, the more likely you are to learn and remember, and the longer you can stay focused. Since working one side of the brain often means giving the other side a chance to rest, you can be more productive at learning for a longer period of time.

And we included ***stories*** and exercises that present ***more than one point of view,*** because your brain is tuned to learn more deeply when it's forced to make evaluations and judgments.

We included ***challenges***, with exercises, and by asking ***questions*** that don't always have a straight answer, because your brain is tuned to learn and remember when it has to *work* at something. Think about it—you can't get your *body* in shape just by *watching* people at the gym. But we did our best to make sure that when you're working hard, it's on the *right* things. That ***you're not spending one extra dendrite*** processing a hard-to-understand example, or parsing difficult, jargon-laden, or overly terse text.

We used ***people***. In stories, examples, pictures, etc., because, well, *you're* a person. And your brain pays more attention to *people* than it does to *things*.



Cut this out and stick it on your refrigerator.

Here's what YOU can do to bend your brain into submission

So, we did our part. The rest is up to you. These tips are a starting point; listen to your brain and figure out what works for you and what doesn't. Try new things.

1. **Slow down. The more you understand, the less you have to memorize.**

   Don't just *read*. Stop and think. When the book asks you a question, don't just skip to the answer. Imagine that someone really *is* asking the question. The more deeply you force your brain to think, the better chance you have of learning and remembering.

2. **Do the exercises. Write your own notes.**

   We put them in, but if we did them for you, that would be like having someone else do your workouts for you. And don't just *look* at the exercises. **Use a pencil.** There's plenty of evidence that physical activity *while* learning can increase the learning.

3. **Read "There Are No Dumb Questions."**

   That means all of them. They're not optional sidebars, ***they're part of the core content!*** Don't skip them.

4. **Make this the last thing you read before bed. Or at least the last challenging thing.**

   Part of the learning (especially the transfer to long-term memory) happens *after* you put the book down. Your brain needs time on its own, to do more processing. If you put in something new during that processing time, some of what you just learned will be lost.

5. **Talk about it. Out loud.**

   Speaking activates a different part of the brain. If you're trying to understand something, or increase your chance of remembering it later, say it out loud. Better still, try to explain it out loud to someone else. You'll learn more quickly, and you might uncover ideas you hadn't known were there when you were reading about it.

6. **Drink water. Lots of it.**

   Your brain works best in a nice bath of fluid. Dehydration (which can happen before you ever feel thirsty) decreases cognitive function.

7. **Listen to your brain.**

   Pay attention to whether your brain is getting overloaded. If you find yourself starting to skim the surface or forget what you just read, it's time for a break. Once you go past a certain point, you won't learn faster by trying to shove more in, and you might even hurt the process.

8. **Feel something.**

   Your brain needs to know that this *matters*. Get involved with the stories. Make up your own captions for the photos. Groaning over a bad joke is *still* better than feeling nothing at all.

9. **Write a lot of code!**

   There's only one way to learn to develop Go programs: **write a lot of code**. And that's what you're going to do throughout this book. Coding is a skill, and the only way to get good at it is to practice. We're going to give you a lot of practice: every chapter has exercises that pose a problem for you to solve. Don't just skip over them—a lot of the learning happens when you solve the exercises. We included a solution to each exercise—don't be afraid to **peek at the solution** if you get stuck! (It's easy to get snagged on something small.) But try to solve the problem before you look at the solution. And definitely get it working before you move on to the next part of the book.

Read me

This is a learning experience, not a reference book. We deliberately stripped out everything that might get in the way of learning whatever it is we're working on at that point in the book. And the first time through, you need to begin at the beginning, because the book makes assumptions about what you've already seen and learned.

It helps if you've done a *little* programming in some other language.

Most developers discover Go *after* they've learned some other programming language. (They often come seeking refuge from that other language.) We touch on the basics enough that a complete beginner can get by, but we don't go into great detail on what a variable is, or how an `if` statement works. You'll have an easier time if you've done at least a *little* of this before.

We don't cover every type, function, and package ever created.

Go comes with a *lot* of software packages built in. Sure, they're all interesting, but we couldn't cover them all even if this book was *twice* as long. Our focus is on the core types and functions that *matter* to you, the beginner. We make sure you have a deep understanding of them, and confidence that you know how and when to use them. In any case, once you're done with *Head First Go*, you'll be able to pick up any reference book and get up to speed quickly on the packages we left out.

The activities are NOT optional.

The exercises and activities are not add-ons; they're part of the core content of the book. Some of them are to help with memory, some are for understanding, and some will help you apply what you've learned. ***Don't skip the exercises.***

The redundancy is intentional and important.

One distinct difference in a Head First book is that we want you to *really* get it. And we want you to finish the book remembering what you've learned. Most reference books don't have retention and recall as a goal, but this book is about *learning*, so you'll see some of the same concepts come up more than once.

The code examples are as lean as possible.

It's frustrating to wade through 200 lines of code looking for the two lines you need to understand. Most examples in this book are shown in the smallest possible context, so that the part you're trying to learn is clear and simple. So don't expect the code to be robust, or even complete. That's *your* assignment after you finish the book. The book examples are written specifically for *learning*, and aren't always fully functional.

We've placed all the example files on the web so you can download them. You'll find them at *http://headfirstgo.com/*.

Acknowledgments

Series founders:

Huge thanks to the Head First founders, **Kathy Sierra** and **Bert Bates**. I loved the series when I encountered it more than a decade ago, but never imagined I might be writing for it. Thank you for creating this amazing style of teaching!

At O'Reilly:

Thanks to everyone at O'Reilly who made this happen, particularly editor **Jeff Bleiel**, and to **Kristen Brown**, **Rachel Monaghan**, and the rest of the production team.

Technical reviewers:

Everyone makes mistakes, but luckily I have tech reviewers **Tim Heckman**, **Edward Yue Shung Wong**, and **Stefan Pochmann** to catch all of mine. You will never know how many problems they found, because I swiftly destroyed all the evidence. But their help and feedback were definitely necessary and are forever appreciated!

And more thanks:

Thanks to **Leo Richardson** for additional proofreading.

Perhaps most importantly, thanks to **Christine**, **Courtney**, **Bryan**, **Lenny**, and **Jeremy** for their patience and support (for two books now)!

O'Reilly Online Learning

For almost 40 years, O'Reilly Media has provided technology and business training, knowledge, and insight to help companies succeed.

Our unique network of experts and innovators share their knowledge and expertise through books, articles, conferences, and our online learning platform. O'Reilly's online learning platform gives you on-demand access to live training courses, in-depth learning paths, interactive coding environments, and a vast collection of text and video from O'Reilly and 200+ other publishers. For more information, please visit *http://oreilly.com*.

1 let's get going

Syntax Basics



Are you ready to turbo-charge your software? Do you want a **simple** programming language that **compiles fast**? That **runs fast**? That makes it **easy to distribute** your work to users? Then **you're ready for Go**!

Go is a programming language that focuses on **simplicity** and **speed**. It's simpler than other languages, so it's quicker to learn. And it lets you harness the power of today's multicore computer processors, so your programs run faster. This chapter will show you all the Go features that will make **your life as a developer easier**, and make your **users happier**.

Ready, set, Go!

Back in 2007, the search engine Google had a problem. They had to maintain programs with millions of lines of code. Before they could test new changes, they had to compile the code into a runnable form, a process which at the time took the better part of an hour. Needless to say, this was bad for developer productivity.

So Google engineers Robert Griesemer, Rob Pike, and Ken Thompson sketched out some goals for a new language:

- Fast compilation
- Less cumbersome code
- Unused memory freed automatically (garbage collection)
- Easy-to-write software that does several operations simultaneously (concurrency)
- Good support for processors with multiple cores

After a couple years of work, Google had created Go: a language that was fast to write code for and produced programs that were fast to compile and run. The project switched to an open source license in 2009. It's now free for anyone to use. And you should use it! Go is rapidly gaining popularity thanks to its simplicity and power.

If you're writing a command-line tool, Go can produce executable files for Windows, macOS, and Linux, all from the same source code. If you're writing a web server, it can help you handle many users connecting at once. And no matter *what* you're writing, it will help you ensure that your code is easier to maintain and add to.

Ready to learn more? Let's Go!



The Go Playground



The easiest way to try Go is to visit *https://play.golang.org* in your web browser. There, the Go team has set up a simple editor where you can enter Go code and run it on their servers. The result is displayed right there in your browser.

(Of course, this only works if you have a stable internet connection. If you don't, see page 25 to learn how to download and run the Go compiler directly on your computer. Then run the following examples using the compiler instead.)

Let's try it out now!



1. Open *https://play.golang.org* in your browser. (Don't worry if what you see doesn't quite match the screenshot; it just means they've improved the site since this book was printed!)

2. Delete any code that's in the editing area, and type this instead:

```
package main

import "fmt"

func main() {
      fmt.Println("Hello, Go!")
}
```

Don't worry, we'll explain what all this means on the next page!

3. Click the Format button, which will automatically reformat your code according to Go conventions.

4. Click the Run button.

You should see "Hello, Go!" displayed at the bottom of the screen. Congratulations, you've just run your first Go program!

Output → `Hello, Go!`

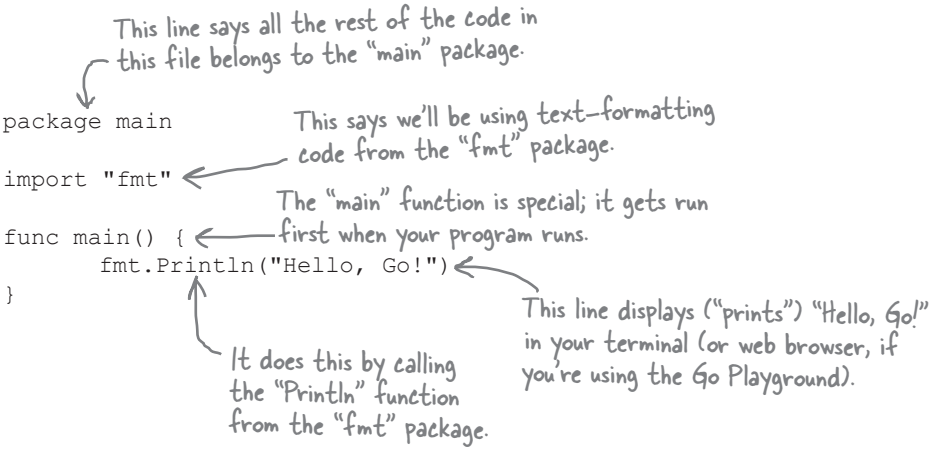
Turn the page, and we'll explain what we just did...

What does it all mean?

You've just run your first Go program! Now let's look at the code and figure out what it actually means...

Every Go file starts with a `package` clause. A **package** is a collection of code that all does similar things, like formatting strings or drawing images. The `package` clause gives the name of the package that this file's code will become a part of. In this case, we use the special package `main`, which is required if this code is going to be run directly (usually from the terminal).

Next, Go files almost always have one or more `import` statements. Each file needs to **import** other packages before its code can use the code those other packages contain. Loading all the Go code on your computer at once would result in a big, slow program, so instead you specify only the packages you need by importing them.



The last part of every Go file is the actual code, which is often split up into one or more functions. A **function** is a group of one or more lines of code that you can **call** (run) from other places in your program. When a Go program is run, it looks for a function named `main` and runs that first, which is why we named this function `main`.



Don't worry if you don't understand all this right now!

We'll look at everything in more detail in the next few pages.

The typical Go file layout

You'll quickly get used to seeing these three sections, in this order, in almost every Go file you work with:

1. The package clause
2. Any `import` statements
3. The actual code

The package clause
```
package main
```
The imports section
```
import "fmt"
```
The actual code
```
func main() {
        fmt.Println("Hello, Go!")
}
```

The saying goes, "a place for everything, and everything in its place." Go is a very *consistent* language. This is a good thing: you'll often find you just *know* where to look in your project for a given piece of code, without having to think about it!

there are no Dumb Questions

Q: My other programming language requires that each statement end with a semicolon. Doesn't Go?

A: You *can* use semicolons to separate statements in Go, but it's not required (in fact, it's generally frowned upon).

Q: What's this Format button? Why did we click that before running our code?

A: The Go compiler comes with a standard formatting tool, called `go fmt`. The Format button is the web version of `go fmt`.

Whenever you share your code, other Go developers will expect it to be in the standard Go format. That means that things like indentation and spacing will be formatted in a standard way, making it easier for everyone to read. Where other languages achieve this by relying on people manually reformatting their code to conform to a style guide, with Go all you have to do is run `go fmt`, and it will automatically fix everything for you.

We ran the formatter on every example we created for this book, and you should run it on all your code, too!

What if something goes wrong?

Go programs have to follow certain rules to avoid confusing the compiler. If we break one of these rules, we'll get an error message.

Suppose we forgot to add parentheses on our call to the `Println` function on line 6.

If we try to run this version of the program, we get an error:

```
Line 1 package main
   2
   3 import "fmt"
   4
   5 func main() {
   6         fmt.Println "Hello, Go!"
   7 }
```

Suppose we forgot the parentheses that used to be here...

```
prog.go:6:14: syntax error: unexpected literal "Hello, Go!" at end of statement
```

Name of file used by Go Playground

Line number where the error occurred

Description of the error

Character number within the line where the error occurred

Go tells us which source code file and line number we need to go to so we can fix the problem. (The Go Playground saves your code to a temporary file before running it, which is where the *prog.go* filename comes from.) Then it gives a description of the error. In this case, because we deleted the parentheses, Go can't tell we're trying to call the `Println` function, so it can't understand why we're putting `"Hello, Go"` at the end of line 6.



Breaking Stuff is Educational!

We can get a feel for the rules Go programs have to follow by intentionally breaking our program in various ways. Take this code sample, try making one of the changes below, and run it. Then undo your change and try the next one. See what happens!

```
package main

import "fmt"

func main() {
        fmt.Println("Hello, Go!")
}
```

Try breaking our code sample and see what happens!

If you do this...	...it will fail because...
Delete the package clause... ~~`package main`~~	Every Go file has to begin with a package clause.
Delete the `import` statement... ~~`import "fmt"`~~	Every Go file has to import every package it references.
Import a second (unused) package... `import "fmt"` `import "strings"`	Go files must import *only* the packages they reference. (This helps keep your code compiling fast!)
Rename the `main` function... `func` ~~`main`~~`hello`	Go looks for a function named `main` to run first.
Change the `Println` call to lowercase... `fmt.`~~`P`~~`println("Hello, Go!")`	Everything in Go is case-sensitive, so although `fmt.Println` is valid, there's no such thing as `fmt.println`.
Delete the package name before `Println`... ~~`fmt.`~~`Println("Hello, Go!")`	The `Println` function isn't part of the `main` package, so Go needs the package name before the function call.

Let's try the first one as an example...

Delete the package clause...

```
import "fmt"

func main() {
        fmt.Println("Hello, Go!")
}
```

You'll get an error!

```
can't load package: package main:
prog.go:1:1: expected 'package', found 'import'
```

Calling functions

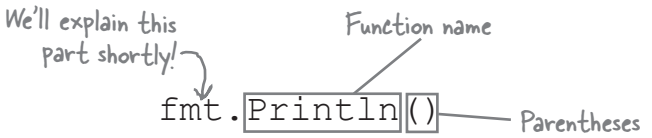
Our example includes a call to the `fmt` package's `Println` function. To call a function, type the function name (`Println` in this case), and a pair of parentheses.

```
package main

import "fmt"

func main() {
        fmt.Println("Hello, Go!")
}
```

A call to the Println function



Like many functions, `Println` can take one or more **arguments**: values you want the function to work with. The arguments appear in parentheses after the function name.

Inside the parentheses are one or more arguments, separated by commas.

```
fmt.Println("First argument", "Second argument")
```

Output →

```
First argument Second argument
```

`Println` can be called with no arguments, or you can provide several arguments. When we look at other functions later, however, you'll find that most require a specific number of arguments. If you provide too few or too many, you'll get an error message saying how many arguments were expected, and you'll need to fix your code.

The Println function

Use the `Println` function when you need to see what your program is doing. Any arguments you pass to it will be printed (displayed) in your terminal, with each argument separated by a space.

After printing all its arguments, `Println` will skip to a new terminal line. (That's why "ln" is at the end of its name.)

```
fmt.Println("First argument", "Second argument")
fmt.Println("Another line")
```

Output →

```
First argument Second argument
Another line
```

Using functions from other packages

The code in our first program is all part of the `main` package, but the `Println` function is in the `fmt` package. (The `fmt` stands for "format.") To be able to call `Println`, we first have to import the package containing it.

```
package main

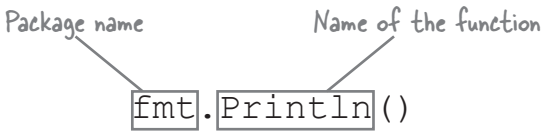
import "fmt"

func main() {
        fmt.Println("Hello, Go!")
}
```

We have to import the "fmt" package before we can access its Println function.

This specifies that we're calling a function that's part of the "fmt" package.

Once we've imported the package, we can access any functions it offers by typing the package name, a dot, and the name of the function we want.



Here's a code sample that calls functions from a couple other packages. Because we need to import multiple packages, we switch to an alternate format for the `import` statement that lets you list multiple packages within parentheses, one package name per line.

```
package main

import (
        "math"
        "strings"
)

func main() {
        math.Floor(2.75)
        strings.Title("head first go")
}
```

This alternate format for the "import" statement lets you import multiple packages at once.

Import the "math" package so we can use math.Floor.

Import the "strings" package so we can use strings.Title.

Call the Floor function from the "math" package.

Call the Title function from the "strings" package.

This program has no output. (We'll explain why in a moment!)

Once we've imported the `math` and `strings` packages, we can access the `math` package's `Floor` function with `math.Floor`, and the `strings` package's `Title` function with `strings.Title`.

You may have noticed that in spite of including those two function calls in our code, the above sample doesn't display any output. We'll look at how to fix that next.

Function return values

In our previous code sample, we tried calling the `math.Floor` and `strings.Title` functions, but they didn't produce any output:

```
package main

import (
       "math"
       "strings"
)

func main() {
       math.Floor(2.75)
       strings.Title("head first go")
}
```

This program produces no output!

When we call the `fmt.Println` function, we don't need to communicate with it any further after that. We pass one or more values for `Println` to print, and we trust that it printed them. But sometimes a program needs to be able to call a function and get data back from it. For this reason, functions in most programming languages can have **return values**: a value that the function computes and returns to its caller.

The `math.Floor` and `strings.Title` functions are both examples of functions that use return values. The `math.Floor` function takes a floating-point number, rounds it down to the nearest whole number, and returns that whole number. And the `strings.Title` function takes a string, capitalizes the first letter of each word it contains (converting it to "title case"), and returns the capitalized string.

To actually see the results of these function calls, we need to take their return values and pass those to `fmt.Println`:

```
package main

import (
       "fmt"
       "math"
       "strings"
)

func main() {
       fmt.Println(math.Floor(2.75))
       fmt.Println(strings.Title("head first go"))
}
```

Import the "fmt" package as well.

Call fmt.Println with the return value from math.Floor.

Takes a number, rounds it down, and returns that value

Output

```
2
Head First Go
```

Call fmt.Println with the return value from strings.Title.

Takes a string, and returns a new string with each word capitalized

Once this change is made, the return values get printed, and we can see the results.

Pool Puzzle



Your **job** is to take code snippets from the pool and place them into the blank lines in the code. **Don't** use the same snippet more than once, and you won't need to use all the snippets. Your **goal** is to make code that will run and produce the output shown.

```
package main          <- We've done the first one for you!

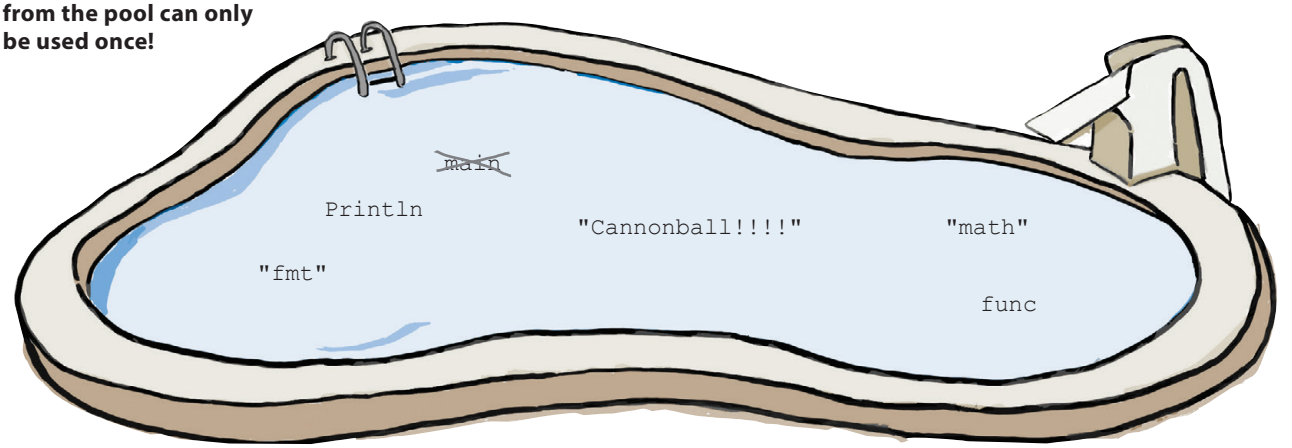
import (
        _____
)

____ main() {
        fmt.Println(_____________)
}
```

Output

```
Cannonball!!!!
```

Note: each snippet from the pool can only be used once!



Answers on page 29.

A Go program template

For the code snippets that follow, just imagine inserting them into this full Go program:

Better yet, try typing this program into the Go Playground, and then insert the snippets one at a time to see for yourself what they do!

```
package main

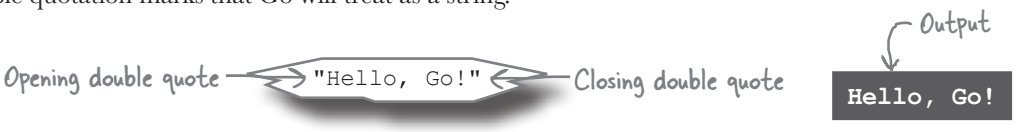
import "fmt"

func main() {
        fmt.Println(
}
```

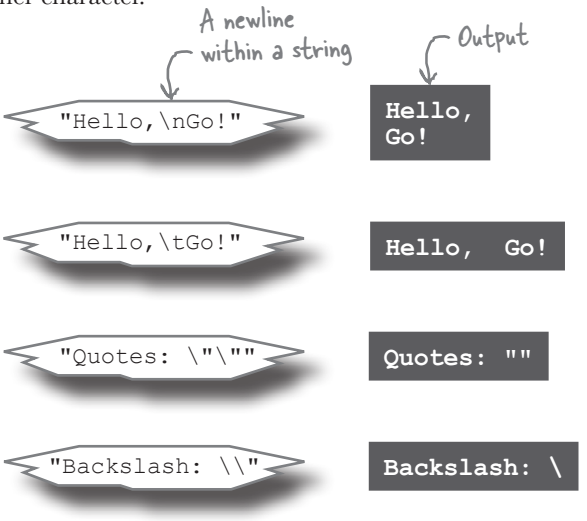
Insert your code here!

Strings

We've been passing **strings** as arguments to `Println`. A string is a series of bytes that usually represent text characters. You can define strings directly within your code using **string literals**: text between double quotation marks that Go will treat as a string.



Within strings, characters like newlines, tabs, and other characters that would be hard to include in program code can be represented with **escape sequences**: a backslash followed by characters that represent another character.



Escape sequence	Value
\n	A newline character.
\t	A tab character.
\"	Double quotation marks.
\\	A backslash.

Runes

Whereas strings are usually used to represent a whole series of text characters, Go's **runes** are used to represent single characters.

String literals are written surrounded by double quotation marks ("), but **rune literals** are written with single quotation marks (').

```
package main

import "fmt"

func main() {
        fmt.Println(          )
}
```

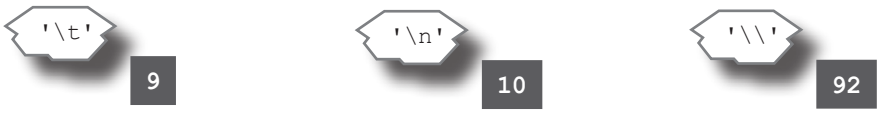
Here's our template again...

Insert your code here!

Go programs can use almost any character from almost any language on earth, because Go uses the Unicode standard for storing runes. Runes are kept as numeric codes, not the characters themselves, and if you pass a rune to `fmt.Println`, you'll see that numeric code in the output, not the original character.



Just as with string literals, escape sequences can be used in a rune literal to represent characters that would be hard to include in program code:



Booleans

Boolean values can be one of only two values: `true` or `false`. They're especially useful with conditional statements, which cause sections of code to run only if a condition is true or false. (We'll look at conditionals in the next chapter.)



Numbers

You can also define numbers directly within your code, and it's even simpler than string literals: just type the number.

```
package main

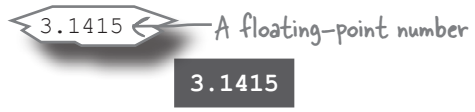
import "fmt"

func main() {
        fmt.Println(          )
}
```

Here's our template again...

Insert your code here!





As we'll see shortly, Go treats integer and floating-point numbers as different types, so remember that a decimal point can be used to distinguish an integer from a floating-point number.

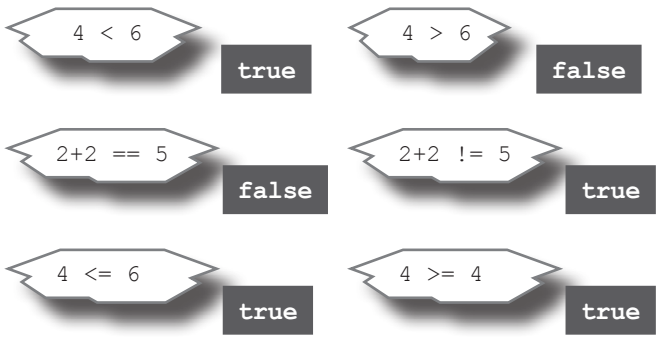
Math operations and comparisons

Go's basic math operators work just like they do in most other languages. The + symbol is for addition, - for subtraction, * for multiplication, and / for division.



You can use < and > to compare two values and see if one is less than or greater than another. You can use == (that's *two* equals signs) to see if two values are equal, and != (that's an exclamation point and an equals sign, read aloud as "not equal") to see if two values are not equal. <= tests whether the first value is less than *or* equal to the second, and >= tests whether the first value is greater than or equal to the second.

The result of a comparison is a Boolean value, either `true` or `false`.



Types

In a previous code sample, we saw the `math.Floor` function, which rounds a floating-point number down to the nearest whole number, and the `strings.Title` function, which converts a string to title case. It makes sense that you would pass a number as an argument to the `Floor` function, and a string as an argument to the `Title` function. But what would happen if you passed a string to `Floor` and a number to `Title`?

```
package main

import (
        "fmt"
        "math"
        "strings"
)

func main() {
        fmt.Println(math.Floor("head first go"))
        fmt.Println(strings.Title(2.75))
}
```

Normally takes a floating-point number!

Normally takes a string!

Errors

```
cannot use "head first go" (type string) as type float64 in argument to math.Floor
cannot use 2.75 (type float64) as type string in argument to strings.Title
```

Go prints two error messages, one for each function call, and the program doesn't even run!

Things in the world around you can often be classified into different types based on what they can be used for. You don't eat a car or truck for breakfast (because they're vehicles), and you don't drive an omelet or bowl of cereal to work (because they're breakfast foods).

Likewise, values in Go are all classified into different **types**, which specify what the values can be used for. Integers can be used in math operations, but strings can't. Strings can be capitalized, but numbers can't. And so on.

Go is **statically typed**, which means that it knows what the types of your values are even before your program runs. Functions expect their arguments to be of particular types, and their return values have types as well (which may or may not be the same as the argument types). If you accidentally use the wrong type of value in the wrong place, Go will give you an error message. This is a good thing: it lets you find out there's a problem before your users do!

Go is statically typed. If you use the wrong type of value in the wrong place, Go will let you know.

Types (continued)

You can view the type of any value by passing it to the `reflect` package's `TypeOf` function. Let's find out what the types are for some of the values we've already seen:

```
package main

import (
       "fmt"
       "reflect"
)

func main() {
       fmt.Println(reflect.TypeOf(42))
       fmt.Println(reflect.TypeOf(3.1415))
       fmt.Println(reflect.TypeOf(true))
       fmt.Println(reflect.TypeOf("Hello, Go!"))
}
```

Import the "reflect" package so we can use its TypeOf function.

Returns the type of its argument

Output

```
int
float64
bool
string
```

Here's what those types are used for:

Type	Description
`int`	An integer. Holds whole numbers.
`float64`	A floating-point number. Holds numbers with a fractional part. (The `64` in the type name is because 64 bits of data are used to hold the number. This means that `float64` values can be fairly, but not infinitely, precise before being rounded off.)
`bool`	A Boolean value. Can only be `true` or `false`.
`string`	A string. A series of data that usually represents text characters.



Exercise

Draw lines to match each code snippet below to a type.
Some types will have more than one snippet that matches with them.

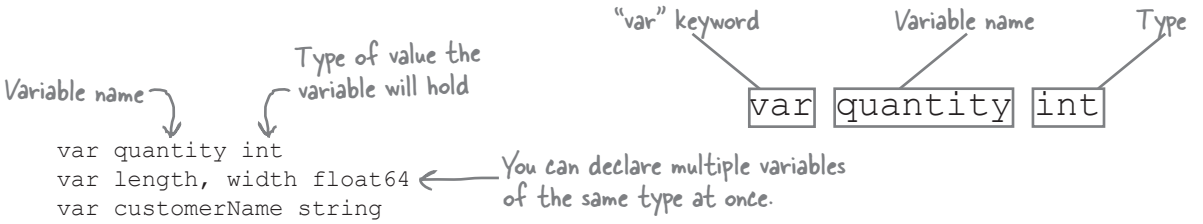
```
reflect.TypeOf(25)
reflect.TypeOf(true)
reflect.TypeOf(5.2)
reflect.TypeOf(1)
reflect.TypeOf(false)
reflect.TypeOf(1.0)
reflect.TypeOf("hello")
```

```
int
float64
bool
string
```

Answers on page 29.

Declaring variables

In Go, a **variable** is a piece of storage containing a value. You can give a variable a name by using a **variable declaration**. Just use the `var` keyword followed by the desired name and the type of values the variable will hold.



Once you declare a variable, you can assign any value of that type to it with = (that's a *single* equals sign):

```
quantity = 2
customerName = "Damon Cole"
```

You can assign values to multiple variables in the same statement. Just place multiple variable names on the left side of the =, and the same number of values on the right side, separated with commas.

```
length, width = 1.2, 2.4
```
← Assigning multiple variables at once.

Once you've assigned values to variables, you can use them in any context where you would use the original values:

```
package main

import "fmt"

func main() {
    var quantity int
    var length, width float64
    var customerName string

    quantity = 4
    length, width = 1.2, 2.4
    customerName = "Damon Cole"

    fmt.Println(customerName)
    fmt.Println("has ordered", quantity, "sheets")
    fmt.Println("each with an area of")
    fmt.Println(length*width, "square meters")
}
```

Declaring the variables

Assigning values to the variables

Using the variables

```
Damon Cole
has ordered 4 sheets
each with an area of
2.88 square meters
```

Declaring variables (continued)

If you know beforehand what a variable's value will be, you can declare variables and assign them values on the same line:

Just add an assignment onto the end.

Declaring variables AND assigning values

```
var quantity int = 4
var length, width float64 = 1.2, 2.4
var customerName string = "Damon Cole"
```

If you're declaring multiple variables, provide multiple values.

You can assign new values to existing variables, but they need to be values of the same type. Go's static typing ensures you don't accidentally assign the wrong kind of value to a variable.

Assigned types don't match the declared types!

```
quantity = "Damon Cole"
customerName = 4
```

Errors

```
cannot use "Damon Cole" (type string) as type int in assignment
cannot use 4 (type int) as type string in assignment
```

If you assign a value to a variable at the same time as you declare it, you can usually omit the variable type from the declaration. The type of the value assigned to the variable will be used as the type of that variable.

Omit variable types.

```
var quantity = 4
var length, width = 1.2, 2.4
var customerName = "Damon Cole"
fmt.Println(reflect.TypeOf(quantity))
fmt.Println(reflect.TypeOf(length))
fmt.Println(reflect.TypeOf(width))
fmt.Println(reflect.TypeOf(customerName))
```

```
int
float64
float64
string
```

Zero values

If you declare a variable without assigning it a value, that variable will contain the **zero value** for its type. For numeric types, the zero value is actually `0`:

```
var myInt int
var myFloat float64
fmt.Println(myInt, myFloat)
```

The zero value for "int" variables is 0.

```
0 0
```

The zero value for "float64" variables is 0.

But for other types, a value of `0` would be invalid, so the zero value for that type may be something else. The zero value for `string` variables is an empty string, for example, and the zero value for `bool` variables is `false`.

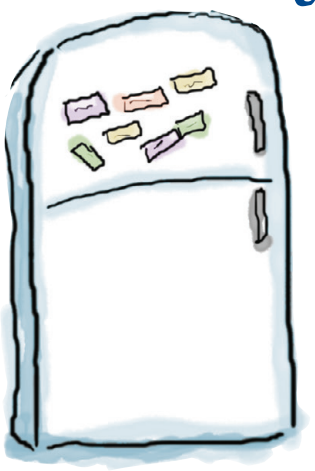
```
var myString string
var myBool bool
fmt.Println(myString, myBool)
```

The zero value for "string" variables is an empty string.

```
false
```

The zero value for "bool" variables is false.

Code Magnets



A Go program is all scrambled up on the fridge. Can you reconstruct the code snippets to make a working program that will produce the given output?

Output

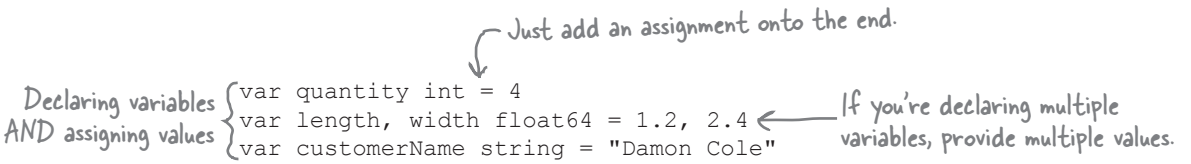
```
I started with 10 apples.
Some jerk ate 4 apples.
There are 6 apples left.
```

```
, "apples.")
, "apples.")
, "apples left.")
var
var
int
originalCount
func main() {
}
int
originalCount
fmt.Println("I started with",
eatenCount
fmt.Println("Some jerk ate",
=
=
fmt.Println("There are",
10
4
eatenCount
import (
    "fmt"
)
package main
originalCount-eatenCount
```

Answers on page 30.

Short variable declarations

We mentioned that you can declare variables and assign them values on the same line:



But if you know what the initial value of a variable is going to be as soon as you declare it, it's more typical to use a **short variable declaration**. Instead of explicitly declaring the type of the variable and later assigning to it with =, you do both at once using `:=`.

Let's update the previous example to use short variable declarations:

```
package main

import "fmt"

func main() {
    quantity := 4
    length, width := 1.2, 2.4
    customerName := "Damon Cole"

    fmt.Println(customerName)
    fmt.Println("has ordered", quantity, "sheets")
    fmt.Println("each with an area of")
    fmt.Println(length*width, "square meters")
}
```

Declaring variables AND assigning values

```
Damon Cole
has ordered 4 sheets
each with an area of
2.88 square meters
```

There's no need to explicitly declare the variable's type; the type of the value assigned to the variable becomes the type of that variable.

Because short variable declarations are so convenient and concise, they're used more often than regular declarations. You'll still see both forms occasionally, though, so it's important to be familiar with both.



Breaking Stuff is Educational!

Take our program that uses variables, try making one of the changes below, and run it. Then undo your change and try the next one. See what happens!

```
package main

import "fmt"

func main() {
        quantity := 4
        length, width := 1.2, 2.4
        customerName := "Damon Cole"

        fmt.Println(customerName)
        fmt.Println("has ordered", quantity, "sheets")
        fmt.Println("each with an area of")
        fmt.Println(length*width, "square meters")
}
```

```
Damon Cole
has ordered 4 sheets
each with an area of
2.88 square meters
```

If you do this...	...it will fail because...
Add a second declaration for the same variable `quantity := 4` `quantity := 4`	You can only declare a variable once. (Although you can assign new values to it as often as you want. You can also declare other variables with the same name, as long as they're in a different scope. We'll learn about scopes in the next chapter.)
Delete the `:` from a short variable declaration `quantity = 4`	If you forget the `:`, it's treated as an assignment, not a declaration, and you can't assign to a variable that hasn't been declared.

Assign a `string` to an `int` variable
`quantity := 4`
`quantity = "a"`
	Variables can only be assigned values of the same type.

Mismatch number of variables and values `length, width := 1.2`	You're required to provide a value for every variable you're assigning, and a variable for every value.
Remove code that uses a variable ~~`fmt.Println(customerName)`~~	All declared variables must be used in your program. If you remove the code that uses a variable, you must also remove the declaration.

Naming rules

Go has one simple set of rules that apply to the names of variables, functions, and types:

- A name must begin with a letter, and can have any number of additional letters and numbers.
- If the name of a variable, function, or type begins with a capital letter, it is considered **exported** and can be accessed from packages outside the current one. (This is why the `P` in `fmt.Println` is capitalized: so it can be used from the `main` package or any other.) If a variable/function/type name begins with a lowercase letter, it is considered **unexported** and can only be accessed within the current package.

OK: `length`, `stack2`, `sales.Total`

Illegal: `2stack` ← Can't start with a number!; `sales.total` ← Can't access anything in another package unless its name is capitalized!

Those are the only rules enforced by the language. But the Go community follows some additional conventions as well:

- If a name consists of multiple words, each word after the first should be capitalized, and they should be attached together without spaces between them, like this: `topPrice`, `RetryConnection`, and so on. (The first letter of the name should only be capitalized if you want to export it from the package.) This style is often called *camel case* because the capitalized letters look like the humps on a camel.
- When the meaning of a name is obvious from the context, the Go community's convention is to abbreviate it: to use `i` instead of `index`, `max` instead of `maximum`, and so on. (However, we at Head First believe that nothing is obvious when you're learning a new language, so we will *not* be following that convention in this book.)

OK: `sheetLength`, `TotalUnits`, `i`

Breaks conventions: `sheetlength` ← Subsequent words should be capitalized!; `Total_Units` ← This is legal, but words should be joined directly!; `index` ← Consider replacing with an abbreviation!

Only variables, functions, or types whose names begin with a capital letter are considered exported: accessible from packages outside the current package.

Conversions

Math and comparison operations in Go require that the included values be of the same type. If they're not, you'll get an error when trying to run your code.

Set up a float64 variable.

Set up an int variable.

```
var length float64 = 1.2
var width int = 2
fmt.Println("Area is", length*width)
fmt.Println("length > width?", length > width)
```

If we use both the float64 and the int in a math operation...

Or a comparison...

...we'll get errors!

Errors

```
invalid operation: length * width (mismatched types float64 and int)
invalid operation: length > width (mismatched types float64 and int)
```

The same is true of assigning new values to variables. If the type of value being assigned doesn't match the declared type of the variable, you'll get an error.

Set up a float64 variable.

Set up an int variable.

```
var length float64 = 1.2
var width int = 2
length = width
fmt.Println(length)
```

If we assign the int value to the float64 variable...

...we'll get an error!

Error

```
cannot use width (type int) as type float64 in assignment
```

The solution is to use **conversions**, which let you convert a value from one type to another type. You just provide the type you want to convert a value to, immediately followed by the value you want to convert in parentheses.

```
var myInt int = 2
float64(myInt)
```

Type to convert to

Value to convert

The result is a new value of the desired type. Here's what we get when we call `TypeOf` on the value in an integer variable, and again on that same value after conversion to a `float64`:

Without a conversion...

```
var myInt int = 2
fmt.Println(reflect.TypeOf(myInt))
fmt.Println(reflect.TypeOf(float64(myInt)))
```

With a conversion...

```
int
float64
```

Type is changed.

Conversions (continued)

Let's update our failing code example to convert the `int` value to a `float64` before using it in any math operations or comparisons with other `float64` values.

```
var length float64 = 1.2
var width int = 2
fmt.Println("Area is", length*float64(width))
fmt.Println("length > width?", length > float64(width))
```

Convert the int to a float64 before multiplying it with another float64.

Convert the int to a float64 before comparing it with another float64.

```
Area is 2.4
length > width? false
```

The math operation and comparison both work correctly now!

Now let's try converting an `int` to a `float64` before assigning it to a `float64` variable:

```
var length float64 = 1.2
var width int = 2
length = float64(width)
fmt.Println(length)
```

Convert the int to a float64 before assigning it to the float64 variable.

```
2
```

Again, with the conversion in place, the assignment is successful.

When making conversions, be aware of how they might change the resulting values. For example, `float64` variables can store fractional values, but `int` variables can't. When you convert a `float64` to an `int`, the fractional portion is simply dropped! This can throw off any operations you do with the resulting value.

```
var length float64 = 3.75
var width int = 5
width = int(length)
fmt.Println(width)
```

This conversion causes the fractional portion to be dropped!

```
3
```

The resulting value is 0.75 lower!

As long as you're cautious, though, you'll find conversions essential to working with Go. They allow otherwise-incompatible types to work together.



We've written the Go code below to calculate a total price with tax and determine if we have enough funds to make a purchase. But we're getting errors when we try to include it in a full program!

```
var price int = 100
fmt.Println("Price is", price, "dollars.")

var taxRate float64 = 0.08
var tax float64 = price * taxRate
fmt.Println("Tax is", tax, "dollars.")

var total float64 = price + tax
fmt.Println("Total cost is", total, "dollars.")

var availableFunds int = 120
fmt.Println(availableFunds, "dollars available.")
fmt.Println("Within budget?", total <= availableFunds)
```

Errors

```
invalid operation: price * taxRate (mismatched types int and float64)
invalid operation: price + tax (mismatched types int and float64)
invalid operation: total <= availableFunds (mismatched types float64 and int)
```

Fill in the blanks below to update this code. Fix the errors so that it produces the expected output. (Hint: Before doing math operations or comparisons, you'll need to use conversions to make the types compatible.)

```
var price int = 100
fmt.Println("Price is", price, "dollars.")

var taxRate float64 = 0.08
var tax float64 = _____________________________
fmt.Println("Tax is", tax, "dollars.")

var total float64 = _____________________
fmt.Println("Total cost is", total, "dollars.")

var availableFunds int = 120
fmt.Println(availableFunds, "dollars available.")
fmt.Println("Within budget?", _____________________________________)
```

Expected output

```
Price is 100 dollars.
Tax is 8 dollars.
Total cost is 108 dollars.
120 dollars available.
Within budget? true
```

Answers on page 30.

Installing Go on your computer

The Go Playground is a great way to try out the language. But its practical uses are limited. You can't use it to work with files, for example. And it doesn't have a way to take user input from the terminal, which we're going to need for an upcoming program.

So, to wrap up this chapter, let's download and install Go on your computer. Don't worry, the Go team has made it really easy! On most operating systems, you just have to run an installer program, and you'll be done.



1. Visit *https://golang.org* in your web browser.
2. Click the download link.
3. Select the installation package for your operating system (OS). The download should begin automatically.
4. Visit the installation instructions page for your OS (you may be taken there automatically after the download starts), and follow the directions there.
5. Open a new terminal or command prompt window.
6. Confirm Go was installed by typing **`go version`** at the prompt and hitting the Return or Enter key. You should see a message with the version of Go that's installed.



Websites are always changing.

It's possible that golang.org *or the Go installer will be updated after this book is published, and these directions will no longer be completely accurate. In that case, visit:*

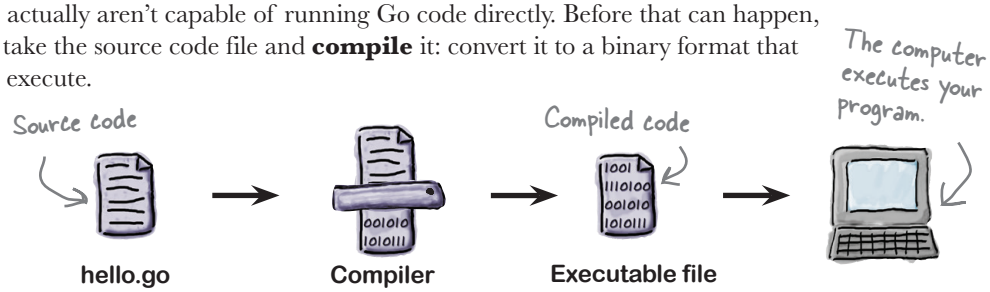
http://headfirstgo.com

for help and troubleshooting tips!

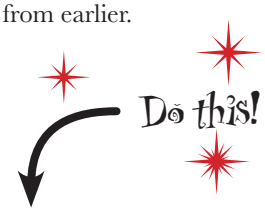
Compiling Go code

Our interaction with the Go Playground has consisted of typing in code and having it mysteriously run. Now that we've actually installed Go on your computer, it's time to take a closer look at how this works.

Computers actually aren't capable of running Go code directly. Before that can happen, we need to take the source code file and **compile** it: convert it to a binary format that a CPU can execute.



Let's try using our new Go installation to compile and run our "Hello, Go!" example from earlier.



Save this to a file.

hello.go

```
package main

import "fmt"

func main() {
        fmt.Println("Hello, Go!")
}
```

1. Using your favorite text editor, save our "Hello, Go!" code from earlier in a plain-text file named *hello.go*.
2. Open a new terminal or command prompt window.
3. In the terminal, change to the directory where you saved *hello.go*.
4. Run **`go fmt hello.go`** to clean up the code formatting. (This step isn't required, but it's a good idea anyway.)
5. Run **`go build hello.go`** to compile the source code. This will add an executable file to the current directory. On macOS or Linux, the executable will be named just *hello*. On Windows, the executable will be named *hello.exe*.
6. Run the executable file. On macOS or Linux, do this by typing **`./hello`** (which means "run a program named `hello` in the current directory"). On Windows, just type **`hello.exe`**.

Change to whatever directory you saved hello.go in.
Format code.
Compile code.
Run executable.

Shell Edit View Window Help

```
$ cd try_go
$ go fmt hello.go
$ go build hello.go
$ ./hello
Hello, Go!
$
```

Compiling and running *hello.go* on macOS or Linux

Change to whatever directory you saved hello.go in.
Format code.
Compile code.
Run executable.

Command Prompt

```
>cd try_go
>go fmt hello.go
>go build hello.go
>hello.exe
Hello, Go!
>
```

Compiling and running *hello.go* on Windows

Go tools

When you install Go, it adds an executable named *go* to your command prompt. The *go* executable gives you access to various commands, including:

Command	Description
`go build`	Compiles source code files into binary files.
`go run`	Compiles and runs a program, without saving an executable file.
`go fmt`	Reformats source files using Go standard formatting.
`go version`	Displays the current Go version.

We just tried the `go fmt` command, which reformats your code in the standard Go format. It's equivalent to the Format button on the Go Playground site. We recommend running `go fmt` on every source file you create.

Most editors can be set up to automatically run go fmt every time you save a file! See https://blog.golang.org/go-fmt-your-code.

We also used the `go build` command to compile code into an executable file. Executable files like this can be distributed to users, and they'll be able to run them even if they don't have Go installed.

But we haven't tried the `go run` command yet. Let's do that now.

Try out code quickly with "go run"

The `go run` command compiles and runs a source file, without saving an executable file to the current directory. It's great for quickly trying out simple programs. Let's use it to run our *hello.go* sample.

1. Open a new terminal or command prompt window.
2. In the terminal, change to the directory where you saved *hello.go*.
3. Type **`go run hello.go`** and hit Enter/Return. (The command is the same on all operating systems.)

```
package main

import "fmt"

func main() {
        fmt.Println("Hello, Go!")
}
```

hello.go

Change to whatever directory you saved hello.go in.

Run source file.

```
Shell Edit View Window Help
$ cd try_go
$ go run hello.go
Hello, Go!
$
```

Running *hello.go* with `go run` (works on any OS)

You'll immediately see the program output. If you make changes to the source code, you don't have to do a separate compilation step; just run your code with `go run` and you'll be able to see the results right away. When you're working on small programs, `go run` is a handy tool to have!



Your Go Toolbox

That's it for Chapter 1! You've added function calls and types to your toolbox.

Function calls

A function is a chunk of code that you can call from other places in your program.

When calling a function, you can use arguments to provide the function with data.

Types

Values in Go are classified into different types, which specify what the values can be used for.

Math operations and comparisons between different types are not allowed, but you can convert a value to a new type if needed.

Go variables can only store values of their declared type.

BULLET POINTS

- A **package** is a group of related functions and other code.
- Before you can use a package's functions within a Go file, you need to **import** that package.
- A `string` is a series of bytes that usually represent text characters.
- A `rune` represents a single text character.
- Go's two most common numeric types are `int`, which holds integers, and `float64`, which holds floating-point numbers.
- The `bool` type holds Boolean values, which are either `true` or `false`.
- A **variable** is a piece of storage that can contain values of a specified type.
- If no value has been assigned to a variable, it will contain the **zero value** for its type. Examples of zero values include `0` for `int` or `float64` variables, or `""` for `string` variables.
- You can declare a variable and assign it a value at the same time using a `:=` **short variable declaration**.
- A variable, function, or type can only be accessed from code in other packages if its name begins with a capital letter.
- The `go fmt` command automatically reformats source files to use Go standard formatting. You should run `go fmt` on any code that you plan to share with others.
- The `go build` command **compiles** Go source code into a binary format that computers can execute.
- The `go run` command compiles and runs a program without saving an executable file in the current directory.

Pool Puzzle Solution

```
package main

import (
        "fmt"
)

func main() {
        fmt.Println("Cannonball!!!!")
}
```

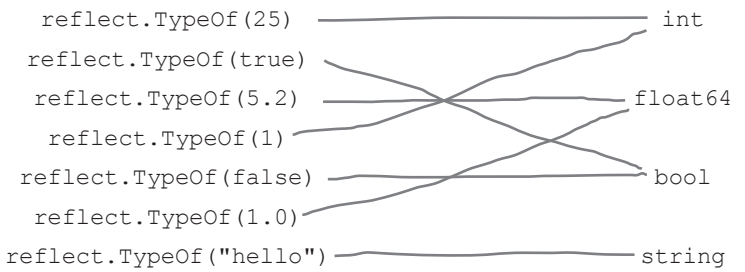
Output

```
Cannonball!!!!
```



Exercise Solution

Draw lines to match each code snippet below to a type.
Some types will have more than one snippet that matches with them.



Code Magnets Solution

```
package main

import (
        "fmt"
)

func main() {

        var originalCount int = 10
        fmt.Println("I started with", originalCount, "apples.")
        var eatenCount int = 4
        fmt.Println("Some jerk ate", eatenCount, "apples.")
        fmt.Println("There are", originalcount-eatenCount, "apples left.")
}
```

Output

```
I started with 10 apples.
Some jerk ate 4 apples.
There are 6 apples left.
```



Exercise Solution

Fill in the blanks below to update this code. Fix the errors so that it produces the expected output. (Hint: Before doing math operations or comparisons, you'll need to use conversions to make the types compatible.)

```
var price int = 100
fmt.Println("Price is", price, "dollars.")

var taxRate float64 = 0.08
var tax float64 = float64(price) * taxRate
fmt.Println("Tax is", tax, "dollars.")

var total float64 = float64(price) + tax
fmt.Println("Total cost is", total, "dollars.")

var availableFunds int = 120
fmt.Println(availableFunds, "dollars available.")
fmt.Println("Within budget?", total <= float64(availableFunds))
```

Expected output

```
Price is 100 dollars.
Tax is 8 dollars.
Total cost is 108 dollars.
120 dollars available.
Within budget? true
```

Conditionals and Loops



Every program has parts that apply only in certain situations.

"This code should run *if* there's an error. Otherwise, that other code should run." Almost every program contains code that should be run only when a certain *condition* is true. So almost every programming language provides **conditional statements** that let you determine whether to run segments of code. Go is no exception.

You may also need some parts of your code to run *repeatedly*. Like most languages, Go provides **loops** that run sections of code more than once. We'll learn to use both conditionals and loops in this chapter!

Calling methods

In Go, it's possible to define **methods**: functions that are associated with values of a given type. Go methods are kind of like the methods that you may have seen attached to "objects" in other languages, but they're a bit simpler.

We'll be taking a detailed look at how methods work in Chapter 9. But we need to use a couple methods to make our examples for this chapter work, so let's look at some brief examples of calling methods now.

The `time` package has a `Time` type that represents a date (year, month, and day) and time (hour, minute, second, etc.). Each `time.Time` value has a `Year` method that returns the year. The code below uses this method to print the current year:

```
package main

import (
        "fmt"
        "time"
)

func main() {
        var now time.Time = time.Now()
        var year int = now.Year()
        fmt.Println(year)
}
```

We need to import the "time" package so we can use the time.Time type.

time.Now returns a time.Time value representing the current date and time.

time.Time values have a Year method that returns the year.

```
2019
```

(Or whatever year your computer's clock is set for.)

The `time.Now` function returns a new `Time` value for the current date and time, which we store in the `now` variable. Then, we call the `Year` method on the value that `now` refers to:

Holds a time.Time value

Call the Year method on the time.Time value.

```
now.Year()
```

The `Year` method returns an integer with the year, which we then print.

Methods are functions that are associated with values of a particular type.

Calling methods (continued)

The `strings` package has a `Replacer` type that can search through a string for a substring, and replace each occurrence of that substring with another string. The code below replaces every # symbol in a string with the letter o:

```
package main

import (
        "fmt"
        "strings"
)

func main() {
        broken := "G# r#cks!"
        replacer := strings.NewReplacer("#", "o")
        fixed := replacer.Replace(broken)
        fmt.Println(fixed)
}
```

Import packages used in the "main" function.

This returns a strings.Replacer that's set up to replace every "#" with "o".

Call the Replace method on the strings.Replacer, and pass it a string to do the replacements on.

Print the string returned from the Replace method.

```
Go rocks!
```

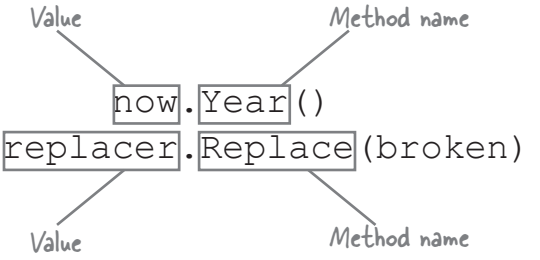
The `strings.NewReplacer` function takes arguments with a string to replace (`"#"`), and a string to replace it with (`"o"`), and returns a `strings.Replacer`. When we pass a string to the `Replacer` value's `Replace` method, it returns a string with those replacements made.



The syntax for calling a method looks a lot like the syntax for calling a function in a different package. Are the two related?

The dot indicates that the thing on its right belongs to the thing on its left.

Whereas the functions we saw earlier belonged to a *package*, the methods belong to an individual *value*. That value is what appears to the left of the dot.



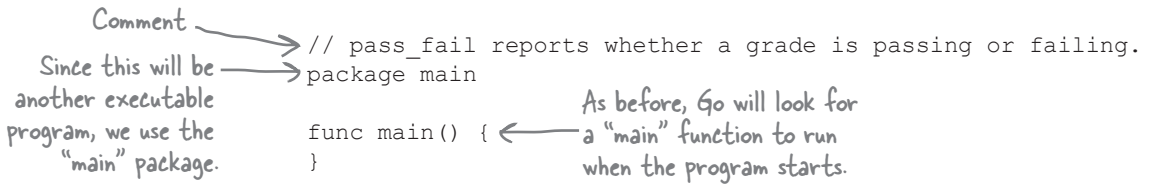
Making the grade

In this chapter, we're going to look at features of Go that let you decide whether to run some code or not, based on a condition. Let's look at a situation where we might need that ability...

We need to write a program that allows a student to type in their percentage grade and tells them whether they passed or not. Passing or failing follows a simple formula: a grade of 60% or more is passing, and less than 60% is failing. So our program will need to give one response if the percentage users enter is 60 or greater, and a different response otherwise.

Comments

Let's create a new file, *pass_fail.go*, to hold our program. We're going to take care of a detail we omitted in our previous programs, and add a description of what the program does at the top.



Most Go programs include descriptions in their source code of what they do, intended for people maintaining the program to read. These **comments** are ignored by the compiler.

The most common form of comment is marked with two slash characters (//). Everything from the slashes to the end of the line is treated as part of the comment. A // comment can appear on a line by itself, or following a line of code.

```
// The total number of widgets in the system.
var TotalCount int // Can only be a whole number.
```

The less frequently used form of comments, **block comments**, spans multiple lines. Block comments start with /* and end with */, and everything between those markers (including newlines) is part of the comment.

```
/*
Package widget includes all the functions used
for processing widgets.
*/
```

Getting a grade from the user

Now let's add some actual code to our *pass_fail.go* program. The first thing it needs to do is allow the user to input a percentage grade. We want them to type a number and press Enter, and we'll store the number they typed in a variable. Let's add code to handle this. *(Note: this code will not actually compile as shown; we'll talk about the reason in a moment!)*

```
// pass_fail reports whether a grade is passing or failing.
package main

import (
        "bufio"
        "fmt"
        "os"
)

func main() {
        fmt.Print("Enter a grade: ")
        reader := bufio.NewReader(os.Stdin)
        input := reader.ReadString('\n')
        fmt.Println(input)
}
```

Import packages used in the "main" function.

Prompt the user to enter a grade.

Set up a "buffered reader" that gets text from the keyboard.

Return everything the user has typed, up to where they pressed the Enter key.

Print what the user typed.

First, we need to let the user know to enter something, so we use the `fmt.Print` function to display a prompt. (Unlike the `Println` function, `Print` doesn't skip to a new terminal line after printing a message, which lets us keep the prompt and the user's entry on the same line.)

Next, we need a way to read (receive and store) input from the program's *standard input*, which all keyboard input goes to. The line `reader := bufio.NewReader(os.Stdin)` stores a `bufio.Reader` in the `reader` variable that can do that for us.

Returns a new bufio.Reader

```
reader := bufio.NewReader(os.Stdin)
```

The Reader will read from standard input (the keyboard).

To actually get the user's input, we call the `ReadString` method on the `Reader`. The `ReadString` method requires an argument with a rune (character) that marks the end of the input. We want to read everything the user types up until they press Enter, so we give `ReadString` a newline rune.

Returns what the user typed, as a string

```
input := reader.ReadString('\n')
```

Everything up until the newline rune will be read.

Once we have the user input, we simply print it.

That's the plan, anyway. But if we try to compile or run this program, we'll get an error:

Error →

```
multiple-value
reader.ReadString()
in single-value context
```



Don't worry too much about the details of how *bufio.Reader* works.

All you really need to know at this point is that it lets us read input from the keyboard.

Multiple return values from a function or method

We're trying to read the user's keyboard input, but we're getting an error. The compiler is reporting a problem in this line of code:

```
input := reader.ReadString('\n')
```

Error → `multiple-value reader.ReadString() in single-value context`

The problem is that the `ReadString` method is trying to return *two* values, and we've only provided *one* variable to assign a value to.

In most programming languages, functions and methods can only have a single return value, but in Go, they can return any number of values. The most common use of multiple return values in Go is to return an additional error value that can be consulted to find out if anything went wrong while the function or method was running. A few examples:

```
bool, err := strconv.ParseBool("true")
file, err := os.Open("myfile.txt")
response, err := http.Get("http://golang.org")
```

Returns an error if the string can't be converted to a boolean

Returns an error if the file can't be opened

Returns an error if the page can't be retrieved



So what's the big deal? Just add a variable to hold that error, and then ignore it!

Go doesn't allow us to declare a variable unless we use it.

Go requires that every variable that gets *declared* must also get *used* somewhere in your program. If we add an `err` variable and then don't check it, our code won't compile. Unused variables often indicate a bug, so this is an example of Go helping you detect and fix bugs!

```
// pass_fail reports whether a grade is...
package main

import (
        "bufio"
        "fmt"
        "os"
)

func main() {
        fmt.Print("Enter a grade: ")
        reader := bufio.NewReader(os.Stdin)
        input, err := reader.ReadString('\n')
        fmt.Println(input)
}
```

If we just add a variable without using it...

...we'll get an error!

Error → `err declared and not used`

Option 1: Ignore the error return value with the blank identifier

The `ReadString` method returns a second value along with the user's input, and we need to do something with that second value. We've tried just adding a second variable and ignoring it, but our code still won't compile.

```
input, err := reader.ReadString('\n')
```

Error → `err declared and not used`

When we have a value that would normally be assigned to a variable, but that we don't intend to use, we can use Go's **blank identifier**. Assigning a value to the blank identifier essentially discards it (while making it obvious to others reading your code that you are doing so). To use the blank identifier, simply type a single underscore (_) character in an assignment statement where you would normally type a variable name.

Let's try using the blank identifier in place of our old `err` variable:

```
// pass_fail reports whether a grade is passing or failing.
package main

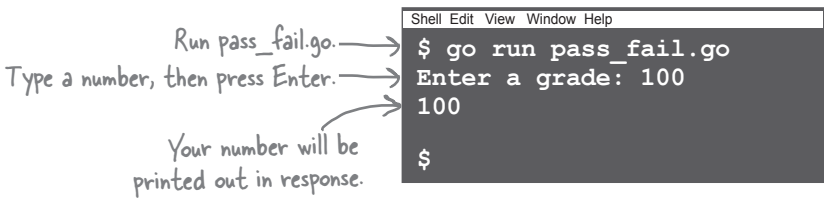
import (
        "bufio"
        "fmt"
        "os"
)

func main() {
        fmt.Print("Enter a grade: ")
        reader := bufio.NewReader(os.Stdin)
        input, _ := reader.ReadString('\n')
        fmt.Println(input)
}
```

Use the blank identifier as a placeholder for the error value.

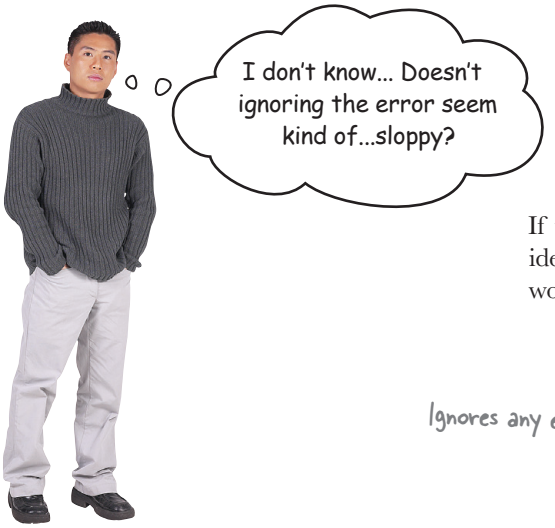
Now we'll try the change out. In your terminal, change to the directory where you saved *pass_fail.go*, and run the program with:

```
go run pass_fail.go
```



When you type a grade (or any other string) at the prompt and press Enter, your entry will be echoed back to you. Our program is working!

Option 2: Handle the error



That's true. If an error actually occurred, this program wouldn't tell us!

If we got an error back from the `ReadString` method, the blank identifier would just cause the error to be ignored, and our program would proceed anyway, possibly with invalid data.

```
func main() {
        fmt.Print("Enter a grade: ")
        reader := bufio.NewReader(os.Stdin)
        input, _ := reader.ReadString('\n')
        fmt.Println(input)
}
```

Ignores any error return value!

Prints what may be an invalid value!

In this case, it would be more appropriate to alert the user and stop the program if there was an error.

The `log` package has a `Fatal` function that can do both of these operations for us at once: log a message to the terminal *and* stop the program. ("Fatal" in this context means reporting an error that "kills" your program.)

Let's get rid of the blank identifier and replace it with an `err` variable so that we're recording the error again. Then, we'll use the `Fatal` function to log the error and halt the program.

```
// pass_fail reports whether a grade is passing or failing.
package main

import (
        "bufio"
        "fmt"
        "log"
        "os"
)

func main() {
        fmt.Print("Enter a grade: ")
        reader := bufio.NewReader(os.Stdin)
        input, err := reader.ReadString('\n')
        log.Fatal(err)
        fmt.Println(input)
}
```

Add the "log" package.

Go back to storing the error return value in a variable.

Report the error and stop the program.

But if we try running this updated program, we'll see there's a new problem...

Conditionals

If our program encounters a problem reading input from the keyboard, we've set it up to report the error and stop running. But now, it stops running even when everything's working correctly!

Store the error return value in a variable.

```
input, err := reader.ReadString('\n')
log.Fatal(err)
```

Log the error return value.

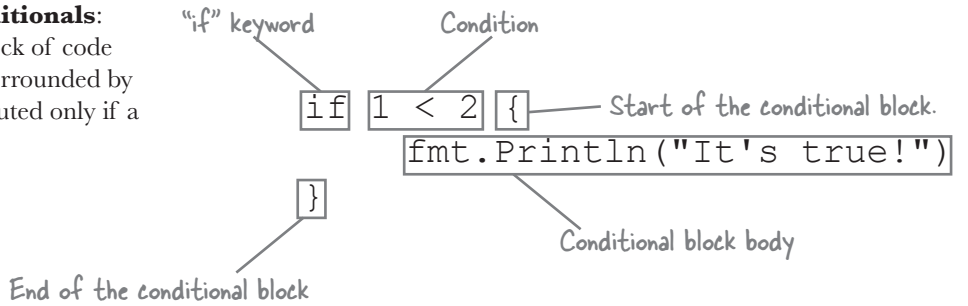
An error gets logged even if everything's working correctly!

```
Shell Edit View Window Help
$ go run pass_fail.go
Enter a grade: 100
2018/03/11 18:27:08 <nil>
exit status 1
$
```

The error value is "nil".

Functions and methods like `ReadString` return an error value of **nil**, which basically means "there's nothing there." In other words, if `err` is `nil`, it means there was no error. But our program is set up to simply report the `nil` error! What we *should* do is exit the program only *if* the `err` variable has a value other than `nil`.

We can do this using **conditionals**: statements that cause a block of code (one or more statements surrounded by { } curly braces) to be executed only if a condition is met.



An expression is evaluated, and if its result is `true`, the code in the conditional block body is executed. If it's `false`, the conditional block is skipped.

```
if true {
        fmt.Println("I'll be printed!")
}
```

```
if false {
        fmt.Println("I won't!")
}
```

As with most other languages, Go supports multiple branches in the conditional. These statements take the form `if...else if...else`.

```
if grade == 100 {
        fmt.Println("Perfect!")
} else if grade >= 60 {
        fmt.Println("You pass.")
} else {
        fmt.Println("You fail!")
}
```

Conditionals (continued)

Conditionals rely on a Boolean expression (one that evaluates to `true` or `false`) to decide whether the code they contain should be executed.

```
if 1 == 1 {
        fmt.Println("I'll be printed!")
}

if 1 > 2 {
        fmt.Println("I won't!")
}

if 1 < 2 {
        fmt.Println("I'll be printed!")
}

if 1 >= 2 {
        fmt.Println("I won't!")
}

if 2 <= 2 {
        fmt.Println("I'll be printed!")
}

if 2 != 2 {
        fmt.Println("I won't!")
}
```

When you need to execute code only if a condition is *false*, you can use `!`, the Boolean negation operator, which lets you take a `true` value and make it `false`, or a `false` value and make it `true`.

```
if !true {
        fmt.Println("I won't be printed!")
}

if !false {
        fmt.Println("I will!")
}
```

If you want to run some code only if two conditions are *both* true, you can use the `&&` ("and") operator. If you want it to run if *either* of two conditions is true, you can use the `||` ("or") operator.

```
if true && true {
        fmt.Println("I'll be printed!")
}

if true && false {
        fmt.Println("I won't!")
}

if false || true {
        fmt.Println("I'll be printed!")
}

if false || false {
        fmt.Println("I won't!")
}
```

there are no Dumb Questions

Q: My other programming language requires that an `if` statement's condition be surrounded with parentheses. Doesn't Go?

A: No, and in fact the `go fmt` tool will remove any parentheses you add, unless you're using them to set order of operations.



Because they're in conditional blocks, only some of the `Println` calls in the code below will be executed. Write down what the output would be.

(We've done the first two lines for you.)

```
if true {
        fmt.Println("true")
}
if false {
        fmt.Println("false")
}
if !false {
        fmt.Println("!false")
}
if true {
        fmt.Println("if true")
} else {
        fmt.Println("else")
}
if false {
        fmt.Println("if false")
} else if true {
        fmt.Println("else if true")
}
if 12 == 12 {
        fmt.Println("12 == 12")
}
if 12 != 12 {
        fmt.Println("12 != 12")
}
if 12 > 12 {
        fmt.Println("12 > 12")
}
if 12 >= 12 {
        fmt.Println("12 >= 12")
}
if 12 == 12 && 5.9 == 5.9 {
        fmt.Println("12 == 12 && 5.9 == 5.9")
}
if 12 == 12 && 5.9 == 6.4 {
        fmt.Println("12 == 12 && 5.9 == 6.4")
}
if 12 == 12 || 5.9 == 6.4 {
        fmt.Println("12 == 12 || 5.9 == 6.4")
}
```

Output:

true

!false

Answers on page 75.

Logging a fatal error, conditionally

Our grading program is reporting an error and exiting, even if it reads input from the keyboard successfully.

Store the error return value in a variable.

```
input, err := reader.ReadString('\n')
log.Fatal(err)
```

Log the error return value.

An error gets logged even if everything's working correctly!

```
Shell Edit View Window Help
$ go run pass_fail.go
Enter a grade: 100
2018/03/11 18:27:08 <nil>
exit status 1
$
```

The error value is "nil".

We know that if the value in our `err` variable is `nil`, it means reading from the keyboard was successful. Now that we know about `if` statements, let's try updating our code to log an error and exit only if `err` is *not* `nil`.

```
// pass_fail reports whether a grade is passing or failing.
package main

import (
        "bufio"
        "fmt"
        "log"
        "os"
)

func main() {
        fmt.Print("Enter a grade: ")
        reader := bufio.NewReader(os.Stdin)
        input, err := reader.ReadString('\n')
        if err != nil {
                log.Fatal(err)
        }
        fmt.Println(input)
}
```

If "error" is not nil...

Report the error and stop the program.

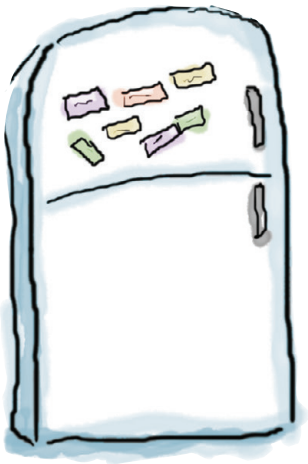
If we rerun our program, we'll see that it's working again. And now, if there are any errors when reading user input, we'll see those as well!

Run pass_fail.go.

```
Shell Edit View Window Help
$ go run pass_fail.go
Enter a grade: 100
100

$
```

Your number will be printed out in response.



Code Magnets

A Go program that prints the size of a file is on the fridge. It calls the `os.Stat` function, which returns an `os.FileInfo` value, and possibly an error value. Then it calls the `Size` method on the `FileInfo` value to get the file size.

But the original program uses the _ blank identifier to ignore the error value from `os.Stat`. If an error occurs (which could happen if the file doesn't exist), this will cause the program to fail.

Reconstruct the extra code snippets to make a program that works just like the original one, but also checks for an error from `os.Stat`. If the error from `os.Stat` is not `nil`, the error should be reported, and the program should exit. Discard the magnet with the _ blank identifier; it won't be used in the finished program.

```
package main
```

This is already a complete program! But it ignores any errors that might happen....

```
import (
        "fmt"
        "log"
        "os"
)
```

```
func main() {
```

The blank identifier ignores any error value. Discard this magnet and replace it with one of the magnets below!

Get a FileInfo value with data regarding the my.txt file.

Holds the file size, date it was changed, etc.

```
fileInfo,  _  :=  os.Stat("my.txt")
```

Add your code here! If the error is not nil, pass it to log.Fatal.

```
fmt.Println(fileInfo.Size())
```

```
}
```

Returns the size of the file

Here are the extra magnets. Add them to the program above!

```
{   !=   }   nil   err   err   if   log.Fatal(err)
```

Answers on page 76.

Avoid shadowing names



Something else is bothering me. You said before that you were trying to avoid abbreviations in this book. Yet here you are, naming variables `err` instead of `error`!

```
fmt.Print("Enter a grade: ")
reader := bufio.NewReader(os.Stdin)
input, err := reader.ReadString('\n')
if err != nil {
       log.Fatal(err)
}
```

Naming a variable `error` would be a bad idea, because it would shadow the name of a type called `error`.



When you declare a variable, you should make sure it doesn't have the same name as any existing functions, packages, types, or other variables. If something by the same name exists in the enclosing scope (we'll talk about scopes shortly), your variable will **shadow** it—that is, take precedence over it. And all too often, that's a bad thing.

Here, we declare a variable named `int` that shadows a type name, a variable named `append` that shadows a built-in function name (we'll see the `append` function in Chapter 6), and a variable named `fmt` that shadows an imported package name. Those names are awkward, but they don't cause any errors by themselves...

```
package main

import "fmt"

func main() {
       var int int = 12
       var append string = "minutes of bonus footage"
       var fmt string = "DVD"
}
```

Naming this variable "int" shadows the name of the built-in "int" type!

Naming this variable "append" shadows the name of the built-in "append" function!

Naming this variable "fmt" shadows the name of the imported "fmt" package!

Avoid shadowing names (continued)

...But if we try to access the type, function, or package the variables are shadowing, we'll get the value in the variable instead. In this case, it results in compile errors:

```
func main() {
        var int int = 12
        var append string = "minutes of bonus footage"
        var fmt string = "DVD"
        var count int
        var languages = append([]string{}, "Español")
        fmt.Println(int, append, "on", fmt, languages)
}
```

"int" now refers to the variable declared above, not the numeric type!

"append" now refers to a variable, not a function!

"fmt" now refers to a variable, not a package!

Compile errors

```
imported and not used: "fmt"
int is not a type
cannot call non-function append (type string), declared at prog.go:7:6
fmt.Println undefined (type string has no field or method Println)
```

To avoid confusion for yourself and your fellow developers, you should avoid shadowing names wherever possible. In this case, fixing the issue is as simple as choosing nonconflicting names for the variables:

```
func main() {
        var count int = 12
        var suffix string = "minutes of bonus footage"
        var format string = "DVD"
        var languages = append([]string{}, "Español")
        fmt.Println(count, suffix, "on", format, languages)
}
```

Rename the "int" variable.

Rename the "append" variable.

Rename the "fmt" variable.

```
12 minutes of bonus footage on DVD [Español]
```

As we'll see in Chapter 3, Go has a built-in type named `error`. So that's why, when declaring variables meant to hold errors, we've been naming them `err` instead of `error`—we want to avoid shadowing the name of the `error` type with our variable name.

"err", not "error"!

```
fmt.Print("Enter a grade: ")
reader := bufio.NewReader(os.Stdin)
input, err := reader.ReadString('\n')
if err != nil {
        log.Fatal(err)
}
```

If you *do* name your variables `error`, your code will *probably* still work. That is, *until* you forget that the `error` type name is shadowed, you try to use the type, and you get the variable instead. Don't take that chance; use the name `err` for your error variables!

Converting strings to numbers

Conditional statements will also let us evaluate the entered grade. Let's add an `if/else` statement to determine whether the grade is passing or failing. If the entered percentage grade is 60 or greater, we'll set the status to `"passing"`. Otherwise, we'll set it to `"failing"`.

```
// package and import statements omitted
func main() {
        fmt.Print("Enter a grade: ")
        reader := bufio.NewReader(os.Stdin)
        input, err := reader.ReadString('\n')
        if err != nil {
                log.Fatal(err)
        }

        if input >= 60 {
                status := "passing"
        } else {
                status := "failing"
        }
}
```

In its current form, though, this gets us a compilation error.

Error →
```
cannot convert 60 to type string
invalid operation: input >= 60 (mismatched types string and int)
```

Here's the problem: input from the keyboard is read in as a string. Go can only compare numbers to other numbers; we can't compare a number with a string. And there's no direct type conversion from `string` to a number:

```
float64("2.6")
```

Error →
```
cannot convert "2.6" (type string) to type float64
```

We have a pair of issues to address here:

- The `input` string still has a newline character on the end, from when the user pressed the Enter key while entering it. We need to strip that off.
- The remainder of the string needs to be converted to a floating-point number.

Converting strings to numbers (continued)

Removing the newline character from the end of the `input` string will be easy. The `strings` package has a `TrimSpace` function that will remove all whitespace characters (newlines, tabs, and regular spaces) from the start and end of a string.

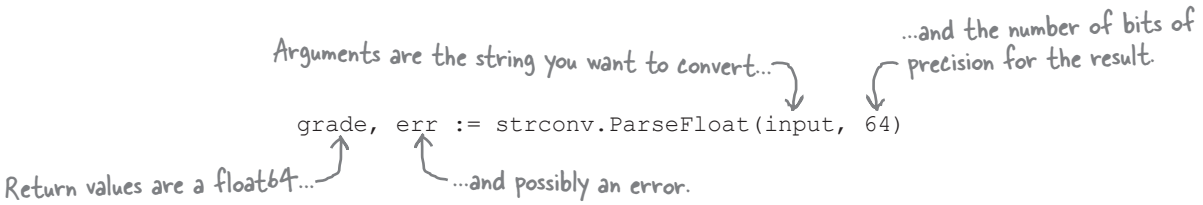
```
s := "\t formerly surrounded by space \n"
fmt.Println(strings.TrimSpace(s))
```

```
formerly surrounded by space
```

So, we can get rid of the newline on `input` by passing it to `TrimSpace`, and assigning the return value back to the `input` variable.

```
input = strings.TrimSpace(input)
```

All that should remain in the `input` string now is the number the user entered. We can use the `strconv` package's `ParseFloat` function to convert it to a `float64` value.



You pass `ParseFloat` a string that you want to convert to a number, as well as the number of bits of precision the result should have. Since we're converting to a `float64` value, we pass the number `64`. (In addition to `float64`, Go offers a less precise `float32` type, but you shouldn't use that unless you have a good reason.)

`ParseFloat` converts the string to a number, and returns it as a `float64` value. Like `ReadString`, it also has a second return value, an error, which will be `nil` unless there was some problem converting the string. (For example, a string that *can't* be converted to a number. We don't know of a numeric equivalent to `"hello"`...)



This whole "bits of precision" thing isn't that important right now.

It's basically just a measure of how much computer memory a floating-point number takes up. As long as you know that you want a `float64`, and so you should pass `64` as the second argument to `ParseFloat`, you'll be fine.

Converting strings to numbers (continued)

Let's update *pass_fail.go* with calls to `TrimSpace` and `ParseFloat`:

```
// pass_fail reports whether a grade is passing or failing.
package main

import (
        "bufio"
        "fmt"
        "log"
        "os"
        "strconv"
        "strings"
)

func main() {
        fmt.Print("Enter a grade: ")
        reader := bufio.NewReader(os.Stdin)
        input, err := reader.ReadString('\n')
        if err != nil {
                log.Fatal(err)
        }

        input = strings.TrimSpace(input)
        grade, err := strconv.ParseFloat(input, 64)
        if err != nil {
                log.Fatal(err)
        }

        if grade >= 60 {
                status := "passing"
        } else {
                status := "failing"
        }
}
```

Add "strconv" so we can use ParseFloat.

Add "strings" so we can use the TrimSpace function.

Trim the newline character from the input string.

Convert the string to a float64 value.

Just as with ReadString, report any error when converting.

Compare to the float64 in "grade", not the string in "input".

First, we add the appropriate packages to the `import` section. We add code to remove the newline character from the `input` string. Then we pass `input` to `ParseFloat`, and store the resulting `float64` value in a new variable, `grade`.

Just as we did with `ReadString`, we test whether `ParseFloat` returns an error value. If it does, we report it and stop the program.

Finally, we update the conditional statement to test the number in `grade`, rather than the string in `input`. That should fix the error stemming from comparing a string to a number.

If we try to run the updated program, we no longer get the `mismatched types string and int` error. So it looks like we've fixed that issue. But we've got a couple more errors to address. We'll look at those next.

Errors

```
status declared
and not used
status declared
and not used
```

Blocks

We've converted the user's grade input to a `float64` value, and added it to a conditional to determine if it's passing or failing. But we're getting a couple more compile errors:

```
if grade >= 60 {
        status := "passing"
} else {
        status := "failing"
}
```

Errors

```
status declared and not used
status declared and not used
```

As we've seen previously, declaring a variable like `status` without using it afterward is an error in Go. It seems a little strange that we're getting the error twice, but let's disregard that for now. We'll add a call to `Println` to print the percentage grade we were given, and the value of `status`.

```
func main() {
        // Omitting code up here...
        if grade >= 60 {
                status := "passing"
        } else {
                status := "failing"
        }
        fmt.Println("A grade of", grade, "is", status)
}
```

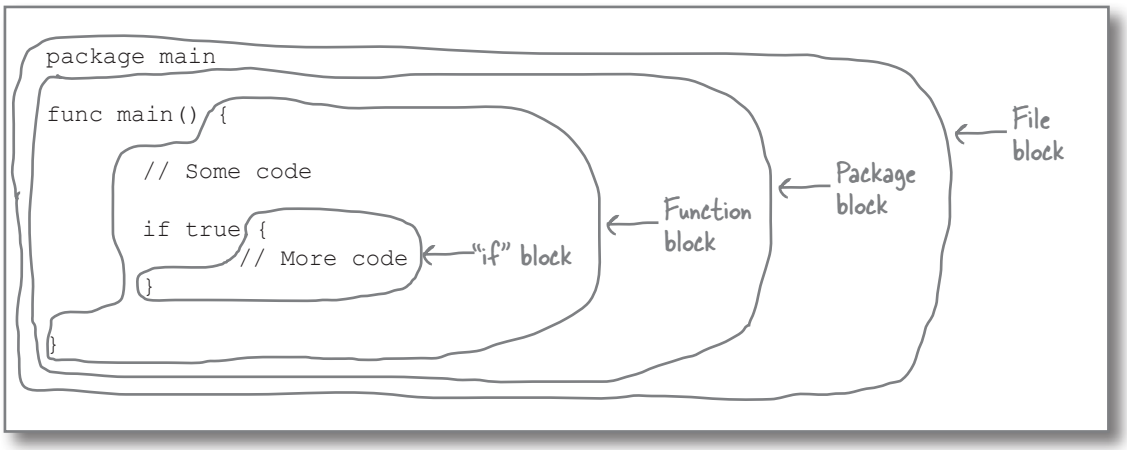
Print status variable

Error

```
undefined: status
```

But now we get a *new* error, saying that the `status` variable is undefined when we attempt to use it in our `Println` statement! What's going on?

Go code can be divided up into **blocks**, segments of code. Blocks are usually surrounded by curly braces ({ }), although there are also blocks at the source code file and package levels. Blocks can be nested inside one another.



The bodies of functions and conditionals are both blocks as well. Understanding this will be key to solving our problem with the `status` variable...

Blocks and variable scope

Each variable you declare has a **scope**: a portion of your code that it's "visible" within. A declared variable can be accessed anywhere within its scope, but if you try to access it outside that scope, you'll get an error.

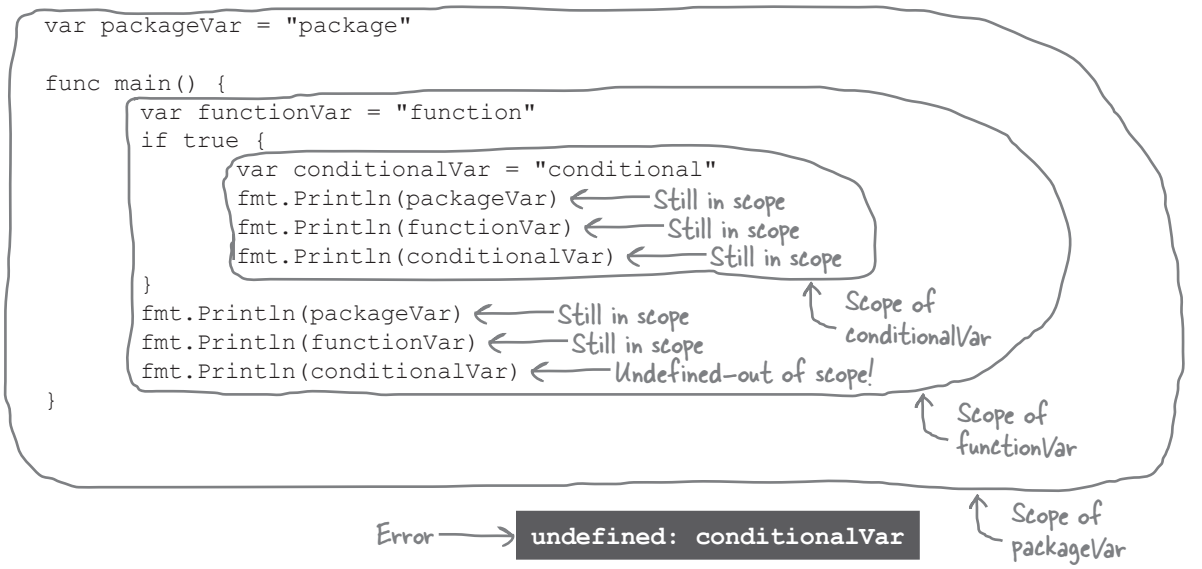
A variable's scope consists of the block it's declared in and any blocks nested within that block.

```
package main

import "fmt"

var packageVar = "package"

func main() {
        var functionVar = "function"
        if true {
                var conditionalVar = "conditional"
                fmt.Println(packageVar)
                fmt.Println(functionVar)
                fmt.Println(conditionalVar)
        }
        fmt.Println(packageVar)
        fmt.Println(functionVar)
        fmt.Println(conditionalVar)
}
```

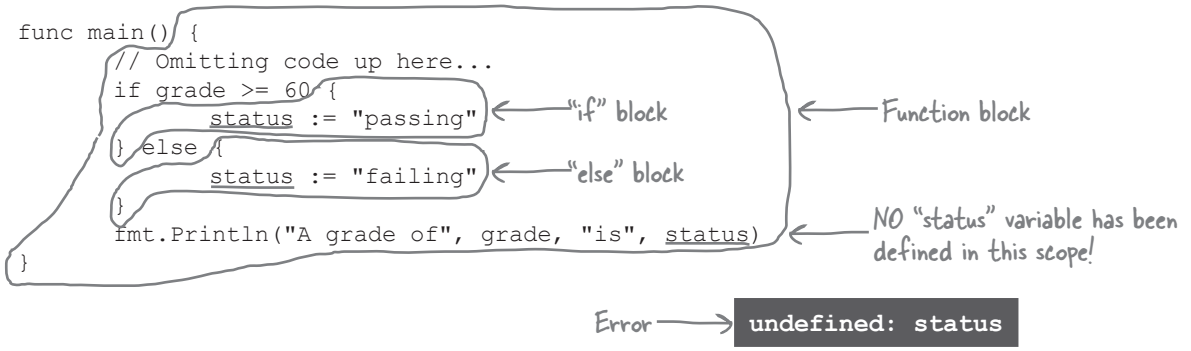


Here are the scopes of the variables in the code above:

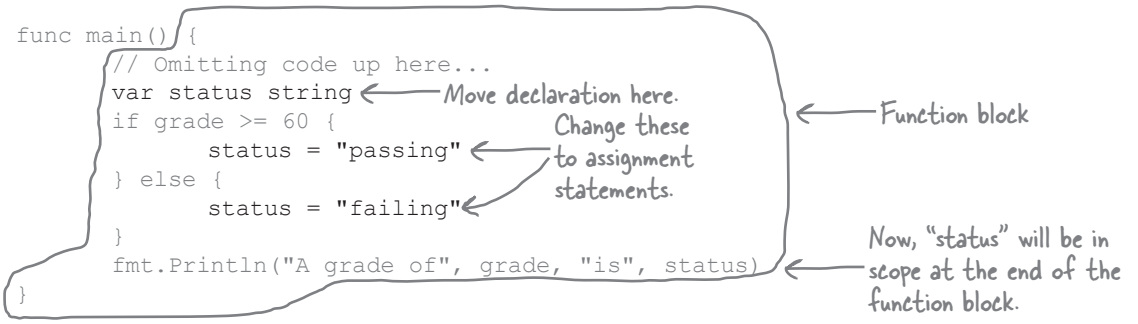
- The scope of `packageVar` is the entire `main` package. You can access `packageVar` anywhere within any function you define in the package.
- The scope of `functionVar` is the entire function it's declared in, including the `if` block nested within that function.
- The scope of `conditionalVar` is limited to the `if` block. When we try to access `conditionalVar` *after* the closing `}` brace of the `if` block, we'll get an error saying that `conditionalVar` is undefined!

Blocks and variable scope (continued)

Now that we understand variable scope, we can explain why our `status` variable was undefined in the grading program. We declared `status` in our conditional blocks. (In fact, we declared it twice, since there are two separate blocks. That's why we got two `status declared and not used` errors.) But then we tried to access `status` *outside* those blocks, where it was no longer in scope.



The solution is to move the declaration of the `status` variable out of the conditional blocks, and up to the function block. Once we do that, the `status` variable will be in scope both within the nested conditional blocks *and* at the end of the function block.





Don't forget to change the short variable declarations within the nested blocks to assignment statements!

If you don't change both occurrences of `:=` to `=`, you'll accidentally create new variables named `status` within the nested conditional blocks, which will then be out of scope at the end of the enclosing function block!

We've finished the grading program!

That was it! Our *pass_fail.go* program is ready for action! Let's take one more look at the complete code:

```
// pass_fail reports whether a grade is passing or failing.
package main

import (
        "bufio"
        "fmt"
        "log"
        "os"
        "strings"
        "strconv"
)

func main() {
        fmt.Print("Enter a grade: ")
        reader := bufio.NewReader(os.Stdin)
        input, err := reader.ReadString('\n')
        if err != nil {
                log.Fatal(err)
        }

        input = strings.TrimSpace(input)
        grade, err := strconv.ParseFloat(input, 64)
        if err != nil {
                log.Fatal(err)
        }

        var status string
        if grade >= 60 {
                status = "passing"
        } else {
                status = "failing"
        }
        fmt.Println("A grade of", grade, "is", status)
}
```

The "main" function gets invoked when the program launches.

Prompt the user to enter a percentage grade.

Create a bufio.Reader, which lets us read keyboard input.

Read what the user types, up until they press Enter.

If there's an error, print the message and exit.

Trim the newline character off the input.

Convert the input string to a float64 (numeric) value.

If there's an error, print the message and exit.

Declare the "status" variable here, so it's in scope for the rest of the function.

If the grade is 60 or over, set the status to "passing". Otherwise, set it to "failing".

Print the entered grade...

...and the pass/fail status.

You can try running the finished program as many times as you like. Enter a percentage grade under 60, and it will report a failing status. Enter a grade over 60, and it will report that it's passing. Looks like everything's working!

```
Shell Edit View Window Help
$ go run pass_fail.go
Enter a grade: 56
A grade of 56 is failing
$ go run pass_fail.go
Enter a grade: 84.5
A grade of 84.5 is passing
$
```



Some of the lines of code below will result in a compile error, because they refer to a variable that is out of scope. Cross out the lines that have errors.

```
package main

import (
	"fmt"
)

var a = "a"

func main() {
	a = "a"
	b := "b"
	if true {
		c := "c"
		if true {
			d := "d"
			fmt.Println(a)
			fmt.Println(b)
			fmt.Println(c)
			fmt.Println(d)
		}
		fmt.Println(a)
		fmt.Println(b)
		fmt.Println(c)
		fmt.Println(d)
	}
	fmt.Println(a)
	fmt.Println(b)
	fmt.Println(c)
	fmt.Println(d)
}
```

Answers on page 77.

Only one variable in a short variable declaration has to be new



One last thing! There's something weird about that grading program code. You said in Chapter 1 that we can't declare a variable twice. And yet the `err` variable appears in two different short variable declarations!

The "err" variable is declared here.

```
input, err := reader.ReadString('\n')
// Code omitted...
grade, err := strconv.ParseFloat(input, 64)
```

But this looks like we're declaring "err" a second time!

It's true that when the same variable name is declared twice in the same scope, we get a compile error:

Attempt to declare "a" again →

```
a := 1
a := 2
```

Compile error

```
no new variables on left side of :=
```

But as long as at least one variable name in a short variable declaration is new, it's allowed. The new variable names are treated as a declaration, and the existing names are treated as an assignment.

```
a := 1          // Declare "a".
b, a := 2, 3    // Declare "b", assign to "a".
a, c := 4, 5    // Assign to "a", declare "c".
fmt.Println(a, b, c)
```

```
4 2 5
```

There's a reason for this special handling: a lot of Go functions return multiple values. It would be a pain if you had to declare all the variables separately just because you want to reuse one of them.

Declaring each variable separately works, but thankfully we don't have to do this...

```
var a, b float64
var err error
a, err = strconv.ParseFloat("1.23", 64)
b, err = strconv.ParseFloat("4.56", 64)
```

Instead, Go lets you use a short variable declaration for everything, even if for *one* of the variables, it's actually an assignment.

...We can just use short variable declaration syntax for everything.

```
a, err := strconv.ParseFloat("1.23", 64)  // Declare "a" and "err".
b, err := strconv.ParseFloat("4.56", 64)  // Declare "b" and assign "err".
fmt.Println(a, b, err)
```

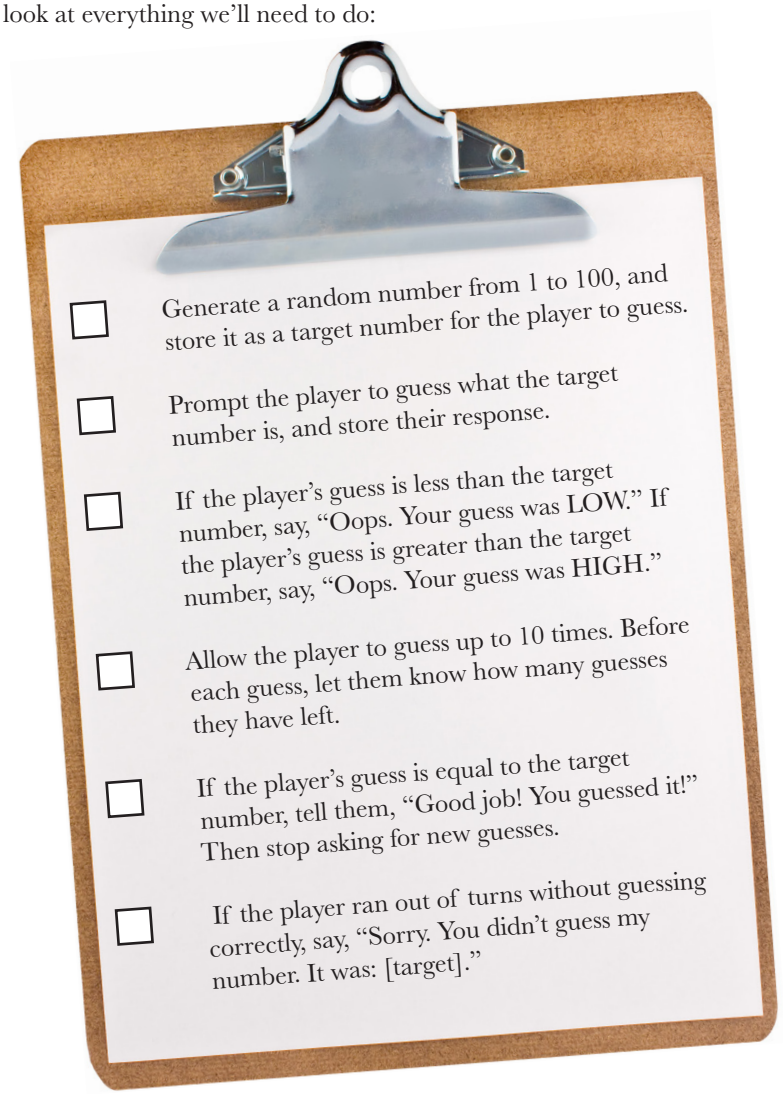
```
1.23 4.56 <nil>
```

Let's build a game

We're going to wrap up this chapter by building a simple game. If that sounds daunting, don't worry; you've already learned most of the skills you're going to need! Along the way, we'll learn about *loops*, which will allow the player to take multiple turns.

This example debuted in Head First Ruby. (Another fine book that you should also buy!) It worked so well that we're using it again here.

Let's look at everything we'll need to do:





Gary Richardott
Game Designer

Let's create a new source file, named *guess.go*.

It looks like our first requirement is to generate a random number. Let's get started!

Package names vs. import paths

The `math/rand` package has a `Intn` function that can generate a random number for us, so we'll need to import `math/rand`. Then we'll call `rand.Intn` to generate the random number.

```
package main

import (
        "fmt"
        "math/rand"
)

func main() {
        target := rand.Intn(100) + 1
        fmt.Println(target)
}
```

Import the "math/rand" package.

Call rand.Intn to generate a random number.



Hang on a second! You said `Intn` came from the `math/rand` package. So why are you just typing `rand.Intn` and not `math/rand.Intn`?

One is the package's import path, and the other is the package's name.

When we say `math/rand` we're referring to the package's *import path*, not its *name*. An **import path** is just a unique string that identifies a package and that you use in an `import` statement. Once you've imported the package, you can refer to it by its package name.

For every package we've used so far, the import path has been identical to the package name. Here are a few examples:

Import path	Package name
`"fmt"`	`fmt`
`"log"`	`log`
`"strings"`	`strings`

But the import path and package name don't have to be identical. Many Go packages fall into similar categories, like compression or complex math. So they're grouped together under similar import path prefixes, such as `"archive/"` or `"math/"`. (Think of them as being similar to the paths of directories on your hard drive.)

Import path	Package name
`"archive"`	`archive`
`"archive/tar"`	`tar`
`"archive/zip"`	`zip`
`"math"`	`math`
`"math/cmplx"`	`cmplx`
`"math/rand"`	`rand`

Package names vs. import paths (continued)

The Go language doesn't require that a package name have anything to do with its import path. But by convention, the last (or only) segment of the import path is also used as the package name. So if the import path is `"archive"`, the package name will be `archive`, and if the import path is `"archive/zip"`, the package name will be `zip`.

Import path	Package name
"archive"	archive
"archive/tar"	tar
"archive/zip"	zip
"math"	math
"math/cmplx"	cmplx
"math/rand"	rand

So, that's why our `import` statement uses a path of `"math/rand"`, but our `main` function just uses the package name: `rand`.

```
package main

import (
        "fmt"
        "math/rand"
)

func main() {
        target := rand.Intn(100) + 1
        fmt.Println(target)
}
```

Use the full import path for "math/rand".

Use the package name: "rand".

Generating a random number

Pass a number to `rand.Intn`, and it will return a random integer between `0` and the number you provided. In other words, if we pass an argument of `100`, we'll get a random number in the range 0–99. Since we need a number in the range 1–100, we'll just add `1` to whatever random value we get. We'll store the result in a variable, `target`. We'll do more with `target` later, but for now we'll just print it.

```
package main

import (
        "fmt"
        "math/rand"
)

func main() {
        target := rand.Intn(100) + 1
        fmt.Println(target)
}
```

Generate an integer from 0 to 99.

Add 1 to make it an integer from 1 to 100.

If we try running our program right now, we'll get a random number. But we just get the *same* random number over and over! The problem is, random numbers generated by computers aren't really that random. But there's a way to increase that randomness...

We get the same random number each time we run the program!

```
Shell Edit View Window Help
$ go run guess.go
82
$ go run guess.go
82
$ go run guess.go
82
$
```

Generating a random number (continued)

To get different random numbers, we need to pass a value to the `rand.Seed` function. That will "seed" the random number generator—that is, give it a value that it will use to generate other random values. But if we keep giving it the same seed value, it will keep giving us the same random values, and we'll be back where we started.

We saw earlier that the `time.Now` function will give us a `Time` value representing the current date and time. We can use that to get a different seed value every time we run our program.

```
package main

import (
        "fmt"
        "math/rand"
        "time"
)

func main() {
        seconds := time.Now().Unix()
        rand.Seed(seconds)
        target := rand.Intn(100) + 1
        fmt.Println("I've chosen a random number between 1 and 100.")
        fmt.Println("Can you guess it?")
        fmt.Println(target)
}
```

Import the "time" package as well.

Get the current date and time, as an integer.

Seed the random number generator.

Now, the generated numbers should be different each time!

Let the player know we've chosen a number.

The `rand.Seed` function expects an integer, so we can't pass it a `Time` value directly. Instead, we call the `Unix` method on the `Time`, which will convert it to an integer. (Specifically, it will convert it to Unix time format, which is an integer with the number of seconds since January 1, 1970. But you don't really need to remember that.) We pass that integer to `rand.Seed`.

We also add a couple `Println` calls to let the user know we've chosen a random number. But aside from that, we can leave the rest of our code, including the call to `rand.Intn`, as is. Seeding the generator should be the only change we need to make.

Now, each time we run our program, we'll see our message, along with a random number. It looks like our changes are successful!

```
Shell Edit View Window Help
$ go run guess.go
I've chosen a random number between 1 and 100.
Can you guess it?
73
$ go run guess.go
I've chosen a random number between 1 and 100.
Can you guess it?
18
$
```

A different number each time we run the program!

Getting an integer from the keyboard

Our first requirement is complete! Next we need to get the user's guess via the keyboard.

That should work in much the same way as when we read in a percentage grade from the keyboard for our grading program.

There will be only one difference: instead of converting the input to a `float64`, we need to convert it to an `int` (since our guessing game uses only whole numbers). So we'll pass the string read from the keyboard to the `strconv` package's `Atoi` (string to integer) function instead of its `ParseFloat` function. `Atoi` will give us an integer as its return value. (Just like `ParseFloat`, `Atoi` might also give us an error if it can't convert the string. If that happens, we again report the error and exit.)




```
package main

import (
        "bufio"
        "fmt"
        "log"
        "math/rand"
        "os"
        "strconv"
        "strings"
        "time"
)

func main() {
        seconds := time.Now().Unix()
        rand.Seed(seconds)
        target := rand.Intn(100) + 1
        fmt.Println("I've chosen a random number between 1 and 100.")
        fmt.Println("Can you guess it?")
        fmt.Println(target)

        reader := bufio.NewReader(os.Stdin)

        fmt.Print("Make a guess: ")
        input, err := reader.ReadString('\n')
        if err != nil {
                log.Fatal(err)
        }
        input = strings.TrimSpace(input)
        guess, err := strconv.Atoi(input)
        if err != nil {
                log.Fatal(err)
        }
}
```

Import these additional packages. (We used all of these in the grading program!)

Create a bufio.Reader, which lets us read keyboard input.

Ask for a number.

Read what the user types, up until they press Enter.

If there's an error, print the message and exit.

Remove the newline.

Convert the input string to an integer.

If there's an error, print the message and exit.

Comparing the guess to the target

Another requirement finished. And this next one will be easy... We just need to compare the user's guess to the randomly generated number, and tell them whether it was higher or lower.

- [x] Prompt the player to guess what the target number is, and store their response.
- [ ] If the player's guess is less than the target number, say, "Oops. Your guess was LOW." If the player's guess is greater than the target number, say, "Oops. Your guess was HIGH."

If `guess` is less than `target`, we need to print a message saying the guess was low. *Otherwise, if* `guess` is greater than `target`, we should print a message saying the guess was high. Sounds like we need an `if...else if` statement. We'll add it below the other code in our `main` function.

```
// No changes to package and import statements; omitting

func main() {
        // No changes to previous code; omitting

        if guess < target {
                fmt.Println("Oops. Your guess was LOW.")
        } else if guess > target {
                fmt.Println("Oops. Your guess was HIGH.")
        }
}
```

If the player's guess was too low, say so.

If the player's guess was too high, say so.

Now try running our updated program from the terminal. It's still set up to print `target` each time it runs, which will be useful for debugging. Just enter a number lower than `target`, and you should be told your guess was low. If you rerun the program, you'll get a new `target` value. Enter a number higher than that, and you'll be told your guess was high.

```
Shell Edit View Window Help
$ go run guess.go
81
I've chosen a random number between 1 and 100.
Can you guess it?
Make a guess: 1
Oops. Your guess was LOW.
$ go run guess.go
54
I've chosen a random number between 1 and 100.
Can you guess it?
Make a guess: 100
Oops. Your guess was HIGH.
$
```

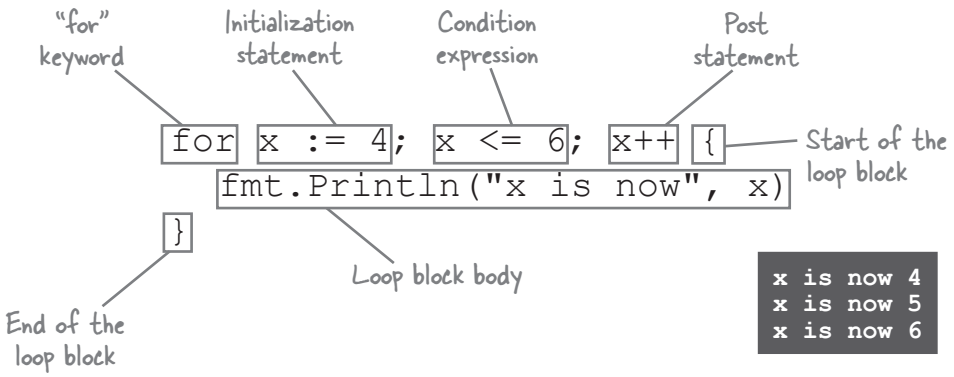
Loops

Another requirement down! Let's look at the next one.

Currently, the player only gets to guess once, but we need to allow them to guess up to 10 times.

The code to prompt for a guess is already in place. We just need to run it *more than once*. We can use a **loop** to execute a block of code repeatedly. If you've used other programming languages, you've probably encountered loops. When you need one or more statements executed over and over, you place them inside a loop.

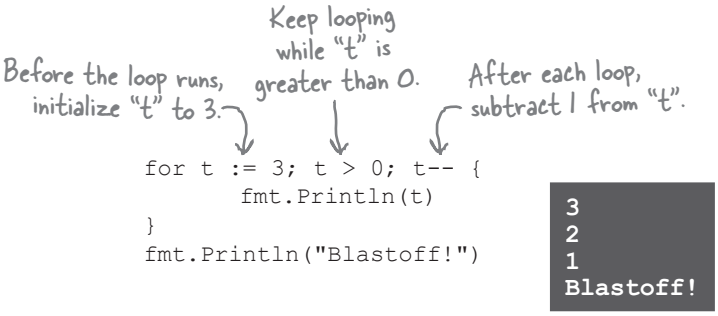
- [x] If the player's guess is less than the target number, say, "Oops. Your guess was LOW." If the player's guess is greater than the target number, say, "Oops. Your guess was HIGH."
- [ ] Allow the player to guess up to 10 times. Before each guess, let them know how many guesses they have left.



Loops always begin with the `for` keyword. In one common kind of loop, `for` is followed by three segments of code that control the loop:

- An initialization (or init) statement that is usually used to initialize a variable
- A condition expression that determines when to break out of the loop
- A post statement that runs after each iteration of the loop

Often, the initialization statement is used to initialize a variable, the condition expression keeps the loop running until that variable reaches a certain value, and the post statement is used to update the value of that variable. For example, in this snippet, the `t` variable is initialized to `3`, the condition keeps the loop going while `t > 0`, and the post statement subtracts `1` from `t` each time the loop runs. Eventually, `t` reaches `0` and the loop ends.



Loops (continued)

The ++ and -- statements are frequently used in loop post statements. Each time they're evaluated, ++ adds 1 to a variable's value, and -- subtracts 1.

```
x := 0
x++
fmt.Println(x)
x++
fmt.Println(x)
x--
fmt.Println(x)
```

```
1
2
1
```

Used in a loop, ++ and -- are convenient for counting up or down.

```
for x := 1; x <= 3; x++ {
        fmt.Println(x)
}
```

```
1
2
3
```

```
for x := 3; x >= 1; x-- {
        fmt.Println(x)
}
```

```
3
2
1
```

Go also includes the assignment operators += and -=. They take the value in a variable, add or subtract another value, and then assign the result back to the variable.

```
x := 0
x += 2
fmt.Println(x)
x += 5
fmt.Println(x)
x -= 3
fmt.Println(x)
```

```
2
7
4
```

+= and -= can be used in a loop to count in increments other than 1.

```
for x := 1; x <= 5; x += 2 {
        fmt.Println(x)
}
```

```
1
3
5
```

```
for x := 15; x >= 5; x -= 5 {
        fmt.Println(x)
}
```

```
15
10
5
```

When the loop finishes, execution will resume with whatever statement follows the loop block. But the loop will keep going as long as the condition expression evaluates to `true`. It's possible to abuse this; here are examples of a loop that will run forever, and a loop that will never run at all:

Infinite loop!

```
for x := 1; true; x++ {
        fmt.Println(x)
}
```

Loop that never runs!

```
for x := 1; false; x++ {
        fmt.Println(x)
}
```



Watch it!

It's possible for a loop to run forever, in which case your program will never stop on its own.

If this happens, with the terminal active, hold the Control key and press C to halt your program.

Init and post statements are optional

If you want, you can leave out the init and post statements from a `for` loop, leaving only the condition expression (although you still need to make sure the condition eventually evaluates to `false`, or you could have an infinite loop on your hands).

```
x := 1                          // Declare x in a separate statement.
for x <= 3 {                    // Use only the condition expression.
        fmt.Println(x)
        x++                     // Increment x in a separate statement.
}
```

```
1
2
3
```

```
x := 3                          // Declare x in a separate statement.
for x >= 1 {                    // Use only the condition expression.
        fmt.Println(x)
        x--                     // Decrement x in a separate statement.
}
```

```
3
2
1
```

Loops and scope

Just like with conditionals, the scope of any variables declared within a loop's block is limited to that block (although the init statement, condition expression, and post statement can be considered part of that scope as well).

```
for x := 1; x <= 3; x++ {
        y := x + 1
        fmt.Println(y)          // Still in scope...
}
fmt.Println(y)                  // Error: out of scope!
```

```
undefined: y                    // Error
```

```
for x := 1; x <= 3; x++ {
        fmt.Println(x)          // Still in scope...
}
fmt.Println(x)                  // Error: out of scope!
```

```
undefined: x                    // Error
```

Also as with conditionals, any variable declared *before* the loop will still be in scope within the loop's control statements and block, *and* will still be in scope after the loop exits.

```
var x int                       // Declared outside loop...
for x = 1; x <= 3; x++ {        // No need to declare x here, just assign to it!
        fmt.Println(x)          // Still in scope
}
fmt.Println(x)                  // Still in scope
```

```
1
2
3
4
```



Breaking Stuff is Educational!

Here's a program that uses a loop to count to 3. Try making one of the changes below and run it. Then undo your change and try the next one. See what happens!

```
package main

import "fmt"

func main() {
        for x := 1; x <= 3; x++ {
                fmt.Println(x)
        }
}
```

```
1
2
3
```

If you do this...	...it will break because...
Add parentheses after the `for` keyword `for (x := 1; x <= 3; x++)`	Some other languages *require* parentheses around a `for` loop's control statements, but not only does Go not require them, it doesn't *allow* them.
Delete the `:` from the init statement `x = 1`	Unless you're assigning to a variable that's already been declared in the enclosing scope (which you usually won't be), the init statement needs to be a *declaration*, not an *assignment*.
Remove the = from the condition expression `x < 3`	The expression `x < 3` becomes `false` when `x` reaches `3` (whereas `x <= 3` would still be `true`). So the loop would only count to `2`.
Reverse the comparison in the condition expression `x >= 3`	Because the condition is already `false` when the loop begins (`x` is initialized to `1`, which is *less* than `3`), the loop will never run.
Change the post statement from `x++` to `x--` `x--`	The `x` variable will start counting *down* from `1` (`1`, `0`, `-1`, `-2`, etc.), and since it will never be greater than `3`, the loop will never end.
Move the `fmt.Println(x)` statement outside the loop's block	Variables declared in the init statement or within the loop block are only in scope within the loop block.



Look carefully at the init statement, condition expression, and post statement for each of these loops. Then write what you think the output will be for each one.

(We've done the first one for you.)

```
for x := 1; x <= 3; x++ {
        fmt.Print(x)
}
```

123

```
for x := 3; x >= 1; x-- {
        fmt.Print(x)
}
```

.....

```
for x := 2; x <= 3; x++ {
        fmt.Print(x)
}
```

.....

```
for x := 1; x < 3; x++ {
        fmt.Print(x)
}
```

.....

```
for x := 1; x <= 3; x+= 2 {
        fmt.Print(x)
}
```

.....

```
for x := 1; x >= 3; x++ {
        fmt.Print(x)
}
```

.....

→ Answers on page 78.

Using a loop in our guessing game

Our game still only prompts the user for a guess once. Let's add a loop around the code that prompts the user for a guess and tells them if it was low or high, so that the user can guess 10 times.

We'll use an `int` variable named `guesses` to track the number of guesses the player has made. In our loop's init statement, we'll initialize `guesses` to `0`. We'll add `1` to `guesses` with each iteration of the loop, and we'll stop the loop when `guesses` reaches `10`.

We'll also add a `Println` statement at the top of the loop's block to tell the user how many guesses they have left.

```
// No changes to package and import statements; omitting

func main() {
        seconds := time.Now().Unix()
        rand.Seed(seconds)
        target := rand.Intn(100) + 1
        fmt.Println("I've chosen a random number between 1 and 100.")
        fmt.Println("Can you guess it?")
        fmt.Println(target)

        reader := bufio.NewReader(os.Stdin)

        for guesses := 0; guesses < 10; guesses++ {
                fmt.Println("You have", 10-guesses, "guesses left.")

                fmt.Print("Make a guess: ")
                input, err := reader.ReadString('\n')
                if err != nil {
                        log.Fatal(err)
                }
                input = strings.TrimSpace(input)
                guess, err := strconv.Atoi(input)
                if err != nil {
                        log.Fatal(err)
                }

                if guess < target {
                        fmt.Println("Oops. Your guess was LOW.")
                } else if guess > target {
                        fmt.Println("Oops. Your guess was HIGH.")
                }
        }
}
```

Use the "guesses" variable to track the number of guesses so far.

Subtract the number of guesses from 10 to tell the player how many they have left.

The existing code, which prompts the user for a guess and tells them if it's low or high, will be run 10 times.

End of the for loop

Using a loop in our guessing game (continued)

Now that our loop is in place, if we run our game again, we'll get asked 10 times what our guess is!

We're still set up to print the target number when the game starts.

Inside the loop, we say how many guesses are left, get the player's guess, and tell them if it was low or high.

Right now, players don't get told when their guess is correct, and the loop doesn't stop.

```
Shell  Edit  View  Window  Help
$ go run guess.go
68
I've chosen a random number between 1 and 100.
Can you guess it?
You have 10 guesses left.
Make a guess: 50
Oops. Your guess was LOW.
You have 9 guesses left.
Make a guess: 75
Oops. Your guess was HIGH.
You have 8 guesses left.
Make a guess: 68
You have 7 guesses left.
Make a guess:
```

Since the code to prompt for a guess and state whether it was high or low is inside the loop, it gets run repeatedly. After 10 guesses, the loop (and the game) will end.

But the loop always runs 10 times, even if the player guesses correctly! Fixing that will be our next requirement.

Skipping parts of a loop with "continue" and "break"

The hard part is done! We only have a couple requirements left to go.

Right now, the loop that prompts the user for a guess always runs 10 times. Even if the player guesses correctly, we don't tell them so, and we don't stop the loop. Our next task is to fix that.

- [x] Allow the player to guess up to 10 times. Before each guess, let them know how many guesses they have left.
- [ ] If the player's guess is equal to the target number, tell them, "Good job! You guessed it!" Then stop asking for new guesses.

Go provides two keywords that control the flow of a loop. The first, `continue`, immediately skips to the next iteration of a loop, without running any further code in the loop block.

Skip directly back to the top of the loop.

```
for x := 1; x <= 3; x++ {
        fmt.Println("before continue")
        continue
        fmt.Println("after continue")
}
```

```
before continue
before continue
before continue
```

In the above example, the string `"after continue"` never gets printed, because the `continue` keyword always skips back to the top of the loop before the second call to `Println` can be run.

The second keyword, `break`, immediately breaks out of a loop. No further code within the loop block is executed, and no further iterations are run. Execution moves to the first statement following the loop.

This WOULD loop three times, but the break prevents that.

```
for x := 1; x <= 3; x++ {
        fmt.Println("before break")
        break
        fmt.Println("after break")
}
fmt.Println("after loop")
```

Immediately break out of the loop.

```
before break
after loop
```

Here, in the first iteration of the loop, the string `"before break"` gets printed, but then the `break` statement immediately breaks out of the loop, without printing the `"after break"` string, and without running the loop again (even though it normally would have run two more times). Execution instead moves to the statement following the loop.

The `break` keyword seems like it would be applicable to our current problem: we need to break out of our loop when the player guesses correctly. Let's try using it in our game...

Breaking out of our guessing loop

We're using an `if...else if` conditional to tell the player the status of their guess. If the player guesses a number too high or too low, we currently print a message telling them so.

It stands to reason that if the guess is neither too high *nor* too low, it must be correct. So let's add an `else` branch onto the conditional, that will run in the event of a correct guess. Inside the block for the `else` branch, we'll tell the player they were right, and then use the `break` statement to stop the guessing loop.

```
// No changes to package and import statements; omitting

func main() {
        // No changes to previous code; omitting

        for guesses := 0; guesses < 10; guesses++ {
                // No changes to previous code; omitting

                if guess < target {
                        fmt.Println("Oops. Your guess was LOW.")
                } else if guess > target {
                        fmt.Println("Oops. Your guess was HIGH.")
                } else {
                        fmt.Println("Good job! You guessed it!")
                        break
                }
        }
}
```

Congratulate the player.

Break out of the loop.

Now, when the player guesses correctly, they'll see a congratulatory message, and the loop will exit without repeating the full 10 times.

Here's the target; we'll cheat and make a correct guess immediately.

We get congratulated, and the loop exits!

```
Shell Edit View Window Help
$ go run guess.go
48
I've chosen a random number between 1 and 100.
Can you guess it?
You have 10 guesses left.
Make a guess: 48
Good job! You guessed it!
$
```

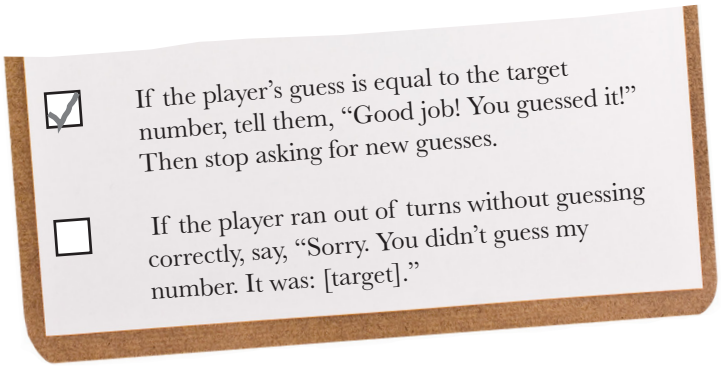
That's another requirement complete!

Revealing the target

We're *so* close! Just one more requirement left!

If the player makes 10 guesses without finding the target number, the loop will exit. In that event, we need to print a message saying they lost, and tell them what the target was.

But we *also* exit the loop if the player guesses correctly. We don't want to say the player has lost when they've already won!



So, before our guessing loop, we'll declare a `success` variable that holds a `bool`. (We need to declare it *before* the loop so that it's still in scope after the loop ends.) We'll initialize `success` to a default value of `false`. Then, if the player guesses correctly, we'll set `success` to `true`, indicating we don't need to print the failure message.

```
// No changes to package and import statements; omitting

func main() {
	// No changes to previous code; omitting

	success := false
	for guesses := 0; guesses < 10; guesses++ {
		// No changes to previous code; omitting

		if guess < target {
			fmt.Println("Oops. Your guess was LOW.")
		} else if guess > target {
			fmt.Println("Oops. Your guess was HIGH.")
		} else {
			success = true
			fmt.Println("Good job! You guessed it!")
			break
		}
	}

	if !success {
		fmt.Println("Sorry, you didn't guess my number. It was:", target)
	}
}
```

Declare "success" before the loop, so it's still in scope after the loop exits.

If the player guesses correctly, indicate we don't need to print the failure message.

If the player was NOT successful (if "success" is false)...

...print the failure message.

After the loop, we add an `if` block that prints the failure message. But an `if` block only runs if its condition evaluates to `true`, and we only want to print the failure message if `success` is `false`. So we add the Boolean negation operator (`!`). As we saw earlier, `!` turns `true` values `false` and `false` values `true`.

The result is that the failure message will be printed if `success` is `false`, but *won't* be printed if `success` is `true`.

The finishing touches

Congratulations, that's the last requirement!

- ☑ If the player ran out of turns without guessing correctly, say, "Sorry. You didn't guess my number. It was: [target]."

Let's take care of a couple final issues with our code, and then try out our game!

First, as we mentioned, it's typical to add a comment at the top of each Go program describing what it does. Let's add one now.

```
// guess challenges players to guess a random number.
package main
...
```

Add a program description comment, above the package clause.

Our program is also encouraging cheaters by printing the target number at the start of every game. Let's remove the `Println` call that does that.

```
fmt.Println("I've chosen a random number between 1 and 100.")
fmt.Println("Can you guess it?")
~~fmt.Println(target)~~
```

Don't reveal the target at the start of each game.

We're finally ready to try running our complete code!

First, we'll run out of guesses on purpose to ensure the target number gets displayed...

```
Shell Edit View Window Help
$ go run guess.go
I've chosen a random number between 1 and 100.
Can you guess it?
You have 10 guesses left.
Make a guess: 10
Oops. Your guess was LOW.
You have 9 guesses left.
Make a guess: 20
Oops. Your guess was LOW.
...
You have 1 guesses left.
Make a guess: 62
Oops. Your guess was LOW.
Sorry, you didn't guess my number. It was: 63
```

Other incorrect guesses omitted...

If we run out of guesses, the correct number is revealed.

Then we'll try guessing successfully.

Our game is working great!

```
Shell Edit View Window Help Cheats
$ go run guess.go
I've chosen a random number between 1 and 100.
Can you guess it?
You have 10 guesses left.
Make a guess: 50
Oops. Your guess was HIGH.
You have 9 guesses left.
Make a guess: 40
Oops. Your guess was LOW.
You have 8 guesses left.
Make a guess: 45
Good job! You guessed it!
```

If we guess correctly, we see the victory message!

Congratulations, your game is complete!



Using conditionals and loops, you've written a complete game in Go! Pour yourself a cold drink—you've earned it!

Here's our complete guess.go source code!

```
// guess challenges players to guess a random number.
package main

import (
        "bufio"
        "fmt"
        "log"
        "math/rand"
        "os"
        "strconv"
        "strings"
        "time"
)

func main() {
        seconds := time.Now().Unix()
        rand.Seed(seconds)
        target := rand.Intn(100) + 1
        fmt.Println("I've chosen a random number between 1 and 100.")
        fmt.Println("Can you guess it?")

        reader := bufio.NewReader(os.Stdin)
        success := false
        for guesses := 0; guesses < 10; guesses++ {
                fmt.Println("You have", 10-guesses, "guesses left.")
                fmt.Print("Make a guess: ")
                input, err := reader.ReadString('\n')
                if err != nil {
                        log.Fatal(err)
                }
                input = strings.TrimSpace(input)
                guess, err := strconv.Atoi(input)
                if err != nil {
                        log.Fatal(err)
                }

                if guess < target {
                        fmt.Println("Oops. Your guess was LOW.")
                } else if guess > target {
                        fmt.Println("Oops. Your guess was HIGH.")
                } else {
                        success = true
                        fmt.Println("Good job! You guessed it!")
                        break
                }
        }

        if !success {
                fmt.Println("Sorry, you didn't guess my number. It was:", target)
        }
}
```

Import all the packages that we use in the code below.

Get the current date and time, as an integer.

Seed the random number generator.

Generate an integer between 1 and 100.

Create a bufio.Reader, which lets us read keyboard input.

Set up to print a failure message by default.

Ask for a number.

Read what the user types, up until they press Enter.

If there's an error, print the message and exit.

Remove the newline.

Convert the input string to an integer.

If there's an error, print the message and exit.

If the guess was too low, say so.

If the guess was too high, say so.

Otherwise, the guess must be correct...

Prevent the failure message from displaying.

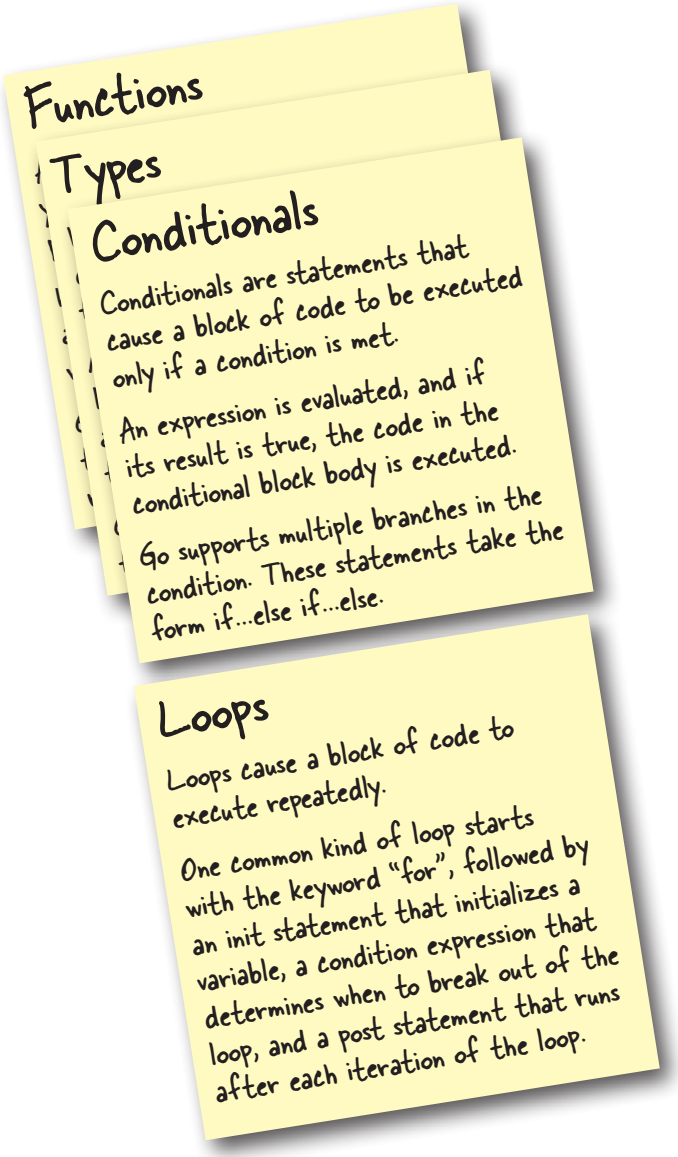
Exit the loop.

If "success" is false, tell player what the target was.



Your Go Toolbox

That's it for Chapter 2! You've added conditionals and loops to your toolbox.



BULLET POINTS

- A **method** is a kind of function that's associated with values of a given type.
- Go treats everything from a `//` marker to the end of the line as a **comment**—and ignores it.
- Multiline comments start with `/*` and end with `*/`. Everything in between, including newlines, is ignored.
- It's conventional to include a comment at the top of every program, explaining what it does.
- Unlike most programming languages, Go allows *multiple* return values from a function or method call.
- One common use of multiple return values is to return the function's main result, and then a second value indicating whether there was an error.
- To discard a value without using it, use the `_` **blank identifier**. The blank identifier can be used in place of any variable in any assignment statement.
- Avoid giving variables the same name as types, functions, or packages; it causes the variable to **shadow** (override) the item with the same name.
- Functions, conditionals, and loops all have **blocks** of code that appear within `{}` braces.
- Their code doesn't appear within `{}` braces, but files and packages also comprise blocks.
- The **scope** of a variable is limited to the block it is defined within, and all blocks nested within that block.
- In addition to a name, a package may have an **import path** that is required when it is imported.
- The `continue` keyword skips to the next iteration of a loop.
- The `break` keyword exits out of a loop entirely.



Because they're in conditional blocks, only some of the `Println` calls in the code below will be executed. Write down what the output would be.

```
if true {                          "if" blocks run if the condition results in true (or if it IS true).
        fmt.Println("true")
}
if false {                         If the condition is false, the block doesn't run.
        fmt.Println("false")
}
if !false {                        The Boolean negation operator turns false into true.
        fmt.Println("!false")
}
if true {                          The "if" branch runs...
        fmt.Println("if true")
} else {                           ...so the "else" branch doesn't.
        fmt.Println("else")
}
if false {                         The "if" branch doesn't run...
        fmt.Println("if false")
} else if true {                   ...so the "else if" branch MIGHT run.
        fmt.Println("else if true")
}
if 12 == 12 {                      12 == 12 is true.
        fmt.Println("12 == 12")
}
if 12 != 12 {                      The values ARE equal, so this is false.
        fmt.Println("12 != 12")
}
if 12 > 12 {                       12 is NOT greater than itself...
        fmt.Println("12 > 12")
}
if 12 >= 12 {                      ...But 12 IS equal to itself.
        fmt.Println("12 >= 12")
}
if 12 == 12 && 5.9 == 5.9 {        The && evaluates to true if BOTH expressions are true.
        fmt.Println("12 == 12 && 5.9 == 5.9")
}
if 12 == 12 && 5.9 == 6.4 {        One expression is false.
        fmt.Println("12 == 12 && 5.9 == 6.4")
}
if 12 == 12 || 5.9 == 6.4 {        The || evaluates to true if EITHER expression is true.
        fmt.Println("12 == 12 || 5.9 == 6.4")
}
```

Output:

true

!false

if true

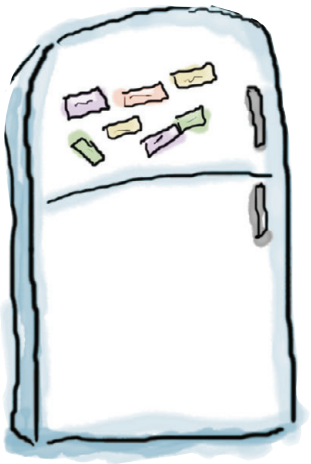
else if true

12 == 12

12 >= 12

12 == 12 && 5.9 == 5.9

12 == 12 || 5.9 == 6.4



Code Magnets Solution

A Go program that prints the size of a file is on the fridge. It calls the `os.Stat` function, which returns an `os.FileInfo` value, and possibly an error. Then it calls the `Size` method on the `FileInfo` value to get the file size.

The original program used the _ blank identifier to ignore the error value from `os.Stat`. If an error occurred (which could happen if the file doesn't exist), this would cause the program to fail.

Your job was to reconstruct the extra code snippets to make a program that works just like the original one, but also checks for an error from `os.Stat`. If the error from `os.Stat` is not `nil`, the error should be reported, and the program should exit.

```
package main

import (
        "fmt"
        "log"
        "os"
)

func main() {
    fileInfo, err := os.Stat("my.txt")
    if err != nil {
        log.Fatal(err)
    }
    fmt.Println(fileInfo.Size())
}
```

Store any error we get from os.Stat.

Get a FileInfo value with data regarding the my.txt file.

Holds the file size, date it was changed, etc.

If the error is not nil, pass it to log.Fatal.

Returns the size of the file.

Discard this magnet; we don't need the blank identifier anymore!

```
_
```



Some of the lines of code below will result in a compile error, because they refer to a variable that is out of scope. Cross out the lines that have errors.

```
package main

import (
        "fmt"
)

var a = "a"

func main() {
        a = "a"
        b := "b"
        if true {
                c := "c"
                if true {
                        d := "d"
                        fmt.Println(a)
                        fmt.Println(b)
                        fmt.Println(c)
                        fmt.Println(d)
                }
                fmt.Println(a)
                fmt.Println(c)
                fmt.Println(d)
        }
        fmt.Println(a)
        fmt.Println(b)
        fmt.Println(c)
        fmt.Println(d)
}
```



Exercise Solution

Look carefully at the init statement, condition expression, and post statement for each of these loops. Then write what you think the output will be for each one.

Start at 1. Stop after 3. Count up.

```
for x := 1; x <= 3; x++ {
        fmt.Print(x)
}
```

123

Start at 3. Stop after 1. Count down.

```
for x := 3; x >= 1; x-- {
        fmt.Print(x)
}
```

321

Start at 2. Stop after 3. Count up.

```
for x := 2; x <= 3; x++ {
        fmt.Print(x)
}
```

23

Start at 1. Stop at 3. Count up.

```
for x := 1; x < 3; x++ {
        fmt.Print(x)
}
```

12

Start at 1. Stop after 3. Count up 2 at a time.

```
for x := 1; x <= 3; x+= 2 {
        fmt.Print(x)
}
```

13

Start at 1. Stop when x < 3 (i.e., immediately). Never runs!

```
for x := 1; x >= 3; x++ {
        fmt.Print(x)
}
```

No output; loop never runs!

3 call me

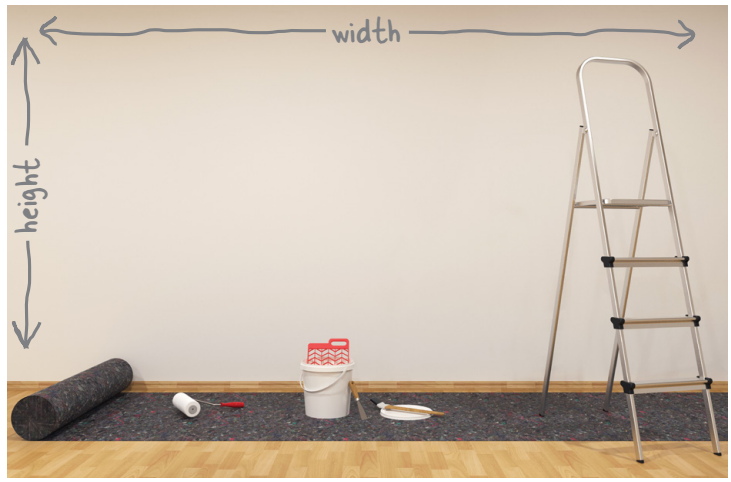
Functions



You've been missing out. You've been calling functions like a pro. But the only functions you could call were the ones Go defined for you. Now, it's your turn. We're going to show you how to create your own functions. We'll learn how to declare functions with and without parameters. We'll declare functions that return a single value, and we'll learn how to return multiple values so that we can indicate when there's been an error. And we'll learn about **pointers**, which allow us to make more memory-efficient function calls.

Some repetitive code

Suppose we need to calculate the amount of paint needed to cover several walls. The manufacturer says each liter of paint covers 10 square meters. So, we'll need to multiply each wall's width (in meters) by its height to get its area, and then divide that by 10 to get the number of liters of paint needed.




```
// package and imports omitted
func main() {
        var width, height, area float64
        width = 4.2
        height = 3.0
        area = width * height
        fmt.Println(area/10.0, "liters needed")
        width = 5.2
        height = 3.5
        area = width * height
        fmt.Println(area/10.0, "liters needed")
}
```

Calculate the amount for a first wall.

Determine the area of the wall.

Calculate how much paint is needed for that area.

Do the same for a second wall.

Determine the area of the wall.

Calculate how much paint is needed for that area.

```
1.2600000000000002 liters needed
1.8199999999999998 liters needed
```

This works, but it has a couple problems:

- The calculations seem to be off by a tiny fraction, and are printing oddly precise floating-point values. We really only need a couple decimal places of precision.
- There's a fair amount of repeated code, even now. This will get worse as we add more walls.

Both items will take a little explanation to address, so let's just look at the first issue for now...

The calculations are slightly off because ordinary floating-point arithmetic on computers is ever-so-slightly inaccurate. (Usually by a few quadrillionths.) The reasons are a little too complicated to get into here, but this problem isn't exclusive to Go.

But as long as we round the numbers to a reasonable degree of precision before displaying them, we should be fine. Let's take a brief detour to look at a function that will help us do that.



Formatting output with Printf and Sprintf



Floating-point numbers in Go are kept with a high degree of precision. This can be cumbersome when you want to display them:

```
fmt.Println("About one-third:", 1.0/3.0)
```

```
About one-third: 0.3333333333333333
```

That's a lot of decimal places!

To deal with these sorts of formatting issues, the `fmt` package provides the `Printf` function. `Printf` stands for "**print**, with **f**ormatting." It takes a string and inserts one or more values into it, formatted in specific ways. Then it prints the resulting string.

```
fmt.Printf("About one-third: %0.2f\n", 1.0/3.0)
```

```
About one-third: 0.33
```

Much more readable!

The `Sprintf` function (also part of the `fmt` package) works just like `Printf`, except that it returns a formatted string instead of printing it.

```
resultString := fmt.Sprintf("About one-third: %0.2f\n", 1.0/3.0)
fmt.Printf(resultString)
```

```
About one-third: 0.33
```

It looks like `Printf` and `Sprintf` *can* help us limit our displayed values to the correct number of places. The question is, *how*? First, to be able to use the `Printf` function effectively, we'll need to learn about two of its features:

- Formatting verbs (the `%0.2f` in the strings above is a verb)
- Value widths (that's the `0.2` in the middle of the verb)



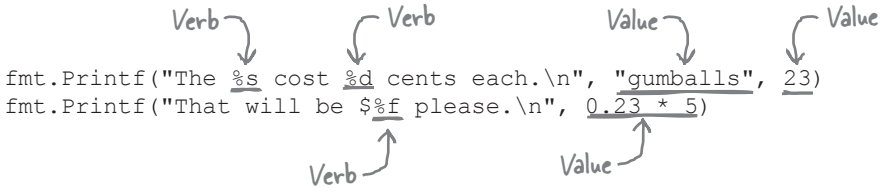
We'll explain exactly what those arguments to `Printf` mean on the next few pages.

We know, those function calls above look a little confusing. We'll show you a ton of examples that should clear that confusion up.

Formatting verbs



The first argument to `Printf` is a string that will be used to format the output. Most of it is formatted exactly as it appears in the string. Any percent signs (`%`), however, will be treated as the start of a **formatting verb**, a section of the string that will be substituted with a value in a particular format. The remaining arguments are used as values with those verbs.




```
The gumballs cost 23 cents each.
That will be $1.150000 please.
```

We'll show how to fix this shortly.

The letter following the percent sign indicates which verb to use. The most common verbs are:

Verb	Output
`%f`	Floating-point number
`%d`	Decimal integer
`%s`	String
`%t`	Boolean (`true` or `false`)
`%v`	Any value (chooses an appropriate format based on the supplied value's type)
`%#v`	Any value, formatted as it would appear in Go program code
`%T`	Type of the supplied value (`int`, `string`, etc.)
`%%`	A literal percent sign

```
fmt.Printf("A float: %f\n", 3.1415)
fmt.Printf("An integer: %d\n", 15)
fmt.Printf("A string: %s\n", "hello")
fmt.Printf("A boolean: %t\n", false)
fmt.Printf("Values: %v %v %v\n", 1.2, "\t", true)
fmt.Printf("Values: %#v %#v %#v\n", 1.2, "\t", true)
fmt.Printf("Types: %T %T %T\n", 1.2, "\t", true)
fmt.Printf("Percent sign: %%\n")
```

```
A float: 3.141500
An integer: 15
A string: hello
A boolean: false
Values: 1.2     true
Values: 1.2 "\t" true
Types: float64 string bool
Percent sign: %
```

Notice, by the way, that we are making sure to add a newline at the end of each formatting string using the `\n` escape sequence. This is because unlike `Println`, `Printf` does not automatically add a newline for us.

Formatting verbs (continued)



We want to point out the `%#v` formatting verb in particular. Because it prints values the way they would appear in Go code, rather than how they normally appear, `%#v` can show you some values that would otherwise be hidden in your output. In this code, for example, `%#v` reveals an empty string, a tab character, and a newline, all of which were invisible when printed with `%v`. We'll use `%#v` more, later in the book!

```
fmt.Printf("%v %v %v", "", "\t", "\n")
fmt.Printf("%#v %#v %#v", "", "\t", "\n")
```

%v prints all the values...

```
 "" "\t" "\n"
```

...but only with %#v can you actually see them!

Formatting value widths

So the `%f` formatting verb is for floating-point numbers. We can use `%f` in our program to format the amount of paint needed.

Insert a floating-point value.

One of the values previously calculated by our program

```
fmt.Printf("%f liters needed\n", 1.8199999999999998)
```

```
1.820000 liters needed
```

Rounded, but still too many digits!

It looks like our value is being rounded to a reasonable number. But it's still showing six places after the decimal point, which is really too much for our current purpose.

For situations like this, formatting verbs let you specify the *width* of the formatted value.

Let's say we want to format some data in a plain-text table. We need to ensure the formatted value fills a minimum number of spaces, so that the columns align properly.

You can specify the minimum width after the percent sign for a formatting verb. If the argument matching that verb is shorter than the minimum width, it will be padded with spaces until the minimum width is reached.

The first field will have a minimum width of 12 characters.

No minimum width for this second field

Print column headings.

```
fmt.Printf("%12s | %s\n", "Product", "Cost in Cents")
fmt.Println("----------------------------") // Print a heading divider.
fmt.Printf("%12s | %2d\n", "Stamps", 50)
fmt.Printf("%12s | %2d\n", "Paper Clips", 5)
fmt.Printf("%12s | %2d\n", "Tape", 99)
```

Minimum width of 12 again

Minimum width of 2

Padding!

```
     Product | Cost in Cents
-----------------------------
      Stamps | 50
 Paper Clips |  5
        Tape | 99
```

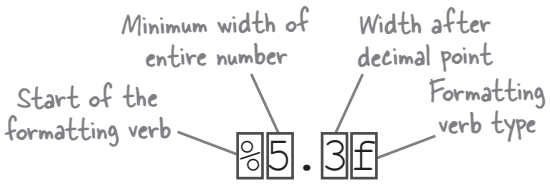
No padding; the value already fills the minimum width.

Padding!

Formatting fractional number widths



And now we come to the part that's important for today's task: you can use value widths to specify the precision (the number of displayed digits) for floating-point numbers. Here's the format:



The minimum width of the entire number includes decimal places and the decimal point. If it's included, shorter numbers will be padded with spaces at the start until this width is reached. If it's omitted, no spaces will ever be added.

The width after the decimal point is the number of decimal places to show. If a more precise number is given, it will be rounded (up or down) to fit in the given number of decimal places.

Here's a quick demonstration of various width values in action:

Doesn't display an actual value; just shows what the verb is

These display the actual values.

```
fmt.Printf("%%7.3f: %7.3f\n", 12.3456)
fmt.Printf("%%7.2f: %7.2f\n", 12.3456)
fmt.Printf("%%7.1f: %7.1f\n", 12.3456)
fmt.Printf("%%.1f: %.1f\n", 12.3456)
fmt.Printf("%%.2f: %.2f\n", 12.3456)
```

```
%7.3f:  12.346
%7.2f:   12.35
%7.1f:    12.3
%.1f: 12.3
%.2f: 12.35
```

Rounded to three places
Rounded to two places
Rounded to one place
Rounded to one place, no padding
Rounded to two places, no padding

That last format, `"%.2f"`, will let us take floating-point numbers of any precision and round them to two decimal places. (It also won't do any unnecessary padding.) Let's try it with the overly precise values from our program to calculate paint volumes.

Rounded to two places!

```
fmt.Printf("%.2f\n", 1.2600000000000002)
fmt.Printf("%.2f\n", 1.8199999999999998)
```

```
1.26
1.82
```

That's much more readable. It looks like the `Printf` function can format our numbers for us. Let's get back to our paint calculator program, and apply what we've learned there.



Using Printf in our paint calculator

Now we have a `Printf` verb, `"%.2f"`, that will let us round a floating-point number to two decimal places. Let's update our paint quantity calculation program to use it.

```
// package and imports omitted
func main() {
        var width, height, area float64
        width = 4.2
        height = 3.0
        area = width * height
        fmt.Printf("%.2f liters needed\n", area/10.0)
        width = 5.2
        height = 3.5
        area = width * height
        fmt.Printf("%.2f liters needed\n", area/10.0)
}
```

Format the value and insert it into the string.

Do the same here!

```
1.26 liters needed
1.82 liters needed
```

Rounded to two places

At last, we have reasonable-looking output! The tiny imprecisions introduced by floating-point arithmetic have been rounded away.



Wasn't it kind of a pain to update that code in two places, though? If you change it, will you remember to update both lines? And what happens when we add more walls?

Good point. Go lets us declare our own functions, so perhaps we should move this code into a function.

As we mentioned way back at the start of Chapter 1, a function is a group of one or more lines of code that you can call from other places in your program. And our program has two groups of lines that look very similar:

```
var width, height, area float64
width = 4.2
height = 3.0
area = width * height
fmt.Printf("%.2f liters needed\n", area/10.0)
```

Calculate the paint needed for the first wall.

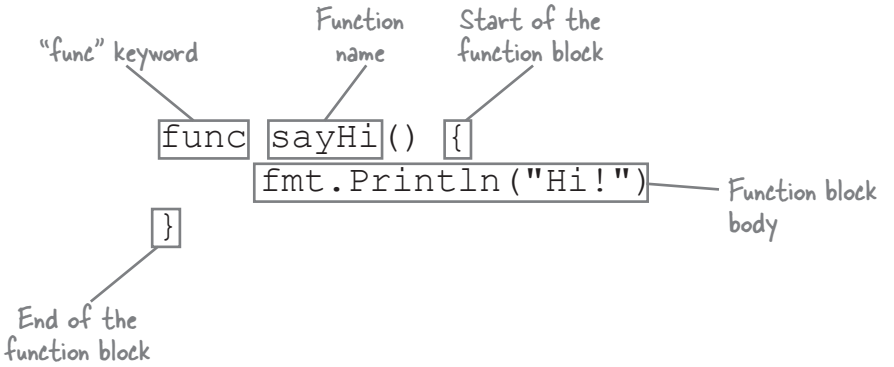
```
width = 5.2
height = 3.5
area = width * height
fmt.Printf("%.2f liters needed\n", area/10.0)
```

Calculate the paint needed for the second wall.

Let's see if we can convert these two sections of code into a single function.

Declaring functions

A simple function declaration might look like this:



A declaration begins with the `func` keyword, followed by the name you want the function to have, a pair of parentheses `()`, and then a block containing the function's code.

Once you've declared a function, you can call it elsewhere in your package simply by typing its name, followed by a pair of parentheses. When you do, the code in the function's block will be run.

Notice that when we call `sayHi`, we're not typing the package name and a dot before the function name. When you call a function that's defined in the current package, you should not specify the package name. (Typing `main.sayHi()` would result in a compile error.)

```
package main

import "fmt"

func sayHi() {
        fmt.Println("Hi!")
}

func main() {
        sayHi()
}
```

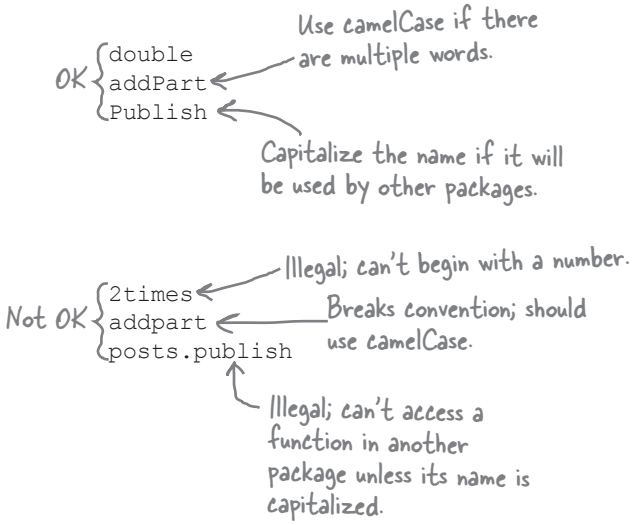
Declare a "sayHi" function.

Call "sayHi".

```
Hi!
```

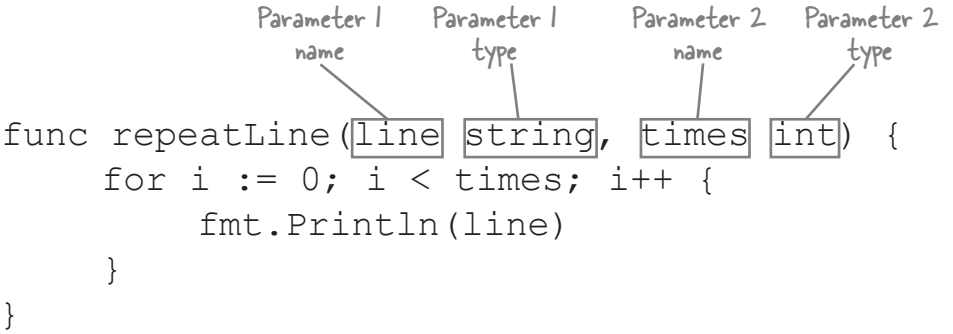
The rules for function names are the same as the rules for variable names:

- A name must begin with a letter, followed by any number of additional letters and numbers. (You'll get a compile error if you break this rule.)
- Functions whose name begins with a capital letter are *exported*, and can be used outside the current package. If you only need to use a function inside the current package, you should start its name with a lowercase letter.
- Names with multiple words should use `camelCase`.



Declaring function parameters

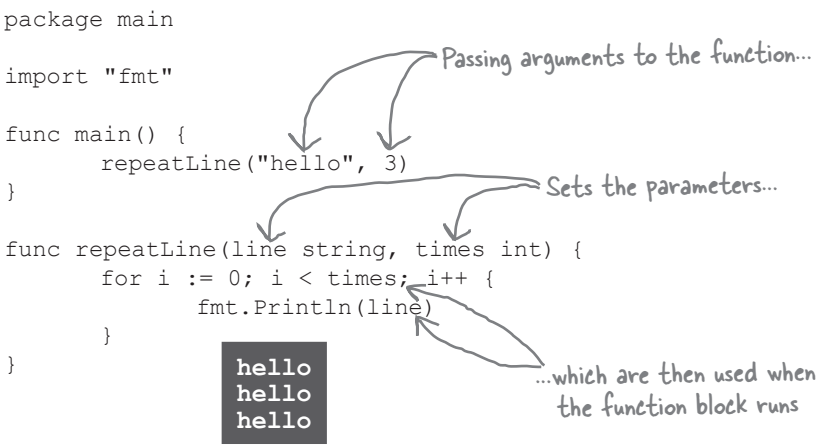
If you want calls to your function to include arguments, you'll need to declare one or more parameters. A **parameter** is a variable, local to a function, whose value is set when the function is called.



You can declare one or more parameters between the parentheses in the function declaration, separated by commas. As with any variable, you'll need to provide a name followed by a type (`float64`, `bool`, etc.) for each parameter you declare.

If a function has parameters defined, then you'll need to pass a matching set of arguments when calling it. When the function is run, each parameter will be set to a copy of the value in the corresponding argument. Those parameter values are then used within the code in the function block.

A parameter is a variable, local to a function, whose value is set when the function is called.



Using functions in our paint calculator

Now that we know how to declare our own functions, let's see if we can get rid of the repetition in our paint calculator.

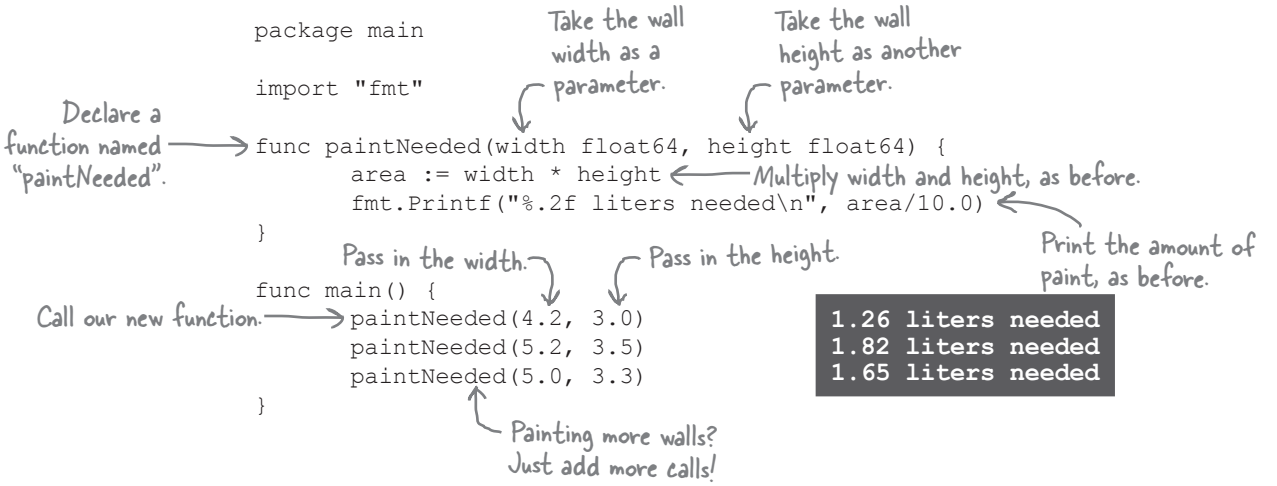
```
// package and imports omitted
func main() {
    var width, height, area float64
    width = 4.2
    height = 3.0
    area = width * height
    fmt.Printf("%.2f liters needed\n", area/10.0)
    width = 5.2
    height = 3.5
    area = width * height
    fmt.Printf("%.2f liters needed\n", area/10.0)
}
```

Repeated code!

Repeated code!

```
1.26 liters needed
1.82 liters needed
```

We'll move the code to calculate the amount of paint to a function named `paintNeeded`. We'll get rid of the separate `width` and `height` variables, and instead take those as function parameters. Then, in our `main` function, we'll just call `paintNeeded` for each wall we need to paint.



No more repeated code, and if we want to calculate the paint needed for additional walls, we just add more calls to `paintNeeded`. This is much cleaner!



Below is a program that declares several functions, then calls those functions within `main`. Write down what the program output would be.

(We've done the first line for you.)

```
package main

import "fmt"

func functionA(a int, b int) {
       fmt.Println(a + b)
}
func functionB(a int, b int) {
       fmt.Println(a * b)
}
func functionC(a bool) {
       fmt.Println(!a)
}
func functionD(a string, b int) {
       for i := 0; i < b; i++ {
              fmt.Print(a)
       }
       fmt.Println()
}

func main() {
       functionA(2, 3)
       functionB(2, 3)
       functionC(true)
       functionD("$", 4)
       functionA(5, 6)
       functionB(5, 6)
       functionC(false)
       functionD("ha", 3)
}
```

Output:

5

→ Answers on page 111.

Functions and variable scope

Our `paintNeeded` function declares an `area` variable within its function block:

```
func paintNeeded(width float64, height float64) {
       area := width * height
       fmt.Printf("%.2f liters needed\n", area/10.0)
}
```

Declare an "area" variable.

Access the variable.

As with conditional and loop blocks, variables declared within a function block are only in scope within that function block. So if we were to try to access the `area` variable outside of the `paintNeeded` function, we'd get a compile error:

```
func paintNeeded(width float64, height float64) {
       area := width * height
       fmt.Printf("%.2f liters needed\n", area/10.0)
}

func main() {
       paintNeeded(4.2, 3.0)
       fmt.Println(area)
}
```

Error → `undefined: area`

Out of scope!

But, also as with conditional and loop blocks, variables declared *outside* a function block will be in scope within that block. That means we can declare a variable at the package level, and access it within any function in that package.

```
package main

import "fmt"

var metersPerLiter float64

func paintNeeded(width, height float64) float64 {
       area := width * height
       return area / metersPerLiter
}

func main() {
       metersPerLiter = 10.0
       fmt.Printf("%.2f", paintNeeded(4.2, 3.0))
}
```

If we declared a variable at the package level...

...still in scope here

...still in scope here

```
1.26
```

Function return values

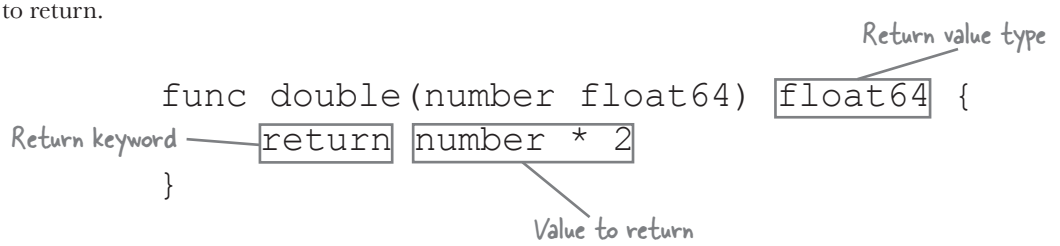
Suppose we wanted to total the amount of paint needed for all the walls we're going to paint. We can't do that with our current `paintNeeded` function; it just prints the amount and then discards it!

```
func paintNeeded(width float64, height float64) {
      area := width * height
      fmt.Printf("%.2f liters needed\n", area/10.0)
}
```

Prints the amount of paint, but then we can't do anything further with it!

So instead, let's revise the `paintNeeded` function to return a value. Then, whoever calls it can print the amount, do additional calculations with it, or do whatever else they need.

Functions always return values of a specific type (and only that type). To declare that a function returns a value, add the type of that return value following the parameters in the function declaration. Then use the `return` keyword in the function block, followed by the value you want to return.



Callers of the function can then assign the return value to a variable, pass it directly to another function, or do whatever else they need to do with it.

```
package main

import "fmt"

func double(number float64) float64 {
      return number * 2
}

func main() {
      dozen := double(6.0)
      fmt.Println(dozen)
      fmt.Println(double(4.2))
}
```

Assign return value to a variable.

Pass return value to another function.

```
12
8.4
```

Function return values (continued)

When a `return` statement runs, the function exits immediately, without running any code that follows it. You can use this together with an `if` statement to exit the function in conditions where there's no point in running the remaining code (due to an error or some other condition).

```
func status(grade float64) string {
        if grade < 60.0 {
                return "failing"
        }
        return "passing"
}

func main() {
        fmt.Println(status(60.1))
        fmt.Println(status(59))
}
```

If grade is failing, return immediately.

Only runs if grade is >= 60

```
passing
failing
```

That means that it's possible to have code that never runs under any circumstances, if you include a `return` statement that isn't part of an `if` block. This almost certainly indicates a bug in the code, so Go helps you detect this situation by requiring that any function that declares a return type must end with a `return` statement. Ending with any other statement will cause a compile error.

```
func double(number float64) float64 {
        return number * 2
        fmt.Println(number * 2)
}
```

Function would always exit here...

This line would never run!

Error →

```
missing return at end of function
```

You'll also get a compile error if the type of your return value doesn't match the declared return type.

Expects a floating-point number...

```
func double(number float64) float64 {
        return int(number * 2)
}
```

...returns an integer!

Error →

```
cannot use int(number * 2) (type int)
as type float64 in return argument
```

Using a return value in our paint calculator

Now that we know how to use function return values, let's see if we can update our paint program to print the total amount of paint needed in addition to the amount needed for each wall.

We'll update the `paintNeeded` function to return the amount needed. We'll use that return value in the `main` function, both to print the amount for the current wall, and to add to a `total` variable that tracks the total amount of paint needed.

```
package main

import "fmt"

func paintNeeded(width float64, height float64) float64 {
        area := width * height
        return area / 10.0
}

func main() {
        var amount, total float64
        amount = paintNeeded(4.2, 3.0)
        fmt.Printf("%0.2f liters needed\n", amount)
        total += amount
        amount = paintNeeded(5.2, 3.5)
        fmt.Printf("%0.2f liters needed\n", amount)
        total += amount
        fmt.Printf("Total: %0.2f liters\n", total)
}
```

Declare that paintNeeded will return a floating-point number.

Return the area instead of printing it.

Declare variables to hold the amount for the current wall, as well as the total for all walls.

Call paintNeeded, and store the return value.

Print the amount for this wall.

Add the amount for this wall to the total.

Repeat the above steps for a second wall.

Print the total for all walls.

```
1.26 liters needed
1.82 liters needed
Total: 3.08 liters
```

It works! Returning the value allowed our `main` function to decide what to do with the calculated amount, rather than relying on the `paintNeeded` function to print it.



Breaking Stuff is Educational!

Here's our updated version of the `paintNeeded` function that returns a value. Try making one of the changes below and try to compile it. Then undo your change and try the next one. See what happens!

```
func paintNeeded(width float64, height float64) float64 {
        area := width * height
        return area / 10.0
}
```

If you do this...	...it will break because...

Remove the `return` statement:
`func paintNeeded(width float64, height float64) float64 {`
`  area := width * height`
~~`  return area / 10.0`~~
`}`
	If your function declares a return type, Go requires that it include a `return` statement.

Add a line *after* the `return` statement:
`func paintNeeded(width float64, height float64) float64 {`
`  area := width * height`
`  return area / 10.0`
`  fmt.Println(area / 10.0)`
`}`
	If your function declares a return type, Go requires that its last statement be a `return` statement.

Remove the return type declaration:
`func paintNeeded(width float64, height float64)` ~~`float64`~~ `{`
`  area := width * height`
`  return area / 10.0`
`}`
	Go doesn't allow you to return a value you haven't declared.

Change the type of value being returned:
`func paintNeeded(width float64, height float64) float64 {`
`  area := width * height`
`  return int(area / 10.0)`
`}`
	Go requires that the type of the returned value match the declared type.

The paintNeeded function needs error handling



Your `paintNeeded` function works great, most of the time. But one of our users recently passed it a negative number by accident, and got a **negative** amount of paint back!

If we accidentally pass in a negative number...

```
func main() {
        amount := paintNeeded(4.2, -3.0)
        fmt.Printf("%0.2f liters needed\n", amount)
}

func paintNeeded(width float64, height float64) float64 {
        area := width * height
        return area / 10.0
}
```

4.2 * -3.0 is -12.6!

-12.6 / 10.0 is -1.26!

```
-1.26 liters needed
```

It looks like the `paintNeeded` function had no idea the argument passed to it was invalid. It went right ahead and used that invalid argument in its calculations, and returned an invalid result. This is a problem—even if you knew a store where you could purchase a negative number of liters of paint, would you really want to apply that to your house? We need a way of detecting invalid arguments and reporting an error.

In Chapter 2, we saw a couple different functions that, in addition to their main return value, also return a second value indicating whether there was an error. The `strconv.Atoi` function, for example, attempted to convert a string to an integer. If the conversion was successful, it returned an error value of `nil`, meaning our program could proceed. But if the error value *wasn't* `nil`, it meant the string couldn't be converted to a number. In that event, we chose to print the error value and exit the program.

```
guess, err := strconv.Atoi(input)
if err != nil {
        log.Fatal(err)
}
```

Convert the input string to an integer.

If there was an error, print the message and exit.

If we want to do the same when calling the `paintNeeded` function, we're going to need two things:

- The ability to create a value representing an error
- The ability to return an additional value from `paintNeeded`

Let's get started figuring this out!

Error values

Before we can return an error value from our `paintNeeded` function, we need an error value to return. An error value is any value with a method named `Error` that returns a string. The simplest way to create one is to pass a string to the `errors` package's `New` function, which will return a new error value. If you call the `Error` method on that error value, you'll get the string you passed to `errors.New`.

```
package main

import (
       "errors"
       "fmt"
)

func main() {
       err := errors.New("height can't be negative")
       fmt.Println(err.Error())
}
```

Create a new error value

Returns the error message

```
height can't be negative
```

But if you're passing the error value to a function in the `fmt` or `log` packages, you probably don't need to call its `Error` method. Functions in `fmt` and `log` have been written to check whether the values passed to them have `Error` methods, and print the return value of `Error` if they do.

```
err := errors.New("height can't be negative")
fmt.Println(err)
log.Fatal(err)
```

Prints the error message

Prints the error message again, then exits the program

```
height can't be negative
2018/03/12 19:49:27 height can't be negative
```

If you need to format numbers or other values for use in your error message, you can use the `fmt.Errorf` function. It inserts values into a format string just like `fmt.Printf` or `fmt.Sprintf`, but instead of printing or returning a `string`, it returns an error value.

Insert a floating-point number, rounded to two decimal places.

Returns an error value

Prints the error message

Also prints the error message

```
err := fmt.Errorf("a height of %0.2f is invalid", -2.33333)
fmt.Println(err.Error())
fmt.Println(err)
```

```
a height of -2.33 is invalid
a height of -2.33 is invalid
```

Declaring multiple return values

Now we need a way to specify that our `paintNeeded` function will return an error value along with the amount of paint needed.

To declare multiple return values for a function, place the return value types in a *second* set of parentheses in the function declaration (after the parentheses for the function parameters), separated with commas. (The parentheses around the return values are optional when there's only one return value, but are required if there's more than one return value.)

From then on, when calling that function, you'll need to account for the additional return values, usually by assigning them to additional variables.

```
package main

import "fmt"

func manyReturns() (int, bool, string) {
       return 1, true, "hello"
}

func main() {
       myInt, myBool, myString := manyReturns()
       fmt.Println(myInt, myBool, myString)
}
```

This function returns an integer, a boolean, and a string.

Store each return value in a variable.

```
1 true hello
```

If it makes the purpose of the return values clearer, you can supply names for each one, similar to parameter names. The main purpose of named return values is as documentation for programmers reading the code.

```
package main

import (
       "fmt"
       "math"
)

func floatParts(number float64) (integerPart int, fractionalPart float64) {
       wholeNumber := math.Floor(number)
       return int(wholeNumber), number - wholeNumber
}

func main() {
       cans, remainder := floatParts(1.26)
       fmt.Println(cans, remainder)
}
```

Name for the first return value

Name for the second return value

```
1 0.26
```

Using multiple return values with our paintNeeded function

As we saw on the previous page, it's possible to return multiple values of any type. But the most common use for multiple return values is to return a primary return value, followed by an additional value indicating whether the function encountered an error. The additional value is usually set to `nil` if there were no problems, or an error value if an error occurred.

We'll follow that convention with our `paintNeeded` function as well. We'll declare that it returns two values, a `float64` and an `error`. (Error values have a type of `error`.) The first thing we'll do in the function block is to check whether the parameters are valid. If either the `width` or `height` parameter is less than `0`, we'll return a paint amount of `0` (which is meaningless, but we do have to return something), and an error value that we generate by calling `fmt.Errorf`. Checking for errors at the start of the function allows us to easily skip the rest of the function's code by calling `return` if there's a problem.

If there were no problems with the parameters, we proceed to calculate and return the paint amount just like before. The only other difference in the function code is that we return a second value of `nil` along with the paint amount, to indicate there were no errors.

```
package main

import "fmt"

func paintNeeded(width float64, height float64) (float64, error) {
        if width < 0 {
                return 0, fmt.Errorf("a width of %0.2f is invalid", width)
        }
        if height < 0 {
                return 0, fmt.Errorf("a height of %0.2f is invalid", height)
        }
        area := width * height
        return area / 10.0, nil
}

func main() {
        amount, err := paintNeeded(4.2, -3.0)
        fmt.Println(err)
        fmt.Printf("%0.2f liters needed\n", amount)
}
```

Here's the return value with the amount of paint, just like before.

Here's a second return value that will indicate whether there were any errors.

If width is invalid, return 0 and an error.

If height is invalid, return 0 and an error.

Return the amount of paint, along with "nil", indicating there was no error.

Add a second variable to hold the second return value.

Prints the error (or "nil" if there was none)

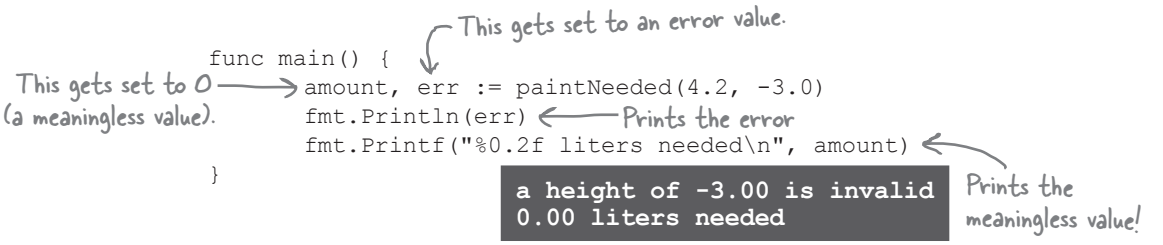
```
a height of -3.00 is invalid
0.00 liters needed
```

In the `main` function, we add a second variable to record the error value from `paintNeeded`. We print the error (if any), and then print the paint amount.

If we pass an invalid argument to `paintNeeded`, we'll get an error return value, and print that error. But we also get `0` as the amount of paint. (As we said, this value is meaningless when there's an error, but we had to use *something* for the first return value.) So we wind up printing the message "0.00 liters needed"! We'll need to fix that...

Always handle errors!

When we pass an invalid argument to `paintNeeded`, we get an error value back, which we print for the user to see. But we also get an (invalid) amount of paint, which we print as well!




```
func main() {
    amount, err := paintNeeded(4.2, -3.0)
    fmt.Println(err)
    fmt.Printf("%0.2f liters needed\n", amount)
}
```

```
a height of -3.00 is invalid
0.00 liters needed
```

When a function returns an error value, it usually has to return a primary return value as well. But any other return values that accompany an error value should be considered unreliable, and ignored.

When you call a function that returns an error value, it's important to test whether that value is `nil` before proceeding. If it's anything other than `nil`, it means there's an error that must be handled.

How the error should be handled depends on the situation. In the case of our `paintNeeded` function, it might be best to simply skip the current calculation and proceed with the rest of the program:

```
func main() {
    amount, err := paintNeeded(4.2, -3.0)
    if err != nil {
        fmt.Println(err)
    } else {
        fmt.Printf("%0.2f liters needed\n", amount)
    }
    // Additional calculations here...
}
```

If the error value is not nil, there must be a problem...
...so print the error.
Otherwise, the error value would be nil...
...so it would be okay to print the amount we got back.

```
a height of -3.00 is invalid
```

But since this is such a short program, you could instead call `log.Fatal` to display the error message and exit the program.

```
func main() {
    amount, err := paintNeeded(4.2, -3.0)
    if err != nil {
        log.Fatal(err)
    }
    fmt.Printf("%0.2f liters needed\n", amount)
}
```

If the error value is not nil, there must be a problem...
...so print the error and exit the program.
This code will never be reached if there's an error.

```
2018/03/12 19:49:27 a height of -3.00 is invalid
```

The important thing to remember is that you should always check the return values to see whether there *is* an error. What you do with the error at that point is up to you!



Breaking Stuff is Educational!

Here's a program that calculates the square root of a number. But if a negative number is passed to the `squareRoot` function, it will return an error value. Make one of the changes below and try to compile it. Then undo your change and try the next one. See what happens!

```
package main

import (
	"fmt"
	"math"
)

func squareRoot(number float64) (float64, error) {
	if number < 0 {
		return 0, fmt.Errorf("can't get square root of negative number")
	}
	return math.Sqrt(number), nil
}

func main() {
	root, err := squareRoot(-9.3)
	if err != nil {
		fmt.Println(err)
	} else {
		fmt.Printf("%0.3f", root)
	}
}
```

If you do this...	...it will break because...
Remove one of the arguments to `return`: `return math.Sqrt(number)~~, nil~~`	The number of arguments to `return` must always match the number of return values in the function declaration.
Remove one of the variables the return values are assigned to: `root~~, err~~ := squareRoot(-9.3)`	If you use any of the return values from a function, Go requires you to use all of them.

Remove the code that uses one of the return values:
`root, err := squareRoot(-9.3)`
`~~if err != nil {~~`
`~~fmt.Println(err)~~`
`~~} else {~~`
`fmt.Printf("%0.3f", root)`
`~~}~~`
	Go requires that you use every variable you declare. This is actually a really useful feature when it comes to error return values, because it helps keep you from accidentally ignoring an error.

Pool Puzzle



Your **job** is to take code snippets from the pool and place them into the blank lines in the code. **Don't** use the same snippet more than once, and you won't need to use all the snippets. Your **goal** is to make code that will run and produce the output shown.

```
package main

import (
        "errors"
        "fmt"
)

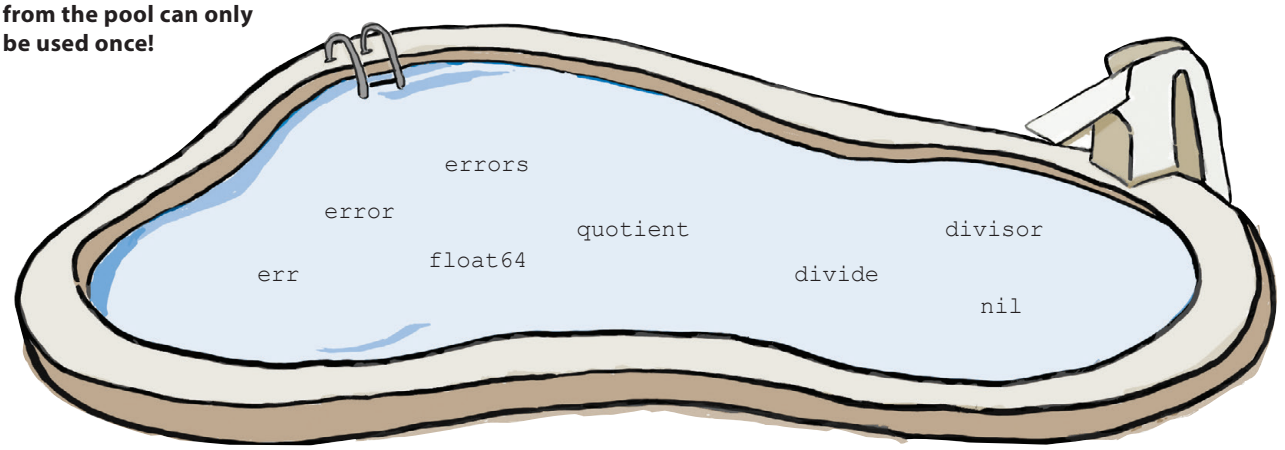
func divide(dividend float64, divisor float64) (float64, _____) {
        if divisor == 0.0 {
                return 0, _________.New("can't divide by 0")
        }
        return dividend / divisor, _____
}

func main() {
        __________, _____ := divide(5.6, 0.0)
        if err != nil {
                fmt.Println(err)
        } else {
                fmt.Printf("%0.2f\n", quotient)
        }
}
```

Output

```
can't divide by 0
```

Note: each snippet from the pool can only be used once!



Answers on page 112.

Function parameters receive copies of the arguments

As we mentioned, when you call a function that has parameters declared, you need to provide arguments to the call. The value in each argument is *copied* to the corresponding parameter variable. (Programming languages that do this are sometimes called "pass-by-value.")

This is fine in most cases. But if you want to pass a variable's value to a function and have it *change* the value in some way, you'll run into trouble. The function can only change the *copy* of the value in its parameter, not the original. So any changes you make within the function won't be visible outside it!

Go is a "pass-by-value" language; function parameters receive a copy of the arguments from the function call.

Here's an updated version of the `double` function we showed earlier. It takes a number, multiplies it by 2, and prints the result. (It uses the `*=` operator, which works just like `+=` , but it multiplies the value the variable holds instead of adding to it.)

```
package main

import "fmt"

func main() {
        amount := 6
        double(amount)
}

func double(number int) {
        number *= 2
        fmt.Println(number)
}
```

Pass an argument to the function.

Parameter is set to a copy of the argument.

```
12
```

Prints the doubled amount

Suppose we wanted to move the statement that prints the doubled value from the `double` function back to the function that calls it, though. It won't work, because `double` only alters its *copy* of the value. Back in the calling function, when we try to print, we'll get the original value, not the doubled one!

```
func main() {
        amount := 6
        double(amount)
        fmt.Println(amount)
}

func double(number int) {
        number *= 2
}
```

Pass an argument to the function.

Prints the original value!

Parameter is set to a copy of the argument.

```
6
```

Prints the unchanged amount!

Alters the copied value, not the original!

We need a way to allow a function to alter the original value a variable holds, rather than a copy. To learn how to do that, we'll need to make one more detour away from functions, to learn about *pointers*.



Go on a Detour

Pointers



You can get the *address* of a variable using `&` (an ampersand), which is Go's "address of" operator. For example, this code initializes a variable, prints its value, and then prints the variable's address...

```
amount := 6
fmt.Println(amount)
fmt.Println(&amount)
```

Retrieve the variable's value

Retrieve the variable's address

```
6
0x1040a124
```

Variable's value

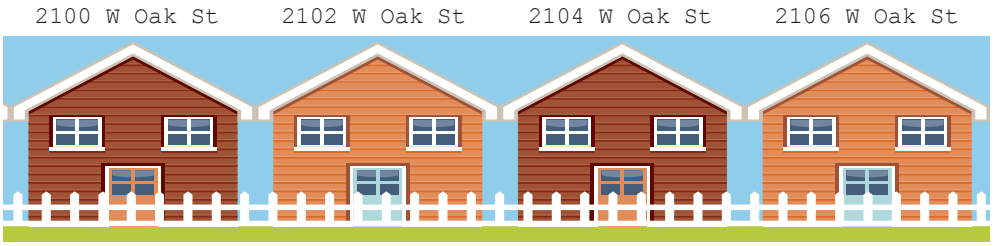
Variable's address

We can get addresses for variables of any type. Notice that the address differs for each variable.

```
var myInt int
fmt.Println(&myInt)
var myFloat float64
fmt.Println(&myFloat)
var myBool bool
fmt.Println(&myBool)
```

```
0x1040a128
0x1040a140
0x1040a148
```

And what are these "addresses," exactly? Well, if you want to find a particular house in a crowded city, you use its address...



Just like a city, the memory your computer sets aside for your program is a crowded place. It's full of variable values: booleans, integers, strings, and more. Just like the address of a house, if you have the address of a variable, you can use it to find the value that variable contains.

Here's the address "0x1040a108"...

0x1040a100	0x1040a108	0x1040a110	0x1040a118	0x1040a120	0x1040a128
true	6	...	...	...	3.1415

...and here's the value at that address!

Values that represent the address of a variable are known as **pointers**, because they *point* to the location where the variable can be found.



Pointer types



The type of a pointer is written with a * symbol, followed by the type of the variable the pointer points to. The type of a pointer to an int variable, for example, would be written *int (you can read that aloud as "pointer to int").

We can use the reflect.TypeOf function to show us the types of our pointers from the previous program:

```
package main

import (
        "fmt"
        "reflect"
)

func main() {
        var myInt int
        fmt.Println(reflect.TypeOf(&myInt))
        var myFloat float64
        fmt.Println(reflect.TypeOf(&myFloat))
        var myBool bool
        fmt.Println(reflect.TypeOf(&myBool))
}
```

Get a pointer to myInt and print the pointer's type.

Get a pointer to myFloat and print the pointer's type.

Get a pointer to myBool and print the pointer's type.

Here are the pointer types.

```
*int
*float64
*bool
```

We can declare variables that hold pointers. A pointer variable can only hold pointers to one type of value, so a variable might only hold *int pointers, only *float64 pointers, and so on.

```
var myInt int
var myIntPointer *int
myIntPointer = &myInt
fmt.Println(myIntPointer)

var myFloat float64
var myFloatPointer *float64
myFloatPointer = &myFloat
fmt.Println(myFloatPointer)
```

Declare a variable that holds a pointer to an int.

Assign a pointer to the variable.

Declare a variable that holds a pointer to a float64.

Assign a pointer to the variable.

```
0x1040a128
0x1040a140
```

As with other types, if you'll be assigning a value to the pointer variable right away, you can use a short variable declaration instead:

```
var myBool bool
myBoolPointer := &myBool
fmt.Println(myBoolPointer)
```

A short declaration for a pointer variable

```
0x1040a148
```

Getting or changing the value at a pointer



You can get the value of the variable a pointer refers to by typing the * operator right before the pointer in your code. To get the value at `myIntPointer`, for example, you'd type `*myIntPointer`. (There's no official consensus on how to read * aloud, but we like to pronounce it as "value at," so `*myIntPointer` is "value at `myIntPointer`.")

```
myInt := 4
myIntPointer := &myInt
fmt.Println(myIntPointer)
fmt.Println(*myIntPointer)

myFloat := 98.6
myFloatPointer := &myFloat
fmt.Println(myFloatPointer)
fmt.Println(*myFloatPointer)

myBool := true
myBoolPointer := &myBool
fmt.Println(myBoolPointer)
fmt.Println(*myBoolPointer)
```

Print the pointer itself.

Print the value at the pointer.

Print the pointer itself.

Print the value at the pointer.

Print the pointer itself.

Print the value at the pointer.

```
0x1040a124
4
0x1040a140
98.6
0x1040a150
true
```

The * operator can also be used to update the value at a pointer:

```
myInt := 4
fmt.Println(myInt)
myIntPointer := &myInt
*myIntPointer = 8
fmt.Println(*myIntPointer)
fmt.Println(myInt)
```

Assign a new value to the variable at the pointer (myInt).

Print the value of the variable at the pointer.

Print the variable's value directly.

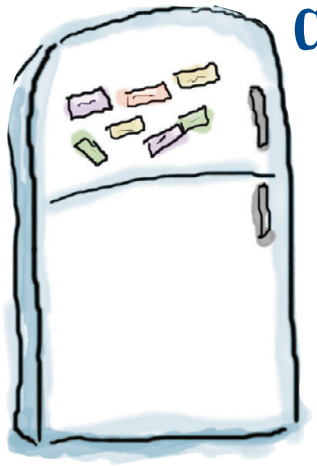
```
4
8
8
```

Initial value of myInt

Result of updating *myIntPointer

Updated value of myInt (same as *myIntPointer)

In the code above, `*myIntPointer = 8` accesses the variable at `myIntPointer` (that is, the `myInt` variable) and assigns a new value to it. So not only is the value of `*myIntPointer` updated, but `myInt` is as well.



Code Magnets

A Go program that uses a pointer variable is scrambled up on the fridge. Can you reconstruct the code snippets to make a working program that will produce the given output?

The program should declare `myInt` as an integer variable, and `myIntPointer` as a variable that holds an integer pointer. Then it should assign a value to `myInt`, and assign a pointer to `myInt` as the value of `myIntPointer`. Finally, it should print the value at `myIntPointer`.

```
package main

import "fmt"

func main() {

Add your code here!

}
```

Output

```
42
```

Here are the extra magnets. Add them to the program above!

```
var     var     myInt     myInt     myInt     42
int     int     myIntPointer     myIntPointer     myIntPointer
=     =     &     *     *     fmt.Println(     )
```

Answers on page 112.

Using pointers with functions



It's possible to return pointers from functions; just declare that the function's return type is a pointer type.

```
func createPointer() *float64 {
        var myFloat = 98.5
        return &myFloat
}

func main() {
        var myFloatPointer *float64 = createPointer()
        fmt.Println(*myFloatPointer)
}
```

Declare that the function returns a float64 pointer.

Return a pointer of the specified type.

Assign the returned pointer to a variable.

Print the value at the pointer.

```
98.5
```

(By the way, unlike in some other languages in Go, it's okay to return a pointer to a variable that's local to a function. Even though that variable is no longer in scope, as long as you still have the pointer, Go will ensure you can still access the value.)

You can also pass pointers to functions as arguments. Just specify that the type of one or more parameters should be a pointer.

```
func printPointer(myBoolPointer *bool) {
        fmt.Println(*myBoolPointer)
}

func main() {
        var myBool bool = true
        printPointer(&myBool)
}
```

Use a pointer type for this parameter.

Print the value at the pointer that gets passed in.

```
true
```

Pass a pointer to the function.

Make sure you only use pointers as arguments, if that's what the function declares it will take. If you try to pass a value directly to a function that's expecting a pointer, you'll get a compile error.

```
func main() {
        var myBool bool = true
        printPointer(myBool)
}
```

Error

```
cannot use myBool (type bool)
as type *bool in argument
to printPointer
```

Now you know the basics of using pointers in Go. We're ready to end our detour, and fix our `double` function!



Fixing our "double" function using pointers

We have a `double` function that takes an `int` value and multiplies it by 2. We want to be able to pass a value in and have that value doubled. But, as we learned, Go is a pass-by-value language, meaning that function parameters receive a *copy* of any arguments from the caller. Our function is doubling its copy of the value and leaving the original untouched!

```
func main() {
        amount := 6
        double(amount)
        fmt.Println(amount)
}

func double(number int) {
        number *= 2
}
```

Pass an argument to the function.

Prints the original value!

Parameter is set to a copy of the argument.

Alters the copied value, not the original!

6

Prints the unchanged amount!

Here's where our detour to learn about pointers is going to be useful. If we pass a pointer to the function and then alter the value at that pointer, the changes will still be effective outside the function!

We only need to make a few small changes to get this working. In the `double` function, we need to update the type of the `number` parameter to take a `*int` rather than an `int`. Then we'll need to change the function code to update the value at the `number` pointer, rather than updating a variable directly. Finally, in the `main` function, we just need to update our call to `double` to pass a pointer rather than a direct value.

```
func main() {
        amount := 6
        double(&amount)
        fmt.Println(amount)
}

func double(number *int) {
        *number *= 2
}
```

Pass a pointer instead of the variable value.

Accept a pointer instead of an int value.

Update the value at the pointer.

12

Prints the doubled amount.

When we run this updated code, a pointer to the `amount` variable will be passed to the `double` function. The `double` function will take the value at that pointer and double it, thereby changing the value in the `amount` variable. When we return to the `main` function and print the `amount` variable, we'll see our doubled value!

You've learned a lot about writing your own functions in this chapter. The benefits of some of these features may not be clear right now. Don't worry—as our programs get more complex in later chapters, we'll be making good use of everything you've learned!



We've written the `negate` function below, which is *supposed* to update the value of the `truth` variable to its opposite (`false`), and update the value of the `lies` variable to its opposite (`true`). But when we call `negate` on the `truth` and `lies` variables and then print their values, we see that they're unchanged!

```
package main

import "fmt"

func negate(myBoolean bool) bool {
        return !myBoolean
}

func main() {
        truth := true
        negate(truth)
        fmt.Println(truth)
        lies := false
        negate(lies)
        fmt.Println(lies)
}
```

Actual output

```
true
false
```

Fill in the blanks below so that `negate` takes a pointer to a Boolean value instead of taking a Boolean value directly, then updates the value at that pointer to the opposite value. Be sure to change the calls to `negate` to pass a pointer instead of passing the value directly!

```
package main

import "fmt"

func negate(myBoolean _______) {
        ___________ = !___________
}

func main() {
        truth := true
        negate(_________)
        fmt.Println(truth)
        lies := false
        negate(_________)
        fmt.Println(lies)
}
```

Output we want

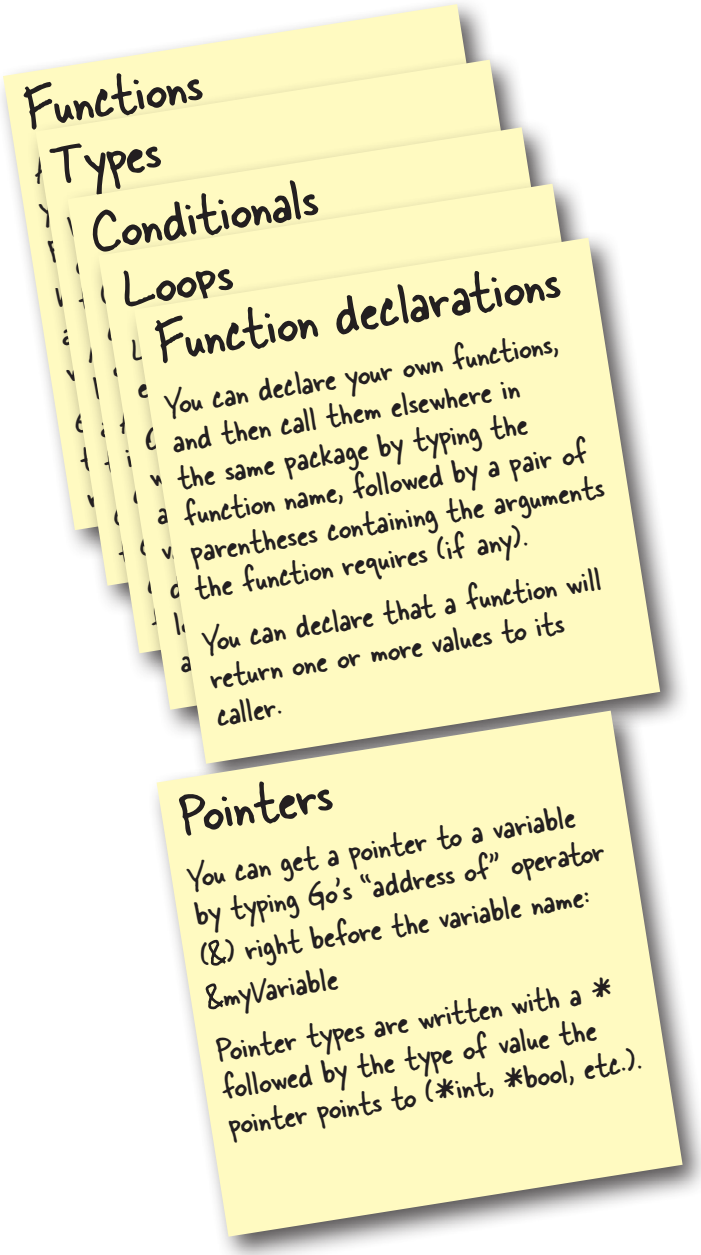
```
false
true
```

Answers on page 112.



Your Go Toolbox

That's it for Chapter 3! You've added function declarations and pointers to your toolbox.



BULLET POINTS

- The `fmt.Printf` and `fmt.Sprintf` functions format values they're given. The first argument should be a formatting string containing **verbs** (`%d`, `%f`, `%s`, etc.) that values will be substituted for.
- Within a formatting verb, you can include a **width**: a minimum number of characters the formatted value will take up. For example, `%12s` results in a 12-character string (padded with spaces), `%2d` results in a 2-character integer, and `%.3f` results in a floating-point number rounded to 3 decimal places.
- If you want calls to your function to accept arguments, you must declare one or more **parameters**, including types for each, in the function declaration. The number and type of arguments must always match the number and type of parameters, or you'll get a compile error.
- If you want your function to return one or more values, you must declare the return value types in the function declaration.
- You can't access a variable declared within a function outside that function. But you can access a variable declared outside a function (usually at the package level) within that function.
- When a function returns multiple values, the last value usually has a type of `error`. Error values have an `Error()` method that returns a string describing the error.
- By convention, functions return an error value of `nil` to indicate there are no errors.
- You can access the value a pointer holds by putting a `*` right before it: `*myPointer`
- If a function receives a pointer as a parameter, and it updates the value at that pointer, then the updated value will still be visible outside the function.



Below is a program that declares several functions, then calls those functions within `main`. Write down what the program output would be.

```
package main

import "fmt"

func functionA(a int, b int) {
       fmt.Println(a + b)
}
func functionB(a int, b int) {
       fmt.Println(a * b)
}
func functionC(a bool) {
       fmt.Println(!a)
}
func functionD(a string, b int) {
       for i := 0; i < b; i++ {
              fmt.Print(a)
       }
       fmt.Println()
}

func main() {
       functionA(2, 3)
       functionB(2, 3)
       functionC(true)
       functionD("$", 4)
       functionA(5, 6)
       functionB(5, 6)
       functionC(false)
       functionD("ha", 3)
}
```

Output:

```
5
6
false
$$$$
11
30
true
hahaha
```

Pool Puzzle Solution

```
package main

import (
	"errors"
	"fmt"
)

func divide(dividend float64, divisor float64) (float64, error) {
	if divisor == 0.0 {
		return 0, errors.New("can't divide by 0")
	}
	return dividend / divisor, nil
}

func main() {
	quotient, err := divide(5.6, 0.0)
	if err != nil {
		fmt.Println(err)
	} else {
		fmt.Printf("%0.2f\n", quotient)
	}
}
```

Code Magnets Solution

```
package main

import "fmt"

func main() {
	var myInt int
	var myIntPointer * int
	myInt = 42
	myIntPointer = & myInt
	fmt.Println( * myIntPointer )
}
```

Output

```
42
```



Exercise Solution

```
package main

import "fmt"

func negate(myBoolean *bool) {
	*myBoolean = !*myBoolean
}

func main() {
	truth := true
	negate(&truth)
	fmt.Println(truth)
	lies := false
	negate(&lies)
	fmt.Println(lies)
}
```

4 bundles of code

Packages



It's time to get organized. So far, we've been throwing all our code together in a single file. As our programs grow bigger and more complex, that's going to quickly become a mess.

In this chapter, we'll show you how to create your own **packages** to help keep related code together in one place. But packages are good for more than just organization. Packages are an easy way to *share code between your programs*. And they're an easy way to *share code with other developers*.

Different programs, same function

We've written two programs, each with an identical copy of a function, and it's becoming a maintenance headache...

On this page, we've got a new version of our *pass_fail.go* program from Chapter 2. The code that reads a grade from the keyboard has been moved to a new `getFloat` function. `getFloat` returns the floating-point number the user typed, unless there's an error, in which case it returns `0` and an error value. If an error is returned, the program reports it and exits; otherwise, it reports whether the grade is passing or failing, as before.


pass_fail.go

```
// pass_fail reports whether a grade is passing or failing.
package main

import (
        "bufio"
        "fmt"
        "log"
        "os"
        "strconv"
        "strings"
)

func getFloat() (float64, error) {
        reader := bufio.NewReader(os.Stdin)
        input, err := reader.ReadString('\n')
        if err != nil {
                return 0, err
        }
        input = strings.TrimSpace(input)
        number, err := strconv.ParseFloat(input, 64)
        if err != nil {
                return 0, err
        }
        return number, nil
}

func main() {
        fmt.Print("Enter a grade: ")
        grade, err := getFloat()
        if err != nil {
                log.Fatal(err)
        }

        var status string
        if grade >= 60 {
                status = "passing"
        } else {
                status = "failing"
        }
        fmt.Println("A grade of", grade, "is", status)
}
```

Identical to the getFloat function on the next page!

Almost identical to the code in Chapter 2, except...

...if there's an error reading input, we return it from the function.

We also return any error converting the string to a float64.

We call getFloat to get a grade...

If an error's returned, we log it and exit.

Unchanged from Chapter 2 code.

```
Enter a grade: 89.7
A grade of 89.7 is passing
```

Different programs, same function (continued)

On this page, we've got a new *tocelsius.go* program that lets the user type a temperature in the Fahrenheit measurement system and converts it to the Celsius system.

Notice that the `getFloat` function in *tocelsius.go* is identical to the `getFloat` function in *pass_fail.go*.


tocelsius.go

```
// tocelsius converts a temperature from Fahrenheit to Celsius.
package main

import (
	"bufio"
	"fmt"
	"log"
	"os"
	"strconv"
	"strings"
)

func getFloat() (float64, error) {
	reader := bufio.NewReader(os.Stdin)
	input, err := reader.ReadString('\n')
	if err != nil {
		return 0, err
	}

	input = strings.TrimSpace(input)
	number, err := strconv.ParseFloat(input, 64)
	if err != nil {
		return 0, err
	}
	return number, nil
}

func main() {
	fmt.Print("Enter a temperature in Fahrenheit: ")
	fahrenheit, err := getFloat()
	if err != nil {
		log.Fatal(err)
	}
	celsius := (fahrenheit - 32) * 5 / 9
	fmt.Printf("%0.2f degrees Celsius\n", celsius)
}
```

Identical to the getFloat function on the previous page!

We call getFloat to get a temperature.

If an error is returned, we log it and exit.

Convert temperature to Celsius...

...and print it with two decimal places of precision.

```
Enter a temperature in Fahrenheit: 98.6
37.00 degrees Celsius
```

Sharing code between programs using packages



More repeated code... If we ever discover a bug in the `getFloat` function, it'll be a pain to fix it in two places. These are two different programs, though, so I guess it can't be helped...

```
func getFloat() (float64, error) {
        reader := bufio.NewReader(os.Stdin)
        input, err := reader.ReadString('\n')
        if err != nil {
                return 0, err
        }

        input = strings.TrimSpace(input)
        number, err := strconv.ParseFloat(input, 64)
        if err != nil {
                return 0, err
        }
        return number, nil
}
```

Actually, there is something we can do—we can move the shared function to a new package!

Go allows us to define our own packages. As we discussed back in Chapter 1, a package is a group of code that all does similar things. The `fmt` package formats output, the `math` package works with numbers, the `strings` package works with strings, and so on. We've used the functions from each of these packages in multiple programs already.

Being able to use the same code between programs is one of the major reasons packages exist. If parts of your code are shared between multiple programs, you should consider moving them into packages.

If parts of your code are shared between multiple programs, you should consider moving them into packages.

The Go workspace directory holds package code

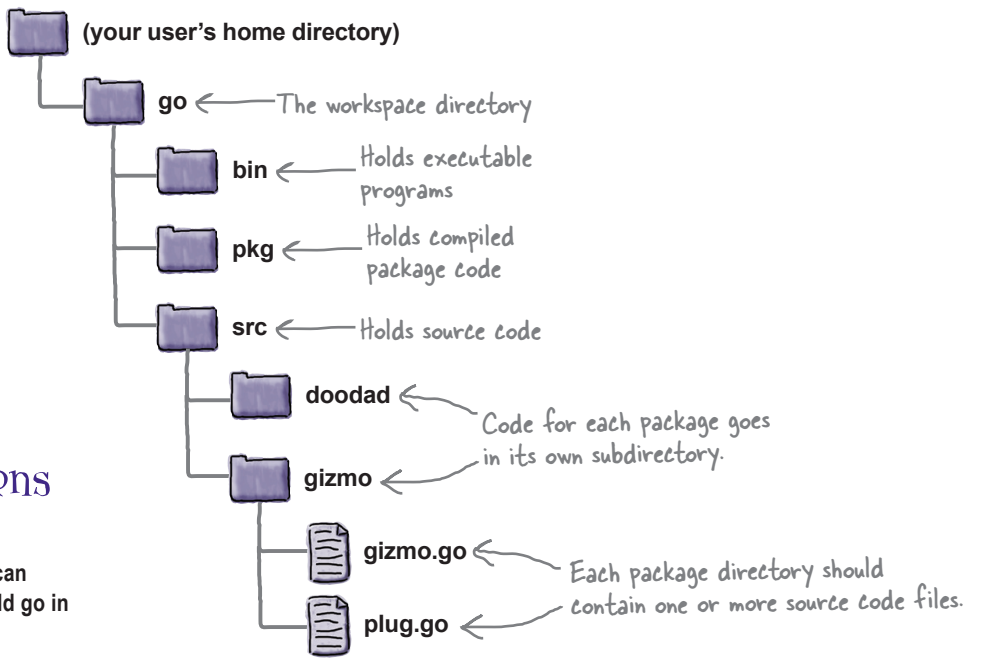
Go tools look for package code in a special directory (folder) on your computer called the **workspace**. By default, the workspace is a directory named *go* in the current user's home directory.

The workspace directory contains three subdirectories:

- *bin*, which holds compiled binary executable programs. (We'll talk more about *bin* later in the chapter.)
- *pkg*, which holds compiled binary package files. (We'll also talk more about *pkg* later in the chapter.)
- *src*, which holds Go source code.

Within *src*, code for each package lives in its own separate subdirectory. By convention, the subdirectory name should be the same as the package name (so code for a `gizmo` package would go in a *gizmo* subdirectory).

Each package directory should contain one or more source code files. The filenames don't matter, but they should end in a *.go* extension.



there are no Dumb Questions

Q: You said a package folder can contain multiple files. What should go in each file?

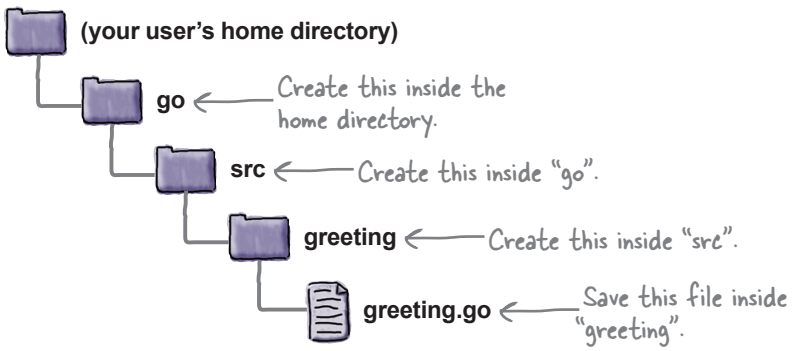
A: Whatever you want! You can keep all of a package's code in one file, or split it between multiple files. Either way, it will all become part of the same package.

Creating a new package

Let's try setting up a package of our own in the workspace. We'll make a simple package named `greeting` that prints greetings in various languages.

The workspace directory isn't created by default when Go is installed, so you'll need to create it yourself. Start by going to your home directory. (The path is *C:\Users\<yourname>* on most Windows systems, */Users/<yourname>* on Macs, and */home/<yourname>* on most Linux systems.) Within the home directory, create a directory named *go*—this will be our new workspace directory. Within the *go* directory, create a directory named *src*.

Finally, we need a directory to hold our package code. By convention, a package's directory should have the same name as a package. Since our package will be named `greeting`, that's the name you should use for the directory.



We know, that seems like a lot of nested directories (and actually, we'll be nesting them even deeper shortly). But trust us, once you've built up a collection of packages of your own as well as packages from others, this structure will help you keep your code organized.

And more importantly, this structure helps Go tools find the code. Because it's always in the *src* directory, Go tools know exactly where to look to find code for the packages you're importing.

Your next step is to create a file within the *greeting* directory, and name it *greeting.go*. The file should include the code below. We'll talk about it more shortly, but for now there's just a couple things we want you to notice...

Like all of our Go source code files thus far, this file starts with a `package` line. But unlike the others, this code isn't part of the `main` package; it's part of a package named `greeting`.

Also notice the two function definitions. They aren't much different from other functions we've seen so far. But because we want these functions to be accessible outside the `greeting` package, notice that we capitalize the first letter of their names so the functions are exported.

```
package greeting

import "fmt"

func Hello() {
        fmt.Println("Hello!")
}

func Hi() {
        fmt.Println("Hi!")
}
```

The package isn't "main", it's "greeting"!

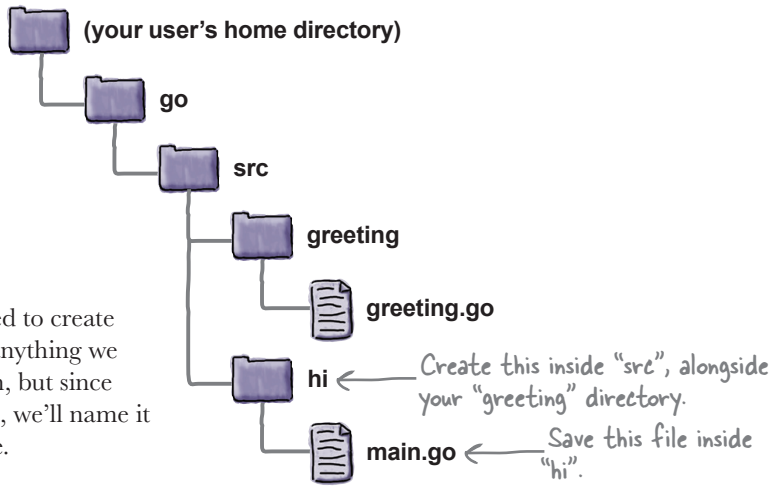
First letters are capitalized so that functions are exported!

greeting.go

Importing our package into a program

Now let's try using our new package within a program.

In your workspace directory, within the *src* subdirectory, create another subdirectory named *hi*. (We don't *have* to store code for executable programs within the workspace, but it's a good idea.)



Then, within your new *hi* directory, we need to create another source file. We can name the file anything we want, as long as it ends with a *.go* extension, but since this is going to be an executable command, we'll name it *main.go*. Save the code below within the file.

Like in every Go source code file, this code starts with a `package` line. But because we intend this to be an executable command, we need to use a package name of `main`. Generally, the package name should match the name of the directory it's kept in, but the `main` package is an exception to that rule.

```
package main

import "greeting"

func main() {
        greeting.Hello()
        greeting.Hi()
}
```

We need to import the package before we can use its functions.

We need the package name and a dot before calls to functions from a different package.

main.go

Next we need to import the `greeting` package so we can use its functions. Go tools look for package code in a folder within the workspace's *src* directory whose name matches the name in the `import` statement. To tell Go to look for code in the *src/greeting* directory within the workspace, we use `import "greeting"`.

Finally, because this is code for an executable, we need a `main` function that will be called when the program runs. In `main` we call both functions that are defined in the `greeting` package. Both calls are preceded by the package name and a dot, so that Go knows which package the functions are a part of.

We're all set; let's try running the program. In your terminal or command prompt window, use the **`cd`** command to change to the *src/hi* directory within your workspace directory. (The path will vary based on the location of your home directory.) Then, use **`go run main.go`** to run the program.

```
Shell  Edit  View  Window  Help
$ cd /Users/jay/go/src/hi
$ go run main.go
Hello!
Hi!
$
```

Functions from the package are called!

When it sees the `import "greeting"` line, Go will look in the *greeting* directory in your workspace's *src* directory for the package source code. That code gets compiled and imported, and we're able to call the `greeting` package's functions!

Packages use the same file layout

Remember back in Chapter 1, we talked about the three sections almost every Go source code file has?

You'll quickly get used to seeing these three sections, in this order, in almost every Go file you work with:

1. The package clause
2. Any `import` statements
3. The actual code

```
The package clause {package main
The imports section {import "fmt"
The actual code    {func main() {
                        fmt.Println("Hello, Go!")
                    }
```

That rule holds true for the `main` package in our *main.go* file, of course. In our code, you can see a `package` clause, followed by an imports section, followed by the actual code for our package.

```
The package clause {package main
The imports section {import "greeting"
The actual code    {func main() {
                        greeting.Hello()
                        greeting.Hi()
                    }
```

Packages other than `main` follow the same format. You can see that our *greeting.go* file also has a package clause, imports section, and the actual package code at the end.

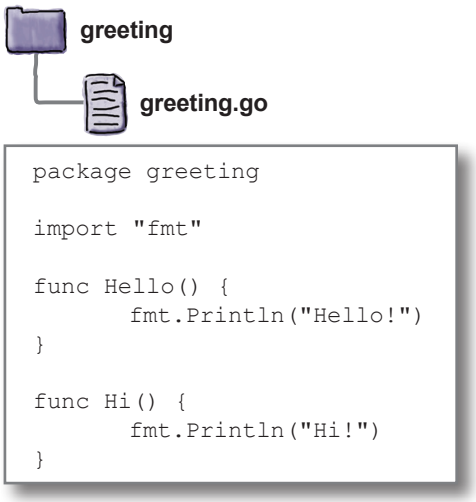
```
The package clause {package greeting
The imports section {import "fmt"
The actual code    {func Hello() {
                        fmt.Println("Hello!")
                    }

                    func Hi() {
                        fmt.Println("Hi!")
                    }
```

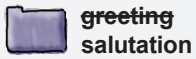


Breaking Stuff is Educational!

Take our code for the `greeting` package, as well as the code for the program that imports it. Try making one of the changes below and run it. Then undo your change and try the next one. See what happens!

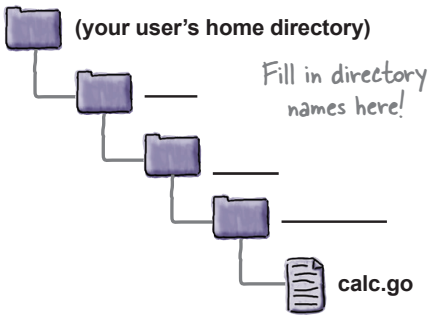




If you do this...	...it will fail because...
Change the name on the *greeting* directory 	The Go tools use the name in the import path as the name of the directory to load the package source code from. If they don't match, the code won't load.
Change the name on the `package` line of *greeting.go* `package salutation`	The contents of the *greeting* directory *will* actually load, as a package named `salutation`. Since the function calls in *main.go* still reference the `greeting` package, though, we'll get errors.
Change the function names in *greeting.go* and *main.go* to all lowercase `func ~~H~~hello()` `func ~~H~~hi()` `greeting.~~H~~hello()` `greeting.~~H~~hi()`	Functions whose names begin with a lowercase letter are unexported, meaning they can only be used within their own package. To use a function from a different package, its name must begin with a capital letter, so it's exported.

Pool Puzzle

Your **job** is to take code snippets from the pool and place them into the blank lines. **Don't** use the same snippet more than once, and you won't need to use all the snippets. Your **goal** is to set up a `calc` package within a Go workspace so `calc`'s functions can be used within *main.go*.




```
package _____

func ____(first float64, second float64) float64 {
        return first + second
}

func _________(first float64, second float64) float64 {
        return first - second
}
```

calc.go

```
package main

import (
        "calc"
        "fmt"
)

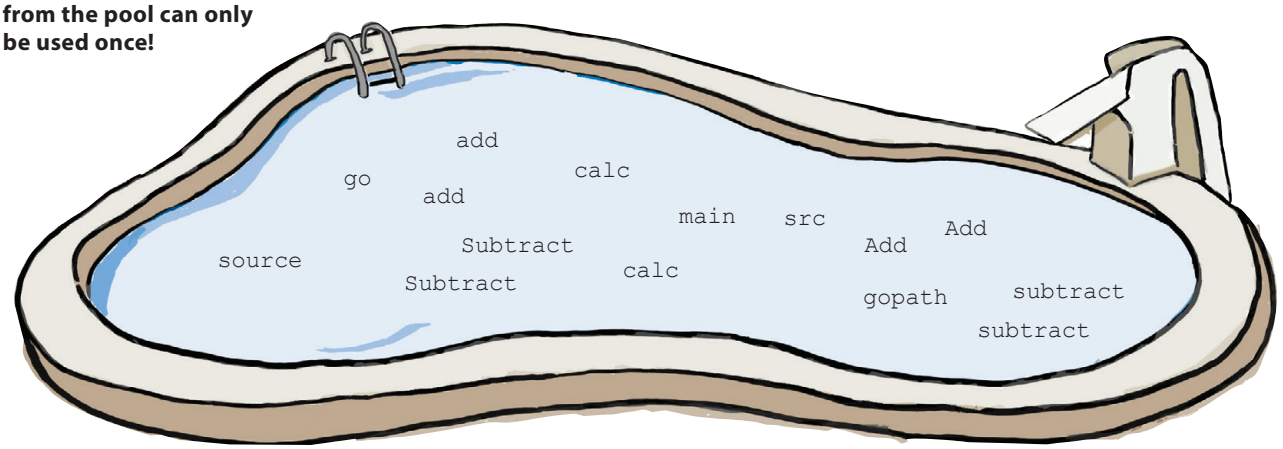
func main() {
        fmt.Println(calc._____(1, 2))
        fmt.Println(calc._____(7, 3))
}
```

main.go

Output

3
4

Note: each snippet from the pool can only be used once!



Answers on page 147.

Package naming conventions

Developers using a package are going to need to type its name each and every time they call a function from that package. (Think of `fmt.Printf`, `fmt.Println`, `fmt.Print`, etc.) To make that as painless as possible, there are a few rules package names should follow:

- A package name should be all lowercase.
- The name should be abbreviated if the meaning is fairly obvious (such as `fmt`).
- It should be one word, if possible. If two words are needed, they should *not* be separated by underscores, and the second word should *not* be capitalized. (The `strconv` package is one example.)
- Imported package names can conflict with local variable names, so don't use a name that package users are likely to want to use as well. (For example, if the `fmt` package were named `format`, anyone who imported that package would risk conflicts if they named a local variable `format`).

Package qualifiers

When accessing a function, variable, or the like that's exported from a different package, you need to qualify the name of the function or variable by typing the package name before it. When you access a function or variable that's defined in the *current* package, however, you should *not* qualify the package name.

In our *main.go* file, since our code is in the `main` package, we need to specify that the `Hello` and `Hi` functions are from the `greeting` package, by typing **`greeting.Hello`** and **`greeting.Hi`**.




```
package main

import "greeting"

func main() {
    greeting.Hello()
    greeting.Hi()
}
```

Package qualifiers




```
package greeting

import "fmt"

func Hello() {
        fmt.Println("Hello!")
}

func Hi() {
        fmt.Println("Hi!")
}

func AllGreetings() {
    Hello()
    Hi()
}
```

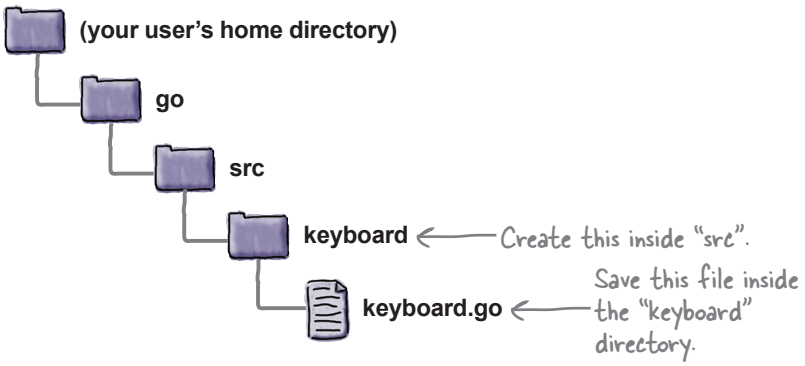
No qualifiers

Suppose that we called the `Hello` and `Hi` functions from another function in the `greeting` package, though. There, we would just type `Hello` and `Hi` (without the package name qualifier) because we'd be calling the functions from the same package where they're defined.

Moving our shared code to a package

Now that we understand how to add packages to the Go workspace, we're finally ready to move our `getFloat` function to a package that our *pass_fail.go* and *tocelsius.go* programs can both use.

Let's name our package `keyboard`, since it reads user input from the keyboard. We'll start by creating a new directory named *keyboard* inside our workspace's *src* directory.



Next, we'll create a source code file within the *keyboard* directory. We can name it anything we want, but we'll just name it after the package: *keyboard.go*.

At the top of the file, we'll need a `package` clause with the package name: `keyboard`.

Then, because this is a separate file, we'll need an `import` statement with all the packages used in our code: `bufio`, `os`, `strconv`, and `strings`. (We need to leave out the `fmt` and `log` packages, as those are only used in the *pass_fail.go* and *tocelsius.go* files.)

Finally, we can copy the code from the old `getFloat` function as is. But we need to be sure to rename the function to `GetFloat`, because it won't be exported unless the first letter of its name is capitalized.

```
package keyboard        <- Add a package clause.

import (
"bufio"                 <- Import only the packages used in this file.
"os"
"strconv"
"strings"
)
                        <- Capitalize the function name, so it's exported.
func GetFloat() (float64, error) {
        reader := bufio.NewReader(os.Stdin)
        input, err := reader.ReadString('\n')
        if err != nil {
                return 0, err
        }

        input = strings.TrimSpace(input)
        number, err := strconv.ParseFloat(input, 64)
        if err != nil {
                return 0, err
        }
        return number, nil
}
```

This code is identical to the old duplicated function code.

keyboard.go

Moving our shared code to a package (continued)

Now the *pass_fail.go* program can be updated to use our new `keyboard` package.

```
// pass_fail reports whether a grade is passing or failing.
package main

import (
	"fmt"
	"keyboard"
	"log"
)

func main() {
	fmt.Print("Enter a grade: ")
	grade, err := keyboard.GetFloat()
	if err != nil {
		log.Fatal(err)
	}

	var status string
	if grade >= 60 {
		status = "passing"
	} else {
		status = "failing"
	}
	fmt.Println("A grade of", grade, "is", status)
}
```

Import only the packages used in this file.

Be sure to import our new package.

We can remove the getFloat function that was here.

Call the "keyboard" package's function instead.

Because we're removing the old `getFloat` function, we need to remove the unused `bufio`, `os`, `strconv`, and `strings` imports. In their place, we'll import the new `keyboard` package.

In our `main` function, in place of the old call to `getFloat`, we'll call the new `keyboard.GetFloat` function. The rest of the code is unchanged.

If we run the updated program, we'll see the same output as before.

```
Enter a grade: 89.7
A grade of 89.7 is passing
```

We can make the same updates to the *tocelsius.go* program.

```
// tocelsius converts a temperature...
package main

import (
	"fmt"
	"keyboard"
	"log"
)

func main() {
	fmt.Print("Enter a temperature in Fahrenheit: ")
	fahrenheit, err := keyboard.GetFloat()
	if err != nil {
		log.Fatal(err)
	}
	celsius := (fahrenheit - 32) * 5 / 9
	fmt.Printf("%0.2f degrees Celsius\n", celsius)
}
```

Import only the packages used in this file.

Be sure to import our new package.

We can remove the getFloat function that was here.

Call the "keyboard" package's function instead.

We update the imports, remove the old `getFloat`, and call `keyboard.GetFloat` instead.

And again, if we run the updated program, we'll get the same output as before. But this time, instead of relying on redundant function code, we're using the shared function in our new package!

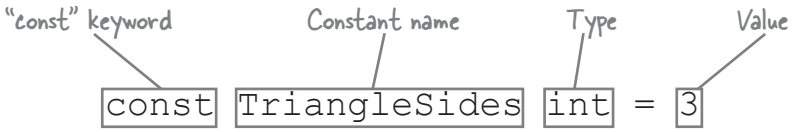
```
Enter a temperature in Fahrenheit: 98.6
37.00 degrees Celsius
```

Constants

Many packages export **constants**: named values that never change.

A constant declaration looks a lot like a variable declaration, with a name, optional type, and value for the constant. But the rules are slightly different:

- Instead of the `var` keyword, you use the `const` keyword.
- You must assign a value at the time the constant is declared; you can't assign a value later as with variables.
- Variables have the `:=` short variable declaration syntax available, but there is no equivalent for constants.



As with variable declarations, you can omit the type, and it will be inferred from the value being assigned:

```
const SquareSides = 4
```

We're assigning an integer, so the type of the constant will be "int".

The value of a *variable* can *vary*, but the value of a *constant* must remain *constant*. Attempting to assign a new value to a constant will result in a compile error. This is a safety feature: constants should be used for values that *shouldn't* ever change.

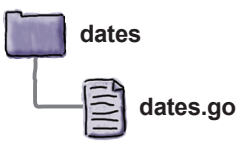
```
const PentagonSides = 5
PentagonSides = 6
```

Attempt to assign a new value to a constant!

Compile error

```
cannot assign to PentagonSides
```

If your program includes "hardcoded" literal values, especially if those values are used in multiple places, you should consider replacing them with constants (even if the program isn't broken up into multiple packages). Here's a package with two functions, both featuring the integer literal 7 representing the number of days in a week:




```
package dates

Accept a number of weeks.
func WeeksToDays(weeks int) int {
        return weeks * 7
}
Accept a number of days.
func DaysToWeeks(days int) float64 {
        return float64(days) / float64(7)
}
```

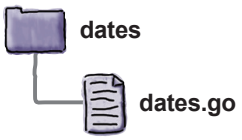
Multiply that by the number of days in a week to get a total number of days.

Divide that by the number of days in a week to get a number of weeks.

Constants (continued)

By replacing the literal values with a constant, `DaysInWeek`, we can document what they mean. (Other developers will see the name `DaysInWeek`, and immediately know we didn't randomly choose the number 7 to use in our functions.) Also, if we add more functions later, we can avoid inconsistencies by having them refer to `DaysInWeek` as well.

Notice that we declare the constant outside of any function, at the package level. Although it's possible to declare a constant inside a function, that would limit its scope to the block for that function. It's much more typical to declare constants at the package level, so they can be accessed by all functions in that package.




```
package dates

const DaysInWeek int = 7

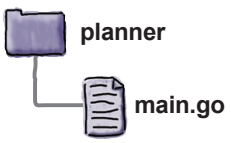
func WeeksToDays(weeks int) int {
  return weeks * DaysInWeek
}
func DaysToWeeks(days int) float64 {
  return float64(days) / float64(DaysInWeek)
}
```

Declare a constant.

Use the constant in place of the integer literal.

Use the constant in place of the integer literal.

As with variables and functions, a constant whose name begins with a capital letter is exported, and we can access it from other packages by qualifying its name. Here, a program makes use of the `DaysInWeek` constant from the `main` package by importing the `dates` package and qualifying the constant name as `dates.DaysInWeek`.




```
package main

import (
  "dates"
  "fmt"
)

func main() {
  days := 3
  fmt.Println("Your appointment is in", days, "days")
  fmt.Println("with a follow-up in", days + dates.DaysInWeek, "days")
}
```

Import the package the constant is declared in.

Qualify the package name.

Use the constant from the "dates" package.

```
Your appointment is in 3 days
with a follow-up in 10 days
```

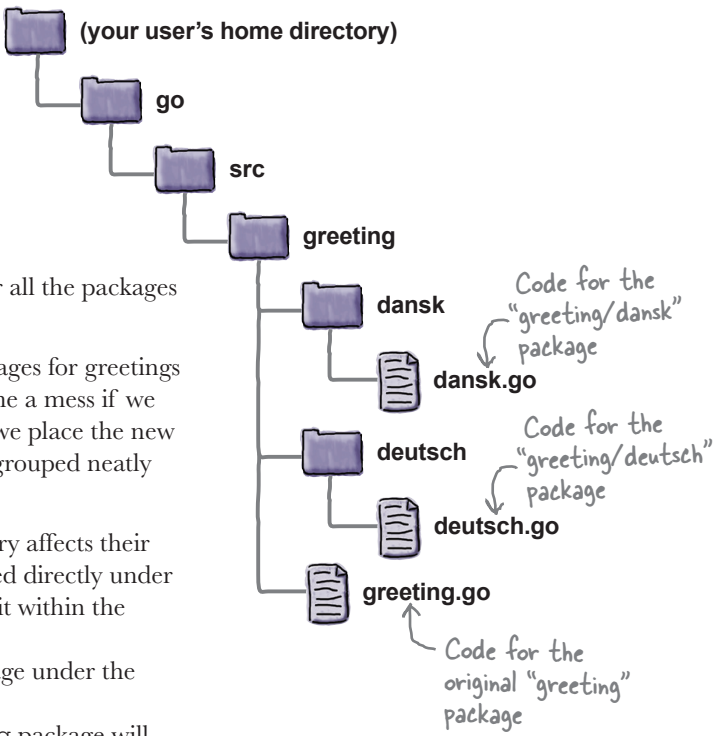
Nested package directories and import paths

When you're working with the packages that come with Go, like `fmt` and `strconv`, the package name is usually the same as its import path (the string you use in an `import` statement to import the package). But as we saw in Chapter 2, that's not always the case...

But the import path and package name don't have to be identical. Many Go packages fall into similar categories, like compression or complex math. So they're grouped together under similar import path prefixes, such as `"archive/"` or `"math/"`. (Think of them as being similar to the paths of directories on your hard drive.)

Import path	Package name
`"archive"`	`archive`
`"archive/tar"`	`tar`
`"archive/zip"`	`zip`
`"math"`	`math`
`"math/cmplx"`	`cmplx`
`"math/rand"`	`rand`

Some sets of packages are grouped together by import path prefixes like `"archive/"` and `"math/"`. We said to think of these prefixes as being similar to the paths of directories on your hard drive...and that wasn't a coincidence. These import path prefixes *are* created using directories!



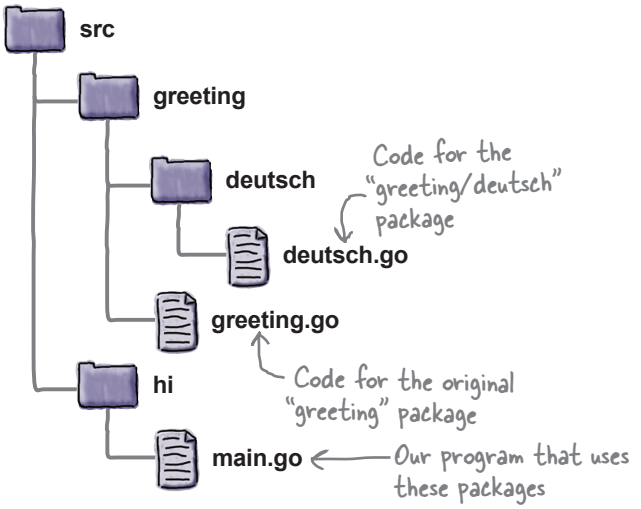
You can nest groups of similar packages together in a directory in your Go workspace. That directory then becomes part of the import path for all the packages it contains.

Suppose, for example, that we wanted to add packages for greetings in additional languages. That would quickly become a mess if we placed them all directly in the *src* directory. But if we place the new packages under the *greeting* directory, they'll all be grouped neatly together.

And placing the packages under the *greeting* directory affects their import path, too. If the `dansk` package were stored directly under *src*, its import path would be `"dansk"`. But place it within the *greeting* directory, and its import path becomes `"greeting/dansk"`. Move the `deutsch` package under the *greeting* directory, and its import path becomes `"greeting/deutsch"`. The original `greeting` package will still be available at an import path of `"greeting"`, as long as its source code file is stored directly under the *greeting* directory (not a subdirectory).

Nested package directories and import paths (continued)

Suppose that we had a `deutsch` package nested under our *greeting* package directory, and that its code looked like this:




```
package deutsch

import "fmt"

func Hallo() {
        fmt.Println("Hallo!")
}

func GutenTag() {
        fmt.Println("Guten Tag!")
}
```

deutsch.go

Let's update our *hi/main.go* code to use the `deutsch` package as well. Since it's nested under the *greeting* directory, we'll need to use an import path of `"greeting/deutsch"`. But once it's imported, we'll be using just the package name to refer to it: `deutsch`.

```
package main

import (
        "greeting"
        "greeting/deutsch"
)

func main() {
        greeting.Hello()
        greeting.Hi()
        deutsch.Hallo()
        deutsch.GutenTag()
}
```

Import the "greeting" package, as before.

Import the "deutsch" package as well.

Add calls to the new package's functions.

main.go

As before, we run our code by using the **`cd`** command to change to the *src/hi* directory within your workspace directory. Then, we use **`go run main.go`** to run the program. We'll see the results of our calls to the `deutsch` package functions in the output.

```
Shell Edit View Window Help
$ cd /Users/jay/go/src/hi
$ go run main.go
Hello!
Hi!
Hallo!
Guten Tag!
```

Here's the output from the "deutsch" package.

Installing program executables with "go install"

When we use `go run`, Go has to compile the program, as well as all the packages it depends on, before it can execute it. And it throws that compiled code away when it's done.

In Chapter 1, we showed you the `go build` command, which compiles and saves an executable binary file (a file you can execute even without Go installed) in the current directory. But using that too much risks littering your Go workspace with executables in random, inconvenient places.

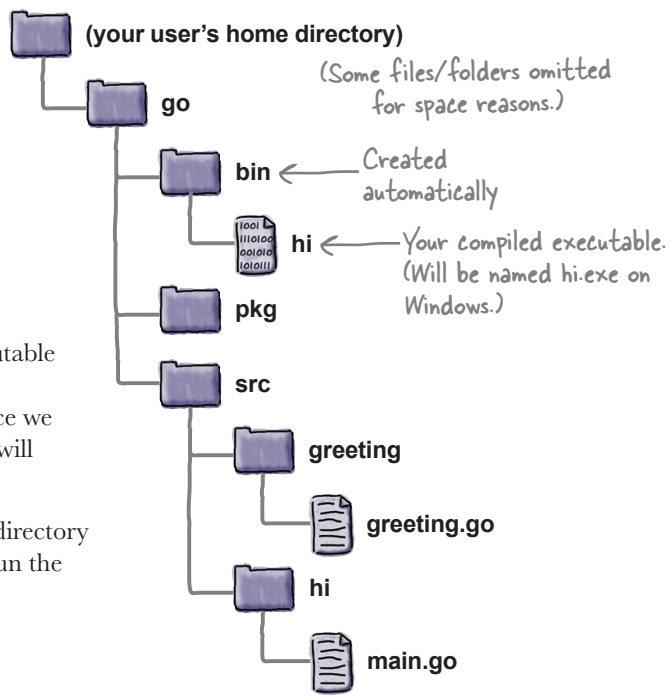
The `go install` command also saves compiled binary versions of executable programs, but in a well-defined, easily accessible place: a *bin* directory in your Go workspace. Just give `go install` the name of a directory within *src* that contains code for an executable program (that is, *.go* files that begin with `package main`). The program will be compiled and an executable will be stored in this standard directory.

(Be sure to pass the name of a directory within "src" to "go install", not the name of a .go file! By default, "go install" isn't set up to handle .go files directly.)

Let's try installing an executable for our *hi/main.go* program. As before, from a terminal, we type **`go install`**, a space, and the name of a folder within our *src* directory (**`hi`**). Again, it doesn't matter what directory you do this from; the `go` tool will look the directory up within the *src* directory.

```
Shell Edit View Window Help
$ go install hi
$
```

When Go sees that the file inside the *hi* directory contains a `package main` declaration, it will know this is code for an executable program. It will compile an executable file, storing it in a directory named *bin* in the Go workspace. (The *bin* directory will be created automatically if it doesn't already exist.)



Unlike the `go build` command, which names an executable after the *.go* file it's based on, `go install` names an executable after the directory that contains the code. Since we compiled the contents of the *hi* directory, the executable will be named `hi` (or `hi.exe` on Windows).

Now, you can use the **`cd`** command to change to the *bin* directory within your Go workspace. Once you're in *bin*, you can run the executable by typing **`./hi`** (or **`hi.exe`** on Windows).

```
Shell Edit View Window Help
$ cd /Users/jay/go/bin
$ ./hi
Hello!
Hi!
Hallo!
Guten Tag!
```

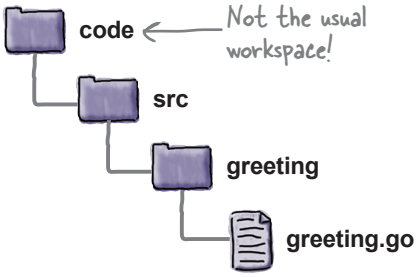
You can also add your workspace's "bin" directory to your system's "PATH" environment variable. Then, you'll be able to run executables in "bin" from anywhere on your system! Recent Go installers for Mac and Windows will update "PATH" for you.

Changing workspaces with the GOPATH environment variable

You may see developers on various websites talking about "setting your `GOPATH`" when discussing the Go workspace. `GOPATH` is an environment variable that Go tools consult to find the location of your workspace. Most Go developers keep all their code in a single workspace, and don't change it from its default location. But if you want, you can use `GOPATH` to move your workspace to a different directory.

An **environment variable** lets you store and retrieve values, kind of like a Go variable, but it's maintained by the operating system, not by Go. You can configure some programs by setting environment variables, and that includes the Go tool.

Suppose that, instead of in your home directory, you had set up your `greeting` package inside a directory named *code* in the root of your hard drive. And now you want to run your *main.go* file, which depends on `greeting`.




```
package main

import "greeting"

func main() {
        greeting.Hello()
        greeting.Hi()
}
```

main.go

But you're getting an error saying the `greeting` package can't be found, because the `go` tool is still looking in the *go* directory in your home directory:

```
Shell Edit View Window Help
$ go run main.go
command.go:3:8: cannot find package "greeting" in any of:
        /usr/local/go/libexec/src/greeting (from $GOROOT)
        /Users/jay/go/src/greeting (from $GOPATH)
```

Setting GOPATH

If your code is stored in a directory other than the default, you'll need to configure the `go` tool to look in the right place. You can do that by setting the `GOPATH` environment variable. How you'll do that depends on your operating system.

On Mac or Linux systems:

You can use the `export` command to set the environment variable. At a terminal prompt, type:

```
export GOPATH="/code"
```

For a directory named *code* in the root of your hard drive, you'll want to use a path of "`/code`". You can substitute a different path if your code is in a different location.

On Windows systems:

You can use the `set` command to set the environment variable. At a command prompt, type:

```
set GOPATH="C:\code"
```

For a directory named *code* in the root of your hard drive, you'll want to use a path of "`C:\code`". You can substitute a different path if your code is in a different location.

Once that's done, `go run` should immediately begin using the directory you specified as its workspace (as should other Go tools). That means the `greeting` library will be found, and the program will run!

On Mac/Linux

```
Shell Edit View Window Help
$ export GOPATH="/code"
$ go run main.go
Hello!
Hi!
```

On Windows

```
Shell Edit View Window Help
C:\Users\jay>set GOPATH="C:\code"
C:\Users\jay>go run main.go
Hello!
Hi!
```

Note that the methods above will only set `GOPATH` for the *current* terminal/command prompt window. You'll need to set it again for each new window you open. But there are ways to set an environment variable permanently, if you want. The methods differ for each operating system, so we don't have space to go into them here. If you type "environment variables" followed by the name of your OS into your favorite search engine, the results should include helpful instructions.

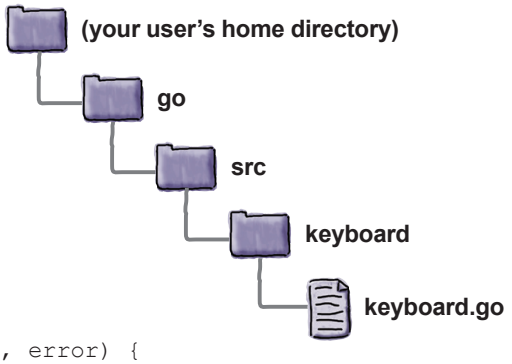
Publishing packages

We're getting so much use out of our `keyboard` package, we wonder if others might find it useful, too.

```
package keyboard

import (
        "bufio"
        "os"
        "strconv"
        "strings"
)

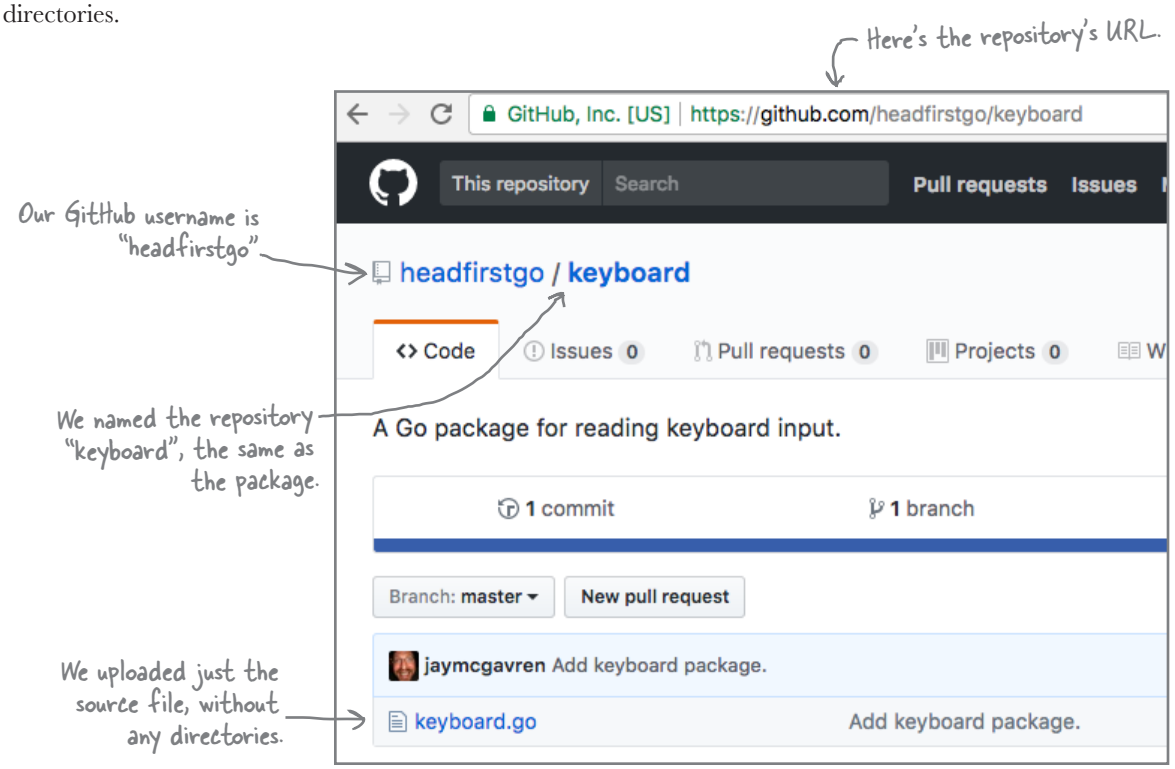
func GetFloat() (float64, error) {
        // GetFloat code here...
}
```



Let's create a repository to hold our code on GitHub, a popular code sharing website. That way, other developers can download it and use it in their own projects! Our GitHub username is `headfirstgo`, and we'll name the repository *keyboard*, so its URL will be:

https://github.com/headfirstgo/keyboard

We'll upload just the *keyboard.go* file to the repository, without nesting it inside any directories.



Publishing packages (continued)



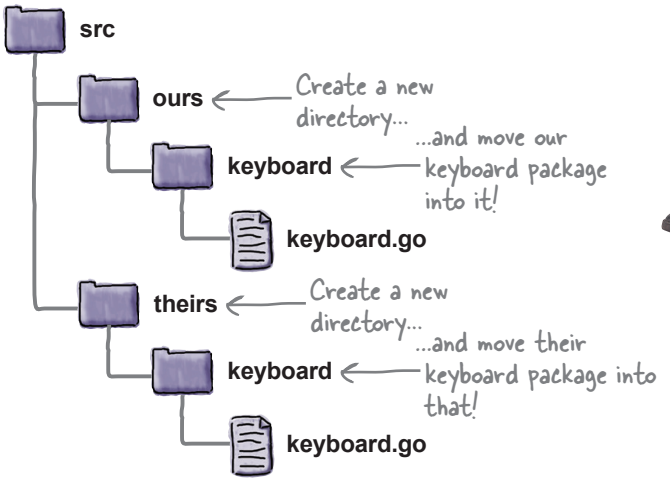
Thanks, but I don't think we can use your package. My music store application already has a `keyboard` package, and if I install **your** `keyboard` package, there will be conflicts!

Hmm, that's a valid concern. There can only be one *keyboard* directory in the Go workspace's *src* directory, and so it *looks* like we can only have one package named `keyboard`!



Wait...what if we nested the directories just like before? We could have one directory to hold **our** `keyboard` package, and another directory to hold **their** `keyboard` package!







Okay, but what do we call the folders that contain the packages? Whose is "ours" and whose is "theirs"?

Publishing packages (continued)

Maybe we need a more universal identifier for the package's author. Our `keyboard` package is the only one available at *http://github.com/headfirstgo/keyboard*, so what if we broke up that URL and used the pieces as directory names?







And then my store can use the URL where **our** `keyboard` package is hosted for **our** directory names. No more conflicts. I like it!

Let's try that: we'll move our package into a directory structure that represents the URL where it's hosted. Inside our *src* directory, we'll create another directory named *github.com*. Inside that, we'll create a directory named after the next segment of the URL, *headfirstgo*. And then we'll move our *keyboard* package directory from the *src* directory into the *headfirstgo* directory.

Although moving the package into a new subdirectory will change its *import path*, it won't change the package *name*. And since the package itself only contains references to the name, we don't have to make any changes to the package code!

```
package keyboard

import (
        "bufio"
        "os"
        "strconv"
        "strings"
)

// More keyboard.go code here...
```

Package name is unchanged, so we don't have to change the package code.



keyboard.go

Publishing packages (continued)

We *will* need to update the programs that rely on our package, though, because the package import path has changed. Because we named each subdirectory after part of the URL where the package is hosted, our new import path looks a lot like that URL:

```
"github.com/headfirstgo/keyboard"
```

We only need to update the `import` statement in each program. Because the package name is the same, references to the package in the rest of the code will be unchanged.

With those changes made, all the programs that rely on our `keyboard` package should resume working normally.

```
// pass_fail reports whether a grade is passing or failing.
package main

import (
        "fmt"
        "github.com/headfirstgo/keyboard"
        "log"
)

func main() {
        fmt.Print("Enter a grade: ")
        grade, err := keyboard.GetFloat()
        if err != nil {
                log.Fatal(err)
        }
        // More code here...
}
```

Update the import path.

No change needed: package name is the same.

```
Enter a grade: 89.7
A grade of 89.7 is passing
```

```
// tocelsius converts a temperature...
package main

import (
        "fmt"
        "github.com/headfirstgo/keyboard"
        "log"
)

func main() {
        fmt.Print("Enter a temperature in Fahrenheit: ")
        fahrenheit, err := keyboard.GetFloat()
        if err != nil {
                log.Fatal(err)
        }
        // More code here...
}
```

Update the import path.

No change needed: package name is the same.

```
Enter a temperature in Fahrenheit: 98.6
37.00 degrees Celsius
```

By the way, we wish we could take credit for this idea of using domain names and paths to ensure a package import path is unique, but we didn't really come up with it. The Go community has been using this as a package naming standard from the beginning. And similar ideas have been used in languages like Java for decades now.

Downloading and installing packages with "go get"

Using a package's hosting URL as an import path has another benefit. The `go` tool has another subcommand named `go get` that can automatically download and install packages for you.

We've set up a Git repository with the `greeting` package that we showed you previously at this URL:

https://github.com/headfirstgo/greeting

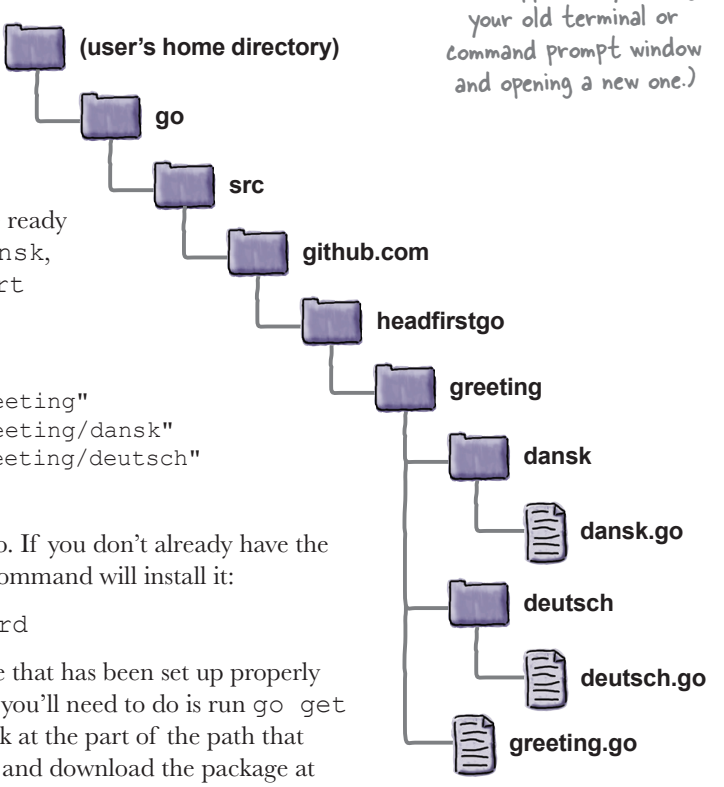
That means that from any computer with Go installed, you can type this in a terminal:

```
go get github.com/headfirstgo/greeting
```

That's `go get` followed by the repository URL, but with the "scheme" portion (the "https://") left off. The `go` tool will connect to *github.com*, download the Git repository at the */headfirstgo/greeting* path, and save it in your Go workspace's *src* directory. (Note: if your system doesn't have Git installed, you'll be prompted to install it when you run the `go get` command. Just follow the instructions on your screen. The `go get` command can also work with Subversion, Mercurial, and Bazaar repositories.)

(Note: "go get" still may not be able to find Git after it's installed. If this happens, try closing your old terminal or command prompt window and opening a new one.)

The `go get` command will automatically create whatever subdirectories are needed to set up the appropriate import path (a *github.com* directory, a *headfirstgo* directory, etc.). The packages saved in the *src* directory will look like this:



With the packages saved in the Go workspace, they're ready for use in programs. You can use the `greeting`, `dansk`, and `deutsch` packages in a program with an `import` statement like this:

```
import (
        "github.com/headfirstgo/greeting"
        "github.com/headfirstgo/greeting/dansk"
        "github.com/headfirstgo/greeting/deutsch"
)
```

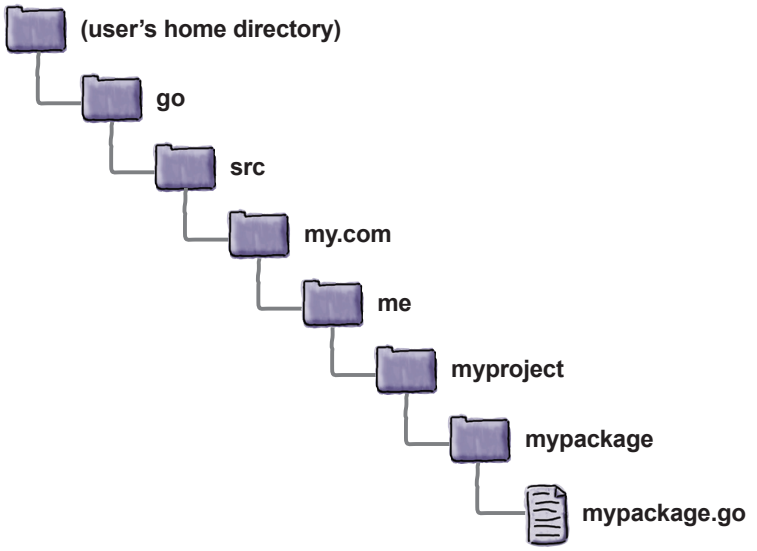
The `go get` command works for other packages, too. If you don't already have the `keyboard` package we showed you previously, this command will install it:

```
go get github.com/headfirstgo/keyboard
```

In fact, the `go get` command works for any package that has been set up properly on a hosting service, no matter who the author is. All you'll need to do is run `go get` and give it the package import path. The tool will look at the part of the path that corresponds to the host address, connect to that host, and download the package at the URL represented by the rest of the import path. It makes using other developers' code really easy!



We've set up a Go workspace with a simple package named `mypackage`. Complete the program below to import `mypackage` and call its `MyFunction` function.




```
package mypackage

func MyFunction() {
}
```

mypackage.go

Your code here:

```
package main

import ____________________________________

func main() {

        ____________________

}
```

Answers on page 147.

Reading package documentation with "go doc"



I installed your `keyboard` package. But I have no idea how to use it! How do I find out what it can do?

You can use the `go doc` command to display documentation on any package or function.

You can get a documentation for a package by passing its import path to `go doc`. For example, we can get info on the `strconv` package by running `go doc strconv`.

Get documentation for strconv package

Package name and import path

(Some output omitted to save space.)

Package description

Included functions

```
Shell Edit View Window Help
$ go doc strconv
package strconv // import "strconv"

Package strconv implements conversions to and from
string representations of basic data types.


Numeric Conversions

The most common numeric conversions are Atoi (string
to int) and Itoa (int to string).

    i, err := strconv.Atoi("-42")
    s := strconv.Itoa(-42)

[...Further description of the package here...]

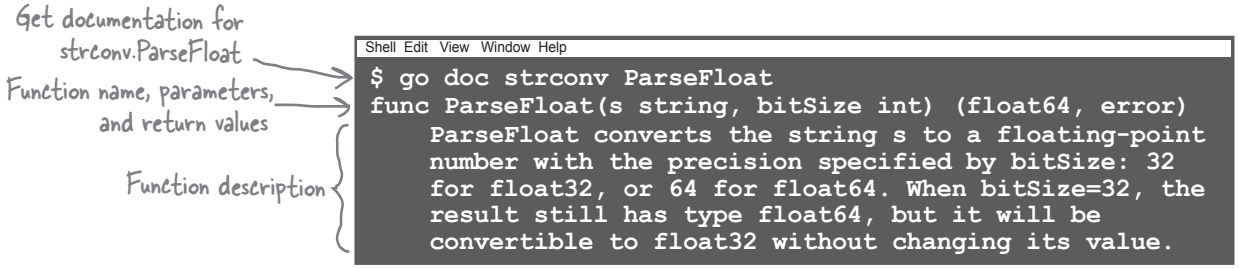
[...Function names...]
func Itoa(i int) string
func ParseBool(str string) (bool, error)
func ParseFloat(s string, bitSize int) (float64, error)
[...More function names...]
```

The output includes the package name and import path (which are one and the same in this case), a description of the package as a whole, and a list of all the functions the package exports.

Reading package documentation with "go doc" (continued)

You can also use `go doc` to get detailed info on specific functions by providing a function name following the package name. Suppose we saw the `ParseFloat` function in the list of the `strconv` package's functions and we wanted to know more about it. We could bring up its documentation with `go doc strconv ParseFloat`.

You'll get back a description of the function and what it does:



The first line looks just like a function declaration would look in code. It includes the function name, followed by parentheses containing the names and types of the parameters it takes (if any). If there are any return values, those will appear after the parameters.

This is followed by a detailed description of what the function does, along with any other information developers need in order to use it.

We can get documentation for our `keyboard` package in the same way, by providing its import path to `go doc`. Let's see if there's anything there that will help our would-be user. From a terminal, run:

```
go doc github.com/headfirstgo/keyboard
```

The `go doc` tool is able to derive basic information like the package name and import path from the code. But there's no package description, so it's not that helpful.



Requesting info on the `GetFloat` function doesn't get us a description either:



Documenting your packages with doc comments

The `go doc` tool works hard to add useful info to its output based on examining the code. Package names and import paths are added for you. So are function names, parameters, and return types.

But `go doc` isn't magic. If you want your users to see documentation of a package or function's intent, you'll need to add it yourself.

Fortunately, that's easy to do: you simply add **doc comments** to your code. Ordinary Go comments that appear immediately before a package clause or function declaration are treated as doc comments, and will be displayed in `go doc`'s output.

Let's try adding doc comments for the `keyboard` package. At the top of the *keyboard.go* file, immediately before the `package` line, we'll add a comment describing what the package does. And immediately before the declaration of `GetFloat`, we'll add a couple comment lines describing that function.

Add ordinary comment lines before the "package" line.

Add ordinary comment lines before a function declaration.

```
// Package keyboard reads user input from the keyboard.
package keyboard

import (
        "bufio"
        "os"
        "strconv"
        "strings"
)

// GetFloat reads a floating-point number from the keyboard.
// It returns the number read and any error encountered.
func GetFloat() (float64, error) {
        // No changes to GetFloat code
}
```

The next time we run `go doc` for the package, it will find the comment before the `package` line and convert it to a package description. And when we run `go doc` for the `GetFloat` function, we'll see a description based on the comment lines we added above `GetFloat`'s declaration.

Package description

```
File  Edit  Window  Help
$ go doc github.com/headfirstgo/keyboard
package keyboard // import "github.com/headfirstgo/keyboard"

Package keyboard reads user input from the keyboard.

func GetFloat() (float64, error)
```

Function description

```
File  Edit  Window  Help
$ go doc github.com/headfirstgo/keyboard GetFloat
func GetFloat() (float64, error)
    GetFloat reads a floating-point number from the
    keyboard. It returns the number read and any error
    encountered.
```

Documenting your packages with doc comments (continued)

Being able to display documentation via `go doc` makes developers that install a package happy.



And doc comments make developers who work on a package's code happy, too! They're ordinary comments, so they're easy to add. And you can easily refer to them while making changes to the code.

Package comment

```
// Package keyboard reads user input from the keyboard.
package keyboard

import (
       "bufio"
       "os"
       "strconv"
       "strings"
)
```

Function comment

```
// GetFloat reads a floating-point number from the keyboard.
// It returns the number read and any error encountered.
func GetFloat() (float64, error) {
       // GetFloat code here
}
```

There are a few conventions to follow when adding doc comments:

- Comments should be complete sentences.
- Package comments should begin with "Package" followed by the package name:

```
// Package mypackage enables widget management.
```

- Function comments should begin with the name of the function they describe:

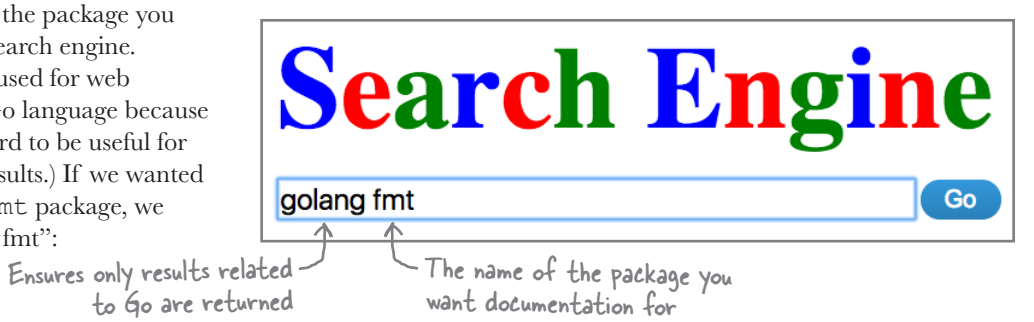
```
// MyFunction converts widgets to gizmos.
```

- You can include code examples in your comments by indenting them.
- Other than indentation for code samples, don't add extra punctuation characters for emphasis or formatting. Doc comments will be displayed as plain text, and should be formatted that way.

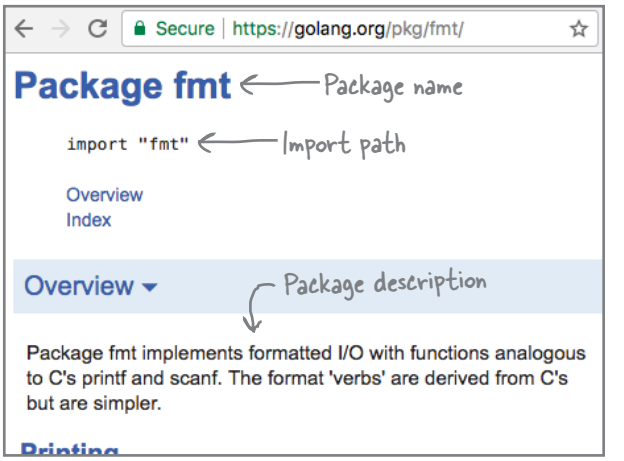
Viewing documentation in a web browser

If you're more comfortable in a web browser than a terminal, there are other ways to view package documentation.

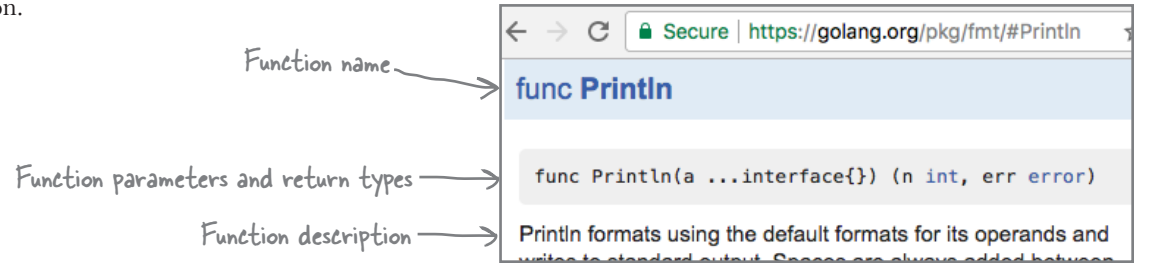
The simplest is to type the word "golang" followed by the name of the package you want into your favorite search engine. ("Golang" is commonly used for web searches regarding the Go language because "go" is too common a word to be useful for filtering out irrelevant results.) If we wanted documentation for the `fmt` package, we could search for "golang fmt":



The results should include sites that offer Go documentation in HTML format. If you're searching for a package in the Go standard library (like `fmt`), one of the top results will probably be from *golang.org*, a site run by the Go development team. The documentation will have much the same contents as the output of the `go doc` tool, with package names, import paths, and descriptions.



One major advantage of the HTML documentation is that each function name in the list of the package's functions will be a handy clickable link leading to the function documentation.

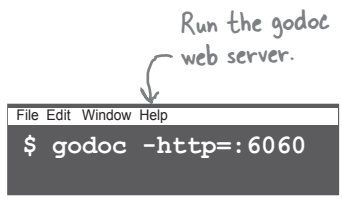


But the content is just the same as what you'd see when running `go doc` in your terminal. It's all based on the same simple doc comments in the code.

Serving HTML documentation to yourself with "godoc"

The same software that powers the *golang.org* site's documentation section is actually available on *your* computer, too. It's a tool called `godoc` (not to be confused with the `go doc` command), and it's automatically installed along with Go. The `godoc` tool generates HTML documentation based on the code in your main Go installation and your workspace. It includes a web server that can share the resulting pages with browsers. (Don't worry, with its default settings `godoc` won't accept connections from any computer other than your own.)

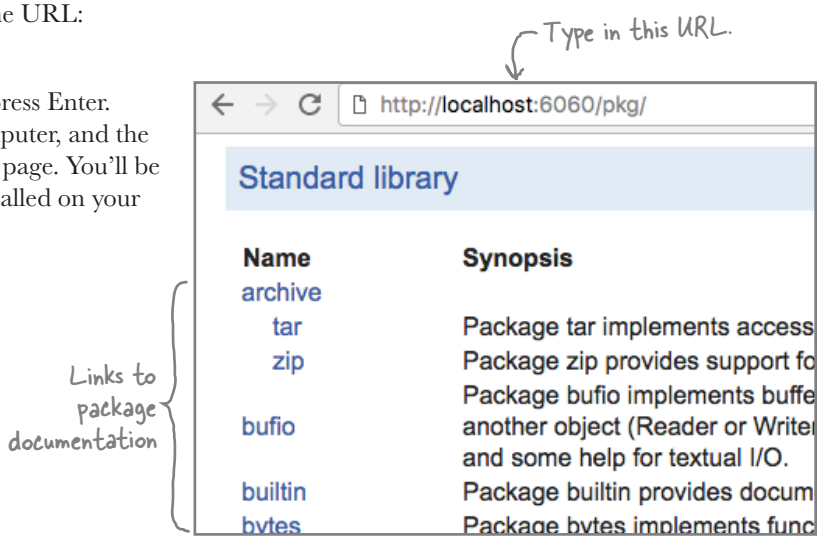
To run `godoc` in web server mode, we'll type the `godoc` command (again, don't confuse that with `go doc`) in a terminal, followed by a special option: `-http=:6060`.



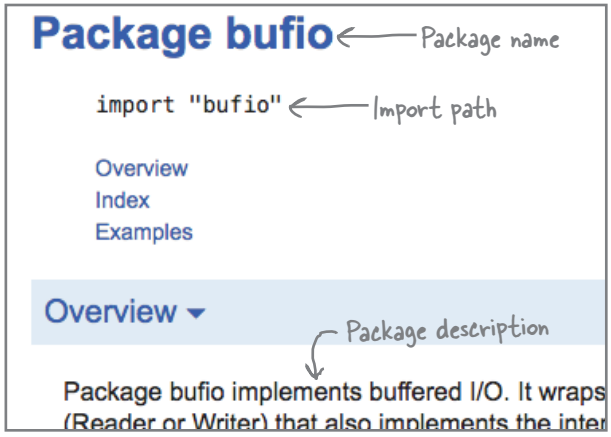
Then with `godoc` running, you can type the URL:

```
http://localhost:6060/pkg
```

...into your web browser's address bar and press Enter. Your browser will connect to your own computer, and the `godoc` server will respond with an HTML page. You'll be presented with a list of all the packages installed on your machine.

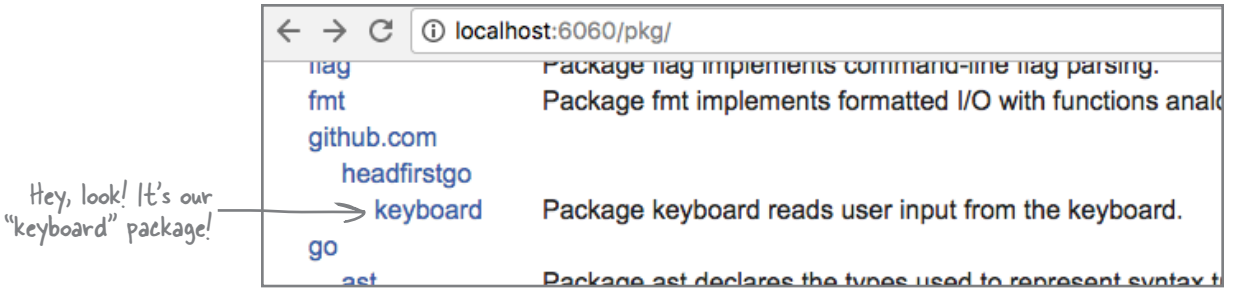


Each package name in the list is a link to that package's documentation. Click it, and you'll see the same package docs that you'd see on *golang.org*.



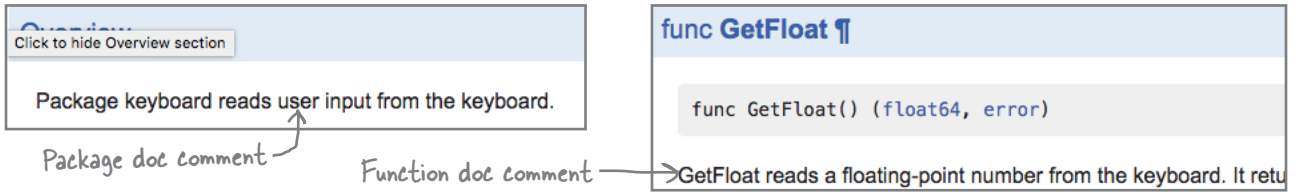
The "godoc" server includes YOUR packages!

If we scroll further through our local `godoc` server's list of packages, we'll see something interesting: our `keyboard` package!

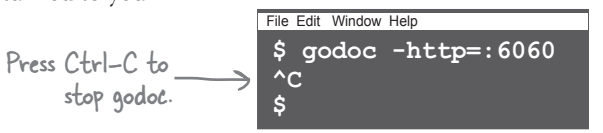


In addition to packages from Go's standard library, the `godoc` tool also builds HTML documentation for any packages in your Go workspace. These could be third-party packages you've installed, or packages you've written yourself.

Click the *keyboard* link, and you'll be taken to the package's documentation. The docs will include any doc comments from our code!



When you're ready to stop the `godoc` server, return to your terminal window, then hold the Ctrl key and press C. You'll be returned to your system prompt.

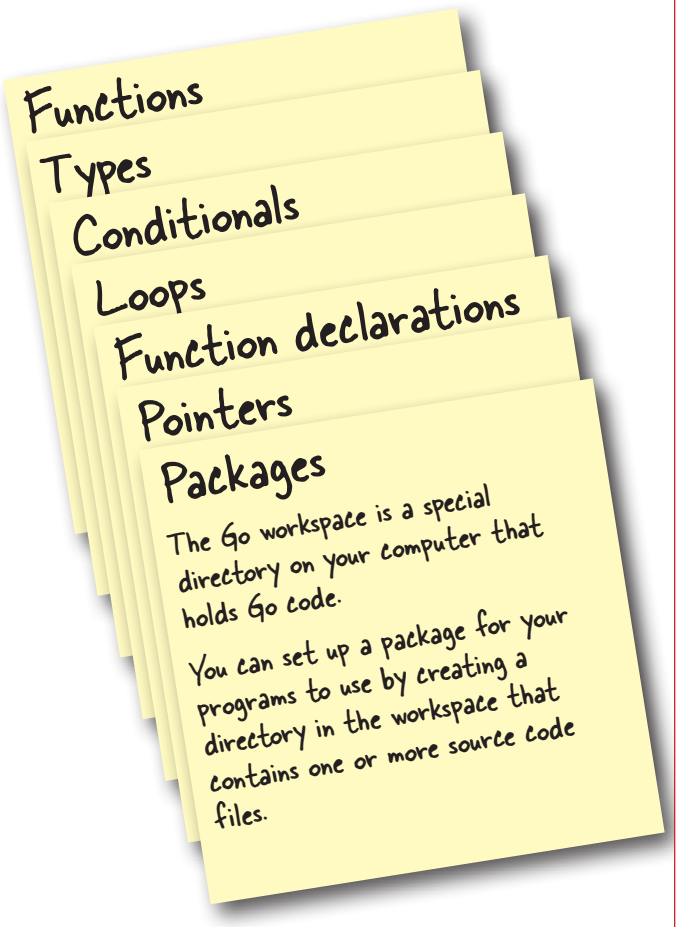


Go makes it easy to document your packages, which makes packages easier to share, which in turn makes them easier for other developers to use. It's just one more feature that makes packages a great way to share code!



Your Go Toolbox

That's it for Chapter 4! You've added packages to your toolbox.

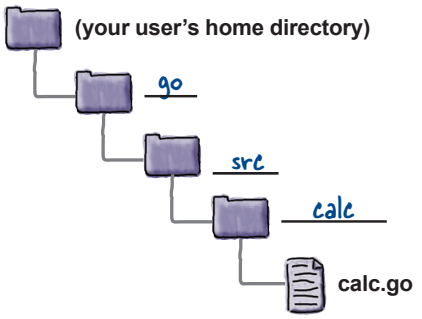


BULLET POINTS

- By default, the workspace directory is a directory named *go* within your user's home directory.
- You can use another directory as your workspace by setting the `GOPATH` environment variable.
- Go uses three subdirectories within the workspace: the *bin* directory holds compiled executable programs, the *pkg* directory holds compiled package code, and the *src* directory holds Go source code.
- The names of the directories within the *src* directory are used to form a package's import path. Names of nested directories are separated by `/` characters in the import path.
- The package's name is determined by the `package` clauses at the top of the source code files within the package directory. Except for the `main` package, the package name should be the same as the name of the directory that contains it.
- Package names should be all lowercase, and ideally consist of a single word.
- A package's functions can only be called from outside that package if they're **exported**. A function is exported if its name begins with a capital letter.
- A **constant** is a name referring to a value that will never change.
- The `go install` command compiles a package's code and stores it in the *pkg* directory for general packages, or the *bin* directory for executable programs.
- A common convention is to use the URL where a package is hosted as its import path. This allows the `go get` tool to find, download, and install packages given only their import path.
- The `go doc` tool displays documentation for packages. Doc comments within the code are included in `go doc`'s output.

Pool Puzzle Solution

Your **job** is to take code snippets from the pool and place them into the blank lines. **Don't** use the same snippet more than once, and you won't need to use all the snippets. Your **goal** is to set up a `calc` package within a Go workspace so `calc`'s functions can be used within *main.go*.




```
package calc
        Make sure the name is capitalized, so the function is exported!
func Add(first float64, second float64) float64 {
        return first + second
}
        Make sure the name is capitalized, so the function is exported!
func Subtract(first float64, second float64) float64 {
        return first - second
}
```

calc.go

```
package main

import (
        "calc"
        "fmt"
)

func main() {
        fmt.Println(calc.Add(1, 2))
        fmt.Println(calc.Subtract(7, 3))
}
```

main.go

Output

```
3
4
```



We've set up a Go workspace with a simple package named `mypackage`. Complete the program below to import `mypackage` and call its `MyFunction` function.

```
package mypackage

func MyFunction() {
}
```

mypackage.go

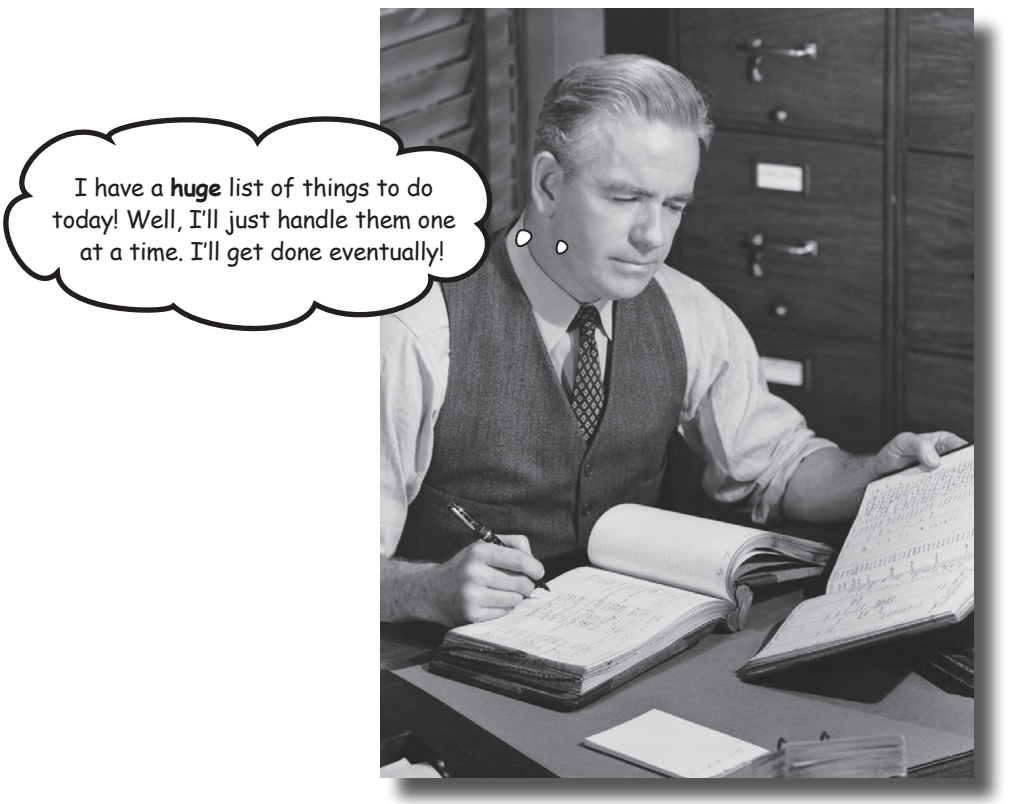
```
package main

import "my.com/me/myproject/mypackage"

func main() {
        mypackage.MyFunction()
}
```

5 on the list

Arrays



A whole lot of programs deal with lists of things. Lists of addresses. Lists of phone numbers. Lists of products. Go has *two* built-in ways of storing lists. This chapter will introduce the first: **arrays**. You'll learn about how to create arrays, how to fill them with data, and how to get that data back out again. Then you'll learn about processing all the elements in array, first the *hard* way with `for` loops, and then the *easy* way with `for...range` loops.

Arrays hold collections of values

A local restaurant owner has a problem. He needs to know how much beef to order for the upcoming week. If he orders too much, the excess will go to waste. If he doesn't order enough, he'll have to tell his customers that he can't make their favorite dishes.

He keeps data on how much meat was used the previous three weeks. He needs a program that will give him some idea of how much to order.



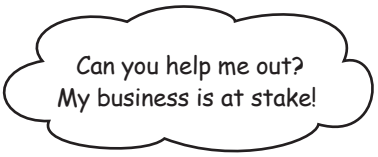
Week A:
71.8 pounds



Week B:
56.2 pounds



Week C:
89.5 pounds



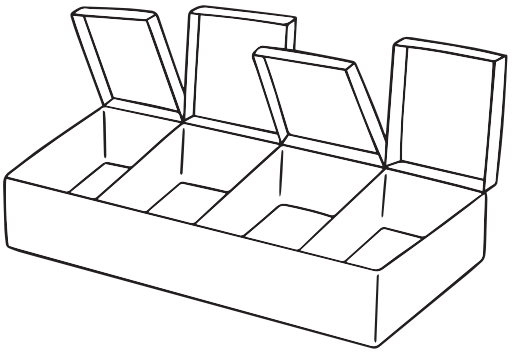


This should be simple enough: we can calculate the average by taking the three amounts, adding them together, and dividing by 3. The average should offer a good estimate of how much to order.

```
(week A + week B + week C) ÷ 3 = average
```

The first issue is going to be storing the sample values. It would be a pain to declare three separate variables, and even more so if we wanted to average more values together later. But, like most programming languages, Go offers a data structure that's perfect for this sort of situation...

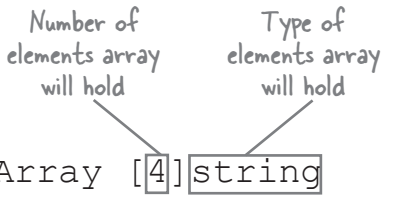
An **array** is a collection of values that all share the same type. Think of it like one of those pill boxes with compartments — you can store and retrieve pills from each compartment separately, but it's also easy to transport the container as a whole.



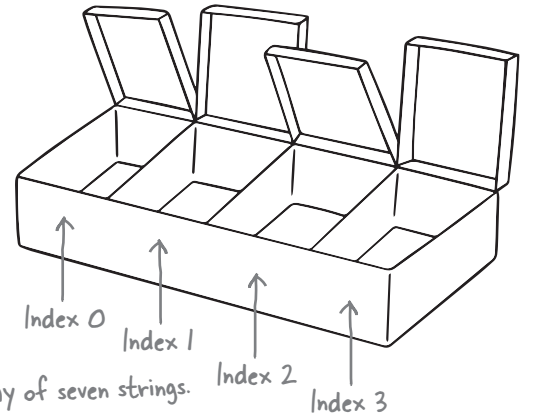
The values an array holds are called its **elements**. You can have an array of strings, an array of booleans, or an array of any other Go type (even an array of arrays). You can store an entire array in a single variable, and then access any element within the array that you need.

Arrays hold collections of values (continued)

An array holds a specific number of elements, and it cannot grow or shrink. To declare a variable that holds an array, you need to specify the number of elements it holds in square brackets (`[]`), followed by the type of elements the array holds.



To set the array elements' values or to retrieve values later, you'll need a way to specify which element you mean. Elements in an array are numbered, starting with 0. An element's number is called its **index**.



If you wanted to make an array with the names of notes on a musical scale, for example, the first note would be assigned to index `0`, the second note would be at index `1`, and so forth. The index is specified in square brackets.

```
var notes [7]string
notes[0] = "do"
notes[1] = "re"
notes[2] = "mi"
fmt.Println(notes[0])
fmt.Println(notes[1])
```

Create an array of seven strings.
Assign a value to the first element.
Assign a value to the second element.
Assign a value to the third element.
Print the first element.
Print the second element.

```
do
re
```

Here's an array of integers:

```
var primes [5]int
primes[0] = 2
primes[1] = 3
fmt.Println(primes[0])
```

Create an array of five integers.
Assign a value to the first element.
Assign a value to the second element.
Print the first element.

```
2
```

And an array of `time.Time` values:

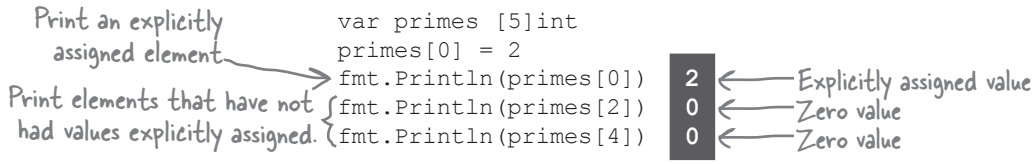
```
var dates [3]time.Time
dates[0] = time.Unix(1257894000, 0)
dates[1] = time.Unix(1447920000, 0)
dates[2] = time.Unix(1508632200, 0)
fmt.Println(dates[1])
```

Create an array of three Time values.
Assign a value to the first element.
Assign a value to the second element.
Assign a value to the third element.
Print the second element.

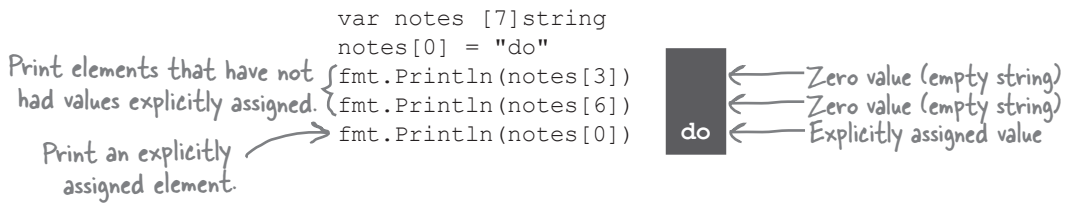
```
2015-11-19 08:00:00 +0000 UTC
```

Zero values in arrays

As with variables, when an array is created, all the values it contains are initialized to the zero value for the type that array holds. So an array of `int` values is filled with zeros by default:

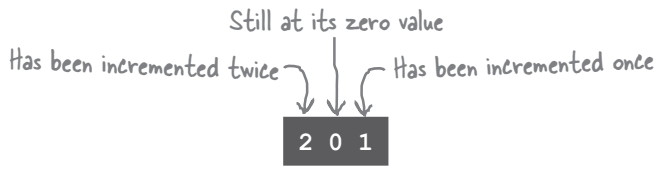


The zero value for strings, however, is an empty string, so an array of `string` values is filled with empty strings by default:



Zero values can make it safe to manipulate an array element even if you haven't explicitly assigned a value to it. For example, here we have an array of integer counters. We can increment any of them without explicitly assigning a value first, because we know they will all start from `0`.

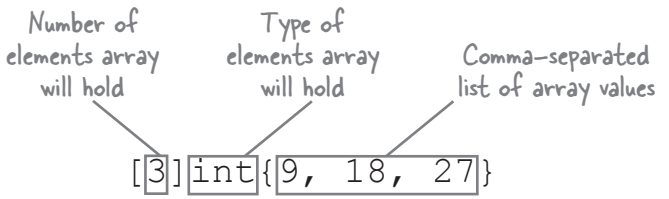
```
var counters [3]int
counters[0]++    Increment the first element from 0 to 1.
counters[0]++    Increment the first element from 1 to 2.
counters[2]++    Increment the third element from 0 to 1.
fmt.Println(counters[0], counters[1], counters[2])
```



When an array is created, all the values it contains are initialized to the zero value for the type the array holds.

Array literals

If you know in advance what values an array should hold, you can initialize the array with those values using an **array literal**. An array literal starts just like an array type, with the number of elements it will hold in square brackets, followed by the type of its elements. This is followed by a list in curly braces of the initial values each element should have. The element values should be separated by commas.



These examples are just like the previous ones we showed, except that instead of assigning values to the array elements one by one, the entire array is initialized using array literals.

```
var notes [7]string = [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}
fmt.Println(notes[3], notes[6], notes[0])
var primes [5]int = [5]int{2, 3, 5, 7, 11}
fmt.Println(primes[0], primes[2], primes[4])
```

Assign values using an array literal.

Assign values using an array literal.

```
fa ti do
2 5 11
```

Using an array literal also allows you to do short variable declarations with `:=`.

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}
primes := [5]int{2, 3, 5, 7, 11}
```

Short variable declaration

Short variable declaration

You can spread array literals over multiple lines, but you're required to use a comma before each newline character in your code. You'll even need a comma following the final entry in the array literal, if it's followed by a newline. (This style looks awkward at first, but it makes it easier to add more elements to the code later.)

```
text := [3]string{
       "This is a series of long strings",
       "which would be awkward to place",
       "together on a single line",
}
```

This is all one array.

This comma at the end is required.



Below is a program that declares a couple arrays and prints out their elements. Write down what the program output would be.

```
package main

import "fmt"

func main() {
        var numbers [3]int
        numbers[0] = 42
        numbers[2] = 108
        var letters = [3]string{"a", "b", "c"}

                                        Output:

        fmt.Println(numbers[0])         .....

        fmt.Println(numbers[1])         .....

        fmt.Println(numbers[2])         .....

        fmt.Println(letters[2])         .....

        fmt.Println(letters[0])         .....

        fmt.Println(letters[1])         .....
}
```

Answers on page 173.

Functions in the "fmt" package know how to handle arrays

When you're just trying to debug code, you don't have to pass array elements to `Println` and other functions in the `fmt` package one by one. Just pass the entire array. There's logic in the `fmt` package to format and print the array for you. (The `fmt` package can also handle slices, maps, and other data structures we'll see later.)

```
var notes [3]string = [3]string{"do", "re", "mi"}
var primes [5]int = [5]int{2, 3, 5, 7, 11}
fmt.Println(notes)
fmt.Println(primes)
```

Pass entire arrays to fmt.Println.

```
[do re mi]
[2 3 5 7 11]
```

You may also remember the `"%#v"` verb used by the `Printf` and `Sprintf` functions, which formats values as they'd appear in Go code. When formatted by `"%#v"`, arrays appear in the result as Go array literals.

```
fmt.Printf("%#v\n", notes)
fmt.Printf("%#v\n", primes)
```

Format arrays as they would appear in Go code.

```
[3]string{"do", "re", "mi"}
[5]int{2, 3, 5, 7, 11}
```

Accessing array elements within a loop

You don't have to explicitly write the integer index of the array element you're accessing in your code. You can also use the value in an integer variable as the array index.

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}
index := 1
fmt.Println(index, notes[index])
index = 3
fmt.Println(index, notes[index])
```

Print the array element at index 1.

Print the array element at index 3.

```
1 re
3 fa
```

That means you can do things like process elements of an array using a `for` loop. You loop through indexes in the array, and use the loop variable to access the element at the current index.

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}
for i := 0; i <= 2; i++ {
        fmt.Println(i, notes[i])
}
```

Loop through indexes 0, 1, and 2.

Print the element at the current index.

```
0 do
1 re
2 mi
```

When accessing array elements using a variable, you need to be careful which index values you use. As we mentioned, arrays hold a specific number of elements. Trying to access an index that is outside the array will cause a **panic**, an error that occurs while your program is running (as opposed to when it's compiling).

The array only has seven elements.

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}

for i := 0; i <= 7; i++ {
        fmt.Println(i, notes[i])
}
```

Loops up through index 7 (the eighth element), which doesn't exist!

Normally, a panic causes your program to crash and display an error message to the user. Needless to say, panics should be avoided whenever possible.

Access indexes 0 through 6.

Accessing index 7 causes a panic!

```
0 do
1 re
2 mi
3 fa
4 so
5 la
6 ti
panic: runtime error: index out of range

goroutine 1 [running]:
main.main()
        /tmp/sandbox732328648/main.go:8 +0x140
```

Checking array length with the "len" function

Writing loops that only access valid array indexes can be somewhat error-prone. Fortunately, there are a couple ways to make the process easier.

The first is to check the actual number of elements in the array before accessing it. You can do this with the built-in `len` function, which returns the length of the array (the number of elements it contains).

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}
fmt.Println(len(notes))
primes := [5]int{2, 3, 5, 7, 11}
fmt.Println(len(primes))
```

Print the length of the "notes" array.

Print the length of the "primes" array.

```
7
5
```

When setting up a loop to process an entire array, you can use `len` to determine which indexes are safe to access.

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}

for i := 0; i < len(notes); i++ {
        fmt.Println(i, notes[i])
}
```

The highest value the "i" variable will reach is 6.

Returns the length of the array, 7

```
0 do
1 re
2 mi
3 fa
4 so
5 la
6 ti
```

This still has the potential for mistakes, though. If `len(notes)` returns `7`, the highest index you can access is `6` (because array indexes start at `0`, not `1`). If you try to access index `7`, you'll get a panic.

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}

for i := 0; i <= len(notes); i++ {
        fmt.Println(i, notes[i])
}
```

The highest value the "i" variable will reach is 7!

Returns the length of the array, 7

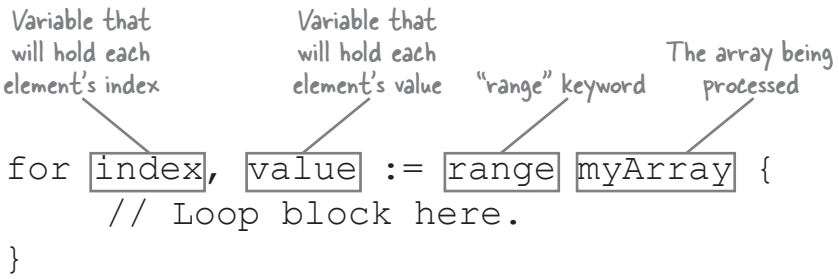
```
0 do
1 re
2 mi
3 fa
4 so
5 la
6 ti
panic: runtime error: index out of range

goroutine 1 [running]:
main.main()
      /tmp/sandbox094804331/main.go:11 +0x140
```

Accessing index 7 causes a panic!

Looping over arrays safely with "for...range"

An even safer way to process each element of an array is to use the special `for...range` loop. In the `range` form, you provide a variable that will hold the integer index of each element, another variable that will hold the value of the element itself, and the array you want to loop over. The loop will run once for each element in the array, assigning the element's index to your first variable and the element's value to your second variable. You can add code to the loop block to process those values.



This form of the `for` loop has no messy init, condition, and post expressions. And because the element value is automatically assigned to a variable for you, there's no risk that you'll accidentally access an invalid array index. Because it's safer and easier to read, you'll see the `for` loop's `range` form used most often when working with arrays and other collections.

Here's our previous code that prints each value in our array of musical notes, updated to use a `for ... range` loop:

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}
```

Variable to hold each index

Variable to hold each string

Process each value in the array.

```
for index, note := range notes {
        fmt.Println(index, note)
}
```

```
0 do
1 re
2 mi
3 fa
4 so
5 la
6 ti
```

The loop runs seven times, once for each element of the `notes` array. For each element, the `index` variable gets set to the element's index, and the `note` variable gets set to the element's value. Then we print the index and value.

Using the blank identifier with "for...range" loops

As always, Go requires that you use every variable you declare. If we stop using the `index` variable from our `for...range` loop, we'll get a compile error:

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}

for index, note := range notes {
        fmt.Println(note)
}
```

The "index" variable has been removed from the output.

Compile error

```
index declared and not used
```

And the same would be true if we didn't use the variable that holds the element value:

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}

for index, note := range notes {
        fmt.Println(index)
}
```

Doesn't use the "note" variable

Compile error

```
note declared and not used
```

Remember in Chapter 2, when we were calling a function with multiple return values, and we wanted to ignore one of them? We assigned that value to the blank identifier (_), which causes Go to discard that value, without giving a compiler error...

We can do the same with values from `for...range` loops. If we don't need the index for each array element, we can just assign it to the blank identifier:

Use the blank identifier as a placeholder for the index value.

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}

for _, note := range notes {
        fmt.Println(note)
}
```

Use only the "note" variable.

```
do
re
mi
fa
so
la
ti
```

And if we don't need the value variable, we can assign that to the blank identifier instead:

Use the blank identifier as a placeholder for the element value.

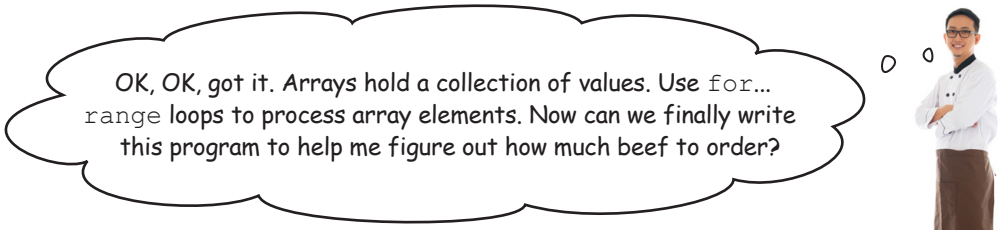
```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}

for index, _ := range notes {
        fmt.Println(index)
}
```

Use only the "index" variable.

```
0
1
2
3
4
5
6
```

Getting the sum of the numbers in an array



We finally know everything we need to create an array of `float64` values and calculate their average. Let's take the amounts of beef that were used in previous weeks, and incorporate them into a program, named `average`.

The first thing we'll need to do is set up a program file. In your Go workspace directory (the *go* directory within your user's home directory, unless you've set the `GOPATH` environment variable), create the following nested directories (if they don't already exist). Within the innermost directory, *average*, save a file named *main.go*.





Now let's write our program code within the *main.go* file. Since this will be an executable program, our code will be part of the `main` package, and will reside in the `main` function.

We'll start by just calculating the total for the three sample values; we can go back later to calculate the average. We use an array literal to create an array of three `float64` values, prepopulated with the sample values from prior weeks. We declare a `float64` variable named `sum` to hold the total, starting with a value of `0`.

Then we use a `for...range` loop to process each number. We don't need the element indexes, so we discard them using the `_` blank identifier. We add each number to the value in `sum`. After we've totaled all the values, we print `sum` before exiting.

```
// average calculates the average of several numbers.
package main

import "fmt"

func main() {
        numbers := [3]float64{71.8, 56.2, 89.5}
        var sum float64 = 0
        for _, number := range numbers {
                sum += number
        }
        fmt.Println(sum)
}
```

This will be an executable program, so we use the "main" package.

Use an array literal to create an array with the three float64 values we're averaging.

Declare a float64 variable to hold the sum of the three numbers.

Discard the element index.

Loop through each number in the array.

Add the current number to the total.

Getting the sum of the numbers in an array (continued)

Let's try compiling and running our program. We'll use the `go install` command to create an executable. We're going to need to provide our executable's import path to `go install`. If we used this directory structure...

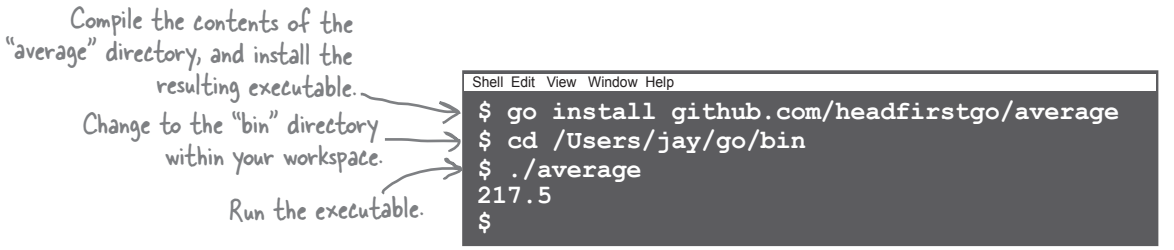


...that means the import path for our package will be `github.com/headfirstgo/average`. So, from your terminal, type:

```
go install github.com/headfirstgo/average
```

You can do so from within any directory. The `go` tool will look for a *github.com/headfirstgo/average* directory within your workspace's *src* directory, and compile any *.go* files it contains. The resulting executable will be named `average`, and will be stored in the *bin* directory within your Go workspace.

Then, you can use the **`cd`** command to change to the *bin* directory within your Go workspace. Once you're in *bin*, you can run the executable by typing **`./average`** (or **`average.exe`** on Windows).



The program will print the total of the three values from our array and exit.

Getting the average of the numbers in an array

We've got our `average` program printing the total of the array's values, so now let's update it to print the actual average. To do that, we'll divide the total by the array's length.

Passing the array to the `len` function returns an `int` value with the array length. But since the total in the `sum` variable is a `float64` value, we'll need to convert the length to a `float64` as well so we can use them together in a math operation. We store the result in the `sampleCount` variable. Once that's done, all we have to do is divide `sum` by `sampleCount`, and print the result.

```
// average calculates the average of several numbers.
package main

import "fmt"

func main() {
        numbers := [3]float64{71.8, 56.2, 89.5}
        var sum float64 = 0
        for _, number := range numbers {
                sum += number
        }
        sampleCount := float64(len(numbers))
        fmt.Printf("Average: %0.2f\n", sum/sampleCount)
}
```

Get the array length as an int and convert it to a float64.

Divide the total of the array's values by the array length to get the average.

Once the code is updated, we can repeat the previous steps to see the new result: run **`go install`** to recompile the code, change to the *bin* directory, and run the updated `average` executable. Instead of the sum of the values in the array, we'll see the average.

```
Shell Edit View Window Help
$ go install github.com/headfirstgo/average
$ cd /Users/jay/go/bin
$ ./average
Average: 72.50
$
```

The average of the array values



Pool Puzzle

Your **job** is to take code snippets from the pool and place them into the blank lines in this code. **Don't** use the same snippet more than once, and you won't need to use all the snippets. Your **goal** is to make a program that will print the index and value of all the array elements that fall between `10` and `20` (it should match the output shown).

```
package main

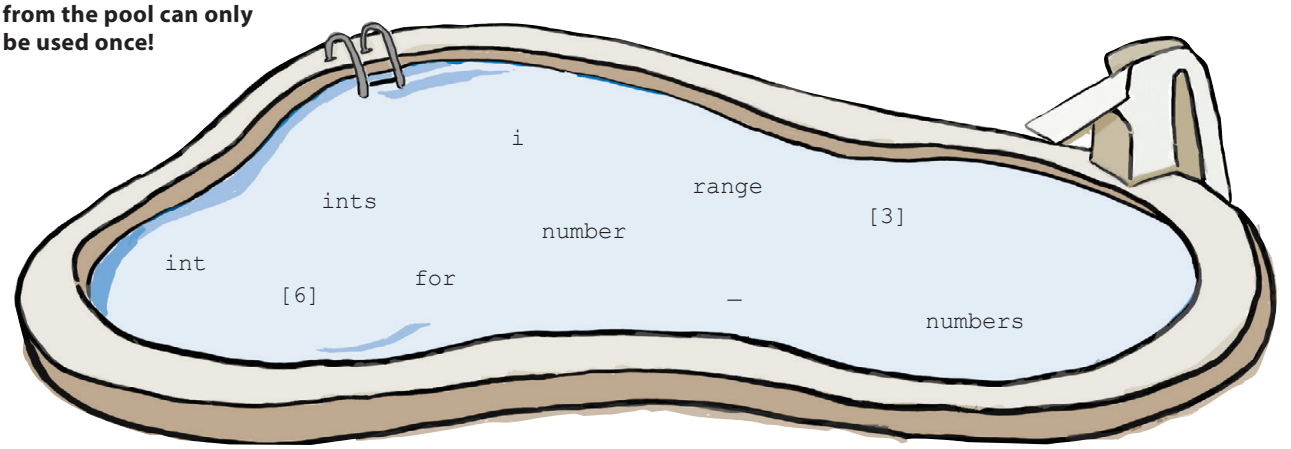
import "fmt"

func main() {
        _________ := ____int{3, 16, -2, 10, 23, 12}
        for i, _________ := ________ numbers {
                if number >= 10 && number <= 20 {
                        fmt.Println(___, number)
                }
        }
}
```

Output

```
1 16
3 10
5 12
```

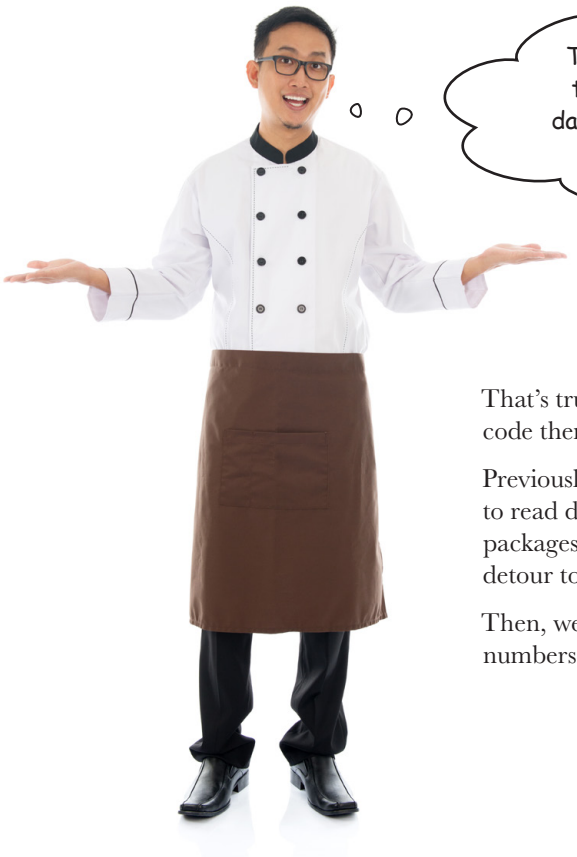
Note: each snippet from the pool can only be used once!



Answers on page 173.

Reading a text file





That's great, but your program only tells me how much to order for **this** week. What should I do when I have data for more weeks? I can't edit the code to change the array values; I don't even have Go installed!

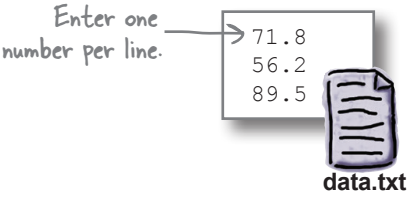
That's true—a program where users have to edit and compile the source code themselves isn't very user-friendly.

Previously, we've used the standard library's `os` and `bufio` packages to read data a line at a time from the keyboard. We can use the same packages to read data a line at a time from text files. Let's go on a brief detour to learn how to do that.

Then, we'll come back and update the `average` program to read its numbers in from a text file.

In your favorite text editor, create a new file named *data.txt*. Save it somewhere *outside* of your Go workspace directory for now.

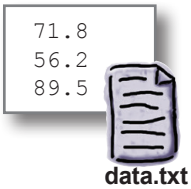
Within the file, enter our three floating-point sample values, one number per line.



Reading a text file (continued)



Before we can update our program to average numbers from a text file, we need to be able to read the file's contents. To start, let's write a program that only reads the file, and then we'll incorporate what we learn into our averaging program.



In the same directory as *data.txt*, create a new program named *readfile.go*. We'll just be running *readfile.go* with `go run`, rather than installing it, so it's okay to save it outside of your Go workspace directory. Save the following code in *readfile.go*. (We'll take a closer look at how this code works on the next page.)




```
package main

import (
        "bufio"
        "fmt"
        "log"
        "os"
)

func main() {
        file, err := os.Open("data.txt")
        if err != nil {
                log.Fatal(err)
        }
        scanner := bufio.NewScanner(file)
        for scanner.Scan() {
                fmt.Println(scanner.Text())
        }
        err = file.Close()
        if err != nil {
                log.Fatal(err)
        }
        if scanner.Err() != nil {
                log.Fatal(scanner.Err())
        }
}
```

Open the data file for reading.

If there was an error opening the file, report it and exit.

Create a new Scanner for the file.

Loops until the end of the file is reached and scanner.Scan returns false

Read a line from the file.

Print the line.

Close the file to free resources.

If there was an error closing the file, report it and exit.

If there was an error scanning the file, report it and exit.

Then, from your terminal, change to the directory where you saved the two files, and run **`go run readfile.go`**. The program will read the contents of *data.txt*, and print them out.

Change to the directory you saved data.txt and readfile.go in.

Run readfile.go.

The contents of data.txt will be printed.

```
Shell  Edit  View  Window  Help
$ cd /Users/jay/code
$ go run readfile.go
71.8
56.2
89.5
```

Reading a text file (continued)



Our test *readfile.go* program is successfully reading the lines of the *data.txt* file and printing them out. Let's take a closer look at how the program works.

We start by passing a string with the name of the file we want to open to the `os.Open` function. Two values are returned from `os.Open`: a pointer to an `os.File` value representing the opened file, and an `error` value. As we've seen with so many other functions, if the `error` value is `nil` it means the file was opened successfully, but any other value means there was an error. (This could happen if the file is missing or unreadable.) If that's the case, we log the error message and exit the program.

Open the data file for reading.

```
file, err := os.Open("data.txt")
if err != nil {
        log.Fatal(err)
}
```

If there was an error opening the file, report it and exit.

Then we pass the `os.File` value to the `bufio.NewScanner` function. That will return a `bufio.Scanner` value that reads from the file.

Create a new Scanner for the file.

```
scanner := bufio.NewScanner(file)
```

The `Scan` method on `bufio.Scanner` is designed to be used as part of a `for` loop. It will read a single line of text from the file, returning `true` if it read data successfully and `false` if it did not. If `Scan` is used as the condition on a `for` loop, the loop will continue running as long as there is more data to be read. Once the end of the file is reached (or there's an error), `Scan` will return `false`, and the loop will exit.

After calling the `Scan` method on the `bufio.Scanner`, calling the `Text` method returns a string with the data that was read. For this program, we simply call `Println` within the loop to print each line out.

Loops until the end of the file is reached and scanner.Scan returns false

```
for scanner.Scan() {
        fmt.Println(scanner.Text())
}
```

Read a line from the file.

Print the line.

Once the loop exits, we're done with the file. Keeping files open consumes resources from the operating system, so files should always be closed when a program is done with them. Calling the `Close` method on the `os.File` will accomplish this. Like the `Open` function, the `Close` method returns an `error` value, which will be `nil` unless there was a problem. (Unlike `Open`, `Close` returns only a *single* value, as there is no useful value for it to return other than the error.)

Close the file to free resources.

```
err = file.Close()
if err != nil {
        log.Fatal(err)
}
```

If there was an error closing the file, report it and exit.

It's also possible that the `bufio.Scanner` encountered an error while scanning through the file. If it did, calling the `Err` method on the scanner will return that error, which we log before exiting.

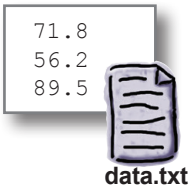
If there was an error scanning the file, report it and exit.

```
if scanner.Err() != nil {
        log.Fatal(scanner.Err())
}
```



Reading a text file into an array

Our *readfile.go* program worked great—we were able to read the lines from our *data.txt* file in as strings, and print them out. Now we need to convert those strings to numbers and store them in an array. Let's create a package named `datafile` that will do this for us.



In your Go workspace directory, create a *datafile* directory within the *headfirstgo* directory. Within the *datafile* directory, save a file named *floats.go*. (We name it *floats.go* because this file will contain code that reads floating-point numbers from files.)



Within *floats.go*, save the following code. A lot of this is based on code from our test *readfile.go* program; we've grayed out the parts where the code is identical. We'll explain the new code in detail on the next page.

```
// Package datafile allows reading data samples from files.
package datafile

import (
        "bufio"
        "os"
        "strconv"
)

// GetFloats reads a float64 from each line of a file.
func GetFloats(fileName string) ([3]float64, error) {
        var numbers [3]float64
        file, err := os.Open(fileName)
        if err != nil {
                return numbers, err
        }
        i := 0
        scanner := bufio.NewScanner(file)
        for scanner.Scan() {
                numbers[i], err = strconv.ParseFloat(scanner.Text(), 64)
                if err != nil {
                        return numbers, err
                }
                i++
        }
        err = file.Close()
        if err != nil {
                return numbers, err
        }
        if scanner.Err() != nil {
                return numbers, scanner.Err()
        }
        return numbers, nil
}
```

Take the file name to read as an argument.

The function will return an array of numbers and any error encountered.

Declare the array we'll be returning.

Open the provided filename.

If there was an error opening the file, return it.

This variable will track which array index we should assign to.

If there was an error converting the line to a number, return it.

Convert the file line string to a float64.

Move to the next array index.

If there was an error closing the file, return it.

If there was an error scanning the file, return it.

If we got this far, there were no errors, so return the array of numbers and a "nil" error.

Reading a text file into an array (continued)

We want to be able to read from files other than *data.txt*, so we accept the name of the file we should open as a parameter. We set the function up to return two values, an array of `float64` values and an `error` value. Like most functions that return an error, the first return value should only be considered usable if the error value is `nil`.

Take the filename to read as an argument.

The function will return an array of numbers and any error encountered.

```
func GetFloats(fileName string) ([3]float64, error) {
```

Next we declare an array of three `float64` values that will hold the numbers we read from the file.

```
var numbers [3]float64
```
← Declare the array we'll be returning.

Just like in *readfile.go*, we open the file for reading. The difference is that instead of a hardcoded string of `"data.txt"`, we open whatever filename was passed to the function. If an error is encountered, we need to return an array along with the error value, so we just return the `numbers` array (even though nothing has been assigned to it yet).

```
file, err := os.Open(fileName)
```
← Open the provided filename.

If there was an error opening the file, return it.

```
if err != nil {
    return numbers, err
}
```

We need to know which array element to assign each line to, so we create a variable to track the current index.

```
i := 0
```
← This variable will track which array index we should assign to.

The code to set up a `bufio.Scanner` and loop over the file's lines is identical to the code from *readfile.go*. The code within the loop is different, however: we need to call `strconv.ParseFloat` on the string read from the file to convert it to a `float64`, and assign the result to the array. If `ParseFloat` results in an error, we need to return that. And if the parsing is successful, we need to increment `i` so that the next number is assigned to the next array element.

```
numbers[i], err = strconv.ParseFloat(scanner.Text(), 64)
```
Convert the file line string to a float64.

If there was an error converting the line to a number, return it.

```
if err != nil {
    return numbers, err
}
```

```
i++
```
← Move to the next array index.

Our code to close the file and report any errors is identical to *readfile.go*, except that we return any errors instead of exiting the program directly. If there are no errors, the end of the `GetFloats` function will be reached, and the array of `float64` values will be returned along with a `nil` error.

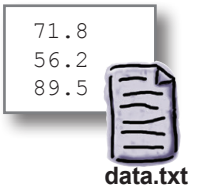
If there was an error scanning the file, return it.

```
if scanner.Err() != nil {
    return numbers, scanner.Err()
}
```

```
return numbers, nil
```
← If we got this far, there were no errors, so return the array of numbers and a "nil" error.

Updating our "average" program to read a text file

We're ready to replace the hardcoded array in our `average` program with an array read in from the *data.txt* file!



Writing our `datafile` package was the hard part. Here in the main program, we only need to do three things:

- Update our `import` declaration to include the `datafile` and `log` packages.
- Replace our array of hardcoded numbers with a call to `datafile.GetFloats("data.txt")`.
- Check whether we got an error back from `GetFloats`, and log it and exit if so.

All the remaining code will be exactly the same.




```
// average calculates the average of several numbers.
package main

import (
        "fmt"
        "github.com/headfirstgo/datafile"
        "log"
)

func main() {
        numbers, err := datafile.GetFloats("data.txt")
        if err != nil {
                log.Fatal(err)
        }
        var sum float64 = 0
        for _, number := range numbers {
                sum += number
        }
        sampleCount := float64(len(numbers))
        fmt.Printf("Average: %0.2f\n", sum/sampleCount)
}
```

Import our package.

Import the "log" package.

Load data.txt, parse the numbers it contains, and store the array.

If there was an error, report it and exit.

Updating our "average" program to read a text file (continued)

We can compile the program using the same terminal command as before:

```
go install github.com/headfirstgo/average
```

Since our program imports the `datafile` package, that will automatically be compiled as well.

Compiles both the "average" program and the "datafile" package it depends on.

```
Shell Edit View Window Help
$ go install github.com/headfirstgo/average
```

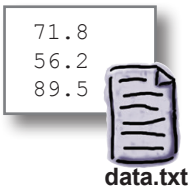
We'll need to move the *data.txt* file to the *bin* subdirectory of the Go workspace. That's because we'll be running the `average` executable from that directory, and it will look for *data.txt* in the same directory. Once you've moved *data.txt*, change into that *bin* subdirectory.

Move the data.txt file to the "bin" subdirectory of the workspace. (Use the appropriate command for your system, or resave it using your text editor.)

```
Shell Edit View Window Help
$ mv data.txt /Users/jay/go/bin
$ cd /Users/jay/go/bin
```

Change to the "bin" subdirectory.

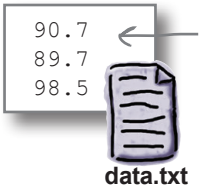
When we run the `average` executable, it will load the values from *data.txt* into an array, and use them to calculate the average.



The average of the data.txt values.

```
Shell Edit View Window Help
$ ./average
Average: 72.50
```

If we change the values in *data.txt*, the average will change as well.



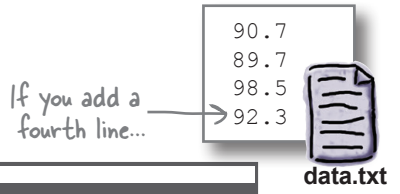
Change the data...

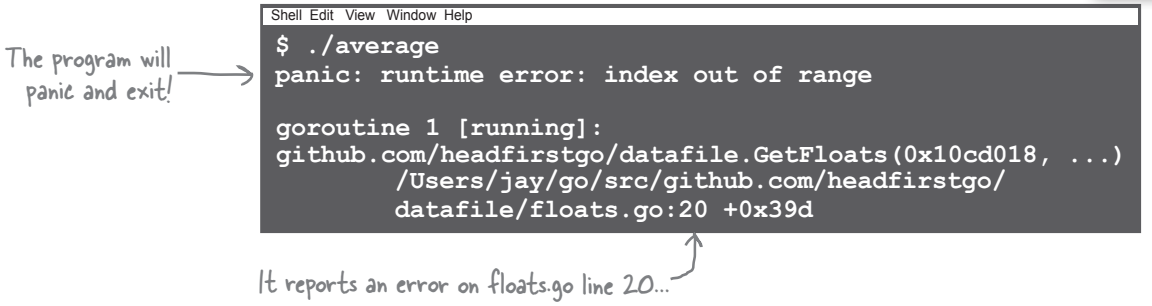
The average is updated.

```
Shell Edit View Window Help
$ ./average
Average: 92.97
```

Our program can only process three values!

But there's a problem—the `average` program only runs if there are three or fewer lines in *data.txt*. If there are four or more, `average` will panic and exit when it's run!





When a Go program panics, it outputs a report with information on the line of code where the problem occurred. In this case, it looks like the problem is on line 20 of the *floats.go* file.

If we look at line 20 of *floats.go*, we'll see that it's the part of the `GetFloats` function where numbers from the file get added to the array!

```
// ...Preceding code omitted...
func GetFloats(fileName string) ([3]float64, error) {
        var numbers [3]float64
        file, err := os.Open(fileName)
        if err != nil {
                return numbers, err
        }
        i := 0
        scanner := bufio.NewScanner(file)
        for scanner.Scan() {
                numbers[i], err = strconv.ParseFloat(scanner.Text(), 64)
                if err != nil {
                        return numbers, err
                }
                i++
        }
        // ...Rest of GetFloats code omitted...
}
```

Here's line 20, where a number is assigned to the array!

Our program can only process three values! (continued)

Remember when a mistake in a previous code sample led a program to attempt to access an eighth element of a seven-element array? That program panicked and exited, too.

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}

for i := 0; i <= 7; i++ {
        fmt.Println(i, notes[i])
}
```

The array only has seven elements.

Loops up through index 7 (the eighth element), which doesn't exist!

```
0 do
1 re
2 mi
3 fa
4 so
5 la
6 ti
panic: runtime error: index out of range
```

Access indexes 0 through 6.

Accessing index 7 causes a panic!

The same problem is happening in our `GetFloats` function. Because we declared that the `numbers` array holds three elements, that's *all* it can hold. When the fourth line of the *data.txt* file is reached, it attempts to assign to a *fourth* element of `numbers`, which results in a panic.

```
func GetFloats(fileName string) ([3]float64, error) {
        var numbers [3]float64
        file, err := os.Open(fileName)
        if err != nil {
                return numbers, err
        }
        i := 0
        scanner := bufio.NewScanner(file)
        for scanner.Scan() {
                numbers[i], err = strconv.ParseFloat(scanner.Text(), 64)
                if err != nil {
                        return numbers, err
                }
                i++
        }
        // ...Rest of GetFloats code omitted...
}
```

The only valid indexes are numbers[0] through numbers[2]...

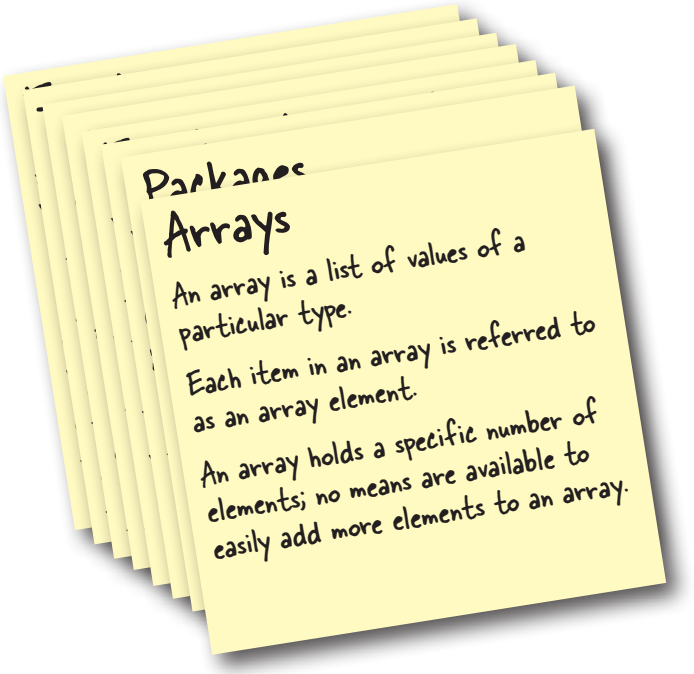
This attempts to assign to numbers[3], which causes a panic!

Go arrays are fixed in size; they can't grow or shrink. But the *data.txt* file can have as many lines as the user wants to add. We'll see a solution for this dilemma in the next chapter!



Your Go Toolbox

That's it for Chapter 5! You've added arrays to your toolbox.



BULLET POINTS

- To declare an array variable, include the array length in square brackets and the type of elements it will hold:
  `var myArray [3]int`
- To assign or access an element of an array, provide its index in square brackets. Indexes start at 0, so the first element of `myArray` is `myArray[0]`.
- As with variables, the default value for all array elements is the zero value for that element's type.
- You can set element values at the time an array is created using an **array literal**:
  `[3]int{4, 9, 6}`
- If you store an index that is not valid for an array in a variable, and then try to access an array element using that variable as an index, you will get a panic—a runtime error.
- You can get the number of elements in an array with the built-in `len` function.
- You can conveniently process all the elements of an array using the special `for...range` loop syntax, which loops through each element and assigns its index and value to variables you provide.
- When using a `for...range` loop, you can ignore either the index or value for each element by assigning it to the `_` blank identifier.
- The `os.Open` function opens a file. It returns a pointer to an `os.File` value representing that opened file.
- Passing an `os.File` value to `bufio.NewScanner` returns a `bufio.Scanner` value whose `Scan` and `Text` methods can be used to read a line at a time from the file as strings.



Below is a program that declares a couple arrays and prints out their elements. Write down what the program output would be.

```
package main

import "fmt"

func main() {
        var numbers [3]int
        numbers[0] = 42
        numbers[2] = 108
        var letters = [3]string{"a", "b", "c"}

                                        Output:

        fmt.Println(numbers[0])         42

        fmt.Println(numbers[1])         0

        fmt.Println(numbers[2])         108

        fmt.Println(letters[2])         c

        fmt.Println(letters[0])         a

        fmt.Println(letters[1])         b
}
```

Pool Puzzle Solution

```
package main

import "fmt"

func main() {
        numbers := [6]int{3, 16, -2, 10, 23, 12}
        for i, number := range numbers {
                if number >= 10 && number <= 20 {
                        fmt.Println(i, number)
                }
        }
}
```

Output

```
1 16
3 10
5 12
```

6 appending issue

Slices

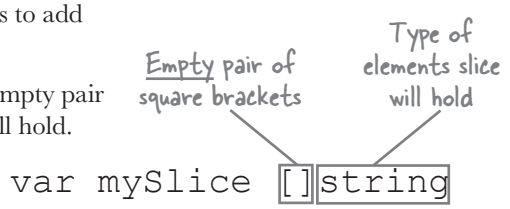


We've learned we can't add more elements to an array. That's a real problem for our program, because we don't know in advance how many pieces of data our file contains. But that's where Go **slices** come in. Slices are a collection type that can grow to hold additional items—just the thing to fix our current program! We'll also see how slices give users an easier way to provide data to *all* your programs, and how they can help you write functions that are more convenient to call.

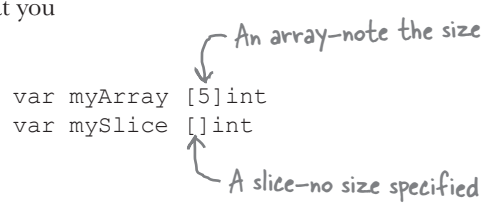
Slices

There actually *is* a Go data structure that we can add more values to—it's called a **slice**. Like arrays, slices are made up of multiple elements, all of the same type. *Unlike* arrays, functions are available that allow us to add extra elements onto the end of a slice.

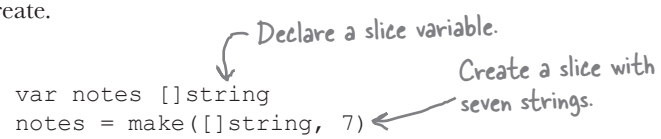
To declare the type for a variable that holds a slice, you use an empty pair of square brackets, followed by the type of elements the slice will hold.



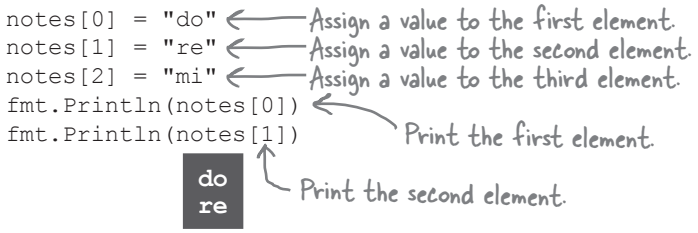
This is just like the syntax for declaring an array variable, except that you don't specify the size.



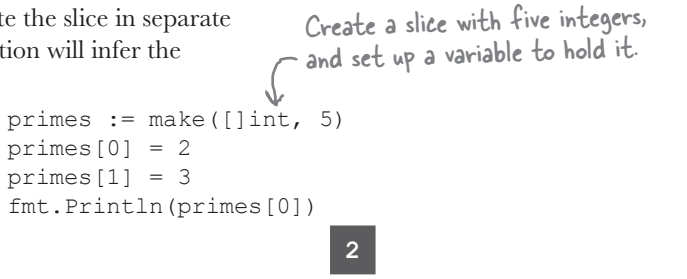
Unlike with array variables, declaring a slice variable doesn't automatically create a slice. For that, you can call the built-in `make` function. You pass `make` the type of the slice you want to create (which should be the same as the type of the variable you're going to assign it to), and the length of slice it should create.



Once the slice is created, you assign and retrieve its elements using the same syntax you would for an array.



You don't have to declare the variable and create the slice in separate steps; using `make` with a short variable declaration will infer the variable's type for you.



Slices (continued)

The built-in `len` function works the same way with slices as it does with arrays. Just pass `len` a slice, and its length will be returned as an integer.

```
notes := make([]string, 7)
primes := make([]int, 5)
fmt.Println(len(notes))
fmt.Println(len(primes))
```

```
7
5
```

Both `for` and `for...range` loops work just the same with slices as they do with arrays, too:

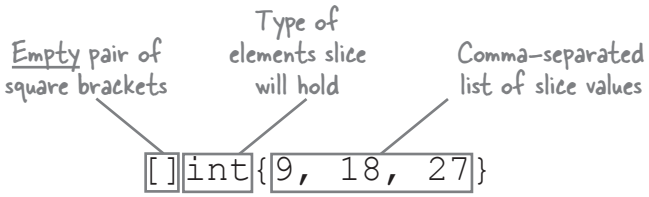
```
letters := []string{"a", "b", "c"}
for i := 0; i < len(letters); i++ {
      fmt.Println(letters[i])
}
for _, letter := range letters {
      fmt.Println(letter)
}
```

```
a
b
c
a
b
c
```

Slice literals

Just like with arrays, if you know in advance what values a slice will start with, you can initialize the slice with those values using a **slice literal**. A slice literal looks a lot like an array literal, but where an array literal has the length of the array in square brackets, a slice literal's square brackets are empty. The empty brackets are then followed by the type of elements the slice will hold, and a list in curly braces of the initial values each element will have.

There's no need to call the `make` function; using a slice literal in your code will create the slice *and* prepopulate it.



These examples are like the previous ones we showed, except that instead of assigning values to the slice elements one by one, the entire slice is initialized using slice literals.

```
notes := []string{"do", "re", "mi", "fa", "so", "la", "ti"}
fmt.Println(notes[3], notes[6], notes[0])
primes := []int{
      2,
      3,
      5,
}
fmt.Println(primes[0], primes[1], primes[2])
```

← Assign values using a slice literal.

← A multi-line slice literal.

```
fa ti do
2 3 5
```

Pool Puzzle



Your **job** is to take code snippets from the pool and place them into the blank lines in this code. **Don't** use the same snippet more than once, and you won't need to use all the snippets. Your **goal** is to make a program that will run and produce the output shown.

```
package main

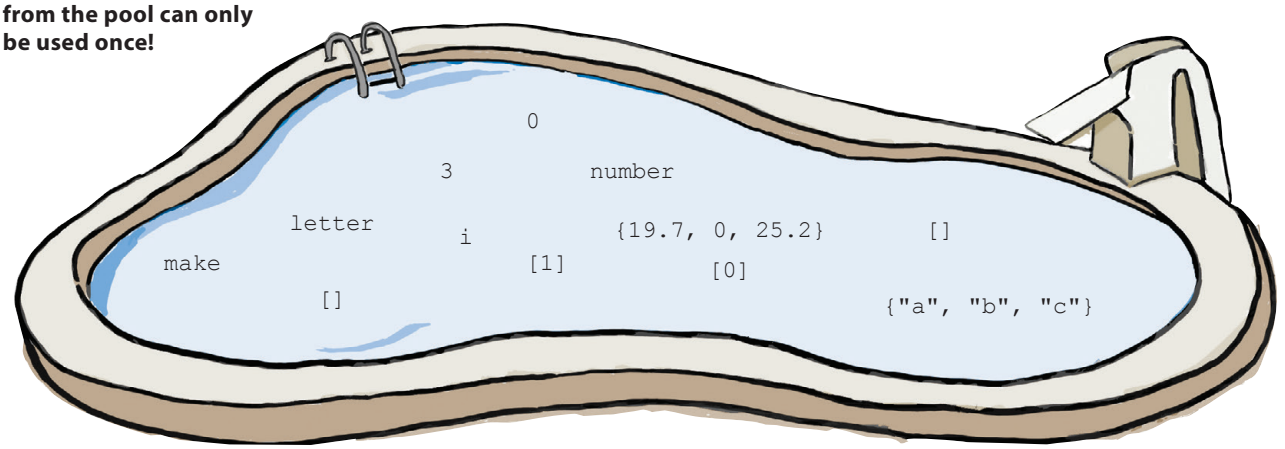
import "fmt"

func main() {
        numbers := ______(__float64, __)
        numbers____ = 19.7
        numbers[2] = 25.2
        for __, _______ := range numbers {
                fmt.Println(i, number)
        }
        var letters = __string_______________
        for i, letter := range letters {
                fmt.Println(i, ________)
        }
}
```

Output

```
0 19.7
1 0
2 25.2
0 a
1 b
2 c
```

Note: each snippet from the pool can only be used once!



Answers on page 203.



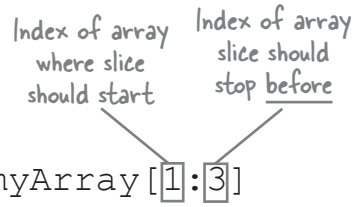
Hold up! It looks like slices can do everything arrays can do, **and** you say we can add additional values to them! Why didn't you just show us slices, and skip that array nonsense?

Because slices are built on top of arrays. You can't understand how slices work without understanding arrays. Here, we'll show you why...

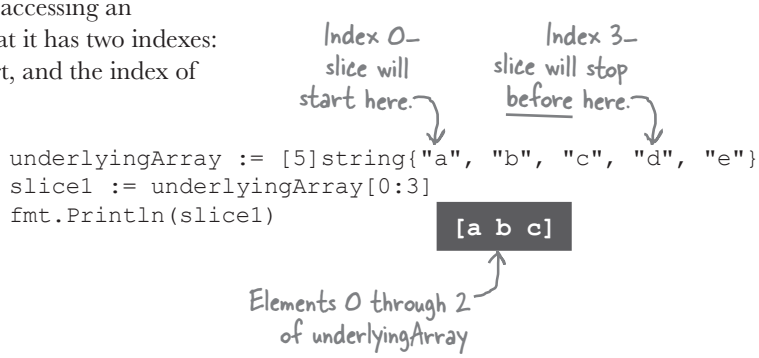
The slice operator

Every slice is built on top of an **underlying array**. It's the underlying array that actually holds the slice's data; the slice is merely a view into some (or all) of the array's elements.

When you use the `make` function or a slice literal to create a slice, the underlying array is created for you automatically (and you can't access it, except through the slice). But you can also create the array yourself, and then create a slice based on it with the **slice operator**.

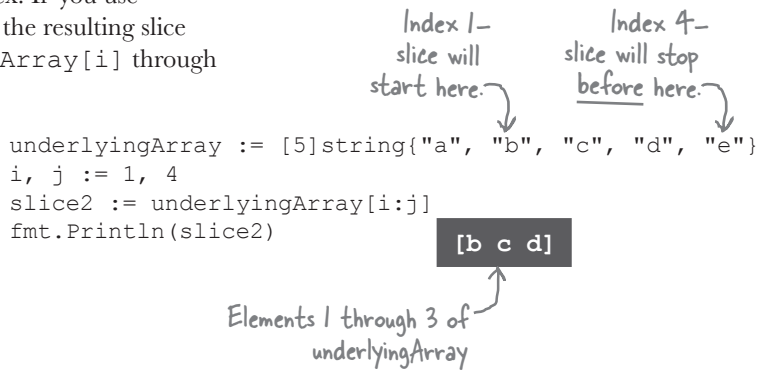


The slice operator looks similar to the syntax for accessing an individual element or slice of an array, except that it has two indexes: the index of the array where the slice should start, and the index of the array that the slice should stop before.



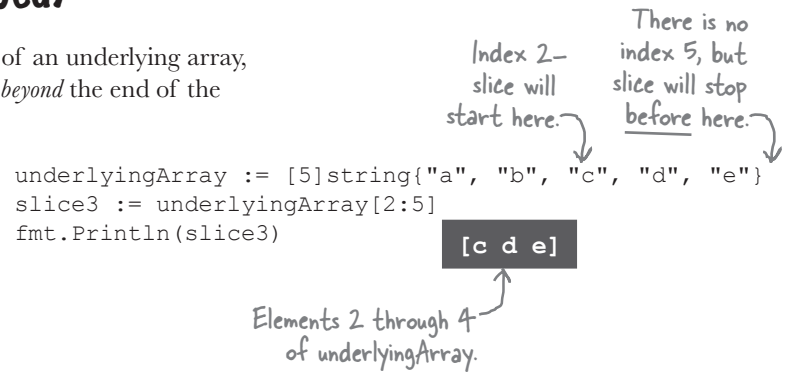
Notice that we emphasize that the second index is the index the slice will stop before. That is, the slice should include the elements up to, but *not* including, the second index. If you use `underlyingArray[i:j]` as a slice operator, the resulting slice will actually contain the elements `underlyingArray[i]` through `underlyingArray[j-1]`.

(We know, it's counterintuitive. But a similar notation has been used in the Python programming language for over 20 years, and it seems to work OK.)



The slice operator (continued)

If you want a slice to include the last element of an underlying array, you actually specify a second index that's one *beyond* the end of the array in your slice operator.

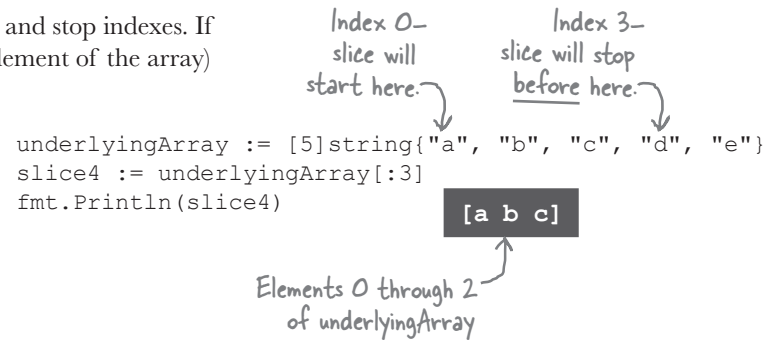


Make sure you don't go any further than that, though, or you'll get an error:

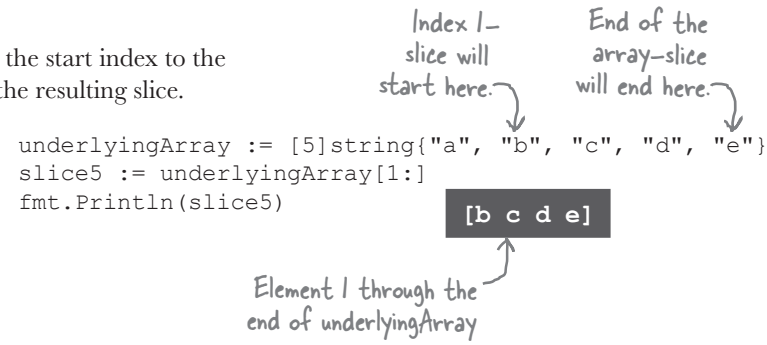
```
underlyingArray := [5]string{"a", "b", "c", "d", "e"}
slice3 := underlyingArray[2:6]
```

```
invalid slice index 6 (out of bounds for 5-element array)
```

The slice operator has defaults for both the start and stop indexes. If you omit the start index, a value of `0` (the first element of the array) will be used.



And if you omit the stop index, everything from the start index to the end of the underlying array will be included in the resulting slice.



Underlying arrays

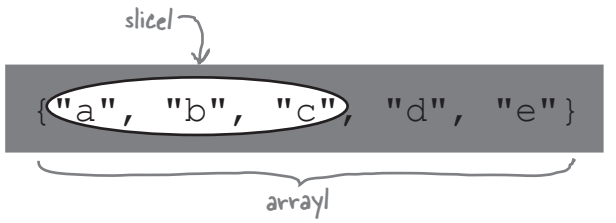
As we mentioned, a slice doesn't hold any data itself; it's merely a view into the elements of an underlying array. You can think of a slice as a microscope, focusing on a particular portion of the contents of a slide (the underlying array).



When you take a slice of an underlying array, you can only "see" the portion of the array's elements that are visible through the slice.

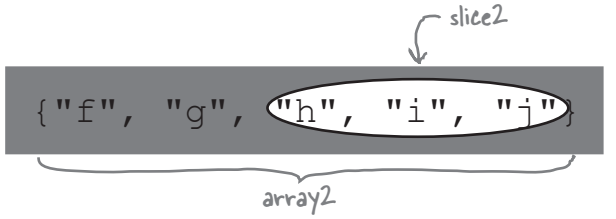
```
array1 := [5]string{"a", "b", "c", "d", "e"}
slice1 := array1[0:3]
fmt.Println(slice1)
```

```
[a b c]
```




```
array2 := [5]string{"f", "g", "h", "i", "j"}
slice2 := array2[2:5]
fmt.Println(slice2)
```

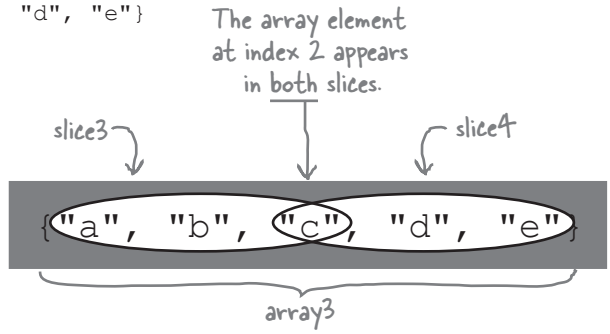
```
[h i j]
```



It's even possible to have multiple slices point to the same underlying array. Each slice will then be a view into its own subset of the array's elements. The slices can even overlap!

```
array3 := [5]string{"a", "b", "c", "d", "e"}
slice3 := array3[0:3]
slice4 := array3[2:5]
fmt.Println(slice3, slice4)
```

```
[a b c] [c d e]
```



Change the underlying array, change the slice

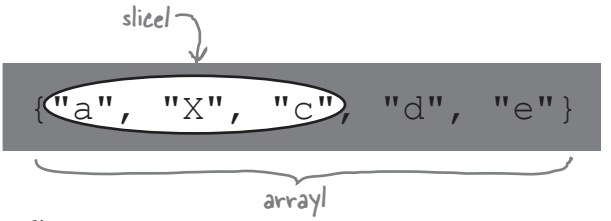
Now, here's something to be careful about: because a slice is just a view into the contents of an array, if you change the underlying array, those changes will *also* be visible within the slice!

```
array1 := [5]string{"a", "b", "c", "d", "e"}
slice1 := array1[0:3]
array1[1] = "X"
fmt.Println(array1)
fmt.Println(slice1)
```

Change an element of the underlying array...

```
[a X c d e]
[a X c]
```

...and the change appears in the slice!



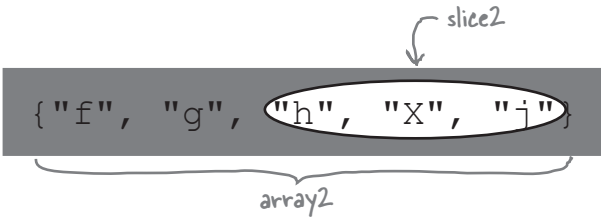
Assigning a new value to a slice element will change the corresponding element in the underlying array.

```
array2 := [5]string{"f", "g", "h", "i", "j"}
slice2 := array2[2:5]
slice2[1] = "X"
fmt.Println(array2)
fmt.Println(slice2)
```

Change an element of the slice...

```
[f g h X j]
[h X j]
```

...and the underlying array is changed!



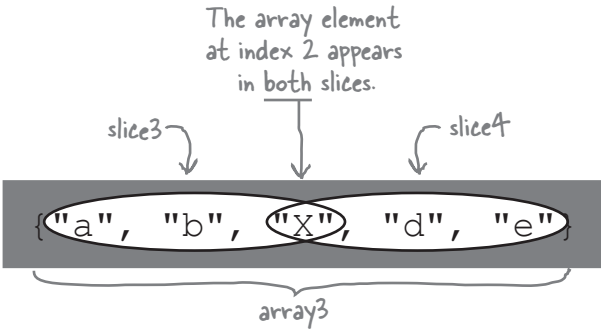
If multiple slices point to the same underlying array, a change to the array's elements will be visible in *all* the slices.

```
array3 := [5]string{"a", "b", "c", "d", "e"}
slice3 := array3[0:3]
slice4 := array3[2:5]
array3[2] = "X"
fmt.Println(array3)
fmt.Println(slice3, slice4)
```

Change an element of the underlying array...

```
[a b X d e]
[a b X] [X d e]
```

...and the change appears in both slices!



Because of these potential issues, you may find it's generally better to create slices using `make` or a slice literal, rather than creating an array and using a slice operator on it. With `make` and with slice literals, you never have to work with the underlying array.

Add onto a slice with the "append" function

Look, all of this info on slices is just super. Really. But I'm still stuck with a program that can only read three lines from a text file, because it uses an array. You said we can add more values onto a slice—I want to hear about that!



Go offers a built-in `append` function that takes a slice, and one or more values you want to append to the end of that slice. It returns a new, larger slice with all the same elements as the original slice, plus the new elements added onto the end.

Assign the return value of "append" back to the same slice variable.

Assign the return value of "append" back to the same slice variable.

```
slice := []string{"a", "b"}
fmt.Println(slice, len(slice))
slice = append(slice, "c")
fmt.Println(slice, len(slice))
slice = append(slice, "d", "e")
fmt.Println(slice, len(slice))
```

Create a slice.

Append an element to the end of the slice.

Append two elements to the end of the slice.

Has one more element, and the length is increased by one.

Has two more elements, and the length is increased by two.

```
[a b] 2
[a b c] 3
[a b c d e] 5
```

You don't have to keep track of what index you want to assign new values to, or anything! Just call `append` with your slice and the value(s) you want added to the end, and you'll get a new, longer slice back. It's really that easy!

Well, with one caution...

Add onto a slice with the "append" function (continued)

Notice that we're making sure to assign the return value of `append` back to the *same* slice variable we passed to `append`. This is to avoid some potentially inconsistent behavior in the slices returned from `append`.

A slice's underlying array can't grow in size. If there isn't room in the array to add elements, all its elements will be copied to a new, larger array, and the slice will be updated to refer to this new array. But since all this happens behind the scenes in the `append` function, there's no easy way to tell whether the slice returned from `append` has the *same* underlying array as the slice you passed in, or a *different* underlying array. If you keep both slices, this can lead to some unpredictable behavior.

Below, for example, we have four slices, the last three created by calls to `append`. Here we are *not* following the convention of assigning `append`'s return value back to the same variable. When we assign a value to an element of the `s4` slice, we can see the change reflected in `s3`, because `s4` and `s3` happen to share the same underlying array. But the change is *not* reflected in `s2` or `s1`, because they have a *different* underlying array.

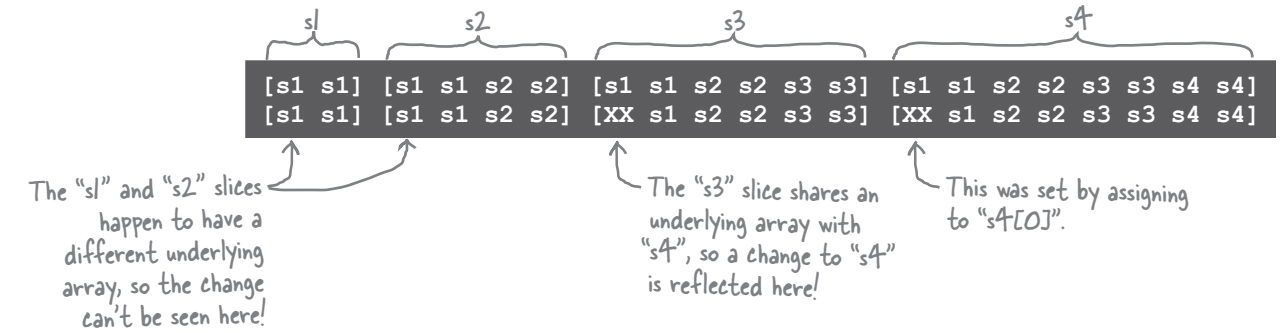
We assign slices returned from "append" to new variables!

```
s1 := []string{"s1", "s1"}
s2 := append(s1, "s2", "s2")
s3 := append(s2, "s3", "s3")
s4 := append(s3, "s4", "s4")
fmt.Println(s1, s2, s3, s4)
s4[0] = "XX"
fmt.Println(s1, s2, s3, s4)
```

Print the slices.

Assign to an element of the fourth slice.

See what's changed.



So when calling `append`, it's conventional to just assign the return value back to the same slice variable you passed to `append`. You don't need to worry about whether two slices have the same underlying array if you're only storing one slice!

We assign slices returned from "append" to the same variable.

```
s1 := []string{"s1", "s1"}
s1 = append(s1, "s2", "s2")
s1 = append(s1, "s3", "s3")
s1 = append(s1, "s4", "s4")
fmt.Println(s1)
```

```
[s1 s1 s2 s2 s3 s3 s4 s4]
```

No nasty surprises here!

Slices and zero values

As with arrays, if you access a slice element that no value has been assigned to, you'll get the zero value for that type back:

Create slices without assigning values to their elements.

```
floatSlice := make([]float64, 10)
boolSlice := make([]bool, 10)
fmt.Println(floatSlice[9], boolSlice[5])
```

```
0 false
```

Unlike arrays, the slice variable itself *also* has a zero value: it's `nil`. That is, a slice variable that no slice has been assigned to will have a value of `nil`.

Declare slice variables without creating slices.

Remember, "%#v" formats a value as it would appear in Go code.

```
var intSlice []int
var stringSlice []string
fmt.Printf("intSlice: %#v, stringSlice: %#v\n", intSlice, stringSlice)
```

The value of both variables is nil.

```
intSlice: []int(nil), stringSlice: []string(nil)
```

In other languages, that might require testing whether a variable actually contains a slice before attempting to use it. But in Go, functions are intentionally written to treat a `nil` slice value as if it were an empty slice. For example, the `len` function will return `0` if it's passed a `nil` slice:

Pass a nil slice to the "len" function.

It will return 0, as if you'd passed an empty slice in!

```
fmt.Println(len(intSlice))
```

```
0
```

The `append` function also treats `nil` slices like empty slices. If you pass an empty slice to `append`, it will add the item you specify to the slice, and return a slice with one item. If you pass a `nil` slice to `append`, you'll *also* get a slice with one item back, even though there technically was no slice to "append" the item to. The `append` function will create the slice behind the scenes.

Pass a nil slice to "append".

It will return a one-item slice, as if you'd appended to an empty slice!

```
intSlice = append(intSlice, 27)
fmt.Printf("intSlice: %#v\n", intSlice)
```

```
intSlice: []int{27}
```

This means you generally don't have to worry about whether you have an empty slice or a `nil` slice. You can treat them both the same, and your code will "just work"!

The variable will contain nil.

The "len" function returns 0.

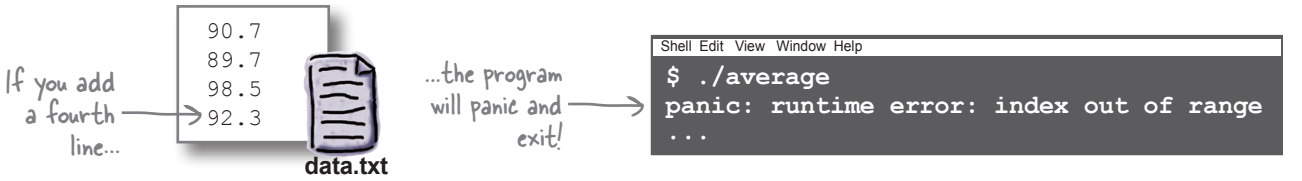
The "append" function returns a one-item slice, as if you'd passed an empty slice in.

```
var slice []string
if len(slice) == 0 {
    slice = append(slice, "first item")
}
fmt.Printf("%#v\n", slice)
```

```
[]string{"first item"}
```

Reading additional file lines using slices and "append"

Now that we know about slices and the `append` function, we can finally fix our `average` program! Remember, `average` was failing as soon as we added a fourth line to the *data.txt* file it reads from:



We traced the problem back to our `datafile` package, which stores the file lines in an array that can't grow beyond three elements:




```
// Package datafile allows reading data samples from files.
package datafile

import (
	"bufio"
	"os"
	"strconv"
)

// GetFloats reads a float64 from each line of a file.
func GetFloats(fileName string) ([3]float64, error) {
	var numbers [3]float64
	file, err := os.Open(fileName)
	if err != nil {
		return numbers, err
	}
	i := 0
	scanner := bufio.NewScanner(file)
	for scanner.Scan() {
		numbers[i], err = strconv.ParseFloat(scanner.Text(), 64)
		if err != nil {
			return numbers, err
		}
		i++
	}
	err = file.Close()
	if err != nil {
		return numbers, err
	}
	if scanner.Err() != nil {
		return numbers, scanner.Err()
	}
	return numbers, nil
}
```

The function returns an array of float64 values.

The only valid indexes are numbers[0] through numbers[2]...

This attempts to assign to numbers[3], which causes a panic!

Reading additional file lines using slices and "append" (continued)

Most of our work with slices has just centered on understanding them. Now that we do, updating the `GetFloats` function to use a slice instead of an array doesn't involve much effort.

First, we update the function declaration to return a slice of `float64` values instead of an array. Previously, we stored the array in a variable called `numbers`; we'll just use that same variable name to hold the slice. We won't assign a value to `numbers`, so at first it will be `nil`.

Instead of assigning values read from the file to a specific array index, we can just call `append` to extend the slice (or create a slice, if it's `nil`) and add new values. That means we can get rid of the code to create and update the `i` variable that tracks the index. We assign the `float64` value returned from `ParseFloat` to a new temporary variable, just to hold it while we check for any errors in parsing. Then we pass the `numbers` slice and the new value from the file to `append`, making sure to assign the return value back to the `numbers` variable.

Aside from that, the code in `GetFloats` can remain the same—the slice is basically a drop-in replacement for the array.




```
// ...Preceding code omitted...
func GetFloats(fileName string) ([]float64, error) {
        var numbers []float64
        file, err := os.Open(fileName)
        if err != nil {
                return numbers, err
        }
        scanner := bufio.NewScanner(file)
        for scanner.Scan() {
                number, err := strconv.ParseFloat(scanner.Text(), 64)
                if err != nil {
                        return numbers, err
                }
                numbers = append(numbers, number)
        }
        err = file.Close()
        if err != nil {
                return numbers, err
        }
        if scanner.Err() != nil {
                return numbers, scanner.Err()
        }
        return numbers, nil
}
```

Switch to returning a slice.

This variable will contain nil by default. (Remember, "append" treats nil just like an empty slice.)

No changes needed for error handling; we can treat the slice the same way we did the array.

Convert the string to a float64 and assign it to a temporary variable.

Append the new number to the slice.

No changes needed here, either.

Trying our improved program

The slice returned from the `GetFloats` function works like a drop-in replacement for an array in our main `average` program, too. In fact, we don't have to make *any* changes to the main program!

Because we used a `:=` short variable declaration to assign the `GetFloats` return value to a variable, the `numbers` variable automatically switches from an inferred type of `[3]float64` (an array) to a type of `[]float64` (a slice). And because the `for...range` loop and the `len` functions work the same way with a slice as they do with an array, no changes are needed to that code, either!



No changes needed anywhere!

```
// average calculates the average of several numbers.
package main

import (
        "fmt"
        "github.com/headfirstgo/datafile"
        "log"
)

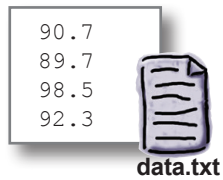
func main() {
        numbers, err := datafile.GetFloats("data.txt")
        if err != nil {
                log.Fatal(err)
        }
        var sum float64 = 0
        for _, number := range numbers {
                sum += number
        }
        sampleCount := float64(len(numbers))
        fmt.Printf("Average: %0.2f\n", sum/sampleCount)
}
```

Automatically gets a type of []float64 instead of [3]float64

Works the same with a slice as it did with an array

Also works the same with a slice

That means we're ready to try the changes out! Ensure the *data.txt* file is still saved in your Go workspace's *bin* subdirectory, and then compile and run the code using the same commands as before. It will read all the lines of *data.txt* and display their average. Then try updating *data.txt* to have more lines, or fewer; it will still work regardless!



Compiles the updated "datafile" package, because "average" depends on it.

Change to the "bin" subdirectory.

Run the program.

The average of the numbers from all four lines of the file!

```
Shell Edit View Window Help
$ go install github.com/headfirstgo/average
$ cd /Users/jay/go/bin
$ ./average
Average: 92.80
```

Returning a nil slice in the event of an error

Let's make one more small improvement to the `GetFloats` function. Currently, we're returning the `numbers` slice even in the event of an error. That means that we could be returning a slice with invalid data:

```
number, err := strconv.ParseFloat(scanner.Text(), 64)
if err != nil {
        return numbers, err
}
```

We're returning invalid data that should not be used!

The code that calls `GetFloats` *should* check the returned error value, see that it's not `nil`, and ignore the contents of the returned slice. But really, why bother to return the slice at all, if the data it contains is invalid? Let's update `GetFloats` to return `nil` instead of a slice in the event of an error.




```
// ...Preceding code omitted...
func GetFloats(fileName string) ([]float64, error) {
        var numbers []float64
        file, err := os.Open(fileName)
        if err != nil {
                return nil, err
        }
        scanner := bufio.NewScanner(file)
        for scanner.Scan() {
                number, err := strconv.ParseFloat(scanner.Text(), 64)
                if err != nil {
                        return nil, err
                }
                numbers = append(numbers, number)
        }
        err = file.Close()
        if err != nil {
                return nil, err
        }
        if scanner.Err() != nil {
                return nil, scanner.Err()
        }
        return numbers, nil
}
```

Return nil instead of the slice. (The slice would be nil at this point anyway, but this change makes it more obvious.)

Return nil instead of the slice.

Return nil instead of the slice.

Return nil instead of the slice.

Let's recompile the program (which will include the updated `datafile` package) and run it. It should work the same as before. But now our error-handling code is a little bit cleaner.

```
Shell Edit View Window Help
$ go install github.com/headfirstgo/average
$ cd /Users/jay/go/bin
$ ./average
Average: 92.80
```



Exercise

Below is a program that takes a slice of an array and then appends elements to the slice. Write down what the program output would be.

```
package main

import "fmt"

func main() {
        array := [5]string{"a", "b", "c", "d", "e"}
        slice := array[1:3]
        slice = append(slice, "x")
        slice = append(slice, "y", "z")
        for _, letter := range slice {
                fmt.Println(letter)
        }
}
```

Output:

.....

.....

.....

.....

.....

.....

.....

We've provided more blanks than you actually need. How many more? That's up to you to figure out!

Answers on page 203.

Command-line arguments

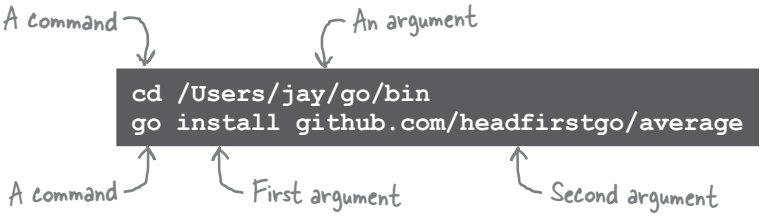
At last! This is working great. I just need one more thing... It's kind of a pain editing *data.txt* every time I need a new average. Is there another way to input the sample values?



There is an alternative—users could pass the values to the program as command-line arguments.

Just as you can control the behavior of many Go functions by passing them arguments, you can pass arguments to many programs you run from the terminal or command prompt. This is known as a program's *command-line interface*.

You've already seen command-line arguments used in this very book. When we run the `cd` ("change directory") command, we pass it the name of the directory we want to change to as an argument. When we run the `go` command, we often pass it multiple arguments: the subcommand (`run`, `install`, etc.) we want to use, and the name of the file or package we want the subcommand to work on.



Getting command-line arguments from the os.Args slice

Let's set up a new version of the `average` program, called `average2`, that takes the values to average as command-line arguments.

The `os` package has a package variable, `os.Args`, that gets set to a slice of strings representing the command-line arguments the currently running program was executed with. We'll start by simply printing the `os.Args` slice to see what it contains.

Create a new *average2* directory alongside the *average* directory in your workspace, and save a *main.go* file within it.



Then, save the following code in *main.go*. It simply imports the `fmt` and `os` packages, and passes the `os.Args` slice to `fmt.Println`.

```
// average2 calculates the average of several numbers.
package main

import (
        "fmt"
        "os"
)

func main() {
        fmt.Println(os.Args)
}
```

Print the os.Args slice.

Let's try it out. From your terminal or command prompt, run this command to compile and install the program:

```
go install github.com/headfirstgo/average2
```

That will install an executable file named *average2* (or *average2.exe* on Windows) to your Go workspace's *bin* subdirectory. Use the `cd` command to change to *bin*, and type **`average2`**, but don't hit the Enter key just yet. Following the program name, type a space, and then type one or more arguments, separated by spaces. *Then* hit Enter. The program will run and print the value of `os.Args`.

Rerun `average2` with different arguments, and you should see different output.

Compile and install the executable.

Change to the "bin" subdirectory.

Run the executable with several arguments.

It will print the value of os.Args.

Run average2 with different arguments to see different results.

```
Shell  Edit  View  Window  Help
$ go install github.com/headfirstgo/average2
$ cd /Users/jay/go/bin
$ ./average2 71.8 56.2 89.5
[./average2 71.8 56.2 89.5]
$ ./average2 do re mi fa so
[./average2 do re mi fa so]
```

The slice operator can be used on other slices

This is working pretty well, but there's one problem: the name of the executable is being included as the first element of `os.Args`.

```
$ ./average2 71.8 56.2 89.5
[./average2 71.8 56.2 89.5]
```

The first element is the name of the program.

That should be easy to remove, though. Remember how we used the slice operator to get a slice that included everything but the first element of an array?

Index 1—slice will start here.

End of the array—slice will end here.

```
underlyingArray := [5]string{"a", "b", "c", "d", "e"}
slice5 := underlyingArray[1:]
fmt.Println(slice5)
```

```
[b c d e]
```

Element 1 through the end of underlyingArray

The slice operator can be used on slices just like it can on arrays. If we use a slice operator of `[1:]` on `os.Args`, it will give us a new slice that omits the first element (whose index is `0`), and includes the second element (index `1`) through the end of the slice.

```
// average2 calculates the average of several numbers.
package main

import (
        "fmt"
        "os"
)

func main() {
        fmt.Println(os.Args[1:])
}
```

Get a new slice that includes the second element (index 1) through the end of os.Args.

If we recompile and rerun `average2`, this time we'll see that the output includes only the actual command-line arguments.

```
Shell Edit View Window Help
$ go install github.com/headfirstgo/average2
$ ./average2 71.8 56.2 89.5
[71.8 56.2 89.5]
$ ./average2 do re mi fa so
[do re mi fa so]
```

Omits the executable name

Omits the executable name

Updating our program to use command-line arguments

Now that we're able to get the command-line arguments as a slice of strings, let's update the `average2` program to convert the arguments to actual numbers, and calculate their average. We'll mostly be able to reuse the concepts we learned about in our original `average` program and the `datafile` package.

We use the slice operator on `os.Args` to omit the program name, and assign the resulting slice to an `arguments` variable. We set up a `sum` variable that will hold the total of all the numbers we're given. Then we use a `for...range` loop to process the elements of the `arguments` slice (using the _ blank identifier to ignore the element index). We use `strconv.ParseFloat` to convert the argument string to a `float64`. If we get an error, we log it and exit, but otherwise we add the current number to `sum`.

When we've looped through all the arguments, we use `len(arguments)` to determine how many data samples we're averaging. We then divide `sum` by this sample count to get the average.




```
// average2 calculates the average of several numbers.
package main

import (
        "fmt"
        "log"
        "os"
        "strconv"
)

func main() {
        arguments := os.Args[1:]
        var sum float64 = 0
        for _, argument := range arguments {
                number, err := strconv.ParseFloat(argument, 64)
                if err != nil {
                        log.Fatal(err)
                }
                sum += number
        }
        sampleCount := float64(len(arguments))
        fmt.Printf("Average: %0.2f\n", sum/sampleCount)
}
```

Import the "log" and "strconv" packages.

Get a slice of strings with all but the first element of os.Args.

Set up a variable to hold the sum of the numbers.

Process each command-line argument.

Convert the string to a float64.

If there was an error converting the string, log it and exit.

Add the number to the total.

The length of the arguments slice can be used as the number of samples.

Calculate the average and print it.

With these changes saved, we can recompile and rerun the program. It will take the numbers you provide as arguments and average them. Give as few or as many arguments as you like; it will still work!

```
Shell Edit View Window Help
$ go install github.com/headfirstgo/average2
$ cd /Users/jay/go/bin
$ ./average2 71.8 56.2 89.5
Average: 72.50
$ ./average2 90.7 89.7 98.5 92.3
Average: 92.80
```

Run the program with several arguments.

Use any number of arguments you like.

Variadic functions

Now that we know about slices, we can cover a feature of Go that we haven't talked about so far. Have you noticed that some function calls can take as few, or as many, arguments as needed? Look at `fmt.Println` or `append`, for example:

```
fmt.Println(1)
fmt.Println(1, 2, 3, 4, 5)
letters := []string{"a"}
letters = append(letters, "b")
letters = append(letters, "c", "d", "e", "f", "g")
```

"Println" can take one argument...

...or five!

"append" can take two arguments...

...or six!

Don't try doing this with just any function, though! With all the functions we've defined so far, there had to be an *exact* match between the number of parameters in the function definition and the number of arguments in the function call. Any difference would result in a compile error.

```
func twoInts(first int, second int) {
      fmt.Println(first, second)
}

func main() {
      twoInts(1)
      twoInts(1, 2, 3)
}
```

If two parameters are expected...

...then we can't pass just one...

...and we can't pass three.

```
tmp/sandbox815038307/main.go:10:9: not enough arguments in call to twoInts
      have (number)
      want (int, int)
tmp/sandbox815038307/main.go:11:9: too many arguments in call to twoInts
      have (number, number, number)
      want (int, int)
```

So how do `Println` and `append` do it? They're declared as variadic functions. A **variadic function** is one that can be called with a *varying* number of arguments. To make a function variadic, use an ellipsis (`...`) before the type of the last (or only) function parameter in the function declaration.

Ellipsis

Type

```
func myFunc(param1 int, param2 ...string) {
    // function code here
}
```

Variadic functions (continued)

The last parameter of a variadic function receives the variadic arguments as a slice, which the function can then process like any other slice.

Here's a variadic version of the `twoInts` function, and it works just fine with any number of arguments:

The "numbers" variable will hold a slice with the arguments.

```
func severalInts(numbers ...int) {
        fmt.Println(numbers)
}

func main() {
        severalInts(1)
        severalInts(1, 2, 3)
}
```

```
[1]
[1 2 3]
```

Here's a similar function that works with strings. Notice that if we provide no variadic arguments, it's not an error; the function just receives an empty slice.

The "strings" variable will hold a slice with the arguments.

```
func severalStrings(strings ...string) {
        fmt.Println(strings)
}

func main() {
        severalStrings("a", "b")
        severalStrings("a", "b", "c", "d", "e")
        severalStrings()
}
```

If there are no arguments, an empty slice is received.

```
[a b]
[a b c d e]
[]
```

A function can take one or more nonvariadic arguments as well. Although a function caller can omit variadic arguments (resulting in an empty slice), nonvariadic arguments are always required; it's a compile error to omit those. Only the *last* parameter in a function definition can be variadic; you can't place it in front of required parameters.

An int argument will be required first.

A Boolean argument will be required second.

Any remaining arguments must be strings and will be stored as a slice here.

```
func mix(num int, flag bool, strings ...string) {
        fmt.Println(num, flag, strings)
}

func main() {
        mix(1, true, "a", "b")
        mix(2, false, "a", "b", "c", "d")
}
```

```
1 true [a b]
2 false [a b c d]
```

Using variadic functions

Here's a `maximum` function that takes any number of `float64` arguments and returns the greatest value out of all of them. The arguments to `maximum` are stored in a slice in the `numbers` parameter. To start, we set the current maximum value to `-Inf`, a special value representing negative infinity, obtained by calling `math.Inf`. (We could start with a current maximum of `0`, but this way `maximum` will work with negative numbers.) Then we use `for...range` to process each argument in the `numbers` slice, comparing it to the current maximum, and setting it as the new maximum if it's greater. Whatever maximum remains after processing all the arguments is the one we return.

```
package main

import (
        "fmt"
        "math"
)

func maximum(numbers ...float64) float64 {
        max := math.Inf(-1)
        for _, number := range numbers {
                if number > max {
                        max = number
                }
        }
        return max
}

func main() {
        fmt.Println(maximum(71.8, 56.2, 89.5))
        fmt.Println(maximum(90.7, 89.7, 98.5, 92.3))
}
```

Take any number of float64 arguments.

Start with a very low value.

Process each variadic argument.

Find the largest value among the arguments.

```
89.5
98.5
```

Here's an `inRange` function that takes a minimum value, a maximum value, and any number of additional `float64` arguments. It will discard any argument that is below the given minimum or above the given maximum, returning a slice containing only the arguments that were in the specified range.

```
package main

import "fmt"

func inRange(min float64, max float64, numbers ...float64) []float64 {
        var result []float64
        for _, number := range numbers {
                if number >= min && number <= max {
                        result = append(result, number)
                }
        }
        return result
}

func main() {
        fmt.Println(inRange(1, 100, -12.5, 3.2, 0, 50, 103.5))
        fmt.Println(inRange(-10, 10, 4.1, 12, -12, -5.2))
}
```

The minimum value in the range

The maximum value in the range

Any number of additional float64 arguments

This slice will hold arguments that were within range.

Process each variadic argument.

If this argument isn't below the minimum or above the maximum...

...add it to the slice to be returned.

Find arguments >= 1 and <= 100.

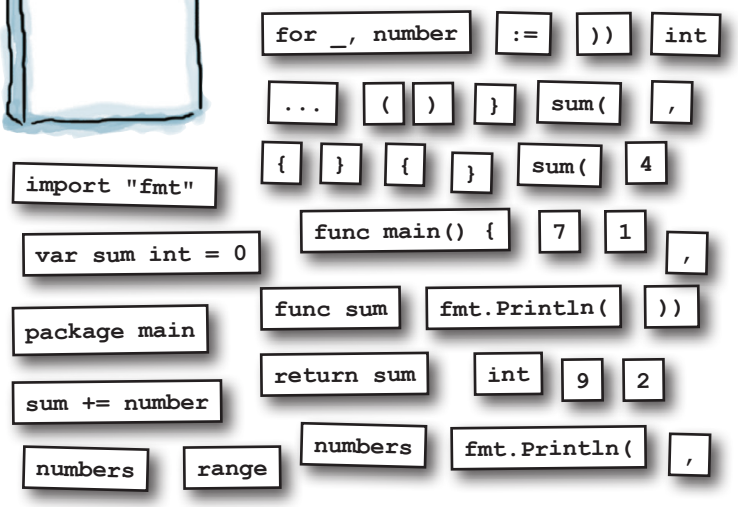
Find arguments >= -10 and <= 10.

```
[3.2 50]
[4.1 -5.2]
```



Code Magnets

A Go program that defines and uses a variadic function is scrambled up on the fridge. Can you reconstruct the code snippets to make a working program that will produce the given output?



Output

```
16
7
```

Answers on page 204.

Using a variadic function to calculate averages

Let's create a variadic `average` function that can take any number of `float64` arguments and return their average. It will look much like the logic from our `average2` program. We'll set up a `sum` variable to hold the total of the argument values. Then we'll loop through the range of arguments, adding each one to the value in `sum`. Finally, we'll divide `sum` by the number of arguments (converted to a `float64`) to get the average. The result is a function that can average as many (or as few) numbers as we want.

```
package main

import "fmt"

func average(numbers ...float64) float64 {
        var sum float64 = 0
        for _, number := range numbers {
                sum += number
        }
        return sum / float64(len(numbers))
}

func main() {
        fmt.Println(average(100, 50))
        fmt.Println(average(90.7, 89.7, 98.5, 92.3))
}
```

Take any number of float64 arguments.

Set up a variable to hold the sum of the arguments.

Process each variadic argument.

Add the argument value to the total.

Divide the total by the number of arguments to get the average.

```
75
92.8
```

Passing slices to variadic functions

Our new `average` variadic function works so well, we should try updating our `average2` program to make use of it. We can paste the `average` function into our `average2` code as is.

In the `main` function, we're still going to need to convert each of the command-line arguments from a `string` to a `float64` value. We'll create a slice to hold the resulting values, and store it in a variable named `numbers`. After each command-line argument is converted, instead of using it to calculate the average directly, we'll just append it to the `numbers` slice.

We then *attempt* to pass the `numbers` slice to the `average` function. But when we go to compile the program, that results in an error...




```
// average2 calculates the average of several numbers.
package main

import (
        "fmt"
        "log"
        "os"
        "strconv"
)

func average(numbers ...float64) float64 {
        var sum float64 = 0
        for _, number := range numbers {
                sum += number
        }
        return sum / float64(len(numbers))
}

func main() {
        arguments := os.Args[1:]
        var numbers []float64
        for _, argument := range arguments {
                number, err := strconv.ParseFloat(argument, 64)
                if err != nil {
                        log.Fatal(err)
                }
                numbers = append(numbers, number)
        }
        fmt.Printf("Average: %0.2f\n", average(numbers))
}
```

Paste in the "average" function as is.

This slice will hold the numbers we're averaging.

Append the converted number to the slice.

Attempt to pass the numbers to the variadic function....

Error

```
cannot use numbers (type []float64)
as type float64 in argument to average
```

The `average` function is expecting one or more `float64` arguments, not a *slice* of `float64` values...

Passing slices to variadic functions (continued)

So what now? Are we forced to choose between making our functions variadic and being able to pass slices to them?

Fortunately, Go provides special syntax for this situation. When calling a variadic function, simply add an ellipsis (`...`) following the slice you want to use in place of variadic arguments.

```
func severalInts(numbers ...int) {
        fmt.Println(numbers)
}

func mix(num int, flag bool, strings ...string) {
        fmt.Println(num, flag, strings)
}

func main() {
        intSlice := []int{1, 2, 3}
        severalInts(intSlice...)
        stringSlice := []string{"a", "b", "c", "d"}
        mix(1, true, stringSlice...)
}
```

Use an int slice in place of the variadic arguments.

Use a string slice in place of the variadic arguments.

```
[1 2 3]
1 true [a b c d]
```

So all we need to do is add an ellipsis following the `numbers` slice in our call to `average`.

```
func main() {
        arguments := os.Args[1:]
        var numbers []float64
        for _, argument := range arguments {
                number, err := strconv.ParseFloat(argument, 64)
                if err != nil {
                        log.Fatal(err)
                }
                numbers = append(numbers, number)
        }
        fmt.Printf("Average: %0.2f\n", average(numbers...))
}
```

Pass the slice to the variadic function.

With that change made, we should be able to compile and run our program again. It will convert our command-line arguments to a slice of `float64` values, then pass that slice to the variadic `average` function.

It works!

```
Shell Edit View Window Help
$ go install github.com/headfirstgo/average2
$ cd /Users/jay/go/bin
$ ./average2 71.8 56.2 89.5
Average: 72.50
$ ./average2 90.7 89.7 98.5 92.3
Average: 92.80
```

Slices have saved the day!



This is great! I can just type in the amount of food I used over the previous weeks, and instantly see the average. And it's so convenient, I can estimate orders for **all** my ingredients this way! I may decide to install Go after all!

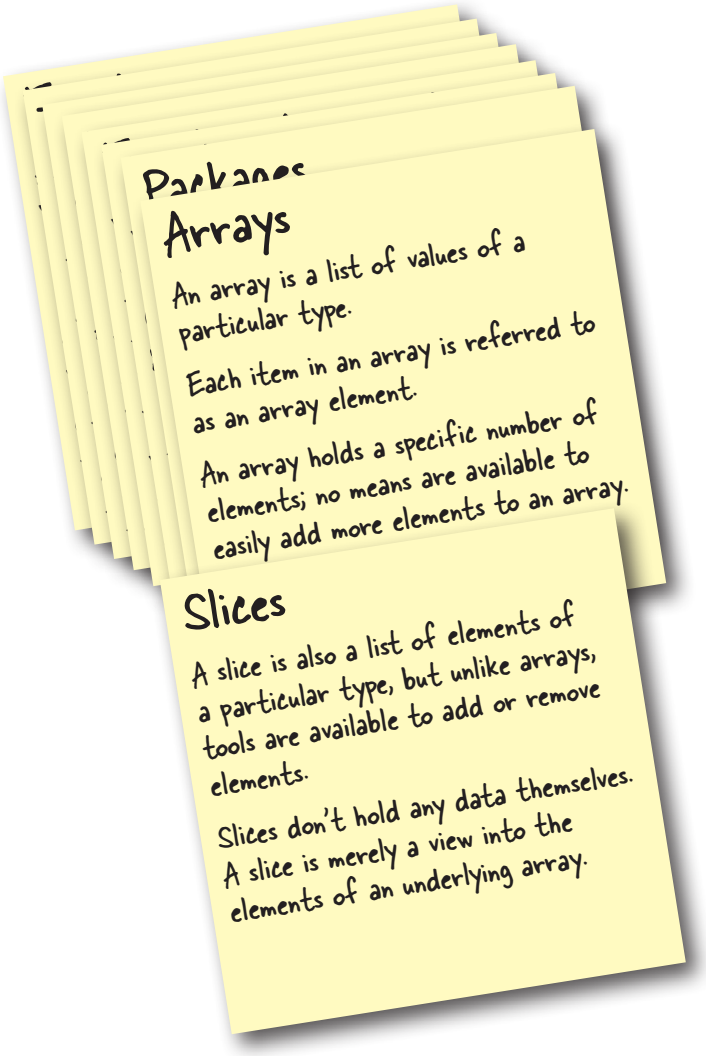
```
Shell  Edit  View  Window  Help
$ go install github.com/headfirstgo/average2
$ cd /Users/jay/go/bin
$ ./average2 71.8 56.2 89.5
Average: 72.50
$ ./average2 90.7 89.7 98.5 92.3
Average: 92.80
```

Working with lists of values is essential for any programming language. With arrays and slices, you can keep your data in collections of whatever size you need. And with features like `for...range` loops, Go makes it easy to process the data in those collections, too!



Your Go Toolbox

That's it for Chapter 6! You've added slices to your toolbox.



BULLET POINTS

- The type for a slice variable is declared just like the type for an array variable, except the length is omitted:
  `var mySlice []int`
- For the most part, code for working with slices is identical to code that works with arrays. This includes: accessing elements, using zero values, passing slices to the `len` function, and `for...range` loops.
- A **slice literal** looks just like an array literal, except the length is omitted:
  `[]int{1, 7, 10}`
- You can get a slice that contains elements `i` through `j - 1` of an array or slice using the **slice operator**: `s[i:j]`
- The `os.Args` package variable contains a slice of strings with the command-line arguments the current program was run with.
- A **variadic function** is one that can be called with a varying number of arguments.
- To declare a variadic function, place an ellipsis (`...`) before the type of the last parameter in the function declaration. That parameter will then receive all the variadic arguments as a slice.
- When calling a variadic function, you can use a slice in place of the variadic arguments by typing an ellipsis after the slice:
  `inRange(1, 10, mySlice...)`

Pool Puzzle Solution

```
package main

import "fmt"

func main() {
        numbers := make([]float64, 3)
        numbers[0] = 19.7
        numbers[2] = 25.2
        for i, number := range numbers {
                fmt.Println(i, number)
        }
        var letters = []string{"a", "b", "c"}
        for i, letter := range letters {
                fmt.Println(i, letter)
        }
}
```



Exercise Solution

Below is a program that takes a slice of an array and then appends elements to the slice. Write down what the program output would be.

```
package main

import "fmt"

func main() {
        array := [5]string{"a", "b", "c", "d", "e"}
        slice := array[1:3]
        slice = append(slice, "x")
        slice = append(slice, "y", "z")
        for _, letter := range slice {
                fmt.Println(letter)
        }
}
```

Output:

```
b
c
x
y
z
```

Code Magnets Solution

```
package main

import "fmt"

func sum ( numbers ... int ) int {
    var sum int = 0
    for _, number := range numbers {
        sum += number
    }
    return sum
}

func main() {
    fmt.Println( sum( 7 , 9 ))
    fmt.Println( sum( 1 , 2 , 4 ))
}
```

Output

```
16
7
```

7 labeling data

Maps



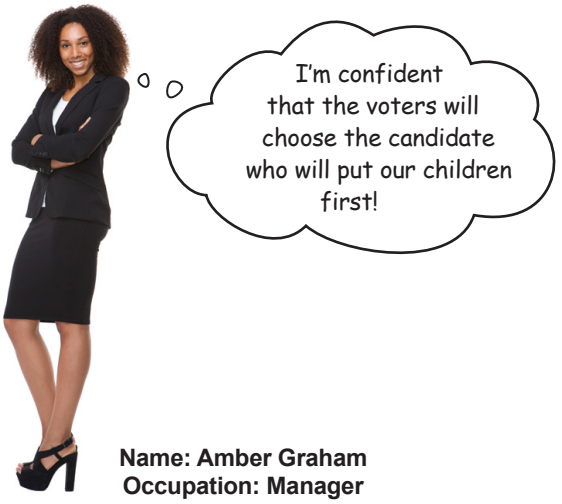
Throwing things in piles is fine, until you need to find something again. You've already seen how to create lists of values using *arrays* and *slices*. You've seen how to apply the same operation to *every value* in an array or slice. But what if you need to work with a *particular* value? To find it, you'll have to start at the beginning of the array or slice, and *look through Every. Single. Value.*

What if there were a kind of collection where every value had a label on it? You could quickly find just the value you needed! In this chapter, we'll look at **maps**, which do just that.

Counting votes

This is another example that debuted in Head First Ruby, in the hashes chapter. Ruby hashes are a lot like Go maps, so this example works great here, too!

A seat on the Sleepy Creek County School Board is up for grabs this year, and polls have been showing that the election is really close. Now that it's election night, the candidates are excitedly watching the votes roll in.



Name: Amber Graham
Occupation: Manager



Name: Brian Martin
Occupation: Accountant

There are two candidates on the ballot, Amber Graham and Brian Martin. Voters also have the option to "write in" a candidate's name (that is, type in a name that doesn't appear on the ballot). Those won't be as common as the main candidates, but we can expect a few such names to appear.

The electronic voting machines in use this year record the votes to text files, one vote per line. (Budgets are tight, so the city council chose the cheap voting machine vendor.)

Here's a file with all the votes for District A:

Each line represents one vote.

```
Amber Graham
Brian Martin
Amber Graham
Brian Martin
Amber Graham
```

votes.txt

We need to process each line of the file and tally the total number of times each name occurs. The name with the most votes will be our winner!

Reading names from a file

Our first order of business is to read the contents of the *votes.txt* file. The `datafile` package from previous chapters already has a `GetFloats` function that reads each line of a file into a slice, but `GetFloats` can only read `float64` values. We're going to need a separate function that can return the file lines as a slice of `string` values.

So let's start by creating a *strings.go* file alongside the *floats.go* file in the *datafile* package directory. In that file, we'll add a `GetStrings` function. The code in `GetStrings` will look much like the code in `GetFloats` (we've grayed out the code that's identical below). But instead of converting each line to a `float64` value, `GetStrings` will just add the line directly to the slice we're returning, as a `string` value.




```
// Package datafile allows reading data samples from files.
package datafile

import (
        "bufio"
        "os"
)

// GetStrings reads a string from each line of a file.
func GetStrings(fileName string) ([]string, error) {
        var lines []string
        file, err := os.Open(fileName)
        if err != nil {
                return nil, err
        }
        scanner := bufio.NewScanner(file)
        for scanner.Scan() {
                line := scanner.Text()
                lines = append(lines, line)
        }
        err = file.Close()
        if err != nil {
                return nil, err
        }
        if scanner.Err() != nil {
                return nil, scanner.Err()
        }
        return lines, nil
}
```

Still part of the same package as GetFloats

Don't import the "strconv" package; we don't need it in this file.

This variable holds a slice of strings.

Return a slice of strings instead of a slice of float64 values.

Instead of converting the file line string to a float64, add it to the slice directly.

Return the slice of strings.

Reading names from a file (continued)

Now let's create the program that will actually count the votes. We'll name it `count`. Within your Go workspace, go into the *src/github.com/headfirstgo* directory and create a new directory named *count*. Then create a file named *main.go* within the *count* directory.

Before writing the full program, let's confirm that our `GetStrings` function is working. At the top of the `main` function, we'll call `datafile.GetStrings`, passing it `"votes.txt"` as the name of the file to read from. We'll store the returned slice of strings in a new variable named `lines`, and any error in a variable named `err`. As usual, if `err` is not `nil`, we'll log the error and exit. Otherwise, we'll simply call `fmt.Println` to print out the contents of the `lines` slice.




```
// count tallies the number of times each line
// occurs within a file.
package main

import (
        "fmt"
        "github.com/headfirstgo/datafile"
        "log"
)

func main() {
        lines, err := datafile.GetStrings("votes.txt")
        if err != nil {
                log.Fatal(err)
        }
        fmt.Println(lines)
}
```

Import the "datafile" package, which now includes the GetStrings function.

Read the "votes.txt" file and return a slice of strings with every line from the file.

If there was an error, log it and exit.

Print the slice of strings.

As we've done with other programs, you can compile this program (plus any packages it depends on, `datafile` in this case) by running `go install` and providing it the package import path. If you used the directory structure shown above, that import path should be `github.com/headfirstgo/count`.

Compile the contents of the "count" directory, and install the resulting executable.

```
Shell Edit View Window Help
$ go install github.com/headfirstgo/count
```

That will save an executable file named *count* (or *count.exe* on Windows) in the *bin* subdirectory of your Go workspace.

Reading names from a file (continued)

As with the *data.txt* file in previous chapters, we need to ensure a *votes.txt* file is saved in the current directory when we run our program. In the *bin* subdirectory of your Go workspace, save a file with the contents shown at right. In your terminal, use the **cd** command to change to that same subdirectory.

Now you should be able to run the executable by typing **./count** (or **count.exe** on Windows). It should read every line of *votes.txt* into a slice of strings, then print that slice out.

```
Amber Graham
Brian Martin
Amber Graham
Brian Martin
Amber Graham
```

votes.txt

Change to the "bin" directory within your workspace.

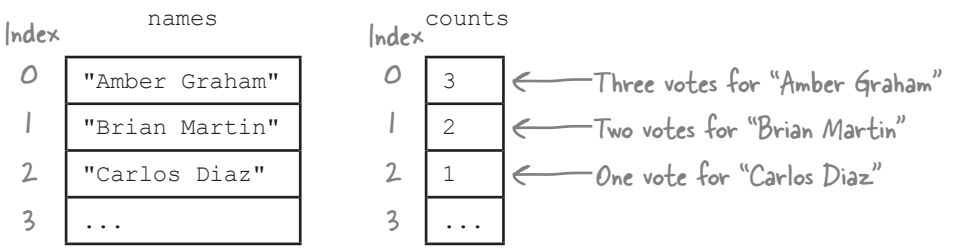
Run the executable.

```
Shell Edit View Window Help
$ cd /Users/jay/go/bin
$ ./count
[Amber Graham Brian Martin Amber Graham Brian Martin
Amber Graham]
$
```

Counting names the hard way, with slices

Reading a slice of names from the file didn't require learning anything new. But now comes the challenge: how do we count the number of times each name occurs? We'll show you two ways, first with slices, and then with a new data structure, *maps*.

For our first solution, we'll create two slices, each with the same number of elements, in a specific order. The first slice would hold the names we found in the file, with each name occurring once. We could call that one `names`. The second slice, `counts`, would hold the number of times each name was found in the file. The element `counts[0]` would hold the count for `names[0]`, `counts[1]` would hold the count for `names[1]`, and so on.



Counting names the hard way, with slices (continued)

Let's update the `count` program to actually count the number of times each name occurs in the file. We'll try this plan of using a `names` slice to hold each unique candidate name, and a corresponding `counts` slice to track the number of times each name occurs.




```
// ...Preceding code omitted...
func main() {
        lines, err := datafile.GetStrings("votes.txt")
        if err != nil {
                log.Fatal(err)
        }
        var names []string
        var counts []int
        for _, line := range lines {
                matched := false
                for i, name := range names {
                        if name == line {
                                counts[i]++
                                matched = true
                        }
                }
                if matched == false {
                        names = append(names, line)
                        counts = append(counts, 1)
                }
        }
        for i, name := range names {
                fmt.Printf("%s: %d\n", name, counts[i])
        }
}
```

This variable will hold a slice of candidate names.

Will hold a slice with the number of times each name occurs

Loop over each value in the names slice.

If this line matches the current name...

...increment the corresponding count.

Mark that we found a match.

If no match was found...

...add it as a new name...

And add a new count (this line will be the first occurrence).

Process each line from the file.

All done; print the results.

Print each element from the names slice...

...and the corresponding element from the counts slice.

As always, we can recompile the program with `go install`. If we run the resulting executable, it will read the *votes.txt* file and print each name it finds, along with the number of times that name occurs!

Compile the program.

Ensure we're in the "bin" subdirectory.

Run the updated program.

Counts for each name will be printed.

```
Shell Edit View Window Help
$ go install github.com/headfirstgo/count
$ cd /Users/jay/go/bin
$ ./count
Amber Graham: 3
Brian Martin: 2
```

Let's take a closer look at how this works...

Counting names the hard way, with slices (continued)

Our `count` program uses a loop nested *inside* another loop to tally the name counts. The outer loop assigns lines of the file to the `line` variable, one at a time.

Process each line of the file.

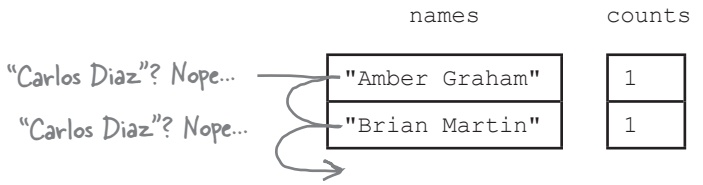
```
for _, line := range lines {
    // ...
}
```

The *inner* loop searches each element of the `names` slice, looking for a name equal to the current line from the file.

Search the "names" slice for one matching the current file line.

```
for i, name := range names {
    if name == line {
        counts[i] += 1
        matched = true
    }
}
```

Say someone adds a write-in candidate to their ballot, causing a line from the text file to be loaded with the string `"Carlos Diaz"`. The program will check the elements of `names`, one by one, to see if any of them equal `"Carlos Diaz"`.



If none matches, the program will append the string `"Carlos Diaz"` to the `names` slice, and a corresponding count of `1` to the `counts` slice (because this line represents the first vote for `"Carlos Diaz"`).

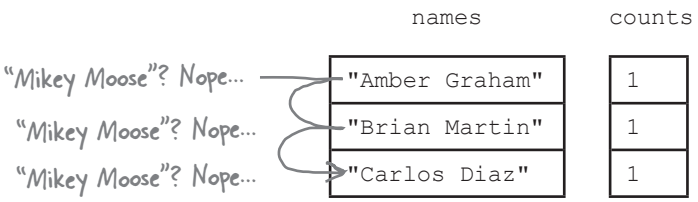


But suppose the next line is the string `"Brian Martin"`. Because that string already exists in the `names` slice, the program will find it and add `1` to the corresponding value in `counts` instead.



Maps

But here's the problem with storing the names in slices: for each and every line of the file, you have to search through many (if not all) of the values in the `names` slice to compare them. That may work okay in a small district like Sleepy Creek County, but in a bigger district with lots of votes, this approach will be way too slow!



Putting data in a slice is like stacking it in a big pile; you can get particular items back out, but you'll have to search through *everything* to find them.

Go has another way of storing collections of data: *maps*. A **map** is a collection where each value is accessed via a *key*. Keys are an easy way to get data back out of your map. It's like having neatly labeled file folders instead of a messy pile.



Slice

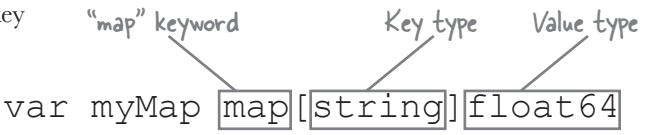


Map

Whereas arrays and slices can only use *integers* as indexes, a map can use *any* type for keys (as long as values of that type can be compared using ==). That includes numbers, strings, and more. The values all have to be of the same type, and the keys all have to be of the same type, but the keys don't have to be the same type as the values.

Maps (continued)

To declare a variable that holds a map, you type the `map` keyword, followed by square brackets (`[]`) containing the key type. Then, following the brackets, provide the value type.



Just as with slices, declaring a map variable doesn't automatically create a map; you need to call the `make` function (the same function you can use to create slices). Instead of a slice type, you can pass `make` the type of the map you want to create (which should be the same as the type of the variable you're going to assign it to).

```
var ranks map[string]int
ranks = make(map[string]int)
```

Declare a map variable.

Actually create the map.

You may find it's easier to just use a short variable declaration, though:

```
ranks := make(map[string]int)
```

Create a map and declare a variable to hold it.

The syntax to assign values to a map and get them back out again looks a lot like the syntax to assign and get values for arrays or slices. But while arrays and slices only let you use integers as element indexes, you can choose almost any type to use for a map's keys. The `ranks` map uses `string` keys:

```
ranks["gold"] = 1
ranks["silver"] = 2
ranks["bronze"] = 3
fmt.Println(ranks["bronze"])
fmt.Println(ranks["gold"])
```

```
3
1
```

Arrays and slices only let you use integer indexes. But you can choose almost any type to use for a map's keys.

Here's another map with strings as keys and strings as values:

```
elements := make(map[string]string)
elements["H"] = "Hydrogen"
elements["Li"] = "Lithium"
fmt.Println(elements["Li"])
fmt.Println(elements["H"])
```

```
Lithium
Hydrogen
```

Here's a map with integers as keys and booleans as values:

```
isPrime := make(map[int]bool)
isPrime[4] = false
isPrime[7] = true
fmt.Println(isPrime[4])
fmt.Println(isPrime[7])
```

```
false
true
```

Map literals

Just as with arrays and slices, if you know keys and values that you want your map to have in advance, you can use a **map literal** to create it. A map literal starts with the map type (in the form `map[KeyType]ValueType`). This is followed by curly braces containing key/value pairs you want the map to start with. For each key/value pair, you include the key, a colon, and then the value. Multiple key/value pairs are separated by commas.



Here are a couple of the preceding map examples, re-created using map literals:

```
ranks := map[string]int{"bronze": 3, "silver": 2, "gold": 1}
fmt.Println(ranks["gold"])
fmt.Println(ranks["bronze"])
elements := map[string]string{
      "H": "Hydrogen",
      "Li": "Lithium",
}
fmt.Println(elements["H"])
fmt.Println(elements["Li"])
```

← Map literal

← Multiline map literal

```
1
3
Hydrogen
Lithium
```

As with slice literals, leaving the curly braces empty creates a map that starts empty.

```
emptyMap := map[string]float64{}
```

Create an empty map.



Fill in the blanks in the program below, so it will produce the output shown.

```
jewelry := _____(map[string]float64)
jewelry["necklace"] = 89.99
jewelry[_____] = 79.99
clothing := ____[string]float64{_____: 59.99, "shirt": 39.99}
fmt.Println("Earrings:", jewelry["earrings"])
fmt.Println("Necklace:", jewelry[_____])
fmt.Println("Shirt:", clothing[_____])
fmt.Println("Pants:", clothing["pants"])
```

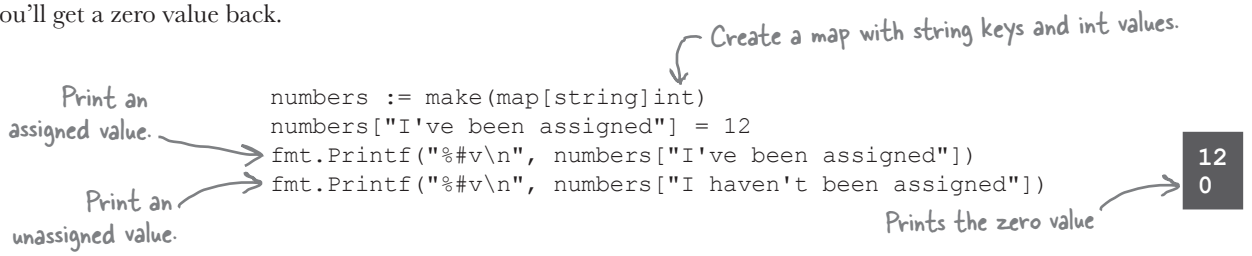
Output

```
Earrings: 79.99
Necklace: 89.99
Shirt: 39.99
Pants: 59.99
```

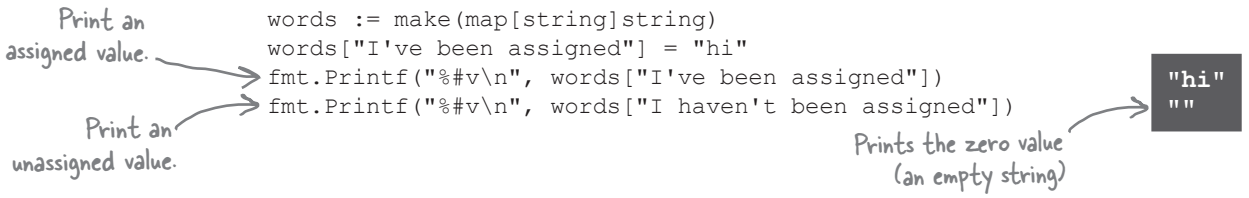
Answers on page 228.

Zero values within maps

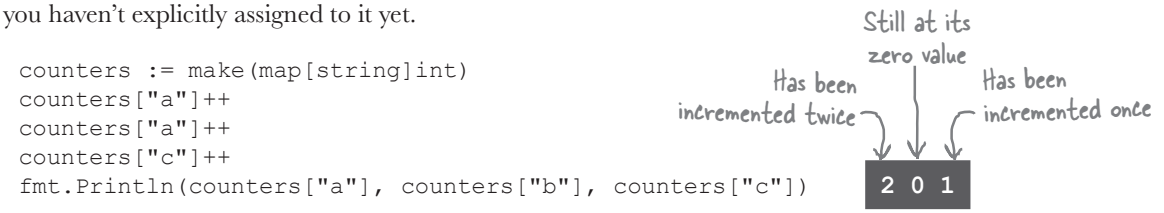
As with arrays and slices, if you access a map key that hasn't been assigned to, you'll get a zero value back.



Depending on the value type, the zero value may not actually be `0`. For maps with a value type of `string`, for example, the zero value will be an empty string.



As with arrays and slices, zero values can make it safe to manipulate a map value even if you haven't explicitly assigned to it yet.

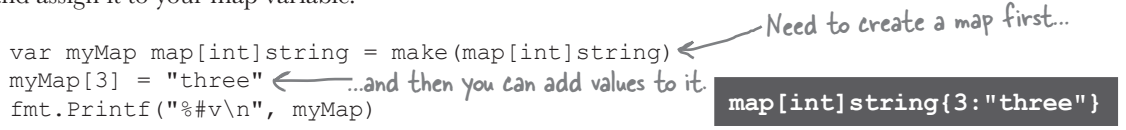


The zero value for a map variable is nil

As with slices, the zero value for the map variable itself is `nil`. If you declare a `map` variable, but don't assign it a value, its value will be `nil`. That means no map exists to add new keys and values to. If you try, you'll get a panic:



Before attempting to add keys and values, create a map using `make` or a map literal, and assign it to your map variable.



How to tell zero values apart from assigned values

Zero values, although useful, can sometimes make it difficult to tell whether a given key has been assigned the zero value, or if it has never been assigned.

Here's an example of a program where this could be an issue. This code erroneously reports that the student `"Carl"` is failing, when in reality he just hasn't had any grades logged:

```
func status(name string) {
      grades := map[string]float64{"Alma": 0, "Rohit": 86.5}
      grade := grades[name]
      if grade < 60 {
            fmt.Printf("%s is failing!\n", name)
      }
}

func main() {
      status("Alma")
      status("Carl")
}
```

A map key with a value of 0 assigned

A map key with no value assigned

```
Alma is failing!
Carl is failing!
```

To address situations like this, accessing a map key optionally returns a second, Boolean value. It will be `true` if the returned value has actually been assigned to the map, or `false` if the returned value just represents the default zero value. Most Go developers assign this Boolean value to a variable named `ok` (because the name is nice and short).

```
counters := map[string]int{"a": 3, "b": 0}
var value int
var ok bool
value, ok = counters["a"]
fmt.Println(value, ok)
value, ok = counters["b"]
fmt.Println(value, ok)
value, ok = counters["c"]
fmt.Println(value, ok)
```

Access an assigned value. "ok" will be true.

Access an assigned value. "ok" will be true.

Access an unassigned value. "ok" will be false.

The Go maintainers refer to this as the "comma ok idiom." We'll see it again with type assertions in Chapter 11.

```
3 true
0 true
0 false
```

If you only want to test whether a value is present, you can have the value itself ignored by assigning it to the _ blank identifier.

Test for the value's presence, but ignore it.

Test for the value's presence, but ignore it.

```
counters := map[string]int{"a": 3, "b": 0}
var ok bool
_, ok = counters["b"]
fmt.Println(ok)
_, ok = counters["c"]
fmt.Println(ok)
```

```
true
false
```

How to tell zero values apart from assigned values (continued)

The second return value can be used to decide whether you should treat the value you got from the map as an assigned value that just happens to match the zero value for that type, or as an unassigned value.

Here's an update to our code that tests whether the requested key has actually had a value assigned before it reports a failing grade:

Get the value, plus a boolean indicating whether this is an assigned value.

If no value was assigned to the specified key...

Otherwise, follow the logic for reporting a failing grade.

...report that no grade has been logged for the student.

```
func status(name string) {
        grades := map[string]float64{"Alma": 0, "Rohit": 86.5}
        grade, ok := grades[name]
        if !ok {
                fmt.Printf("No grade recorded for %s.\n", name)
        } else if grade < 60 {
                fmt.Printf("%s is failing!\n", name)
        }
}

func main() {
        status("Alma")
        status("Carl")
}
```

```
Alma is failing!
No grade recorded for Carl.
```



Write down what the output of this program snippet would be.

```
data := []string{"a", "c", "e", "a", "e"}
counts := make(map[string]int)
for _, item := range data {
        counts[item]++
}
letters := []string{"a", "b", "c", "d", "e"}
for _, letter := range letters {
        count, ok := counts[letter]
        if !ok {
                fmt.Printf("%s: not found\n", letter)
        } else {
                fmt.Printf("%s: %d\n", letter, count)
        }
}
```

Output:

Answers on page 228.

Removing key/value pairs with the "delete" function

At some point after assigning a value to a key, you may want to remove it from your map. Go provides the built-in `delete` function for this purpose. Just pass the `delete` function two things: the map you want to delete a key from, and the key you want deleted. That key and its corresponding value will be removed from the map.

In the code below, we assign values to keys in two different maps, then delete them again. After that, when we try accessing those keys, we get a zero value (which is `0` for the `ranks` map, `false` for the `isPrime` map). The secondary Boolean value is also `false` in each case, which means that the key is not present.

```
var ok bool
ranks := make(map[string]int)
var rank int
ranks["bronze"] = 3
rank, ok = ranks["bronze"]
fmt.Printf("rank: %d, ok: %v\n", rank, ok)
delete(ranks, "bronze")
rank, ok = ranks["bronze"]
fmt.Printf("rank: %d, ok: %v\n", rank, ok)

isPrime := make(map[int]bool)
var prime bool
isPrime[5] = true
prime, ok = isPrime[5]
fmt.Printf("prime: %v, ok: %v\n", prime, ok)
delete(isPrime, 5)
prime, ok = isPrime[5]
fmt.Printf("prime: %v, ok: %v\n", prime, ok)
```

Assign a value to the "bronze" key.

"ok" will be true because a value is present.

Delete the "bronze" key and its corresponding value.

"ok" will be false because the value's been deleted.

Assign a value to the 5 key.

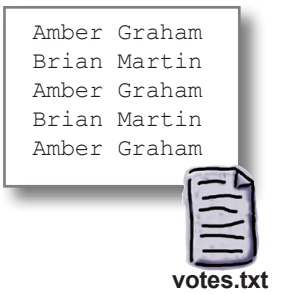
"ok" will be true because a value is present.

Delete the 5 key and its corresponding value.

"ok" will be false because the value's been deleted.

```
rank: 3, ok: true
rank: 0, ok: false
prime: true, ok: true
prime: false, ok: false
```

Updating our vote counting program to use maps



Now that we understand maps a bit better, let's see if we can use what we've learned to simplify our vote counting program.

Previously, we used a pair of slices, one called `names` that held candidate names, and one called `counts` held vote counts for each name. For each name we read from the file, we had to search through the slice of names, one by one, for a match. We then incremented the vote count for that name in the corresponding element of the `counts` slice.

```
// ...
var names []string      <- This variable will hold a slice of candidate names.
var counts []int        <- This variable will hold a slice with the number of times each name occurs.
for _, line := range lines {
        matched := false
        for i, name := range names {      <- Loop over each value in the names slice.
                if name == line {         <- If this line matches the current name...
                        counts[i] += 1    <- ...increment the corresponding count.
// ...
```

Using a map will be much simpler. We can replace the two slices with a single map (which we'll also call `counts`). Our map will use candidate names as its keys, and integers (which will hold the vote counts for that name) as its values. Once that's set up, all we have to do is use each candidate name we read from the file as a map key, and increment the value that key holds.

Here's some simplified code that creates a map and increments the values for some candidate names directly:

```
counts := make(map[string]int)
counts["Amber Graham"]++
counts["Brian Martin"]++
counts["Amber Graham"]++
fmt.Println(counts)
```

```
map[Amber Graham:2 Brian Martin:1]
```

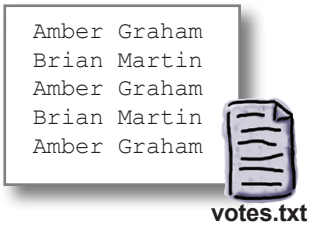
Our previous program needed separate logic to add new elements to both slices if the name wasn't found...

```
if matched == false {      <- If no match was found...
        names = append(names, line)      <- ...add it as a new name...
        counts = append(counts, 1)       <- ...and add a new count (this line will be the first occurrence).
}
```

But we don't need to do that with a map. If the key we're accessing doesn't already exist, we'll get the zero value back (literally `0` in this case, since our values are integers). We then increment that value, giving us `1`, which gets assigned to the map. When we encounter that name again, we'll get the assigned value, which we can then increment as normal.

Updating our vote counting program to use maps (continued)

Next, let's try incorporating our `counts` map into the actual program, so it can tally the votes from the actual file.



We'll be honest; after all that work to learn about maps, the final code looks a little anticlimactic! We replace the two slice declarations with a single map declaration. Next is the code in the loop that processes strings from the file. We replace the original 11 lines of code there with a single line, which increments the count in the map for the current candidate name. And we replace the loop at the end that prints the results with a single line that prints the whole `counts` map.




```
package main

import (
        "fmt"
        "github.com/headfirstgo/datafile"
        "log"
)

func main() {
        lines, err := datafile.GetStrings("votes.txt")
        if err != nil {
                log.Fatal(err)
        }
        counts := make(map[string]int)
        for _, line := range lines {
                counts[line]++
        }
        fmt.Println(counts)
}
```

Declare a map that will use candidate names as keys, and vote counts as values.

Increment the vote count for the current candidate.

Print the populated map.

Trust us, though, the code only *looks* anticlimactic. There are still complex operations going on here. But the map is handling them all for you, which means you don't have to write as much code!

As before, you can recompile the program using the `go install` command. When we rerun the executable, the *votes.txt* file will be loaded and processed. We'll see the `counts` map printed, with the number of times each name was encountered in the file.

```
Shell Edit View Window Help
$ go install github.com/headfirstgo/count
$ cd /Users/jay/go/bin
$ ./count
map[Amber Graham:3 Brian Martin:2]
```

Using for...range loops with maps



This program is really handy. But we can't show the results to the press like this... Can you print them in a more legible format?

The format we have

```
map[Amber Graham:3 Brian Martin:2]
```

Name: Kevin Wagner
Occupation: Election Volunteer

The format we want

```
Amber Graham: 3
Brian Martin: 2
```

That's true. A format of one name and one vote count per line would probably be better:

To format each key and value from the map as a separate line, we're going to need to loop through each entry in the map.

The same `for...range` loop we've been using to process array and slice elements works on maps, too. Instead of assigning an integer index to the first variable you provide, however, the current map key will be assigned.

Variable that will hold each map key

Variable that will hold each corresponding value

"range" keyword

The map being processed

```
for key, value := range myMap {
    // Loop block here.
}
```

Using for...range loops with maps (continued)

The `for...range` loop makes it easy to loop through a map's keys and values. Just provide a variable to hold each key, and another to hold the corresponding value, and it will automatically loop through each entry in the map.

```
package main

import "fmt"

func main() {
      grades := map[string]float64{"Alma": 74.2, "Rohit": 86.5, "Carl": 59.7}
      for name, grade := range grades {
            fmt.Printf("%s has a grade of %0.1f%%\n", name, grade)
      }
}
```

Loop through each key/value pair.

Print each key and its corresponding value.

```
Carl has a grade of 59.7%
Alma has a grade of 74.2%
Rohit has a grade of 86.5%
```

If you only need to loop through the keys, you can omit the variable that holds the values:

Process only the keys.

```
fmt.Println("Class roster:")
for name := range grades {
      fmt.Println(name)
}
```

```
Class roster:
Alma
Rohit
Carl
```

And if you only need the values, you can assign the keys to the _ blank identifier:

Process only the values.

```
fmt.Println("Grades:")
for _, grade := range grades {
      fmt.Println(grade)
}
```

```
Grades:
59.7
74.2
86.5
```

But there's one potential issue with this example... If you save the preceding example to a file and run it with `go run`, you'll find that the map keys and values are printed in a random order. If you run the program multiple times, you'll get a different order each time.

(Note: The same is not true of code run via the online Go Playground site. There, the order will still be random, but it will produce the same output each time it's run.)

The loop follows a different order each time!

```
Shell  Edit  View  Window  Help
$ go run temp.go
Alma has a grade of 74.2%
Rohit has a grade of 86.5%
Carl has a grade of 59.7%
$ go run temp.go
Carl has a grade of 59.7%
Alma has a grade of 74.2%
Rohit has a grade of 86.5%
```

The for...range loop handles maps in random order!

The `for...range` loop processes map keys and values in a random order because a map is an *unordered* collection of keys and values. When you use a `for...range` loop with a map, you never know what order you'll get the map's contents in! Sometimes that's fine, but if you need more consistent ordering, you'll need to write the code for that yourself.

Here's an update to the previous program that always prints the names in alphabetical order. It does using two separate `for` loops. The first loops over each key in the map, ignoring the values, and adds them to a slice of strings. Then, the slice is passed to the `sort` package's `Strings` function to sort it alphabetically, in place.

The second `for` loop doesn't loop over the map, it loops over the sorted slice of names. (Which, thanks to the preceding code, now contains every key from the map in alphabetical order.) It prints the name and then gets the value that matches that name from the map. It still processes every key and value in the map, but it gets the keys from the sorted slice, not the map itself.

```
package main

import (
        "fmt"
        "sort"
)

func main() {
        grades := map[string]float64{"Alma": 74.2, "Rohit": 86.5, "Carl": 59.7}
        var names []string
        for name := range grades {
                names = append(names, name)
        }
        sort.Strings(names)
        for _, name := range names {
                fmt.Printf("%s has a grade of %0.1f%%\n", name, grades[name])
        }
}
```

Build a slice with all the map keys.

Sort the slice alphabetically.

Process the names alphabetically.

Use the current student name to get the grade from the map.

If we save the above code and run it, this time the student names are printed in alphabetical order. This will be true no matter how many times we run the program.

If it doesn't matter what order your map data is processed in, using a `for...range` loop directly on the map will probably work for you. But if order matters, you may want to consider setting up your own code to handle the processing order.

```
Shell Edit View Window Help
$ go run temp.go
Alma has a grade of 74.2%
Carl has a grade of 59.7%
Rohit has a grade of 86.5%
$ go run temp.go
Alma has a grade of 74.2%
Carl has a grade of 59.7%
Rohit has a grade of 86.5%
```

The names are processed in alphabetical order each time.

Updating our vote counting program with a for...range loop

There aren't a lot of candidates in Sleepy Creek County, so we don't see a need to sort the output by name. We'll just use a `for...range` loop to process the keys and values directly from the map.

It's a pretty simple change to make; we just replace the line that prints the entire map with a `for...range` loop. We'll assign each key to a `name` variable, and each value to a `count` variable. Then we'll call `Printf` to print the current candidate name and vote count.

```
Amber Graham
Brian Martin
Amber Graham
Brian Martin
Amber Graham
```

votes.txt

your workspace > src > github.com > headfirstgo > count > main.go

```
package main

import (
	"fmt"
	"github.com/headfirstgo/datafile"
	"log"
)

func main() {
	lines, err := datafile.GetStrings("votes.txt")
	if err != nil {
		log.Fatal(err)
	}
	counts := make(map[string]int)
	for _, line := range lines {
		counts[line]++
	}
	for name, count := range counts {
		fmt.Printf("Votes for %s: %d\n", name, count)
	}
}
```

Process each map key and value.

Print the key (the candidate name).

Print the value (the vote count).

Another compilation via `go install`, another run of the executable, and we'll see our output in its new format. Each candidate name and their vote count is here, neatly formatted on its own line.

```
Shell Edit View Window Help
$ go install github.com/headfirstgo/count
$ cd /Users/jay/go/bin
$ ./count
Votes for Amber Graham: 3
Votes for Brian Martin: 2
```

The vote counting program is complete!



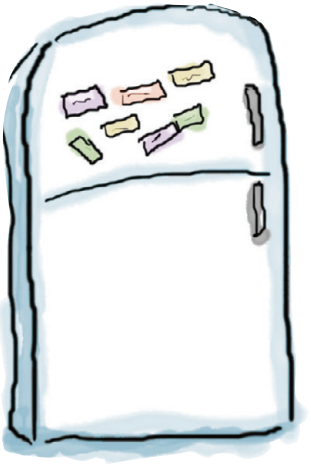
I knew the voters would make the right choice! I'd like to congratulate my opponent on a hard-fought campaign...

```
Shell Edit View Window Help
$ go install github.com/headfirstgo/count
$ cd /Users/jay/go/bin
$ ./count
Votes for Amber Graham: 3
Votes for Brian Martin: 2
```

Our vote counting program is complete!

When the only data collections we had available were arrays and slices, we needed a lot of extra code and processing time to look values up. But maps have made the process easy! Anytime you need to be able to find a collection's values again, you should consider using a map!

Code Magnets



A Go program that uses a `for...range` loop to print out the contents of a map is scrambled up on the fridge. Can you reconstruct the code snippets to make a working program that will produce the given output? (It's okay if the output order differs between runs of the program.)

```
package main
import "fmt"
func main() {
ranks
{
}
}
}
{
"bronze": 3
"silver": 2
"gold": 1
:=
:=
,
,
,
,
,
fmt.Printf(
)
"The %s medal's rank is %d\n"
map
range
ranks
for
[string]
int
medal
rank
medal
rank
```

Output

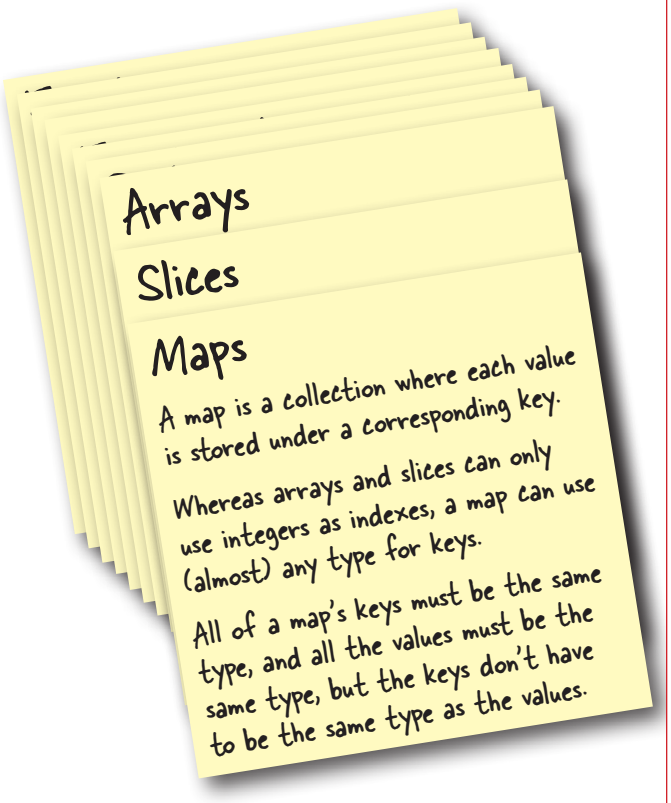
```
The gold medal's rank is 1
The bronze medal's rank is 3
The silver medal's rank is 2
```

Answers on page 229.



Your Go Toolbox

That's it for Chapter 7! You've added maps to your toolbox.



BULLET POINTS

- When declaring a map variable, you must provide the types for its keys and its values:
  `var myMap map[string]int`
- To create a new map, call the `make` function with the type of the map you want:
  `myMap = make(map[string]int)`
- To assign a value to a map, provide the key you want to assign it to in square brackets:
  `myMap["my key"] = 12`
- To get a value, you provide the key as well:
  `fmt.Println(myMap["my key"])`
- You can create a map and initialize it with data at the same time using a **map literal**:
  `map[string]int{"a": 2, "b": 3}`
- As with arrays and slices, if you access a map key that hasn't been assigned a value, you'll get a zero value back.
- Getting a value from a map can return a second, optional Boolean value that indicates whether that value was assigned, or if it represents a default zero value:
  `value, ok := myMap["c"]`
- If you only want to test whether a key has had a value assigned, you can ignore the actual value using the _ blank identifier:
  `_, ok := myMap["c"]`
- You can delete keys and their corresponding values from a map using the `delete` built-in function:
  `delete(myMap, "b")`
- You can use `for...range` loops with maps, much like you can with arrays or slices. You provide one variable that will be assigned each key in turn, and a second variable that will be assigned each value in turn.

```
for key, value := range myMap {
    fmt.Println(key, value)
}
```



Exercise Solution

Fill in the blanks in the program below, so it will produce the output shown.

```
jewelry := make(map[string]float64)
jewelry["necklace"] = 89.99
jewelry[__"earrings"__] = 79.99
clothing := map[string]float64{__"pants"__: 59.99, "shirt": 39.99}
fmt.Println("Earrings:", jewelry["earrings"])
fmt.Println("Necklace:", jewelry[__"necklace"__])
fmt.Println("Shirt:", clothing[__"shirt"__])
fmt.Println("Pants:", clothing["pants"])
```

Make a new, empty map.

Assign values to keys.

Make a new, pre-populated map using a map literal.

Print various values from the maps.

Output

```
Earrings: 79.99
Necklace: 89.99
Shirt: 39.99
Pants: 59.99
```



Exercise Solution

Write down what the output of this program snippet would be.

We'll count the number of times each letter occurs within this slice.

```
data := []string{"a", "c", "e", "a", "e"}
counts := make(map[string]int)
for _, item := range data {
        counts[item]++
}
letters := []string{"a", "b", "c", "d", "e"}
for _, letter := range letters {
        count, ok := counts[letter]
        if !ok {
                fmt.Printf("%s: not found\n", letter)
        } else {
                fmt.Printf("%s: %d\n", letter, count)
        }
}
```

A map to hold the counts

Process each letter.

Increment the count for the current letter.

We'll see if each of these letters exists as a key in the map.

Get the count for the current letter, as well as an indicator of whether it was found at all.

If letter was not found...

...say so.

Otherwise, letter was found...

...so print the letter and the count that was recorded for it.

Output:

a: 2

b: not found

c: 1

d: not found

e: 2

Code Magnets Solution

```
package main

import "fmt"

func main() {

    ranks := map [string] int { "bronze": 3 , "silver": 2 , "gold": 1 }

    for medal , rank := range ranks {
        fmt.Printf( "The %s medal's rank is %d\n" , medal , rank )
    }

}
```

Process each key and value in the map.

Print the key and value.

Output

```
The gold medal's rank is 1
The bronze medal's rank is 3
The silver medal's rank is 2
```

Structs



Sometimes you need to store more than one type of data.

We learned about slices, which store a list of values. Then we learned about maps, which map a list of keys to a list of values. But both of these data structures can only hold values of *one* type. Sometimes, you need to group together values of *several* types. Think of billing receipts, where you have to mix item names (strings) with quantities (integers). Or student records, where you have to mix student names (strings) with grade point averages (floating-point numbers). You can't mix value types in slices or maps. But you *can* if you use another type called a **struct**. We'll learn all about structs in this chapter!

Slices and maps hold values of ONE type

Gopher Fancy is a new magazine devoted to lovable rodents. They're currently working on a system to keep track of their subscriber base.



To start, we need to store the subscriber's name, the monthly rate we're charging them, and whether their subscription is active. But the name is a `string`, the rate is a `float64`, and the active indicator is a `bool`. We can't make one slice hold all those types!

A slice can only be set up to hold one type of value.

```
subscriber := []string{}
subscriber = append(subscriber, "Aman Singh")
subscriber = append(subscriber, 4.99)
subscriber = append(subscriber, true)
```

We can't add this float64!

We can't add this boolean!

```
cannot use 4.99 (type float64) as type string in append
cannot use true (type bool) as type string in append
```



Then we tried maps. We wish that would have worked, because we could have used the keys to label what each value represented. But just like slices, maps can only hold one type of value!

A map can only hold one type of value.

```
subscriber := map[string]float64{}
subscriber["name"] = "Aman Singh"
subscriber["rate"] = 4.99
subscriber["active"] = true
```

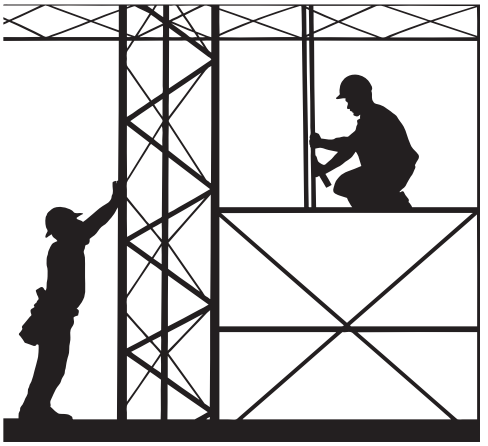
We can't store this string!

We can't store this boolean!

```
cannot use "Aman Singh" (type string)
as type float64 in assignment
cannot use true (type bool)
as type float64 in assignment
```

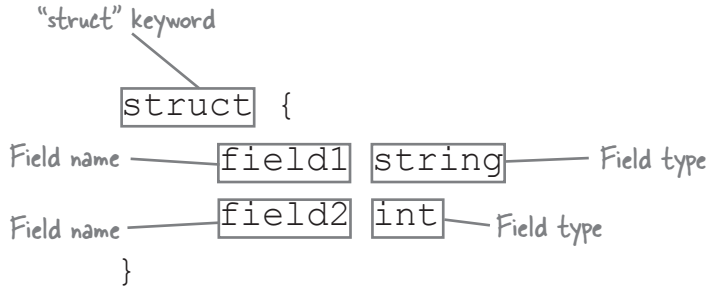
It's true: arrays, slices, and maps are no help if you need to mix values of different types. They can only be set up to hold values of a single type. But Go does have a way to solve this problem...

Structs are built out of values of MANY types



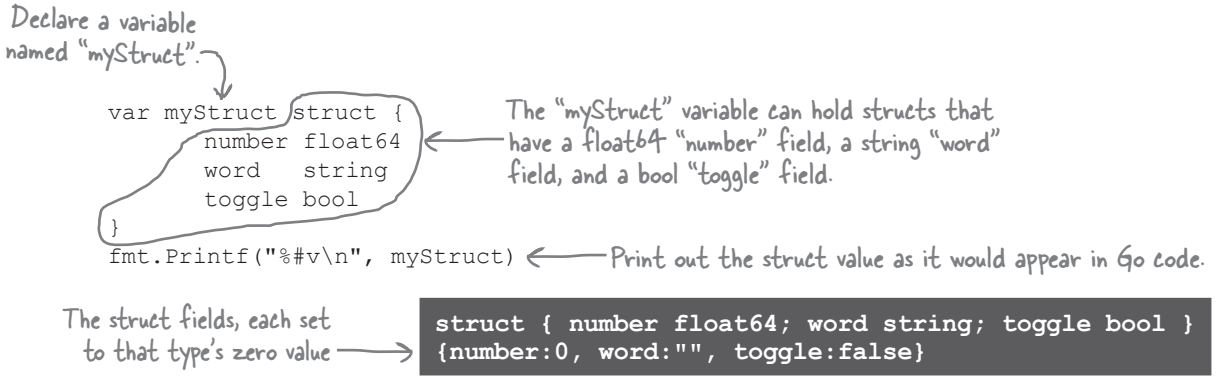
A **struct** (short for "structure") is a value that is constructed out of other values of many different types. Whereas a slice might only be able to hold `string` values or a map might only be able to hold `int` values, you can create a struct that holds `string` values, `int` values, `float64` values, `bool` values, and more—all in one convenient group.

You declare a struct type using the `struct` keyword, followed by curly braces. Within the braces, you can define one or more **fields**: values that the struct groups together. Each field definition appears on a separate line, and consists of a field name, followed by the type of value that field will hold.



You can use a struct type as the type of a variable you're declaring. This code declares a variable named `myStruct` that holds structs that have a `float64` field named `number`, a `string` field named `word`, and a `bool` field named `toggle`:

(It's more common to use a defined type to declare struct variables, but we won't cover type definitions for a few more pages, so we'll write it this way for now.)



When we call `Printf` with the `%#v` verb above, it prints the value in `myStruct` as a struct literal. We'll be covering struct literals later in the chapter, but for now you can see that the struct's `number` field has been set to `0`, the `word` field to an empty string, and the `toggle` field to `false`. Each field has been set to the zero value for its type.



Don't worry about the number of spaces between struct field names and their types.

When you write your struct fields, just insert a single space between the field name and its type. When you run the `go fmt` command on your files (which you should always do), it will insert extra spaces so that all the types align vertically. The alignment just makes the code easier to read; it doesn't change its meaning at all!

```
var aStruct struct {
        shortName   int
        longerName  float64
        longestName string
}
```

Extra spaces align the field types.

Access struct fields using the dot operator

Now we can define a struct, but to actually use it, we need a way to store new values in the struct's fields and retrieve them again.

All along, we've been using the dot operator to indicate functions that "belong to" another package, or methods that "belong to" a value:

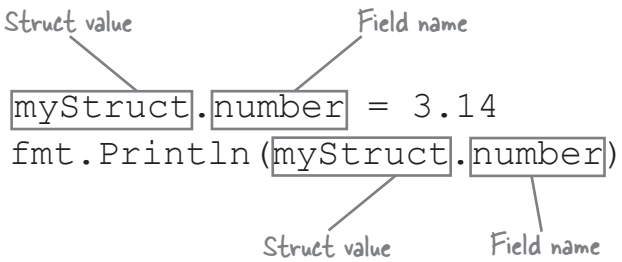
```
fmt.Println("hi")
```

Call a function belonging to the "fmt" package.

```
var myTime time.Time
myTime.Year()
```

Call a method belonging to a "Time" value.

Similarly, we can use a dot operator to indicate fields that "belong to" a struct. This works for both assigning values and retrieving them.



We can use dot operators to assign values to all the fields of `myStruct` and then print them back out:

```
var myStruct struct {
        number float64
        word   string
        toggle bool
}
myStruct.number = 3.14
myStruct.word = "pie"
myStruct.toggle = true
fmt.Println(myStruct.number)
fmt.Println(myStruct.word)
fmt.Println(myStruct.toggle)
```

Assign values to struct fields.

Retrieve values from struct fields.

```
3.14
pie
true
```

Storing subscriber data in a struct

Now that we know how to declare a variable that holds a struct and assign values to its fields, we can create a struct to hold magazine subscriber data.

First, we'll define a variable named `subscriber`. We'll give `subscriber` a `struct` type with `name` (`string`), `rate` (`float64`), and `active` (`bool`) fields.

With the variable and its type declared, we can then use dot operators to access the struct's fields. We assign values of the appropriate type to each field, and then print the values back out again.

Declare a "subscriber" variable...

...that holds structs.

```
var subscriber struct {
        name   string
        rate   float64
        active bool
}
```

The struct will have a "name" field that holds a string...

...a "rate" field that holds a float64...

...and an "active" field that holds a bool.

Assign values to struct fields.

```
subscriber.name = "Aman Singh"
subscriber.rate = 4.99
subscriber.active = true
```

Retrieve values from struct fields.

```
fmt.Println("Name:", subscriber.name)
fmt.Println("Monthly rate:", subscriber.rate)
fmt.Println("Active?", subscriber.active)
```

```
Name: Aman Singh
Monthly rate: 4.99
Active? true
```

Even though the data we have for a subscriber is stored using a variety of types, structs let us keep it all in one convenient package!



Exercise

At the right is a program that creates a struct variable to hold a pet's name (a `string`) and age (an `int`). Fill in the blanks so that the code will produce the output shown.

```
package main

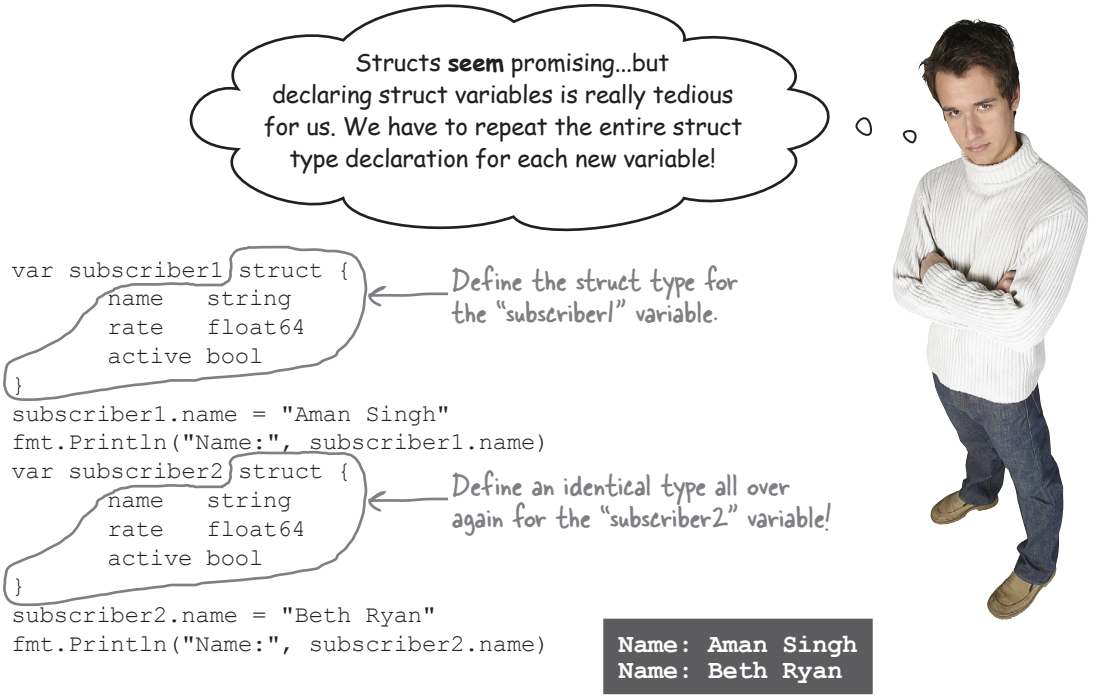
import "fmt"

func main() {
        var pet _______ {
                name   _______
                ____    int
        }
        pet._____ = "Max"
        pet.age = 5
        fmt.Println("Name:", ____.name)
        fmt.Println("Age:", pet.____)
}
```

```
Name: Max
Age: 5
```

Answers on page 262.

Defined types and structs

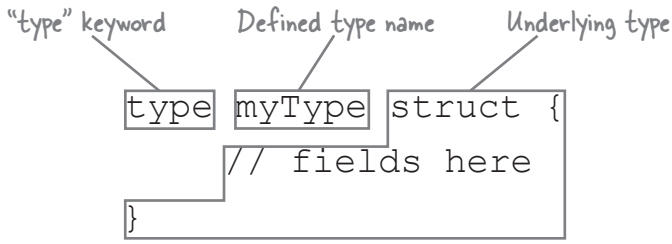


Throughout this book, you've used a variety of types, like `int`, `string`, `bool`, slices, maps, and now structs. But you haven't been able to create completely *new* types.

Type definitions allow you to create types of your own. They let you create a new **defined type** that's based on an **underlying type**.

Although you can use any type as an underlying type, such as `float64`, `string`, or even slices or maps, in this chapter we're going to focus on using struct types as underlying types. We'll try using other underlying types when we take a deeper look at defined types in the next chapter.

To write a type definition, use the `type` keyword, followed by the name for your new defined type, and then the underlying type you want to base it on. If you're using a struct type as your underlying type, you'll use the `struct` keyword followed by a list of field definitions in curly braces, just as you did when declaring struct variables.



Defined types and structs (continued)

Just like variables, type definitions *can* be written within a function. But that will limit its scope to that function's block, meaning you won't be able to use it outside that function. So types are usually defined outside of any functions, at the package level.

As a quick demonstration, the code below defines two types: `part` and `car`. Each defined type uses a struct as its underlying type.

Then, within the `main` function, we declare a `porsche` variable of the `car` type, and a `bolts` variable of the `part` type. There's no need to rewrite the lengthy struct definitions when declaring the variables; we just use the names of the defined types.

```
package main

import "fmt"

type part struct {
        description string
        count       int
}

type car struct {
        name     string
        topSpeed float64
}

func main() {
        var porsche car
        porsche.name = "Porsche 911 R"
        porsche.topSpeed = 323
        fmt.Println("Name:", porsche.name)
        fmt.Println("Top speed:", porsche.topSpeed)

        var bolts part
        bolts.description = "Hex bolts"
        bolts.count = 24
        fmt.Println("Description:", bolts.description)
        fmt.Println("Count:", bolts.count)
}
```

Define a type named "part".

The underlying type for "part" will be a struct with these fields.

Define a type named "car".

The underlying type for "car" will be a struct with these fields.

Declare a variable of type "car".

Access the struct fields.

Declare a variable of type "part".

Access the struct fields.

```
Name: Porsche 911 R
Top speed: 323
Description: Hex bolts
Count: 24
```

With the variables declared, we can set the values of their struct fields and get the values back out, just as we did in previous programs.

Using a defined type for magazine subscribers

Previously, to create more than one variable that stored magazine subscriber data in a struct, we had to write out the full struct type (including all its fields) for each variable.

```
var subscriber1 struct {
        name   string
        rate   float64
        active bool
}
// ...
var subscriber2 struct {
        name   string
        rate   float64
        active bool
}
// ...
```

Define a struct type.

Define an identical type.

But now, we can simply define a `subscriber` type at the package level. We write the struct type just once, as the underlying type for the defined type. When we're ready to declare variables, we don't have to write the struct type again; we simply use `subscriber` as their type. No more need to repeat the entire struct definition!

```
package main

import "fmt"

type subscriber struct {
        name   string
        rate   float64
        active bool
}

func main() {
        var subscriber1 subscriber
        subscriber1.name = "Aman Singh"
        fmt.Println("Name:", subscriber1.name)
        var subscriber2 subscriber
        subscriber2.name = "Beth Ryan"
        fmt.Println("Name:", subscriber2.name)
}
```

Define a type named "subscriber".

Use the struct type that was on the variables as the underlying type for the type definition.

Declare a variable of type "subscriber".

Use the "subscriber" type for the second variable, too.

```
Name: Aman Singh
Name: Beth Ryan
```

Using defined types with functions

Defined types can be used for more than just variable types. They also work for function parameters and return values.

Here's our `part` type again, together with a new `showInfo` function that prints a part's fields. The function takes a single parameter, with `part` as its type. Within `showInfo`, we access the fields via the parameter variable just like any other struct variable's.

```
package main

import "fmt"

type part struct {
        description string
        count       int
}

func showInfo(p part) {
        fmt.Println("Description:", p.description)
        fmt.Println("Count:", p.count)
}

func main() {
        var bolts part
        bolts.description = "Hex bolts"
        bolts.count = 24
        showInfo(bolts)
}
```

Declare one parameter, with "part" as its type.

Access the parameter's fields.

Create a "part" value.

Pass the "part" to the function.

```
Description: Hex bolts
Count: 24
```

And here's a `minimumOrder` function that creates a `part` with a specified description and a predefined value for the `count` field. We declare `minimumOrder`'s return type to be `part` so it can return the new struct.

```
// Package, imports, type definition omitted

func minimumOrder(description string) part {
        var p part
        p.description = description
        p.count = 100
        return p
}

func main() {
        p := minimumOrder("Hex bolts")
        fmt.Println(p.description, p.count)
}
```

Declare one return value, with a type of "part".

Create a new "part" value.

Return the "part".

Call minimumOrder. Use a short variable declaration to store the returned "part".

```
Hex bolts 100
```

Using defined types with functions (continued)

Let's go over a couple functions that work with the magazine's `subscriber` type...

The `printInfo` function takes a `subscriber` as a parameter, and prints the values of its fields.

We also have a `defaultSubscriber` function that sets up a new `subscriber` struct with some default values. It takes a string parameter called `name`, and uses that to set a new `subscriber` value's `name` field. Then it sets the `rate` and `active` fields to default values. Finally, it returns the completed `subscriber` struct to its caller.

```
package main

import "fmt"

type subscriber struct {
        name   string
        rate   float64
        active bool
}
func printInfo(s subscriber) {
        fmt.Println("Name:", s.name)
        fmt.Println("Monthly rate:", s.rate)
        fmt.Println("Active?", s.active)
}

func defaultSubscriber(name string) subscriber {
        var s subscriber
        s.name = name
        s.rate = 5.99
        s.active = true
        return s
}

func main() {
        subscriber1 := defaultSubscriber("Aman Singh")
        subscriber1.rate = 4.99
        printInfo(subscriber1)

        subscriber2 := defaultSubscriber("Beth Ryan")
        printInfo(subscriber2)
}
```

Declare one parameter...

...with a type of "subscriber".

Return a "subscriber" value.

Create a new "subscriber".

Set the struct's fields.

Return the "subscriber".

Set up a subscriber with this name.

Use a custom rate.

Print the field values.

Set up a subscriber with this name.

Print the field values.

```
Name: Aman Singh
Monthly rate: 4.99
Active? true
Name: Beth Ryan
Monthly rate: 5.99
Active? true
```

In our `main` function, we can pass a subscriber name to `defaultSubscriber` to get a new `subscriber` struct. One subscriber gets a discounted `rate`, so we reset that struct field directly. We can pass filled-out `subscriber` structs to `printInfo` to print out their contents.



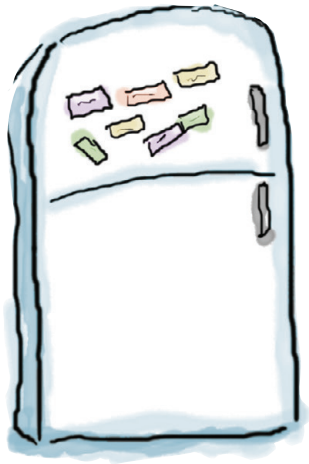
Watch it!

Don't use an existing type name as a variable name!

If you've defined a type named `car` *in the current package, and you declare a variable that's also named* `car`*, the variable name will shadow the type name, making it inaccessible.*

```
Refers to the type
var car car
var car2 car
Refers to the variable,
resulting in an error!
```

This isn't a common problem in practice, because defined types are often exported from their packages (and their names are therefore capitalized), and variables often are not (and their names are therefore lowercase). `Car` *(an exported type name) can't conflict with* `car` *(an unexported variable name). We'll see more about exporting defined types later in the chapter. Still, shadowing is a confusing problem when it occurs, so it's good to be aware that it can happen.*



Code Magnets

A Go program is scrambled up on the fridge. Can you reconstruct the code snippets to make a working program that will produce the given output? The finished program will have a defined struct type named `student`, and a `printInfo` function that accepts a `student` value as a parameter.

```
package main
{
}
}
import "fmt"
) {
}
type
struct
s student
student
var s
fmt.Println("Name:", s.name)
printInfo(s)
fmt.Printf("Grade: %0.1f\n", s.grade)
name string
func printInfo(
func main() {
grade float64
s.name = "Alonzo Cole"
s.grade = 92.3
student
```

Output →

```
Name: Alonzo Cole
Grade: 92.3
```

→ Answers on page 241.

Modifying a struct using a function

We'll be offering this $4.99 discounted rate to a lot of subscribers, so I tried to create an `applyDiscount` function to set the `rate` field for us. But it's not working!



```
func applyDiscount(s subscriber) {
        s.rate = 4.99
}

func main() {
        var s subscriber
        applyDiscount(s)
        fmt.Println(s.rate)
}
```

Accept a "subscriber" parameter.

Set the "rate" field.

Attempt to set a "subscriber" struct's "rate" field to 4.99.

0

But it's still set to 0!

Our friends at *Gopher Fancy* are trying to write a function that takes a struct as a parameter and updates one of the fields in that struct.

Remember way back in Chapter 3, when we were trying to write a `double` function that took a number and doubled it? After `double` returned, the number was back to its original value!

That's when we learned that Go is a "pass-by-value" language, meaning that function parameters receive a *copy* of the arguments the function was called with. If a function changes a parameter value, it's changing the *copy*, not the *original*.

```
func main() {
        amount := 6
        double(amount)
        fmt.Println(amount)
}

func double(number int) {
        number *= 2
}
```

Pass an argument to the function.

Prints the original value!

Parameter is set to a copy of the argument.

Alters the copied value, not the original!

6

Prints the unchanged amount!

The same thing is true for structs. When we pass a `subscriber` struct to `applyDiscount`, the function receives a *copy* of the struct. So when we set the `rate` field on the struct, we're modifying the *copied* struct, not the original.

```
func applyDiscount(s subscriber) {
        s.rate = 4.99
}
```

Receives a copy of the struct!

Modifies the copy, not the original!

Modifying a struct using a function (continued)

Back in Chapter 3, our solution was to update the function parameter to accept a *pointer* to a value, instead of accepting a value directly. When calling the function, we used the address-of operator (&) to pass a pointer to the value we wanted to update. Then, within the function, we used the * operator to update the value at that pointer.

As a result, the updated value was still visible after the function returned.

```
func main() {
        amount := 6
        double(&amount)
        fmt.Println(amount)
}

func double(number *int) {
        *number *= 2
}
```

Pass a pointer instead of the variable value.

Accept a pointer instead of an int value.

Update the value at the pointer.

```
12
```

Prints the doubled amount

We can use pointers to allow a function to update a struct as well.

Here's an updated version of the `applyDiscount` function that should work correctly. We update the `s` parameter to accept a pointer to a `subscriber` struct, rather than the struct itself. Then we update the value in the struct's `rate` field.

In `main`, we call `applyDiscount` with a pointer to a `subscriber` struct. When we print the value in the struct's `rate` field, we can see that it's been updated successfully!

```
package main

import "fmt"

type subscriber struct {
        name   string
        rate   float64
        active bool
}

func applyDiscount(s *subscriber) {
        s.rate = 4.99
}

func main() {
        var s subscriber
        applyDiscount(&s)
        fmt.Println(s.rate)
}
```

Take a pointer to a struct, not the struct itself.

Update the struct field.

Pass a pointer, not a struct.

```
4.99
```



Wait, how does that work? In the `double` function, we had to use the * operator to get the value at the pointer. Don't you need * when you set the `rate` field in `applyDiscount`?

Actually, no! The dot notation to access fields works on struct pointers as well as the structs themselves.

Accessing struct fields through a pointer

If you try to print a pointer variable, what you'll see is the memory address it points to. This is generally not what you want.

```
func main() {
        var value int = 2
        var pointer *int = &value
        fmt.Println(pointer)
}
```

Create a value.

Get a pointer to the value.

Oops! This prints the pointer, not the value!

```
0xc420014100
```

Instead, you need to use the * operator (what we like to call the "value-at operator") to get the value at the pointer.

```
func main() {
        var value int = 2
        var pointer *int = &value
        fmt.Println(*pointer)
}
```

Print the value at the pointer.

```
2
```

So you might think you'd need to use the * operator with pointers to structs as well. But just putting a * before the struct pointer won't work:

```
type myStruct struct {
        myField int
}

func main() {
        var value myStruct
        value.myField = 3
        var pointer *myStruct = &value
        fmt.Println(*pointer.myField)
}
```

Create a struct value.

Get a pointer to the struct value.

Attempt to get the struct value at the pointer.

Error!

```
invalid indirect of
pointer.myField (type int)
```

If you write `*pointer.myField`, Go thinks that `myField` must contain a pointer. It doesn't, though, so an error results. To get this to work, you need to wrap `*pointer` in parentheses. That will cause the `myStruct` value to be retrieved, after which you can access the struct field.

```
func main() {
        var value myStruct
        value.myField = 3
        var pointer *myStruct = &value
        fmt.Println((*pointer).myField)
}
```

Get the struct value at the pointer, then access the struct field.

```
3
```

Accessing struct fields through a pointer (continued)

Having to write `(*pointer).myField` all the time would get tedious quickly, though. For this reason, the dot operator lets you access fields via *pointers* to structs, just as you can access fields directly from struct values. You can leave off the parentheses and the `*` operator.

```
func main() {
        var value myStruct
        value.myField = 3
        var pointer *myStruct = &value
        fmt.Println(pointer.myField)
}
```

Access the struct field through the pointer.

```
3
```

This works for assigning to struct fields through a pointer as well:

```
func main() {
        var value myStruct
        var pointer *myStruct = &value
        pointer.myField = 9
        fmt.Println(pointer.myField)
}
```

Assign to a struct field through the pointer.

```
9
```

And that's how the `applyDiscount` function is able to update the struct field without using the `*` operator. It assigns to the `rate` field *through* the struct pointer.

```
func applyDiscount(s *subscriber) {
        s.rate = 4.99
}

func main() {
        var s subscriber
        applyDiscount(&s)
        fmt.Println(s.rate)
}
```

Assign to the struct field through the pointer.

```
4.99
```

there are no Dumb Questions

Q: You showed a `defaultSubscriber` function before that set a struct's fields, but it didn't need to use any pointers! Why not?

A: The `defaultSubscriber` function *returned* a struct value. If a caller stores the returned value, then the values in its fields will be preserved. Only functions that *modify existing* structs without returning them have to use pointers for those changes to be preserved.

But `defaultSubscriber` *could* have returned a pointer to a struct, if we had wanted it to. In fact, we make just that change in the next section!

Pass large structs using pointers

So function parameters receive a copy of the arguments from the function call, even for structs... If you pass a big struct with a lot of fields, won't that take up a lot of the computer's memory?



Yes, it will. It has to make room for the original struct and the copy.

Functions receive a copy of the arguments they're called with, even if they're a big value like a struct.

That's why, unless your struct has only a couple small fields, it's often a good idea to pass functions a *pointer* to a struct, rather than the struct itself. (This is true even if the function doesn't need to modify the struct.) When you pass a struct pointer, only one copy of the original struct exists in memory. The function just receives the memory address of that single struct, and can read the struct, modify it, or whatever else it needs to do, all without making an extra copy.

Here's our `defaultSubscriber` function, updated to return a pointer, and our `printInfo` function, updated to receive a pointer. Neither of these functions needs to change an existing struct like `applyDiscount` does. But using pointers ensures that only one copy of each struct needs to be kept in memory, while still allowing the program to work as normal.

```
// Code above here omitted
type subscriber struct {
        name   string
        rate   float64
        active bool
}

func printInfo(s *subscriber) {
        fmt.Println("Name:", s.name)
        fmt.Println("Monthly rate:", s.rate)
        fmt.Println("Active?", s.active)
}

func defaultSubscriber(name string) *subscriber {
        var s subscriber
        s.name = name
        s.rate = 5.99
        s.active = true
        return &s
}

func applyDiscount(s *subscriber) {
        s.rate = 4.99
}

func main() {
        subscriber1 := defaultSubscriber("Aman Singh")
        applyDiscount(subscriber1)
        printInfo(subscriber1)

        subscriber2 := defaultSubscriber("Beth Ryan")
        printInfo(subscriber2)
}
```

Update to take a pointer.

Update to return a pointer.

Return a pointer to a struct instead of the struct itself.

This is no longer a struct, it's a struct pointer...

Since this is already a pointer, remove the address-of operator.

```
Name: Aman Singh
Monthly rate: 4.99
Active? true
Name: Beth Ryan
Monthly rate: 5.99
Active? true
```



The two programs below aren't working quite right. The `nitroBoost` function in the lefthand program is supposed to add `50` kilometers/hour to a `car`'s top speed, but it's not. And the `doublePack` function in the righthand program is supposed to double a `part` value's `count` field, but it's not, either.

See if you can fix the programs. Only minimal changes will be necessary; we've left a little extra space in the code so you can make the necessary updates.

```
package main

import "fmt"

type car struct {
       name      string
       topSpeed  float64
}

func nitroBoost( c  car ) {
       c.topSpeed += 50
}

func main() {
       var mustang car
       mustang.name = "Mustang Cobra"
       mustang.topSpeed = 225
       nitroBoost( mustang )
       fmt.Println( mustang.name )
       fmt.Println( mustang.topSpeed )
}
```

This is supposed to be 50 km/h higher!

```
Mustang Cobra
225
```

```
package main

import "fmt"

type part struct {
       description  string
       count        int
}

func doublePack( p  part ) {
       p.count *= 2
}

func main() {
       var fuses part
       fuses.description = "Fuses"
       fuses.count = 5
       doublePack( fuses )
       fmt.Println( fuses.description )
       fmt.Println( fuses.count )
}
```

This is supposed to be doubled!

```
Fuses
5
```

Answers on page 263.

Moving our struct type to a different package



We're definitely starting to appreciate the convenience of this `subscriber` struct type. But the code in our `main` package is getting a little long. Can we move `subscriber` out to another package?

That should be easy to do. Find the *headfirstgo* directory within your Go workspace, and create a new directory in there to hold a package named `magazine`. Within *magazine*, create a file named *magazine.go*.



Be sure to add a `package magazine` declaration at the top of *magazine.go*. Then, copy the `subscriber` struct definition from your existing code and paste it into *magazine.go*.

We'll try pasting the type definition in here without any changes.

```
package magazine

type subscriber struct {
        name   string
        rate   float64
        active bool
}
```

Next, let's create a program to try out the new package. Since we're just experimenting for now, let's not create a separate package folder for this code; we'll just run it using the `go run` command. Create a file named *main.go*. You can save it in any directory you want, but make sure you save it *outside* your Go workspace, so it doesn't interfere with any other packages.



(You can move this code into your Go workspace later, if you want, as long as you create a separate package directory for it.)

Within *main.go*, save this code, which simply creates a new `subscriber` struct and accesses one of its fields.

There are two differences from the previous examples. First, we need to import the `magazine` package at the top of the file. Second, we need to use `magazine.subscriber` as the type name, since it belongs to another package now.

```
package main

import (
        "fmt"
        "github.com/headfirstgo/magazine"
)

func main() {
        var s magazine.subscriber
        s.rate = 4.99
        fmt.Println(s.rate)
}
```

Import packages we need...

...including our new "magazine" package.

Type name needs to be prefixed with the package name now.

A defined type's name must be capitalized to be exported

Let's see if our experimental code can still access the `subscriber` struct type in its new package. In your terminal, change into the directory where you saved *main.go*, then enter **`go run main.go`**.

```
Shell Edit View Window Help
$ cd temp
$ go run main.go
./main.go:9:18: cannot refer to unexported name magazine.subscriber
./main.go:9:18: undefined: magazine.subscriber
```

We get a couple errors, but here's the important one: `cannot refer to unexported name magazine.subscriber`.

Go type names follow the same rule as variable and function names: if the name of a variable, function, or type begins with a capital letter, it is considered *exported* and can be accessed from outside the package it's declared in. But our `subscriber` type name begins with a lowercase letter. That means it can only be used within the `magazine` package.

Well, that seems like an easy fix. We'll just open our *magazine.go* file and capitalize the name of the defined type. Then, we'll open *main.go* and capitalize the names of any references to that type. (There's just one right now.)

For a type to be accessed outside the package it's defined in, it must be exported: its name must begin with a capital letter.

magazine.go

```
package magazine

type Subscriber struct {
        name   string
        rate   float64
        active bool
}
```

Capitalize the type name.

main.go

```
package main

import (
        "fmt"
        "github.com/headfirstgo/magazine"
)

func main() {
        var s magazine.Subscriber
        s.rate = 4.99
        fmt.Println(s.rate)
}
```

Capitalize the type name.

If we try running the updated code with `go run main.go`, we no longer get the error saying that the `magazine.subscriber` type is unexported. So that seems to be fixed. But we get a couple new errors in its place...

```
Shell Edit View Window Help
$ go run main.go
./main.go:10:13: s.rate undefined
(cannot refer to unexported field or method rate)
./main.go:11:25: s.rate undefined
(cannot refer to unexported field or method rate)
```

Struct field names must be capitalized to be exported

With the `Subscriber` type name capitalized, we seem to be able to access it from the `main` package. But now we're getting an error saying that we can't refer to the `rate` *field*, because *that* is unexported.

```
Shell Edit View Window Help
$ go run main.go
./main.go:10:13: s.rate undefined
(cannot refer to unexported field or method rate)
./main.go:11:25: s.rate undefined
(cannot refer to unexported field or method rate)
```

Even if a struct type is exported from a package, its fields will be *unexported* if their names don't begin with a capital letter. Let's try capitalizing `Rate` (in both *magazine.go* and *main.go*)...

Struct field names must also be capitalized if you want to export them from their package.

magazine.go

```
package magazine

type Subscriber struct {
        name   string
Capitalize. → Rate   float64
        active bool
}
```

main.go

```
package main

import (
        "fmt"
        "github.com/headfirstgo/magazine"
)

func main() {
        var s magazine.Subscriber
        s.Rate = 4.99 ← Capitalize.
        fmt.Println(s.Rate) ← Capitalize.
}
```

Run *main.go* again, and you'll see that everything works this time. Now that they're exported, we can access the `Subscriber` type *and* its `Rate` field from the `main` package.

```
Shell Edit View Window Help
$ go run main.go
4.99
```

Notice that the code worked even though the `name` and `active` fields were still unexported. You can have a mixture of exported and unexported fields within a single struct type, if you want.

That's probably not advisable in the case of the `Subscriber` type, though. It wouldn't make sense to be able to access the subscription rate from other packages, but not the name or active status. So let's go back into *magazine.go* and export the other fields as well. Simply capitalize their names: `Name` and `Active`.

magazine.go

```
package magazine

type Subscriber struct {
Capitalize. → Name   string
        Rate   float64
Capitalize. → Active bool
}
```

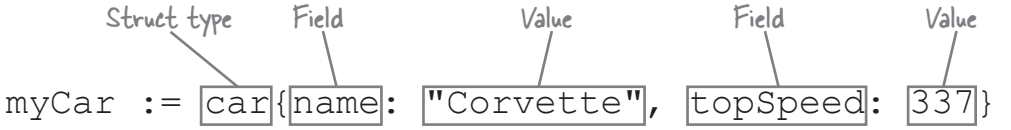
Struct literals

The code to define a struct and then assign values to its fields one by one can get a bit tedious:

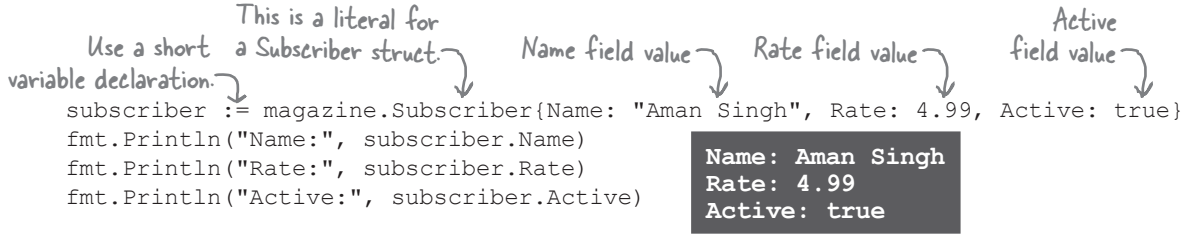
```
var subscriber magazine.Subscriber
subscriber.Name = "Aman Singh"
subscriber.Rate = 4.99
subscriber.Active = true
```

So, just as with slices and maps, Go offers **struct literals** to let you create a struct and set its fields at the same time.

The syntax looks similar to a map literal. The type is listed first, followed by curly braces. Within the braces, you can specify values for some or all of the struct fields, using the field name, a colon, and then the value. If you specify multiple fields, separate them with commas.

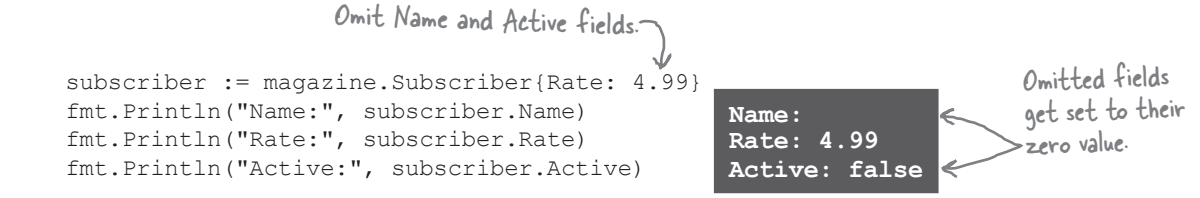


Above, we showed some code that creates a `Subscriber` struct and sets its fields, one by one. This code does the same thing in a single line, using a struct literal:



You may have noticed that for most of the chapter, we've had to use long-form declarations for struct variables (unless the struct was being returned from a function). Struct literals allow us to use short variable declarations for a struct we've just created.

You can omit some or even all of the fields from the curly braces. Omitted fields will be set to the zero value for their type.





Pool Puzzle

Your **job** is to take code snippets from the pool and place them into the blank lines in this code. **Don't** use the same snippet more than once, and you won't need to use all the snippets. Your **goal** is to make a program that will run and produce the output shown.

```
package main

import (
        "fmt"
        "geo"
)

func main() {
        location := geo._____{_____: 37.42, _____: -122.08}
        fmt.Println("Latitude:", location.Latitude)
        fmt.Println("Longitude:", location.Longitude)
}
```

main.go

```
package geo

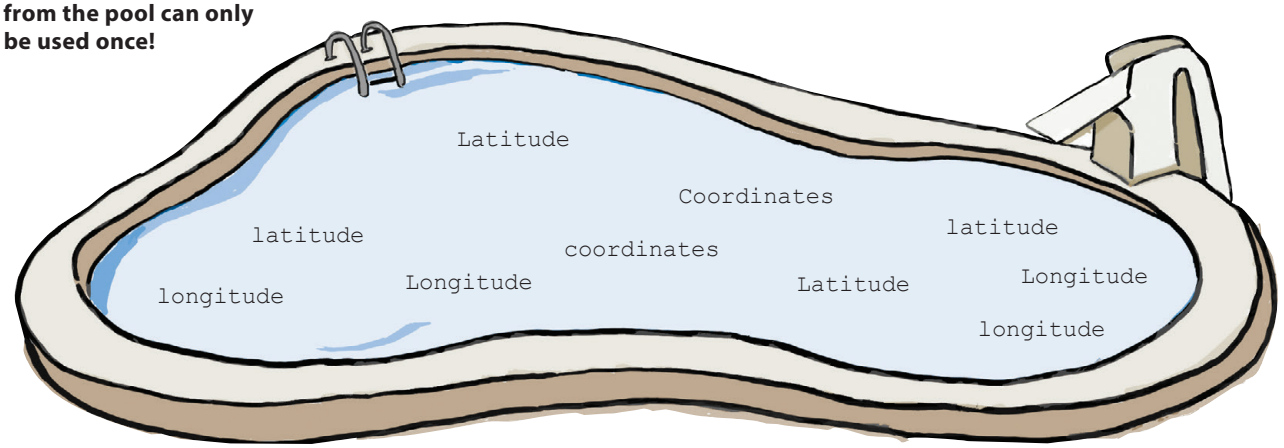
type Coordinates struct {
        _____ float64
        _____ float64
}
```

geo.go

Output

```
Latitude: 37.42
Longitude: -122.08
```

Note: each snippet from the pool can only be used once!



Answers on page 264.

Creating an Employee struct type



This new `magazine` package is working out great! Just a couple more things before we can publish our first issue... We need an `Employee` struct type to track the names and salaries of our employees. And we need to store mailing addresses for both employees and subscribers.

Adding an `Employee` struct type should be pretty easy. We'll just add it to the `magazine` package, alongside the `Subscriber` type. In *magazine.go*, define a new `Employee` type, with a `struct` underlying type. Give the struct type a `Name` field with a type of `string`, and a `Salary` field with a type of `float64`. Be sure to capitalize the type name *and* all the fields, so that they're exported from the `magazine` package.

We can update the `main` function in *main.go* to try the new type out. First, declare a variable with the type `magazine.Employee`. Then assign values of the appropriate type to each of the fields. Finally, print those values out.

magazine.go

```
package magazine

type Subscriber struct {
        Name   string
        Rate   float64
        Active bool
}

type Employee struct {
        Name   string
        Salary float64
}
```

Capitalize the name, so it's exported.

Export field names, too.

main.go

```
package main

import (
        "fmt"
        "github.com/headfirstgo/magazine"
)

func main() {
        var employee magazine.Employee
        employee.Name = "Joy Carr"
        employee.Salary = 60000
        fmt.Println(employee.Name)
        fmt.Println(employee.Salary)
}
```

Try creating an Employee value.

If you execute `go run main.go` from your terminal, it should run, create a new `magazine.Employee` struct, set its field values, and then print those values out.

```
Joy Carr
60000
```

Creating an Address struct type

Next, we need to track mailing addresses for both the `Subscriber` and `Employee` types. We're going to need fields for the street address, city, state, and postal code (zip code).

We *could* add separate fields to both the `Subscriber` and `Employee` types, like this:

```
type Subscriber struct {
        Name   string
        Rate   float64
        Active bool
        Street     string
        City       string
        State      string
        PostalCode string
}
```

If we added fields here...

```
type Employee struct {
        Name   string
        Salary float64
        Street     string
        City       string
        State      string
        PostalCode string
}
```

...we'd have to repeat them here...

But mailing addresses are going to have the same format, no matter what type they belong to. It's a pain to have to repeat all those fields between multiple types.

Struct fields can hold values of any type, *including other structs*. So, instead, let's try building an `Address` struct type, and then adding an `Address` field on the `Subscriber` and `Employee` types. That will save us some effort now, and ensure consistency between the types later if we have to change the address format.

We'll create just the `Address` type first, so we can ensure it's working correctly. Place it in the `magazine` package, alongside the `Subscriber` and `Employee` types. Then, replace the code in *main.go* with a few lines to create an `Address` and ensure its fields are accessible.

magazine.go

```
package magazine

// Subscriber and Employee
// code omitted...

type Address struct {
        Street     string
        City       string
        State      string
        PostalCode string
}
```

Add the new type here.

main.go

```
package main

import (
        "fmt"
        "github.com/headfirstgo/magazine"
)

func main() {
        var address magazine.Address
        address.Street = "123 Oak St"
        address.City = "Omaha"
        address.State = "NE"
        address.PostalCode = "68111"
        fmt.Println(address)
}
```

Try creating an Address value.

Type **`go run main.go`** in your terminal, and it should create an `Address` struct, populate its fields, and then print the whole struct out.

```
{123 Oak St Omaha NE 68111}
```

Adding a struct as a field on another type

Now that we're sure the `Address` struct type works by itself, let's add `HomeAddress` fields to the `Subscriber` and `Employee` types.

Adding a struct field that is itself a struct type is no different than adding a field of any other type. You provide a name for the field, followed by the field's type (which in this case will be a struct type).

Add a field named `HomeAddress` to the `Subscriber` struct. Make sure to capitalize the field name, so that it's accessible from outside the `magazine` package. Then specify the field type, which is `Address`.

Add a `HomeAddress` field to the `Employee` type as well.

magazine.go

```
package magazine

type Subscriber struct {
        Name        string
        Rate        float64
        Active      bool
        HomeAddress Address
}

type Employee struct {
        Name        string
        Salary      float64
        HomeAddress Address
}

type Address struct {
        // Fields omitted
}
```

Capitalized field name

The field type

Capitalized field name

The field type

Setting up a struct within another struct

Now let's see if we can populate the fields of the `Address` struct *within* the `Subscriber` struct. There are a couple ways to go about this.

The first approach is to create an entirely separate `Address` struct and then use it to set the entire `Address` field of the `Subscriber` struct. Here's an update to *main.go* that follows this approach.

main.go

```
package main

import (
        "fmt"
        "github.com/headfirstgo/magazine"
)

func main() {
        address := magazine.Address{Street: "123 Oak St",
                City: "Omaha", State: "NE", PostalCode: "68111"}
        subscriber := magazine.Subscriber{Name: "Aman Singh"}
        subscriber.HomeAddress = address
        fmt.Println(subscriber.HomeAddress)
}
```

Create an Address value and populate its fields.

Create the Subscriber struct that the Address will belong to.

Set the HomeAddress field.

Print the HomeAddress field.

Type **`go run main.go`** in your terminal, and you'll see the subscriber's `HomeAddress` field has been set to the struct you built.

```
{123 Oak St Omaha NE 68111}
```

Setting up a struct within another struct (continued)

Another approach is to set the fields of the inner struct *through* the outer struct.

When a `Subscriber` struct is created, its `HomeAddress` field is already set: it's an `Address` struct with all its fields set to their zero values. If we print `HomeAddress` using the `"%#v"` verb for `fmt.Printf`, it will print the struct as it would appear in Go code — that is, as a struct literal. We'll see that each of the `Address` fields is set to an empty string, which is the zero value for the `string` type.

```
subscriber := magazine.Subscriber{}
fmt.Printf("%#v\n", subscriber.HomeAddress)
```

Each of the Address struct's fields is set to an empty string (which is the zero value for strings).

The field is already set as a new Address struct.

```
magazine.Address{Street:"", City:"", State:"", PostalCode:""}
```

If `subscriber` is a variable that contains a `Subscriber` struct, then when you type `subscriber.HomeAddress`, you'll get an `Address` struct, even if you haven't explicitly set `HomeAddress`.

You can use this fact to "chain" dot operators together so you can access the fields of the `Address` struct. Simply type **`subscriber.HomeAddress`** to access the `Address` struct, followed by *another* dot operator and the name of the field you want to access on that `Address` struct.

```
subscriber.HomeAddress.City
```

This part gives you an Address struct.

This part accesses the City field on that Address struct.

This works both for assigning values to the inner struct's fields...

```
subscriber.HomeAddress.PostalCode = "68111"
```

...and for retrieving those values again later.

```
fmt.Println("Postal Code:", subscriber.HomeAddress.PostalCode)
```

Setting up a struct within another struct (continued)

Here's an update to *main.go* that uses dot operator chaining. First we store a `Subscriber` struct in the `subscriber` variable. That will automatically create an `Address` struct in `subscriber`'s `HomeAddress` field. We set values for `subscriber.HomeAddress.Street`, `subscriber.HomeAddress.City`, and so on, and then print those values out again.

Then we store an `Employee` struct in the `employee` variable, and do the same for its `HomeAddress` struct.

main.go

```
package main

import (
        "fmt"
        "github.com/headfirstgo/magazine"
)

func main() {
        subscriber := magazine.Subscriber{Name: "Aman Singh"}
        subscriber.HomeAddress.Street = "123 Oak St"
        subscriber.HomeAddress.City = "Omaha"
        subscriber.HomeAddress.State = "NE"
        subscriber.HomeAddress.PostalCode = "68111"
        fmt.Println("Subscriber Name:", subscriber.Name)
        fmt.Println("Street:", subscriber.HomeAddress.Street)
        fmt.Println("City:", subscriber.HomeAddress.City)
        fmt.Println("State:", subscriber.HomeAddress.State)
        fmt.Println("Postal Code:", subscriber.HomeAddress.PostalCode)

        employee := magazine.Employee{Name: "Joy Carr"}
        employee.HomeAddress.Street = "456 Elm St"
        employee.HomeAddress.City = "Portland"
        employee.HomeAddress.State = "OR"
        employee.HomeAddress.PostalCode = "97222"
        fmt.Println("Employee Name:", employee.Name)
        fmt.Println("Street:", employee.HomeAddress.Street)
        fmt.Println("City:", employee.HomeAddress.City)
        fmt.Println("State:", employee.HomeAddress.State)
        fmt.Println("Postal Code:", employee.HomeAddress.PostalCode)
}
```

Set the fields of subscriber.HomeAddress.

Retrieve field values from subscriber.HomeAddress.

Set the fields of employee.HomeAddress.

Retrieve field values from employee.HomeAddress.

Type **`go run main.go`** in your terminal, and the program will print out the completed fields of both `subscriber.HomeAddress` and `employee.HomeAddress`.

```
Subscriber Name: Aman Singh
Street: 123 Oak St
City: Omaha
State: NE
Postal Code: 68111
Employee Name: Joy Carr
Street: 456 Elm St
City: Portland
State: OR
Postal Code: 97222
```

Anonymous struct fields

The code to access the fields of an inner struct through its outer struct can be a bit tedious, though. You have to write the field name of the inner struct (`HomeAddress`) each time you want to access any of the fields it contains.

```
subscriber := magazine.Subscriber{Name: "Aman Singh"}
subscriber.HomeAddress.Street = "123 Oak St"
subscriber.HomeAddress.City = "Omaha"
subscriber.HomeAddress.State = "NE"
subscriber.HomeAddress.PostalCode = "68111"
```

You have to write the field name of the inner struct...

...and only then can you access its fields.

Go allows you to define **anonymous fields**: struct fields that have no name of their own, just a type. We can use an anonymous field to make our inner struct easier to access.

Here's an update to the `Subscriber` and `Employee` types to convert their `HomeAddress` fields to an anonymous field. To do this, we simply remove the field name, leaving only the type.

magazine.go

```
package magazine

type Subscriber struct {
        Name   string
        Rate   float64
        Active bool
        Address
}

type Employee struct {
        Name   string
        Salary float64
        Address
}

type Address struct {
        // Fields omitted
}
```

Delete the field name ("HomeAddress"), leaving only the type.

Delete the field name ("HomeAddress"), leaving only the type.

When you declare an anonymous field, you can use the field's type name as if it were the name of the field. So `subscriber.Address` and `employee.Address` in the code below still access the `Address` structs:

```
subscriber := magazine.Subscriber{Name: "Aman Singh"}
subscriber.Address.Street = "123 Oak St"
subscriber.Address.City = "Omaha"
fmt.Println("Street:", subscriber.Address.Street)
fmt.Println("City:", subscriber.Address.City)
employee := magazine.Employee{Name: "Joy Carr"}
employee.Address.State = "OR"
employee.Address.PostalCode = "97222"
fmt.Println("State:", employee.Address.State)
fmt.Println("Postal Code:", employee.Address.PostalCode)
```

Access the inner struct field through its new "name", which is "Address".

```
Street: 123 Oak St
City: Omaha
State: OR
Postal Code: 97222
```

Embedding structs

But anonymous fields offer much more than just the ability to skip providing a name for a field in a struct definition.

An inner struct that is stored within an outer struct using an anonymous field is said to be **embedded** within the outer struct. Fields for an embedded struct are **promoted** to the outer struct, meaning you can access them as if they belong to the outer struct.

So now that the `Address` struct type is embedded within the `Subscriber` and `Employee` struct types, you don't have to write out `subscriber.Address.City` to get at the `City` field; you can just write `subscriber.City`. You don't need to write `employee.Address.State`; you can just write `employee.State`.

Here's one last version of *main.go*, updated to treat `Address` as an embedded type. You can write the code as if there were no `Address` type at all; it's like the `Address` fields belong to the struct type they're embedded within.



main.go

```
package main

import (
        "fmt"
        "github.com/headfirstgo/magazine"
)

func main() {
        subscriber := magazine.Subscriber{Name: "Aman Singh"}
        subscriber.Street = "123 Oak St"
        subscriber.City = "Omaha"
        subscriber.State = "NE"
        subscriber.PostalCode = "68111"
        fmt.Println("Street:", subscriber.Street)
        fmt.Println("City:", subscriber.City)
        fmt.Println("State:", subscriber.State)
        fmt.Println("Postal Code:", subscriber.PostalCode)

        employee := magazine.Employee{Name: "Joy Carr"}
        employee.Street = "456 Elm St"
        employee.City = "Portland"
        employee.State = "OR"
        employee.PostalCode = "97222"
        fmt.Println("Street:", employee.Street)
        fmt.Println("City:", employee.City)
        fmt.Println("State:", employee.State)
        fmt.Println("Postal Code:", employee.PostalCode)
}
```

Set the fields of the Address as if they were defined on Subscriber.

Retrieve Address fields through the Subscriber.

Set the fields of the Address as if they were defined on Employee.

Retrieve Address fields through the Employee.

```
Street: 123 Oak St
City: Omaha
State: NE
Postal Code: 68111
Street: 456 Elm St
City: Portland
State: OR
Postal Code: 97222
```

Keep in mind that you don't *have* to embed inner structs. You don't have to use inner structs at all. Sometimes adding new fields on the outer struct leads to the clearest code. Consider your current situation, and go with the solution that works best for you and your users.

Our defined types are complete!



These struct types you've made for us are great! No more passing around bunches of variables just to represent one subscriber. Everything we need is stored in one convenient bundle. Thanks to you, we're ready to fire up the presses and mail out our first issue!

Nice work! You've defined `Subscriber` and `Employee` struct types, and embedded an `Address` struct in each of them. You've found a way to represent all the data the magazine needed!

You're still missing an important aspect to defined types, though. In previous chapters, you've used types like `time.Time` and `strings.Replacer` that have *methods*: functions that you can call *on* their values. But you haven't learned how to define methods for your own types yet. Don't worry; we'll learn all about it in the next chapter!



Exercise

Here's a source file from the `geo` package, which we saw in a previous exercise. Your goal is to make the code in *main.go* work correctly. But here's the catch: you need to do it by adding just *two* fields to the `Landmark` struct type within *geo.go*.

geo.go

```
package geo

type Coordinates struct {
        Latitude  float64
        Longitude float64
}

type Landmark struct {

}
```

Add two fields here!

main.go

```
package main

import (
        "fmt"
        "geo"
)

func main() {
        location := geo.Landmark{}
        location.Name = "The Googleplex"
        location.Latitude = 37.42
        location.Longitude = -122.08
        fmt.Println(location)
}
```

Output →

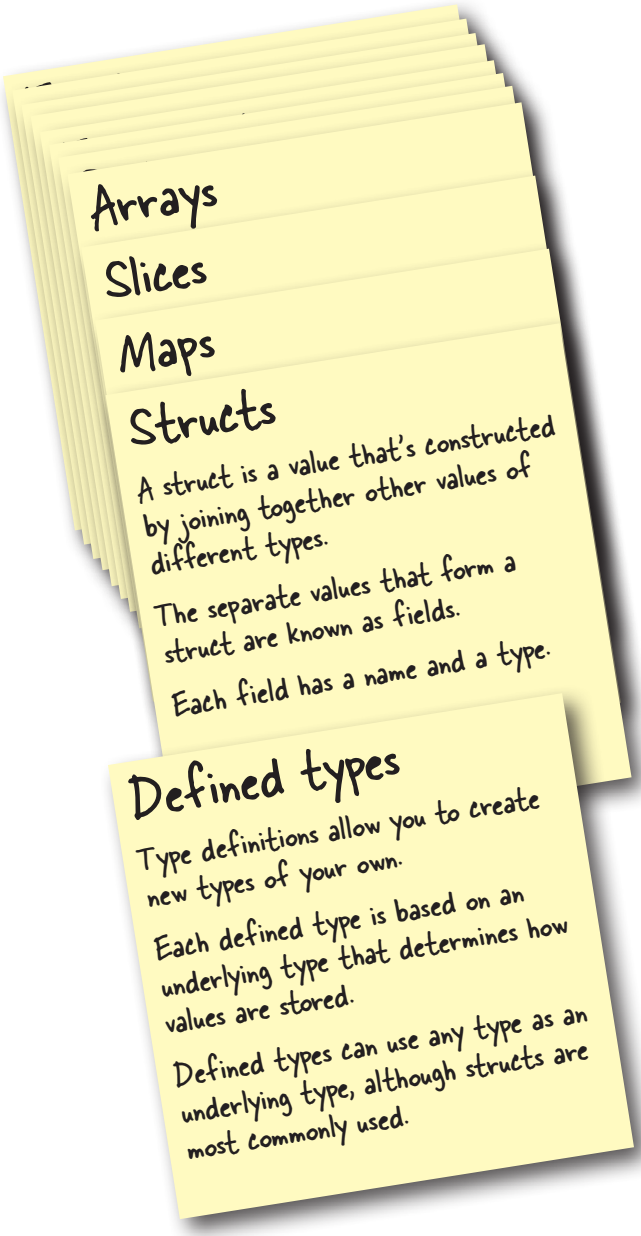
```
{The Googleplex {37.42 -122.08}}
```

Answers on page 264.



Your Go Toolbox

That's it for Chapter 8! You've added structs and defined types to your toolbox.



BULLET POINTS

- You can declare a variable with a struct type. To specify a struct type, use the `struct` keyword, followed by a list of field names and types within curly braces.

```
var myStruct struct {
    field1 string
    field2 int
}
```

- Writing struct types repeatedly can get tedious, so it's usually best to define a type with an underlying struct type. Then the defined type can be used for variables, function parameters or return values, and so on.

```
type myType struct {
    field1 string
}
var myVar myType
```

- Struct fields are accessed via the dot operator.

```
myVar.field1 = "value"
fmt.Println(myVar.field1)
```

- If a function needs to modify a struct or if a struct is large, it should be passed to the function as a pointer.
- Types will only be exported from the package they're defined in if their name begins with a capital letter.
- Likewise, struct fields will not be accessible outside their package unless their name is capitalized.
- Struct literals let you create a struct and set its fields at the same time.

```
myVar := myType{field1: "value"}
```

- Adding a struct field with no name, only a type, defines an anonymous field.
- An inner struct that is added as part of an outer struct using an anonymous field is said to be **embedded** within the outer struct.
- You can access the fields of an embedded struct as if they belong to the outer struct.



Exercise Solution

At the right is a program that creates a struct variable to hold a pet's name (a `string`) and age (an `int`). Fill in the blanks so that the code will produce the output shown.

```
package main

import "fmt"

func main() {
        var pet  struct  {
                name   string
                age    int
        }
        pet. name   = "Max"
        pet.age = 5
        fmt.Println("Name:", pet .name)
        fmt.Println("Age:", pet. age )
}
```

```
Name: Max
Age: 5
```

Code Magnets Solution

```
package main

import "fmt"

type  student  struct  {
        name string
        grade float64
}

func printInfo(  s student  ) {
        fmt.Println("Name:", s.name)
        fmt.Printf("Grade: %0.1f\n", s.grade)
}

func main() {
        var s   student
        s.name = "Alonzo Cole"
        s.grade = 92.3
        printInfo(s)
}
```

Define a "student" struct type.

Define a function that takes a "student" struct as a parameter.

Pass a struct to the function.

Output

```
Name: Alonzo Cole
Grade: 92.3
```



The two programs below weren't working quite right. The `nitroBoost` function in the lefthand program was supposed to add `50` kilometers/hour to a `car`'s top speed, but it wasn't. And the `doublePack` function in the righthand program was supposed to double a `part` value's `count` field, but it wasn't, either.

Fixing both programs was simply a matter of updating the functions to accept pointers, and updating the function calls to pass pointers. The code within the functions that updates the struct fields doesn't need to be changed; the code to access a field through a pointer to a struct is the same as the code to access a field on the struct directly.

```
package main

import "fmt"

type car struct {
        name       string
        topSpeed   float64
}

func nitroBoost( c *car ) {
        c.topSpeed += 50
}

func main() {
        var mustang car
        mustang.name = "Mustang Cobra"
        mustang.topSpeed = 225
        nitroBoost(&mustang )
        fmt.Println( mustang.name )
        fmt.Println( mustang.topSpeed )
}
```

Accept a pointer to a struct instead of a struct.

No change needed; works with a pointer as well as the struct itself.

Pass a pointer.

Fixed; it's 50 km/h higher.

```
Mustang Cobra
275
```

```
package main

import "fmt"

type part struct {
        description  string
        count        int
}

func doublePack( p *part ) {
        p.count *= 2
}

func main() {
        var fuses part
        fuses.description = "Fuses"
        fuses.count = 5
        doublePack(&fuses )
        fmt.Println( fuses.description )
        fmt.Println( fuses.count )
}
```

Accept a pointer to a struct instead of a struct.

No change needed; works with a pointer as well as the struct itself.

Pass a pointer.

Fixed; it's double the original value.

```
Fuses
10
```

Pool Puzzle Solution

```
package main

import (
        "fmt"
        "geo"
)

func main() {
        location := geo.Coordinates{Latitude: 37.42, Longitude: -122.08}
        fmt.Println("Latitude:", location.Latitude)
        fmt.Println("Longitude:", location.Longitude)
}
```

Type name has to be capitalized because it needs to be exported.

Field names need to be capitalized, too.

```
package geo

type Coordinates struct {
        Latitude  float64
        Longitude float64
}
```

geo.go

main.go

Output

```
Latitude: 37.42
Longitude: -122.08
```



The *geo.go* source file is from the `geo` package, which we saw in a previous exercise. Your goal was to make the code in *main.go* work correctly, by adding just *two* fields to the `Landmark` struct type within *geo.go*.

```
package geo

type Coordinates struct {
        Latitude  float64
        Longitude float64
}

type Landmark struct {
        Name string
        Coordinates
}
```

geo.go

Embed Coordinates as an anonymous field, which allows you to access its Latitude and Longitude fields as if they were defined on Landmark.

```
package main

import (
        "fmt"
        "geo"
)

func main() {
        location := geo.Landmark{}
        location.Name = "The Googleplex"
        location.Latitude = 37.42
        location.Longitude = -122.08
        fmt.Println(location)
}
```

main.go

Output

```
{The Googleplex {37.42 -122.08}}
```

9 you're my type

Defined Types



There's more to learn about defined types. In the previous chapter, we showed you how to define a type with a struct underlying type. What we *didn't* show you was that you can use *any* type as an underlying type.

And do you remember methods—the special kind of function that's associated with values of a particular type? We've been calling methods on various values throughout the book, but we haven't shown you how to define your *own* methods. In this chapter, we're going to fix all of that. Let's get started!

Type errors in real life

If you live in the US, you are probably used to the quirky system of measurement used there. At gas stations, for example, fuel is sold by the gallon, a volume nearly four times the size of the liter used in much of the rest of the world.

Steve is an American, renting a car in another country. He pulls into a gas station to refuel. He intends to purchase 10 gallons, figuring that will be enough to reach his hotel in another city.



He gets back on the road, but only gets one-fourth of the way to his destination before running out of fuel.

If Steve had looked at the labels on the gas pump more closely, he would have realized that it was measuring the fuel in liters, not gallons, and that he needed to purchase 37.85 liters to get the equivalent of 10 gallons.

When you have a number, it's best to be certain what that number is measuring. You want to know if it's liters or gallons, kilograms or pounds, dollars or yen.



Defined types with underlying basic types

If you have the following variable:

```
var fuel float64 = 10
```

...does that represent 10 gallons or 10 liters? The person who wrote that declaration knows, but no one else does, not for sure.

You can use Go's defined types to make it clear what a value is to be used for. Although defined types most commonly use structs as their underlying types, they *can* be based on `int`, `float64`, `string`, `bool`, or any other type.

Go defined types most often use structs as their underlying types, but they can also be based on ints, strings, booleans, or any other type.

Here's a program that defines two new types, `Liters` and `Gallons`, both with an underlying type of `float64`. These are defined at the package level, so that they're available within any function in the current package.

Within the `main` function, we declare a variable with a type of `Gallons`, and another with a type of `Liters`. We assign values to each variable, and then print them out.

```
package main

import "fmt"

type Liters float64
type Gallons float64

func main() {
        var carFuel Gallons
        var busFuel Liters
        carFuel = Gallons(10.0)
        busFuel = Liters(240.0)
        fmt.Println(carFuel, busFuel)
}
```

Define two new types, each with an underlying type of float64.

Define a variable with a type of Gallons.

Define a variable with a type of Liters.

Convert a float64 to Gallons.

Convert a float64 to Liters.

```
10 240
```

Once you've defined a type, you can do a conversion to that type from any value of the underlying type. As with any other conversion, you write the type you want to convert to, followed by the value you want to convert in parentheses.

If we had wanted, we could have written short variable declarations in the code above using type conversions:

```
carFuel := Gallons(10.0)
busFuel := Liters(240.0)
```

Use short variable declarations together with type conversions.

Defined types with underlying basic types (continued)

If you have a variable that uses a defined type, you *cannot* assign a value of a different defined type to it, even if the other type has the same underlying type. This helps protect developers from confusing the two types.

```
carFuel = Liters(240.0)
busFuel = Gallons(10.0)
```

Errors →
```
cannot use Liters(240) (type Liters) as type Gallons in assignment
cannot use Gallons(10) (type Gallons) as type Liters in assignment
```

But you can *convert* between types that have the same underlying type. So `Liters` can be converted to `Gallons` and vice versa, because both have an underlying type of `float64`. But Go only considers the value of the underlying type when doing a conversion; there is no difference between `Gallons(Liters(240.0))` and `Gallons(240.0)`. Simply converting raw values from one type to another defeats the protection against conversion errors that types are supposed to provide.

```
carFuel = Gallons(Liters(40.0))   ← 40 liters does NOT equal 40 gallons!
busFuel = Liters(Gallons(63.0))   ← 63 gallons does NOT equal 63 liters!
fmt.Printf("Gallons: %0.1f Liters: %0.1f\n", carFuel, busFuel)
```

Legal, but incorrect! →
```
Gallons: 40.0 Liters: 63.0
```

Instead, you'll want to perform whatever operations are necessary to convert the underlying type value to a value appropriate for the type you're converting to.

A quick web search shows that one liter equals roughly 0.264 gallons, and that one gallon equals roughly 3.785 liters. We can multiply by these conversion rates to convert from `Gallons` to `Liters`, and vice versa.

```
carFuel = Gallons(Liters(40.0) * 0.264)   ← Convert from Liters to Gallons.
busFuel = Liters(Gallons(63.0) * 3.785)   ← Convert from Gallons to Liters.
fmt.Printf("Gallons: %0.1f Liters: %0.1f\n", carFuel, busFuel)
```

Properly converted values →
```
Gallons: 10.6 Liters: 238.5
```

Defined types and operators

A defined type supports all the same operations as its underlying type. Types based on `float64`, for example, support arithmetic operators like +, -, *, and /, as well as comparison operators like ==, >, and <.

```
fmt.Println(Liters(1.2) + Liters(3.4))
fmt.Println(Gallons(5.5) - Gallons(2.2))
fmt.Println(Liters(2.2) / Liters(1.1))
fmt.Println(Gallons(1.2) == Gallons(1.2))
fmt.Println(Liters(1.2) < Liters(3.4))
fmt.Println(Liters(1.2) > Liters(3.4))
```

```
4.6
3.3
2
true
true
false
```

A type based on an underlying type of `string`, however, would support +, ==, >, and <, but not -, because - is not a valid operator for strings.

```
// package and import statements omitted
type Title string
```

Define a type with an underlying type of "string".

```
func main() {
    fmt.Println(Title("Alien") == Title("Alien"))
    fmt.Println(Title("Alien") < Title("Zodiac"))
    fmt.Println(Title("Alien") > Title("Zodiac"))
    fmt.Println(Title("Alien") + "s")
    fmt.Println(Title("Jaws 2") - " 2")
}
```

These work...

This doesn't!

Error

```
invalid operation:
Title("Jaws 2") - " 2"
(operator - not defined
on string)
```

A defined type can be used in operations together with literal values:

```
fmt.Println(Liters(1.2) + 3.4)
fmt.Println(Gallons(5.5) - 2.2)
fmt.Println(Gallons(1.2) == 1.2)
fmt.Println(Liters(1.2) < 3.4)
```

```
4.6
3.3
true
true
```

But defined types *cannot* be used in operations together with values of a different type, even if the other type has the same underlying type. Again, this is to protect developers from accidentally mixing the two types.

```
fmt.Println(Liters(1.2) + Gallons(3.4))
fmt.Println(Gallons(1.2) == Liters(1.2))
```

If you want to add a value in `Liters` to a value in `Gallons`, you'll need to convert one type to match the other first.

Errors

```
invalid operation: Liters(1.2) + Gallons(3.4)
(mismatched types Liters and Gallons)
invalid operation: Gallons(1.2) == Liters(1.2)
(mismatched types Gallons and Liters)
```



Pool Puzzle

Your **job** is to take code snippets from the pool and place them into the blank lines in this code. **Don't** use the same snippet more than once, and you won't need to use all the snippets. Your **goal** is to make a program that will run and produce the output shown.

```
package main

import "fmt"

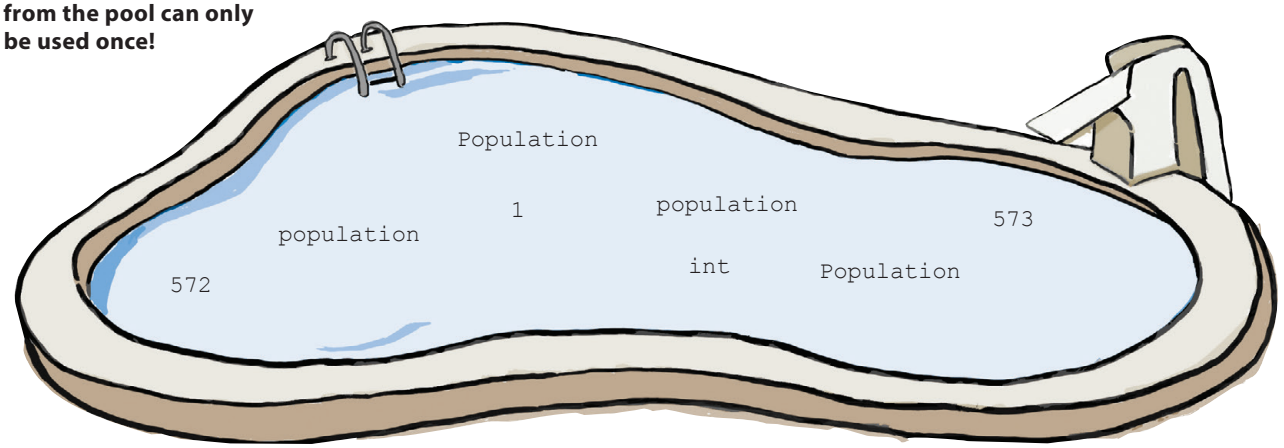
type _____________ int

func main() {
        var _____________ Population
        population = _____________(_____)
        fmt.Println("Sleepy Creek County population:", population)
        fmt.Println("Congratulations, Kevin and Anna! It's a girl!")
        population += ___
        fmt.Println("Sleepy Creek County population:", population)
}
```

Output →

```
Sleepy Creek County population: 572
Congratulations, Kevin and Anna! It's a girl!
Sleepy Creek County population: 573
```

Note: each snippet from the pool can only be used once!



→ Answers on page 286.

Converting between types using functions

Suppose we wanted to take a car whose fuel level is measured in `Gallons` and refill it at a gas pump that measures in `Liters`. Or take a bus whose fuel is measured in `Liters` and refill it at a gas pump that measures in `Gallons`. To protect us from inaccurate measurements, Go will give us a compile error if we try to combine values of different types:

```
package main

import "fmt"

type Liters float64
type Gallons float64

func main() {
        carFuel := Gallons(1.2)
        busFuel := Liters(2.5)
        carFuel += Liters(8.0)
        busFuel += Gallons(30.0)
}
```

Can't add a Liters value to a Gallons value!

Can't add a Gallons value to a Liters value!

Errors →

```
invalid operation: carFuel += Liters(8)
(mismatched types Gallons and Liters)
invalid operation: busFuel += Gallons(20)
(mismatched types Liters and Gallons)
```

In order to do operations with values of different types, we need to convert the types to match first. Previously, we demonstrated multiplying a `Liters` value by 0.264 and converted the result to `Gallons`. We also multiplied a `Gallons` value by 3.785 and converted the result to `Liters`.

```
carFuel = Gallons(Liters(40.0) * 0.264)
busFuel = Liters(Gallons(63.0) * 3.785)
```

Convert from Liters to Gallons.

Convert from Gallons to Liters.

We can create `ToGallons` and `ToLiters` functions that do the same thing, then call them to perform the conversion for us:

```
// Imports, type declarations omitted
func ToGallons(l Liters) Gallons {
        return Gallons(l * 0.264)
}

func ToLiters(g Gallons) Liters {
        return Liters(g * 3.785)
}

func main() {
        carFuel := Gallons(1.2)
        busFuel := Liters(4.5)
        carFuel += ToGallons(Liters(40.0))
        busFuel += ToLiters(Gallons(30.0))
        fmt.Printf("Car fuel: %0.1f gallons\n", carFuel)
        fmt.Printf("Bus fuel: %0.1f liters\n", busFuel)
}
```

The number of Gallons is just over 1/4 the number of Liters.

The number of Liters is just under four times the number of Gallons.

Convert Liters to Gallons before adding.

Convert Gallons to Liters before adding.

```
Car fuel: 11.8 gallons
Bus fuel: 118.1 liters
```

Converting between types using functions (continued)

Gasoline isn't the only liquid we need to measure the volume of. There's cooking oil, bottles of soda, and juice, to name a few. And so there are many more measures of volume than just liters and gallons. In the US there are teaspoons, cups, quarts, and more. The metric system has other units of measure as well, but the milliliter (1/1000 of a liter) is the most commonly used.

Let's add a new type, `Milliliters`. Like the others, it will use `float64` as an underlying type.

Add a new type.

```
type Liters float64
type Milliliters float64
type Gallons float64
```

We're also going to want a way to convert from `Milliliters` to the other types. But if we start adding a function to convert from `Milliliters` to `Gallons`, we run into a problem: we can't have two `ToGallons` functions in the same package!

```
func ToGallons(l Liters) Gallons {
      return Gallons(l * 0.264)
}
func ToGallons(m Milliliters) Gallons {
      return Gallons(m * 0.000264)
}
```

We can't add another function to convert from Milliliters to Gallons if it has the same name!

Error →

```
12:31: ToGallons redeclared in this block
      previous declaration at prog.go:9:26
```

We could rename the two `ToGallons` functions to include the type they're converting from: `LitersToGallons` and `MillilitersToGallons`, respectively. But those names would be a pain to write out all the time, and as we start adding functions to convert between the other types, it becomes clear this isn't sustainable.

```
Eliminates the conflict, but the name is really long!
func LitersToGallons(l Liters) Gallons {
      return Gallons(l * 0.264)
}
Eliminates the conflict, but the name is really long!
func MillilitersToGallons(m Milliliters) Gallons {
      return Gallons(m * 0.000264)
}
Avoids conflict, but the name is really long!
func GallonsToLiters(g Gallons) Liters {
      return Liters(g * 3.785)
}
Avoids conflict, but the name is really long!
func GallonsToMilliliters(g Gallons) Milliliters {
      return Milliliters(g * 3785.41)
}
```

there are no Dumb Questions

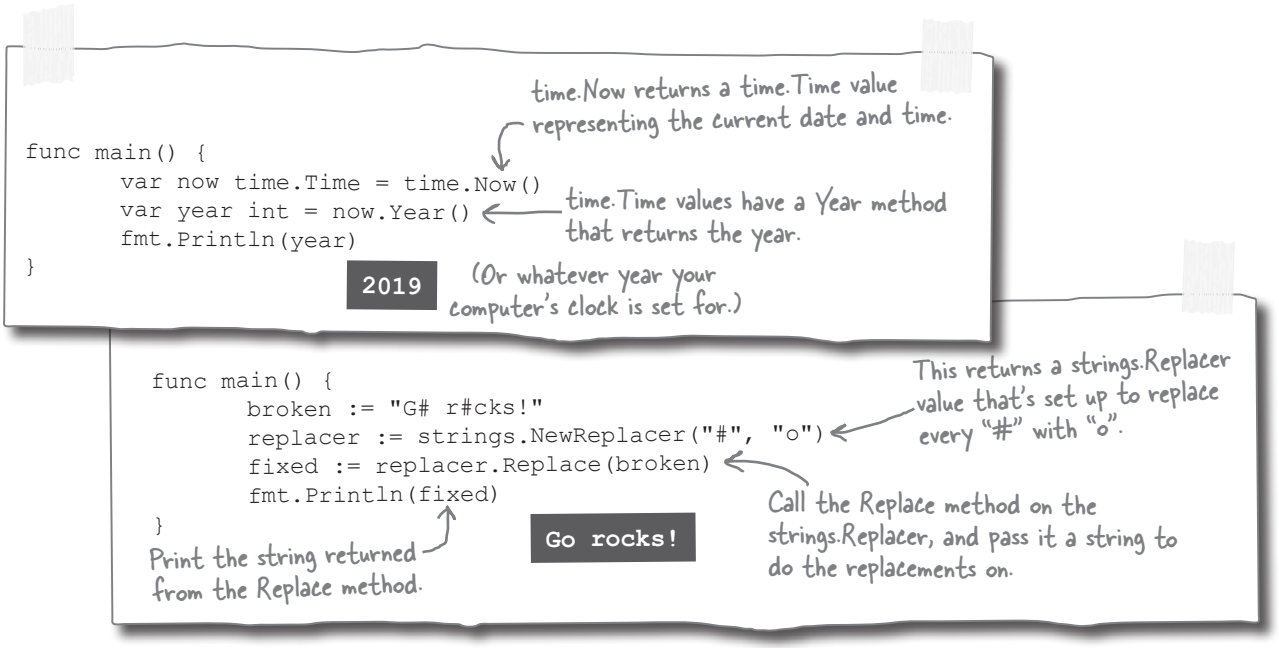
Q: I've seen other languages that support function *overloading*: they allow you to have multiple functions with the same name, as long as their parameter types are different. Doesn't Go support that?

A: The Go maintainers get this question frequently too, and they answer it at *https://golang.org/doc/faq#overloading*: "Experience with other languages told us that having a variety of methods with the same name but different signatures was occasionally useful but that it could also be confusing and fragile in practice." The Go language is simplified by *not* supporting overloading, and so it doesn't support it. As you'll see later in the book, the Go team made similar decisions in other areas of the language, too; when they have to choose between simplicity and adding more features, they generally choose simplicity. But that's okay! As we'll see shortly, there are other ways to get the same benefits...


Wouldn't it be dreamy if you could write a ToGallons function that worked with Liters values, and another ToGallons function that worked with Milliliters values? But I know it's just a fantasy…

Fixing our function name conflict using methods

Remember way back in Chapter 2, we introduced you to *methods*, which are functions associated with values of a given type? Among other things, we created a `time.Time` value and called its `Year` method, and we created a `strings.Replacer` value and called its `Replace` method.



We can define methods of our own to help with our type conversion problem.

We're not allowed to have multiple functions named `ToGallons`, so we had to write long, cumbersome function names that incorporated the type we were converting:

```
LitersToGallons(Liters(2))
MillilitersToGallons(Milliliters(500))
```

But we *can* have multiple *methods* named `ToGallons`, as long as they're defined on separate types. Not having to worry about name conflicts will let us make our method names much shorter.

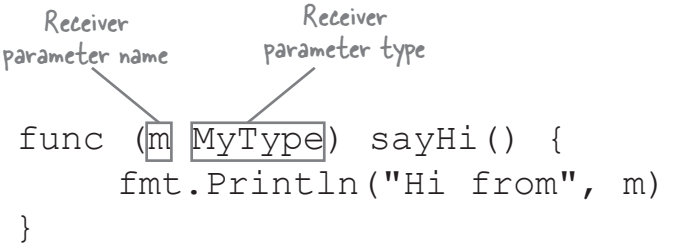
```
Liters(2).ToGallons()
Milliliters(500).ToGallons()
```

But let's not get ahead of ourselves. Before we can do anything else, we need to know how to define a method...

Defining methods

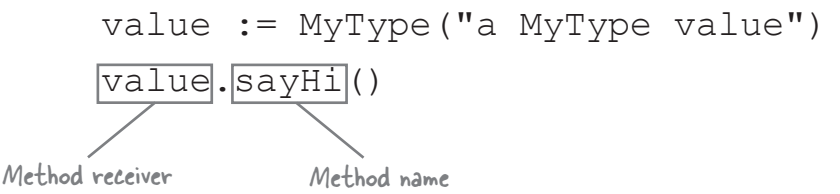
A method definition is very similar to a function definition. In fact, there's really only one difference: you add one extra parameter, a **receiver parameter**, in parentheses *before* the function name.

As with any function parameter, you need to provide a name for the receiver parameter, followed by a type.



To call a method you've defined, you write the value you're calling the method on, a dot, and the name of the method you're calling, followed by parentheses. The value you're calling the method on is known as the method **receiver**.

The similarity between method calls and method definitions can help you remember the syntax: the receiver is listed first when you're *calling* a method, and the receiver parameter is listed first when you're *defining* a method.



The name of the receiver parameter in the method definition isn't important, but its type is; the method you're defining becomes associated with all values of that type.

Below, we define a type named `MyType`, with an underlying type of `string`. Then, we define a method named `sayHi`. Because `sayHi` has a receiver parameter with a type of `MyType`, we'll be able to call the `sayHi` method on any `MyType` value. (Most developers would say that `sayHi` is defined "on" `MyType`.)

Once a method is defined on a type, it can be called on any value of that type.

Here, we create two different `MyType` values, and call `sayHi` on each of them.

```
package main

import "fmt"

type MyType string

func (m MyType) sayHi() {
        fmt.Println("Hi")
}

func main() {
        value := MyType("a MyType value")
        value.sayHi()
        anotherValue := MyType("another value")
        anotherValue.sayHi()
}
```

Define a new type.

Define a receiver parameter.

The method will be defined on MyType.

Create a MyType value.

Call sayHi on that value.

Create another MyType value.

Call sayHi on the new value.

```
Hi
Hi
```

The receiver parameter is (pretty much) just another parameter

The type of the receiver parameter is the type that the method becomes associated with. But aside from that, the receiver parameter doesn't get special treatment from Go. You can access its contents within the method block just like you would any other function parameter.

The code sample below is almost identical to the previous one, except that we've updated it to print the value of the receiver parameter. You can see the receivers in the resulting output.

```
package main

import "fmt"

type MyType string

func (m MyType) sayHi() {
        fmt.Println("Hi from", m)
}

func main() {
        value := MyType("a MyType value")
        value.sayHi()
        anotherValue := MyType("another value")
        anotherValue.sayHi()
}
```

Print the receiver parameter's value.

Value to call method on

Receivers passed to receiver parameter

Value to call method on

```
Hi from a MyType value
Hi from another value
```

See receiver values in the output.

Go lets you name a receiver parameter whatever you want, but it's more readable if all the methods you define for a type have receiver parameters with the same name.

By convention, Go developers usually use a name consisting of a single letter—the first letter of the receiver's type name, in lowercase. (This is why we used `m` as the name for our `MyType` receiver parameter.)

Go uses receiver parameters instead of the "self" or "this" values seen in other languages.

there are no Dumb Questions

Q: Can I define new methods on *any* type?

A: Only types that are defined in the same package where you define the method. That means no defining methods for types from someone else's `security` package from your `hacking` package, and no defining new methods on universal types like `int` or `string`.

Q: But I need to be able to use methods of my own with someone else's type!

A: First you should consider whether a function would work well enough; a function can take any type you want as a parameter. But if you *really* need a value that has some methods of your own, plus some methods from a type in another package, you can make a struct type that embeds the other package's type as an anonymous field. We'll look at how that works in the next chapter.

Q: I've seen other languages where a method receiver was available in a method block in a special variable named `self` or `this`. Does Go do that?

A: Go uses receiver parameters instead of `self` and `this`. The big difference is that `self` and `this` are set *implicitly*, whereas you *explicitly* declare a receiver parameter. Other than that, receiver parameters are used in the same way, and there's no need for Go to reserve `self` or `this` as keywords! (You could even name your receiver parameter `this` if you wanted, but don't do that; the convention is to use the first letter of the receiver's type name instead.)

A method is (pretty much) just like a function

Aside from the fact that they're called on a receiver, methods are otherwise pretty similar to any other function.

As with any other function, you can define additional parameters within parentheses following the method name. These parameter variables can be accessed in the method block, along with the receiver parameter. When you call the method, you'll need to provide an argument for each parameter.

```
func (m MyType) MethodWithParameters(number int, flag bool) {
        fmt.Println(m)
        fmt.Println(number)
        fmt.Println(flag)
}

func main() {
        value := MyType("MyType value")
        value.MethodWithParameters(4, true)
}
```

Receiver parameter

Parameter

Parameter

Receiver

Argument

Argument

```
MyType value
4
true
```

As with any other function, you can declare one or more return values for a method, which will be returned when the method is called:

```
func (m MyType) WithReturn() int {
        return len(m)
}

func main() {
        value := MyType("MyType value")
        fmt.Println(value.WithReturn())
}
```

Return value

Return the length of the receiver's underlying string value.

Print the method's return value.

```
12
```

As with any other function, a method is considered exported from the current package if its name begins with a capital letter, and it's considered unexported if its name begins with a lowercase letter. If you want to use your method outside the current package, be sure its name begins with a capital letter.

Exported; name begins with a capital letter.

```
func (m MyType) ExportedMethod() {
}

func (m MyType) unexportedMethod() {
}
```

Unexported; name begins with a lowercase letter.



Fill in the blanks to define a `Number` type with `Add` and `Subtract` methods that will produce the output shown.

```
type Number int

func (__ _________) ____(_____________ int) {
        fmt.Println(n, "plus", otherNumber, "is", int(n)+otherNumber)
}

func (__ _________) _________(_____________ int) {
        fmt.Println(n, "minus", otherNumber, "is", int(n)-otherNumber)
}

func main() {
        ten := Number(10)
        ten.Add(4)
        ten.Subtract(5)
        four := Number(4)
        four.Add(3)
        four.Subtract(2)
}
```

```
10 plus 4 is 14
10 minus 5 is 5
4 plus 3 is 7
4 minus 2 is 2
```

Answers on page 286.

Pointer receiver parameters

Here's an issue that may look familiar by now. We've defined a new `Number` type with an underlying type of `int`. We've given `Number` a `double` method that is supposed to multiply the underlying value of its receiver by two and then update the receiver. But we can see from the output that the method receiver isn't actually getting updated.

```
package main

import "fmt"

type Number int

func (n Number) Double() {
        n *= 2
}

func main() {
        number := Number(4)
        fmt.Println("Original value of number:", number)
        number.Double()
        fmt.Println("number after calling Double:", number)
}
```

Define a type with an underlying type of "int".

Define a method on the Number type.

Multiply the receiver by two, and attempt to update the receiver.

Create a Number value.

Attempt to double the Number.

```
Original value of number: 4
number after calling Double: 4
```

Number is unchanged!

Back in Chapter 3, we had a `double` *function* with a similar problem. Back then, we learned that function parameters receive a copy of the values the function is called with, not the original values, and that any updates to the copy would be lost when the function exited. To make the `double` function work, we had to pass a *pointer* to the value we wanted to update, and then update the value at that pointer within the function.

```
func main() {
        amount := 6
        double(&amount)
        fmt.Println(amount)
}

func double(number *int) {
        *number *= 2
}
```

Pass a pointer instead of the variable value.

Accept a pointer instead of an int value.

Update the value at the pointer.

```
12
```

Prints the doubled amount

Pointer receiver parameters (continued)

We've said that receiver parameters are treated no differently than ordinary parameters. And like any other parameter, a receiver parameter receives a *copy* of the receiver value. If you make changes to the receiver within a method, you're changing the copy, not the original.

As with the `double` function in Chapter 3, the solution is to update our `Double` method to use a pointer for its receiver parameter. This is done in the same way as any other parameter: we place a `*` in front of the receiver type to indicate it's a pointer type. We'll also need to modify the method block so that it updates the value at the pointer. Once that's done, when we call `Double` on a `Number` value, the `Number` should be updated.

```
// Package, imports, type omitted
func (n *Number) Double() {
        *n *= 2
}
func main() {
        number := Number(4)
        fmt.Println("Original value of number:", number)
        number.Double()
        fmt.Println("number after calling Double:", number)
}
```

Change the receiver parameter to a pointer type.

Update the value at the pointer.

We DON'T have to update the method call!

```
Original value of number: 4
number after calling Double: 8
```

Value at pointer was updated.

Notice that we *didn't* have to change the method call at all. When you call a method that requires a pointer receiver on a variable with a nonpointer type, Go will automatically convert the receiver to a pointer for you. The same is true for variables with pointer types; if you call a method requiring a value receiver, Go will automatically get the value at the pointer for you and pass that to the method.

You can see this at work in the code at right. The method named `method` takes a value receiver, but we can call it using both direct values and pointers, because Go autoconverts if needed. And the method named `pointerMethod` takes a pointer receiver, but we can call it on both direct values and pointers, because Go will autoconvert if needed.

By the way, the code at right breaks a convention: for consistency, all of your type's methods can take value receivers, or they can all take pointer receivers, but you should avoid mixing the two. We're only mixing the two kinds here for demonstration purposes.

```
// Package, imports omitted
type MyType string

func (m MyType) method() {
        fmt.Println("Method with value receiver")
}
func (m *MyType) pointerMethod() {
        fmt.Println("Method with pointer receiver")
}

func main() {
        value := MyType("a value")
        pointer := &value
        value.method()
        value.pointerMethod()
        pointer.method()
        pointer.pointerMethod()
}
```

Value automatically converted to pointer

Value at pointer automatically retrieved

```
Method with value receiver
Method with pointer receiver
Method with value receiver
Method with pointer receiver
```



To call a method that requires a pointer receiver, you have to be able to get a pointer to the value!

You can only get pointers to values that are stored in variables. If you try to get the address of a value that's not stored in a variable, you'll get an error:

```
&MyType("a value")
```

Error →

```
cannot take the address
of MyType("a value")
```

The same limitation applies when calling methods with pointer receivers. Go can automatically convert values to pointers for you, but only if the receiver value is stored in a variable. If you try to call a method on the value itself, Go won't be able to get a pointer, and you'll get a similar error:

```
MyType("a value").pointerMethod()
```

Errors →

```
cannot call pointer method
on MyType("a value")
cannot take the address
of MyType("a value")
```

Instead, you'll need to store the value in a variable, which will then allow Go to get a pointer to it:

```
value := MyType("a value")
value.pointerMethod()
```

Go converts this to a pointer.



Breaking Stuff is Educational!

Here is our `Number` type again, with definitions for a couple methods. Make one of the changes below and try to compile the code. Then undo your change and try the next one. See what happens!

```
package main

import "fmt"

type Number int

func (n *Number) Display() {
       fmt.Println(*n)
}
func (n *Number) Double() {
       *n *= 2
}

func main() {
       number := Number(4)
       number.Double()
       number.Display()
}
```

If you do this...	...the code will break because...

Change a receiver parameter to a type not defined in this package:
`func (n *~~Number~~int) Double() {`
`  *n *= 2`
`}`
	You can only define new methods on types that were declared in the current package. Defining a method on a globally defined type like `int` will result in a compile error.

Change the receiver parameter for `Double` to a nonpointer type:
`func (n *Number) Double() {`
`  *n *= 2`
`}`
	Receiver parameters receive a copy of the value the method was called on. If the `Double` function only modifies the copy, the original value will be unchanged when `Double` exits.

Call a method that requires a pointer receiver on a value that's not in a variable: `Number(4).Double()`	When calling a method that takes a pointer receiver, Go can automatically convert a value to a pointer to a receiver *if* it's stored in a variable. If it's not, you'll get an error.

Change the receiver parameter for `Display` to a nonpointer type:
`func (n *Number) Display() {`
`  fmt.Println(*n)`
`}`
	The code will actually still *work* after making this change, but it breaks convention! Receiver parameters in the methods for a type can be all pointers, or all values, but it's best to avoid mixing the two.

Converting Liters and Milliliters to Gallons using methods

When we added a `Milliliters` type to our defined types for measuring volume, we discovered we couldn't have `ToGallons` functions for both `Liters` and `Milliliters`. To work around this, we had to create functions with lengthy names:

```
func LitersToGallons(l Liters) Gallons {
      return Gallons(l * 0.264)
}
func MillilitersToGallons(m Milliliters) Gallons {
      return Gallons(m * 0.000264)
}
```

But unlike functions, method names don't have to be unique, as long as they're defined on different types.

Let's try implementing a `ToGallons` method on the `Liters` type. The code will be almost identical to the `LitersToGallons` function, but we'll make the `Liters` value a receiver parameter rather than an ordinary parameter. Then we'll do the same for the `Milliliters` type, converting the `MillilitersToGallons` function to a `ToGallons` method.

Notice that we're not using pointer types for the receiver parameters. We're not modifying the receivers, and the values don't consume much memory, so it's fine for the parameter to receive a copy of the value.

```
package main

import "fmt"

type Liters float64
type Milliliters float64
type Gallons float64

func (l Liters) ToGallons() Gallons {
      return Gallons(l * 0.264)
}

func (m Milliliters) ToGallons() Gallons {
      return Gallons(m * 0.000264)
}

func main() {
      soda := Liters(2)
      fmt.Printf("%0.3f liters equals %0.3f gallons\n", soda, soda.ToGallons())
      water := Milliliters(500)
      fmt.Printf("%0.3f milliliters equals %0.3f gallons\n", water, water.ToGallons())
}
```

Method for Liters

Names can be identical, if they're on separate types.

Method block unchanged from function block

Method for Milliliters

Names can be identical, if they're on separate types.

Method block unchanged from function block

Create Liters value.

Convert Liters to Gallons.

Create Milliliters value.

Convert Milliliters to Gallons.

```
2.000 liters equals 0.528 gallons
500.000 milliliters equals 0.132 gallons
```

In our `main` function, we create a `Liters` value, then call `ToGallons` on it. Because the receiver has the type `Liters`, the `ToGallons` method for the `Liters` type is called. Likewise, calling `ToGallons` on a `Milliliters` value causes the `ToGallons` method for the `Milliliters` type to be called.

Converting Gallons to Liters and Milliliters using methods

The process is similar when converting the `GallonsToLiters` and `GallonsToMilliliters` functions to methods. We just move the `Gallons` parameter to a receiver parameter in each.

```
func (g Gallons) ToLiters() Liters {
       return Liters(g * 3.785)
}
func (g Gallons) ToMilliliters() Milliliters {
       return Milliliters(g * 3785.41)
}

func main() {
       milk := Gallons(2)
       fmt.Printf("%0.3f gallons equals %0.3f liters\n", milk, milk.ToLiters())
       fmt.Printf("%0.3f gallons equals %0.3f milliliters\n", milk, milk.ToMilliliters())
}
```

Define a ToLiters method on the Gallons type.

Define a ToMilliliters method on the Gallons type.

Create a Gallons value.

Convert it to Liters.

Convert it to Milliliters.

```
2.000 gallons equals 7.570 liters
2.000 gallons equals 7570.820 milliliters
```



Exercise

The code below should add a `ToMilliliters` method on the `Liters` type, and a `ToLiters` method on the `Milliliters` type. The code in the `main` function should produce the output shown. Fill in the blanks to complete the code.

```
type Liters float64
type Milliliters float64
type Gallons float64

func _____ ToMilliliters() _____ {
       return Milliliters(l * 1000)
}
func _____ ToLiters() _____ {
       return Liters(m / 1000)
}

func main() {
       l := _____(3)
       fmt.Printf("%0.1f liters is %0.1f milliliters\n", l, l._____())
       ml := _____(500)
       fmt.Printf("%0.1f milliliters is %0.1f liters\n", ml, ml._____())
}
```

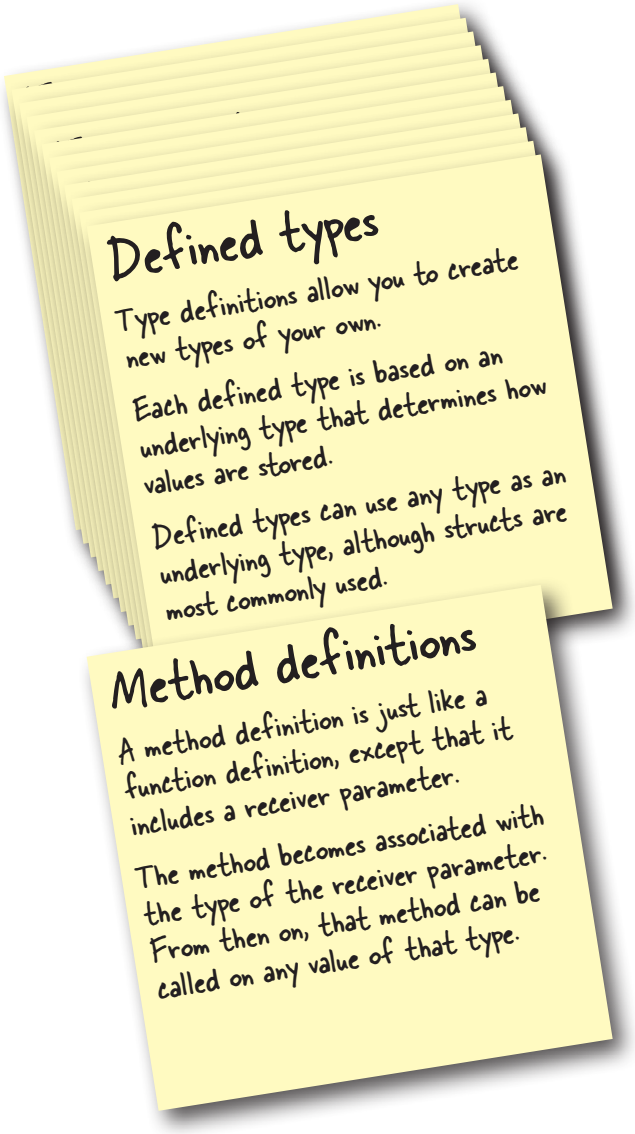
```
3.0 liters is 3000.0 milliliters
500.0 milliliters is 0.5 liters
```

Answers on page 287.



Your Go Toolbox

That's it for Chapter 9! You've added method definitions to your toolbox.



BULLET POINTS

- Once you've defined a type, you can do a conversion to that type from any value of the same underlying type:
  `Gallons(10.0)`
- Once a variable's type is defined, values of other types cannot be assigned to that variable, even if they have the same underlying type.
- A defined type supports all the same operators as its underlying type. A type based on `int`, for example, would support `+`, `-`, `*`, `/`, `==`, `>`, and `<` operators.
- A defined type can be used in operations together with literal values:
  `Gallons(10.0) + 2.3`
- To define a method, provide a receiver parameter in parentheses before the method name:
  ```
  func (m MyType) MyMethod() {
  }
  ```
- The receiver parameter can be used within the method block like any other parameter:
  ```
  func (m MyType) MyMethod() {
      fmt.Println("called on", m)
  }
  ```
- You can define additional parameters or return values on a method, just as you would with any other function.
- Defining multiple functions with the same name in the same package is not allowed, even if they have parameters of different types. But you *can* define multiple *methods* with the same name, as long as each is defined on a different type.
- You can only define methods on types that were defined in the same package.
- As with any other parameter, receiver parameters receive a copy of the original value. If your method needs to modify the receiver, you should use a pointer type for the receiver parameter, and modify the value at that pointer.

Pool Puzzle Solution

```
package main

import "fmt"

type Population int

func main() {
    var population Population
    population = Population(572)
    fmt.Println("Sleepy Creek County population:", population)
    fmt.Println("Congratulations, Kevin and Anna! It's a girl!")
    population += 1
    fmt.Println("Sleepy Creek County population:", population)
}
```

Declare a Population type with an underlying type of "int".

Convert an integer to a Population value.

The underlying type supports the += operator; therefore, Population does too.

Output

```
Sleepy Creek County population: 572
Congratulations, Kevin and Anna! It's a girl!
Sleepy Creek County population: 573
```



Exercise Solution

Fill in the blanks to define a `Number` type with `Add` and `Subtract` methods that will produce the output shown.

```
type Number int

func (n Number) Add(otherNumber int) {
    fmt.Println(n, "plus", otherNumber, "is", int(n)+otherNumber)
}

func (n Number) Subtract(otherNumber int) {
    fmt.Println(n, "minus", otherNumber, "is", int(n)-otherNumber)
}

func main() {
    ten := Number(10)
    ten.Add(4)
    ten.Subtract(5)
    four := Number(4)
    four.Add(3)
    four.Subtract(2)
}
```

The receiver parameter

Method will be defined on the Number type.

Print the receiver.

Print the regular parameter.

Can't add a Number to an int; need to do a conversion

Need to do a conversion

Convert an integer to a Number.

Call the Number's methods.

Convert an integer to a Number.

Call the Number's methods.

```
10 plus 4 is 14
10 minus 5 is 5
4 plus 3 is 7
4 minus 2 is 2
```



The code below should add a `ToMilliliters` method on the `Liters` type, and a `ToLiters` method on the `Milliliters` type. The code in the `main` function should produce the output shown. Fill in the blanks to complete the code.

```
type Liters float64
type Milliliters float64
type Gallons float64

func (l Liters) ToMilliliters() Milliliters {
        return Milliliters(l * 1000)
}
func (m Milliliters) ToLiters() Liters {
        return Liters(m / 1000)
}

func main() {
        l := Liters(3)
        fmt.Printf("%0.1f liters is %0.1f milliliters\n", l, l.ToMilliliters())
        ml := Milliliters(500)
        fmt.Printf("%0.1f milliliters is %0.1f liters\n", ml, ml.ToLiters())
}
```

Multiply the receiver value by 1,000, and convert the result's type to Milliliters.

Divide the receiver value by 1,000, and convert the result's type to Liters.

```
3.0 liters is 3000.0 milliliters
500.0 milliliters is 0.5 liters
```

10 keep it to yourself

Encapsulation and Embedding

I heard that `Paragraph` type of hers stores its data in a simple `string` field! And that fancy `Replace` method? It's just promoted from an embedded `strings.Replacer`! You'd never know it from using `Paragraph`, though!



Mistakes happen. Sometimes, your program will receive invalid data from user input, a file you're reading in, or elsewhere. In this chapter, you'll learn about **encapsulation**: a way to protect your struct type's fields from that invalid data. That way, you'll know your field data is safe to work with!

We'll also show you how to **embed** other types within your struct type. If your struct type needs methods that already exist on another type, you don't have to copy and paste the method code. You can embed the other type within your struct type, and then use the embedded type's methods just as if they were defined on your own type!

Creating a Date struct type

A local startup called Remind Me is developing a calendar application to help users remember birthdays, anniversaries, and more.



The year, month, and day sound like they all need to be grouped together; none of those values would be useful by itself. A struct type would probably be useful for keeping those separate values together in a single bundle.

As we've seen, defined types can use any other type as their underlying type, including structs. In fact, struct types served as our introduction to defined types, back in Chapter 8.

Let's create a `Date` struct type to hold our year, month, and day values. We'll add `Year`, `Month`, and `Day` fields to the struct, each with a type of `int`. In our `main` function, we'll run a quick test of the new type, using a struct literal to create a `Date` value with all its fields populated. We'll just use `Println` to print the `Date` out for now.

```
package main

import "fmt"

type Date struct {
        Year  int
        Month int
        Day   int
}

func main() {
        date := Date{Year: 2019, Month: 5, Day: 27}
        fmt.Println(date)
}
```

Define a new struct type.

Define struct fields.

Use a struct literal to create a Date value.

```
{2019 5 27}
```

If we run the finished program, we'll see the `Year`, `Month`, and `Day` fields of our `Date` struct. It looks like everything's working!

People are setting the Date struct field to invalid values!



Ah, we can see how that might happen. Only year numbers `1` or greater are valid, but we don't have anything preventing users from accidentally setting the `Year` field to `0` or `-999`. Only month numbers from `1` through `12` are valid, but nothing prevents users from setting the `Month` field to `0` or `13`. Only the numbers `1` through `31` are valid for the `Day` field, but users can enter days like `-2` or `50`.

```
date := Date{Year: 2019, Month: 14, Day: 50}
fmt.Println(date)
date = Date{Year: 0, Month: 0, Day: -2}
fmt.Println(date)
date = Date{Year: -999, Month: -1, Day: 0}
fmt.Println(date)
```

Invalid! Invalid! Invalid! Invalid! Invalid! Invalid!

```
{2019 14 50}
{0 0 -2}
{-999 -1 0}
```

What we need is a way for our programs to ensure the user data is valid before accepting it. In computer science, this is known as *data validation*. We need to test that the `Year` is being set to a value of `1` or greater, the `Month` is being set between `1` and `12`, and the `Day` is being set between `1` and `31`.

(Yes, some months have fewer than 31 days, but to keep our code samples a reasonable length, we'll just check that it's between 1 and 31.)

Setter methods

A struct type is just another defined type, and that means you can define methods on it just like any other. We should be able to create `SetYear`, `SetMonth`, and `SetDay` methods on the `Date` type that take a value, check whether it's valid, and if so, set the appropriate struct field.

This kind of method is often called a **setter method**. By convention, Go setter methods are usually named in the form `Set`*X*, where *X* is the thing that you're setting.

Here's our first attempt at a `SetYear` method. The receiver parameter is the `Date` struct you're calling the method on. `SetYear` accepts the year you want to set as a parameter, and sets the `Year` field on the receiver `Date` struct. It doesn't validate the value at all currently, but we'll add validation in a little bit.

In our `main` method, we create a `Date` and call `SetYear` on it. Then we print the struct's `Year` field.

Setter methods are methods used to set fields or other values within a defined type's underlying value.

```
package main

import "fmt"

type Date struct {
        Year  int
        Month int
        Day   int
}

func (d Date) SetYear(year int) {
        d.Year = year
}

func main() {
        date := Date{}
        date.SetYear(2019)
        fmt.Println(date.Year)
}
```

Accepts the value the field should be set to

Set the struct field.

Create a Date.

Set its Year field via the method.

Print the Year field.

```
0
```

Year is still set to its zero value!

When we run the program, though, we'll see that it didn't work quite right. Even though we create a `Date` and call `SetYear` with a new value, the `Year` field is still set to its zero value!

Setter methods need pointer receivers

Remember the `Double` method on the `Number` type we showed you earlier? Originally, we wrote it with a plain value receiver type, `Number`. But we learned that, like any other parameter, receiver parameters receive a *copy* of the original value. The `Double` method was updating the copy, which was lost when the function exited.

```
func (n *Number) Double() {
        *n *= 2
}
```

Change the receiver parameter to a pointer type.

Update the value at the pointer.

We needed to update `Double` to take a pointer receiver type, `*Number`. When we updated the value at the pointer, the changes were preserved after `Double` exited.

The same holds true for `SetYear`. The `Date` receiver gets a *copy* of the original struct. Any updates to the fields of the copy are lost when `SetYear` exits!

```
func (d Date) SetYear(year int) {
        d.Year = year
}
```

Receives a copy of the Date struct

Updates the copy, not the original!

We can fix `SetYear` by updating it to take a pointer receiver: `(d *Date)`. That's the only change that's necessary. We don't have to update the `SetYear` method block, because `d.Year` automatically gets the value at the pointer for us (as if we'd typed `(*d).Year`). The call to `date.SetYear` in `main` doesn't need to be changed either, because the `Date` value is automatically converted to a `*Date` when it's passed to the method.

```
type Date struct {
        Year  int
        Month int
        Day   int
}

func (d *Date) SetYear(year int) {
        d.Year = year
}

func main() {
        date := Date{}
        date.SetYear(2019)
        fmt.Println(date.Year)
}
```

Needs to be a pointer receiver, so original value can be updated

Now updates original value, not a copy

Automatically gets value at pointer

Automatically converted to a pointer

```
2019
```

Year field has been updated.

Now that `SetYear` takes a pointer receiver, if we rerun the code, we'll see that the `Year` field has been updated.

Adding the remaining setter methods

Now it should be easy to follow the same pattern to define `SetMonth` and `SetDay` methods on the `Date` type. We just need to be sure to use a pointer receiver in the method definition. Go will convert the receiver to a pointer when we call each method, and convert the pointer back to a struct value when updating its fields.

```
package main

import "fmt"

type Date struct {
       Year  int
       Month int
       Day   int
}

func (d *Date) SetYear(year int) {
       d.Year = year
}
func (d *Date) SetMonth(month int) {
       d.Month = month
}
func (d *Date) SetDay(day int) {
       d.Day = day
}
```

Be sure to use pointer receivers!

In `main`, we can create a `Date` struct value; set its `Year`, `Month`, and `Day` fields via our new methods; and print the whole struct out to see the results.

```
func main() {
       date := Date{}
       date.SetYear(2019)
       date.SetMonth(5)
       date.SetDay(27)
       fmt.Println(date)
}
```

Set the month.

Set the day of the month.

Print all fields.

```
{2019 5 27}
```

Now we have setter methods for each of our `Date` type's fields. But even if they use the methods, users can still accidentally set the fields to invalid values. We'll look at preventing that next.

```
date := Date{}
date.SetYear(0)
date.SetMonth(14)
date.SetDay(50)
fmt.Println(date)
```

Invalid!

Invalid!

Invalid!

```
{0 14 50}
```



In the Chapter 8 exercises, you saw code for a `Coordinates` struct type. We've moved that type definition to a *coordinates.go* file within the *geo* package directory.

We need to add setter methods to the `Coordinates` type for each of its fields.Fill in the blanks in the *coordinates.go* file below, so that the code in *main.go* will run and produce the output shown.

```
package geo

type Coordinates struct {
        Latitude  float64
        Longitude float64
}

func (c _______________) SetLatitude(__________ float64) {
        _____________ = latitude
}

func (c _______________) SetLongitude(__________ float64) {
        _____________ = longitude
}
```

coordinates.go

```
package main

import (
        "fmt"
        "geo"
)

func main() {
        coordinates := geo.Coordinates{}
        coordinates.SetLatitude(37.42)
        coordinates.SetLongitude(-122.08)
        fmt.Println(coordinates)
}
```

main.go

Output

```
{37.42 -122.08}
```

Answers on page 317.

Adding validation to the setter methods

Adding validation to our setter methods will take a bit of work, but we learned everything we need to do it in Chapter 3.

In each setter method, we'll test whether the value is in a valid range. If it's invalid, we'll return an `error` value. If it's valid, we'll set the `Date` struct field as normal and return `nil` for the error value.

Let's add validation to the `SetYear` method first. We add a declaration that the method will return a value, of type `error`. At the start of the method block, we test whether the `year` parameter provided by the caller is any number less than `1`. If it is, we return an `error` with a message of `"invalid year"`. If not, we set the struct's `Year` field and return `nil`, indicating there was no error.

In `main`, we call `SetYear` and store its return value in a variable named `err`. If `err` is not `nil`, it means the assigned value was invalid, so we log the error and exit. Otherwise, we proceed to print the `Date` struct's `Year` field.

```
package main

import (
        "errors"
        "fmt"
        "log"
)

type Date struct {
        Year  int
        Month int
        Day   int
}

func (d *Date) SetYear(year int) error {
        if year < 1 {
                return errors.New("invalid year")
        }
        d.Year = year
        return nil
}
// SetMonth, SetDay omitted

func main() {
        date := Date{}
        err := date.SetYear(0)
        if err != nil {
                log.Fatal(err)
        }
        fmt.Println(date.Year)
}
```

Lets us create error values

Lets us log an error and exit

Add an error return value.

If the given year is invalid, return an error.

Otherwise, set the field...

...and return an error of "nil".

This value is invalid!

Capture any error.

If the value was invalid, log the error and exit.

```
2018/03/17 19:58:02 invalid year
exit status 1
```

Error gets logged.

Passing an invalid value to `SetYear` causes the program to report the error and exit. But if we pass a valid value, the program will proceed to print it out. Looks like our `SetYear` method is working!

```
date := Date{}
err := date.SetYear(2019)
if err != nil {
        log.Fatal(err)
}
fmt.Println(date.Year)
```

Valid value

```
2019
```

Field gets printed.

Adding validation to the setter methods (continued)

Validation code in the `SetMonth` and `SetDay` methods will be similar to the code in `SetYear`.

In `SetMonth`, we test whether the provided month number is less than `1` or greater than `12`, and return an error if so. Otherwise, we set the field and return `nil`.

And in `SetDay`, we test whether the provided day of the month is less than `1` or greater than `31`. Invalid values result in a returned error, but valid values cause the field to be set and `nil` to be returned.

```
// Package, imports, type declaration omitted
func (d *Date) SetYear(year int) error {
	if year < 1 {
		return errors.New("invalid year")
	}
	d.Year = year
	return nil
}
func (d *Date) SetMonth(month int) error {
	if month < 1 || month > 12 {
		return errors.New("invalid month")
	}
	d.Month = month
	return nil
}
func (d *Date) SetDay(day int) error {
	if day < 1 || day > 31 {
		return errors.New("invalid day")
	}
	d.Day = day
	return nil
}

func main() {
	// Try the below code snippets here
}
```

You can test the setter methods by inserting the code snippets below into the block for `main`...

Passing `14` to `SetMonth` results in an error:

```
date := Date{}
err := date.SetMonth(14)
if err != nil {
	log.Fatal(err)
}
fmt.Println(date.Month)
```

```
2018/03/17 20:17:42
invalid month
exit status 1
```

But passing `5` to `SetMonth` works:

```
date := Date{}
err := date.SetMonth(5)
if err != nil {
	log.Fatal(err)
}
fmt.Println(date.Month)
```

```
5
```

Passing `50` to `SetDay` results in an error:

```
date := Date{}
err := date.SetDay(50)
if err != nil {
	log.Fatal(err)
}
fmt.Println(date.Day)
```

```
2018/03/17 20:30:54
invalid day
exit status 1
```

But passing `27` to `SetDay` works:

```
date := Date{}
err := date.SetDay(27)
if err != nil {
	log.Fatal(err)
}
fmt.Println(date.Day)
```

```
27
```

The fields can still be set to invalid values!

The validation provided by your setter methods is great, when people actually use them. But we've got people setting the struct fields directly, and they're still entering invalid data!



It's true; there's nothing preventing anyone from setting the `Date` struct fields directly. And if they do so, it bypasses the validation code in the setter methods. They can set any value they want!

```
date := Date{}
date.Year = 2019
date.Month = 14
date.Day = 50
fmt.Println(date)
```

```
{2019 14 50}
```

We need a way to protect these fields, so that users of our `Date` type can only update the fields using the setter methods.

Go provides a way of doing this: we can move the `Date` type to another package and make its date fields unexported.

So far, unexported variables, functions, and the like have mostly gotten in our way. The most recent example of this was in Chapter 8, when we discovered that even though our `Subscriber` struct type was exported from the `magazine` package, its fields were *unexported*, making them inaccessible outside the `magazine` package.

With the `Subscriber` type name capitalized, we seem to be able to access it from the `main` package. But now we're getting an error saying that we can't refer to the `rate` *field*, because *that* is unexported.

```
Shell Edit View Window Help
$ go run main.go
./main.go:10:13: s.rate undefined
(cannot refer to unexported field or method rate)
./main.go:11:25: s.rate undefined
(cannot refer to unexported field or method rate)
```

Even if a struct type is exported from a package, its fields will be *unexported* if their names don't begin with a capital letter. Let's try capitalizing `Rate` (in both *magazine.go* and *main.go*)...

But in this case, we don't *want* the fields to be accessible. Unexported struct fields are exactly what we need!

Let's try moving our `Date` type to another package and making its fields unexported, and see if that fixes our problem.

Moving the Date type to another package

In the *headfirstgo* directory within your Go workspace, create a new directory to hold a package named `calendar`. Within *calendar*, create a file named *date.go*. (Remember, you can name the files within a package directory anything you want; they'll all become part of the same package.)



Within *date.go*, add a `package calendar` declaration and import the `"errors"` package. (That's the only package that the code in this file will be using.) Then, copy all your old code for the `Date` type and paste it into this file.

```
package calendar

import "errors"

type Date struct {
       Year  int
       Month int
       Day   int
}

func (d *Date) SetYear(year int) error {
       if year < 1 {
              return errors.New("invalid year")
       }
       d.Year = year
       return nil
}
func (d *Date) SetMonth(month int) error {
       if month < 1 || month > 12 {
              return errors.New("invalid month")
       }
       d.Month = month
       return nil
}
func (d *Date) SetDay(day int) error {
       if day < 1 || day > 31 {
              return errors.New("invalid day")
       }
       d.Day = day
       return nil
}
```

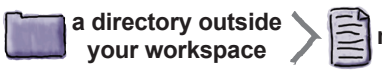
This file is part of the "calendar" package.

This file only uses functions from the "error" package.

Copy and paste all the code for the Date type into the new file.

Moving the Date type to another package (continued)

Next, let's create a program to try out the `calendar` package. Since this is just for experimenting, we'll do as we did in Chapter 8 and save a file *outside* the Go workspace, so it doesn't interfere with any other packages. (We'll just use the `go run` command to run it.) Name the file *main.go*.



(You can move this code into your Go workspace later, if you want, as long as you create a separate package directory for it.)

At this point, code we add in *main.go* will still be able to create an invalid `Date`, either by setting its fields directly or by using a struct literal.

```
package main
```
Use the "main" package, since we'll be running this code as a program.

```
import (
        "fmt"
        "github.com/headfirstgo/calendar"
)
```
Import our new package.

```
func main() {
        date := calendar.Date{}
```
Need to specify package we imported from.

Create a new Date value.

```
        date.Year = 2019
        date.Month = 14
        date.Day = 50
```
Set the Date's fields directly.

```
        fmt.Println(date)
        date = calendar.Date{Year: 0, Month: 0, Day: -2}
        fmt.Println(date)
}
```
Specify package.

Set another Date's fields using a struct literal.

If we run *main.go* from the terminal, we'll see that both ways of setting the fields worked, and two invalid dates are printed.

```
Shell  Edit  View  Window  Help
$ cd temp
$ go run main.go
{2019 14 50}
{0 0 -2}
```
Invalid dates!

Making Date fields unexported

Now let's try updating the `Date` struct so that its fields are unexported. Simply change the field names to begin with lowercase letters in the type definition and everywhere else they occur.

The `Date` type itself needs to remain exported, as do all of the setter methods, because we *will* need to access these from outside the `calendar` package.

date.go

```
package calendar

import "errors"
type Date struct {
        year  int
        month int
        day   int
}
func (d *Date) SetYear(year int) error {
        if year < 1 {
                return errors.New("invalid year")
        }
        d.year = year
        return nil
}
func (d *Date) SetMonth(month int) error {
        if month < 1 || month > 12 {
                return errors.New("invalid month")
        }
        d.month = month
        return nil
}
func (d *Date) SetDay(day int) error {
        if day < 1 || day > 31 {
                return errors.New("invalid day")
        }
        d.day = day
        return nil
}
```

Date type needs to remain exported!

Change field names so they are unexported.

No changes to method names

No changes to method parameters

Update field name to match declaration above.

Update field name to match declaration above.

Update field name to match declaration above.

To test our changes, update the field names in *main.go* to match the field names in *date.go*.

main.go

```
// Package, import statements omitted
func main() {
        date := calendar.Date{}
        date.year = 2019
        date.month = 14
        date.day = 50
        fmt.Println(date)

        date = calendar.Date{year: 0, month: 0, day: -2}
        fmt.Println(date)
}
```

Change field names to match.

Change field names to match.

Accessing unexported fields through exported methods

As you might expect, now that we've converted the fields of `Date` to unexported, trying to access them from the `main` package results in compile errors. This is true both when we're trying to set the field values directly, and when using them in a struct literal.

Can't access fields directly

```
Shell Edit View Window Help
$ cd temp
$ go run main.go
./main.go:10:6: date.year undefined (cannot refer to unexported field or method year)
./main.go:11:6: date.month undefined (cannot refer to unexported field or method month)
./main.go:12:6: date.day undefined (cannot refer to unexported field or method day)
./main.go:15:27: unknown field 'year' in struct literal of type calendar.Date
./main.go:15:37: unknown field 'month' in struct literal of type calendar.Date
./main.go:15:45: unknown field 'day' in struct literal of type calendar.Date
```

But we can still access the fields indirectly. *Unexported* variables, struct fields, functions, methods, and the like can still be accessed by *exported* functions and methods in the same package. So when code in the `main` package calls the exported `SetYear` method on a `Date` value, `SetYear` can update the `Date`'s `year` struct field, even though it's unexported. The exported `SetMonth` method can update the unexported `month` field. And so on.

If we modify *main.go* to use the setter methods, we'll be able to update a `Date` value's fields:

main.go

```
package main

import (
        "fmt"
        "github.com/headfirstgo/calendar"
        "log"
)

func main() {
        date := calendar.Date{}
        err := date.SetYear(2019)
        if err != nil {
                log.Fatal(err)
        }
        err = date.SetMonth(5)
        if err != nil {
                log.Fatal(err)
        }
        err = date.SetDay(27)
        if err != nil {
                log.Fatal(err)
        }
        fmt.Println(date)
}
```

Use the setter method.

Use the setter method.

Use the setter method.

You can update fields via the setter methods!

```
Shell Edit View Window Help
$ cd temp
$ go run main.go
{2019 5 27}
```

Unexported variables, struct fields, functions, and methods can still be accessed by exported functions and methods in the same package.

Accessing unexported fields through exported methods (continued)

If we update *main.go* to call `SetYear` with an invalid value, we'll get an error when we run it:

main.go

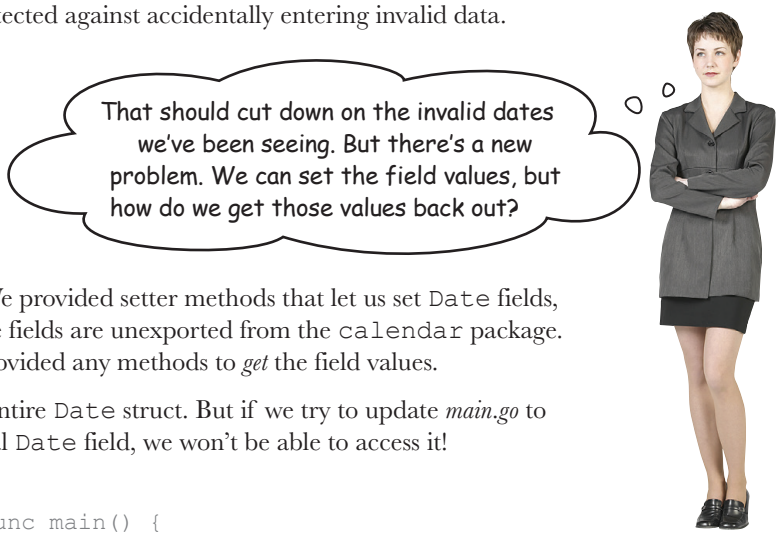
```
func main() {
        date := calendar.Date{}
        err := date.SetYear(0)
        if err != nil {
                log.Fatal(err)
        }
        fmt.Println(date)
}
```

Call the setter method with an invalid value.

Invalid values get reported!

```
Shell Edit View Window Help
$ cd temp
$ go run main.go
2018/03/23 19:20:17 invalid year
exit status 1
```

Now that a `Date` value's fields can only be updated via its setter methods, programs are protected against accidentally entering invalid data.



Ah, that's right. We provided setter methods that let us set `Date` fields, even though those fields are unexported from the `calendar` package. But we haven't provided any methods to *get* the field values.

We can print an entire `Date` struct. But if we try to update *main.go* to print an individual `Date` field, we won't be able to access it!

main.go

```
func main() {
        date := calendar.Date{}
        err := date.SetYear(2019)
        if err != nil {
                log.Fatal(err)
        }
        fmt.Println(date.year)
}
```

Set to a valid year.

Try to print the year field.

Get an error, because the field is unexported!

```
Shell Edit View Window Help
$ cd temp
$ go run main.go
# command-line-arguments
./main.go:16:18: date.year undefined
(cannot refer to unexported field or method year)
```

Getter methods

As we've seen, methods whose main purpose is to *set* the value of a struct field or variable are called *setter methods*. And, as you might expect, methods whose main purpose is to *get* the value of a struct field or variable are called **getter methods**.

Compared to the setter methods, adding getter methods to the `Date` type will be easy. They don't need to do anything except return the field value when they're called.

By convention, a getter method's name should be the same as the name of the field or variable it accesses. (Of course, if you want the method to be exported, its name will need to start with a capital letter.) So `Date` will need a `Year` method to access the `year` field, a `Month` method for the `month` field, and a `Day` method for the `day` field.

Getter methods don't need to modify the receiver at all, so we *could* use a direct `Date` value as a receiver. But if any method on a type takes a pointer receiver, convention says that they *all* should, for consistency's sake. Since we have to use a pointer receiver for our setter methods, we use a pointer for the getter methods as well.

With the changes to *date.go* complete, we can update *main.go* to set all the `Date` fields, then use the getter methods to print them all out.

main.go

```
// Package, import statements omitted
func main() {
	date := calendar.Date{}
	err := date.SetYear(2019)
	if err != nil {
		log.Fatal(err)
	}
	err = date.SetMonth(5)
	if err != nil {
		log.Fatal(err)
	}
	err = date.SetDay(27)
	if err != nil {
		log.Fatal(err)
	}
	fmt.Println(date.Year())
	fmt.Println(date.Month())
	fmt.Println(date.Day())
}
```

date.go

```
package calendar

import "errors"

type Date struct {
	year  int
	month int
	day   int
}

func (d *Date) Year() int {
	return d.year
}
func (d *Date) Month() int {
	return d.month
}
func (d *Date) Day() int {
	return d.day
}
// Setter methods omitted
```

Use a pointer receiver type for consistency with the setter methods.

Same name as the field (but capitalized so it's exported)

Return the field value.

```
Shell Edit View Window Help
$ cd temp
$ go run main.go
2019
5
27
```

Values returned from getter methods

Encapsulation

The practice of hiding data in one part of a program from code in another part is known as **encapsulation**, and it's not unique to Go. Encapsulation is valuable because it can be used to protect against invalid data (as we've seen). Also, you can change an encapsulated portion of a program without worrying about breaking other code that accesses it, because direct access isn't allowed.

Many other programming languages encapsulate data within classes. (Classes are a concept similar, but not identical, to a Go type.) In Go, data is encapsulated within packages, using unexported variables, struct fields, functions, or methods.

Encapsulation is used far more frequently in other languages than it is in Go. In some languages it's conventional to define getters and setters for every field, even when accessing those fields directly would work just as well. Go developers generally only rely on encapsulation when it's necessary, such as when field data needs to be validated by setter methods. In Go, if you don't see a need to encapsulate a field, it's generally okay to export it and allow direct access to it.

there are no Dumb Questions

Q: Many other languages don't allow access to encapsulated values outside of the class where they're defined. Is it safe for Go to allow other code in the same package to access unexported fields?

A: Generally, all the code in a package is the work of a single developer (or group of developers). All the code in a package generally has a similar purpose, as well. The authors of code within the same package are most likely to need access to unexported data, and they're also likely to only use that data in valid ways. So, yes, sharing unexported data with the rest of the package is generally safe.

Code *outside* the package is likely to be written by *other* developers, but that's okay because the unexported fields are hidden from them, so they can't accidentally change their values to something invalid.

Q: I've seen other languages where the name of every getter method started with "`Get`", as in `GetName`, `GetCity`, and so on. Can I do that in Go?

A: The Go language will allow you to do that, but you shouldn't. The Go community has decided on a convention of leaving the `Get` prefix off of getter method names. Including it would only lead to confusion for your fellow developers!

Go still uses a `Set` prefix for setter methods, just like many other languages, because it's needed to distinguish setter method names from getter method names for the same field.



Bear with us; we'll need two pages to fit all the code for this exercise...
Fill in the blanks to make the following changes to the `Coordinates` type:

- Update its fields so they're unexported.
- Add getter methods for each field. (Be sure to follow the convention: a getter method's name should be the same as the name of the field it accesses, with capitalization if the method needs to be exported.)
- Add validation to the setter methods. `SetLatitude` should return an error if the passed-in value is less than `-90` or greater than `90`. `SetLongitude` should return an error if the new value is less than `-180` or greater than `180`.


coordinates.go

```
package geo

import "errors"

type Coordinates struct {
        _________  float64
        _________  float64
}

func (c *Coordinates) _________() _______ {
        return c.latitude
}

func (c *Coordinates) _________() _______ {
        return c.longitude
}

func (c *Coordinates) SetLatitude(latitude float64) _____ {
        if latitude < -90 || latitude > 90 {
                return __________("invalid latitude")
        }
        c.latitude = latitude
        return ___
}
func (c *Coordinates) SetLongitude(longitude float64) _____ {
        if longitude < -180 || longitude > 180 {
                return __________("invalid longitude")
        }
        c.longitude = longitude
        return ___
}
```



Exercise (Continued)

Next, update the `main` package code to make use of the revised `Coordinates` type.

- For each call to a setter method, store the `error` return value.
- If the `error` is not `nil`, use the `log.Fatal` function to log the error message and exit.
- If there were no errors setting the fields, call both getter methods to print the field values.

The completed code should produce the output shown when it runs. (The call to `SetLatitude` should be successful, but we're passing an invalid value to `SetLongitude`, so it should log an error and exit at that point.)




```
package main

import (
        "fmt"
        "geo"
        "log"
)

func main() {
        coordinates := geo.Coordinates{}
        ___ := coordinates.SetLatitude(37.42)
        if err != ___ {
                log.Fatal(err)
        }
        err = coordinates.SetLongitude(-1122.08)
        if err != ___ {
                log.Fatal(err)
        }
        fmt.Println(coordinates._________())
        fmt.Println(coordinates._________())
}
```

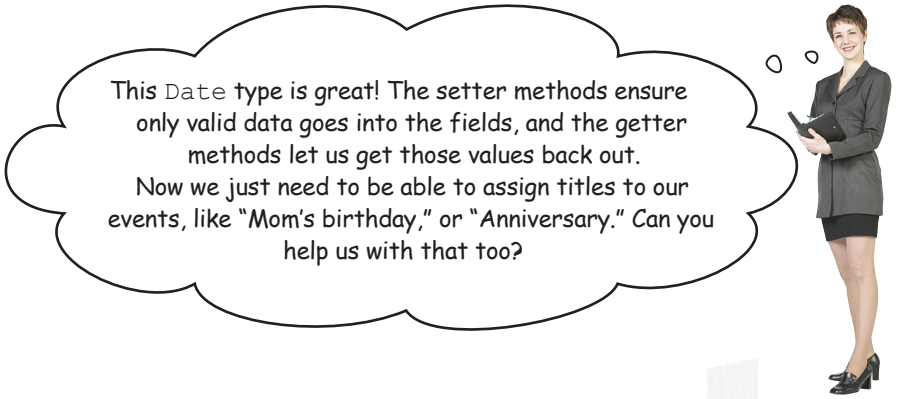
(An invalid value!)

Output

```
2018/03/23 20:12:49 invalid longitude
exit status 1
```

Answers on page 318.

Embedding the Date type in an Event type



That shouldn't take much work. Remember how we embedded an `Address` struct type within two other struct types back in Chapter 8?

The `Address` type was considered "embedded" because we used an anonymous field (a field with no name, just a type) in the outer struct to store it. This caused the fields of `Address` to be promoted to the outer struct, allowing us to access fields of the inner struct as if they belonged to the outer struct.

```
package magazine

type Subscriber struct {
        Name   string
        Rate   float64
        Active bool
        Address
}

type Employee struct {
        Name   string
        Salary float64
        Address
}

type Address struct {
        // Fields omitted
}
```

Set the fields of the Address as if they were defined on Subscriber.

```
subscriber.Street = "123 Oak St"
subscriber.City = "Omaha"
subscriber.State = "NE"
subscriber.PostalCode = "68111"
```

your workspace > src > github.com > headfirstgo > calendar > event.go

Since that strategy worked so well before, let's define an `Event` type that embeds a `Date` with an anonymous field.

Create another file within the `calendar` package folder, named *event.go*. (We could put it within the existing *date.go* field, but this organizes things a bit more neatly.) Within that file, define an `Event` type with two fields: a `Title` field with a type of `string`, and an anonymous `Date` field.

```
package calendar

type Event struct {
        Title string
        Date
}
```

Embed a Date using an anonymous field.

Unexported fields don't get promoted

Embedding a `Date` in the `Event` type will *not* cause the `Date` fields to be promoted to the `Event`, though. The `Date` fields are unexported, and Go doesn't promote unexported fields to the enclosing type. That makes sense; we made sure the fields were encapsulated so they can only be accessed through setter and getter methods, and we don't want that encapsulation to be circumvented through field promotion.

event.go

```
package calendar

type Event struct {
        Title string
        Date
}
```

Embedded using an anonymous field → `Date`

In our `main` package, if we try to set the `month` field of a `Date` through its enclosing `Event`, we'll get an error:

main.go

```
package main

import "github.com/headfirstgo/calendar"

func main() {
        event := calendar.Event{}
        event.month = 5
}
```

Unexported Date fields aren't promoted to the Event!

```
event.month undefined (type calendar.Event has no field or method month)
```

Error

And, of course, using dot operator chaining to retrieve the `Date` field and then access fields on it directly won't work, either. You can't access a `Date` value's unexported fields when it's by itself, and you can't access its unexported fields when it's part of an `Event`, either.

main.go

```
func main() {
        event := calendar.Event{}
        event.Date.year = 2019
}
```

Can't access Date fields directly on the Date!

```
event.Date.year undefined (cannot refer to unexported field or method year)
```

Error

So does that mean we won't be able to access the fields of the `Date` type, if it's embedded within the `Event` type? Don't worry; there's another way!

Exported methods get promoted just like fields

If you embed a type with exported methods within a struct type, its methods will be promoted to the outer type, meaning you can call the methods as if they were defined on the outer type. (Remember how embedding one struct type within another causes the inner struct's fields to be promoted to the outer struct? This is the same idea, but with methods instead of fields.)

Here's a package that defines two types. `MyType` is a struct type and it embeds a second type, `EmbeddedType`, as an anonymous field.

```
package mypackage

import "fmt"

type MyType struct {
        EmbeddedType
}

type EmbeddedType string

func (e EmbeddedType) ExportedMethod() {
        fmt.Println("Hi from ExportedMethod on EmbeddedType")
}

func (e EmbeddedType) unexportedMethod() {
}
```

These types are in their own package.

Declare MyType as a struct type.

EmbeddedType is embedded in MyType.

Declare a type to embed (doesn't matter whether it's a struct).

This method will be promoted to MyType.

This method will not be promoted.

Because `EmbeddedType` defines an exported method (named `ExportedMethod`), that method is promoted to `MyType`, and can be called on `MyType` values.

```
package main

import "mypackage"

func main() {
        value := mypackage.MyType{}
        value.ExportedMethod()
}
```

Call the method that was promoted from EmbeddedType.

```
Hi from ExportedMethod on EmbeddedType
```

As with unexported fields, unexported methods are *not* promoted. You'll get an error if you try to call one.

```
value.unexportedMethod()
```

Attempt to call unexported method.

Error

```
value.unexportedMethod undefined (type mypackage.MyType
has no field or method unexportedMethod)
```

Exported methods get promoted just like fields (continued)

Our `Date` fields weren't promoted to the `Event` type, because they're unexported. But the getter and setter methods on `Date` *are* exported, and they *do* get promoted to the `Event` type!

That means we can create an `Event` value, and then call the getter and setter methods for the `Date` directly on the `Event`. That's just what we do in the updated *main.go* code below. As always, the exported methods are able to access the unexported `Date` fields for us.

main.go

```
package main

import (
	"fmt"
	"github.com/headfirstgo/calendar"
	"log"
)

func main() {
	event := calendar.Event{}
	err := event.SetYear(2019)
	if err != nil {
		log.Fatal(err)
	}
	err = event.SetMonth(5)
	if err != nil {
		log.Fatal(err)
	}
	err = event.SetDay(27)
	if err != nil {
		log.Fatal(err)
	}
	fmt.Println(event.Year())
	fmt.Println(event.Month())
	fmt.Println(event.Day())
}
```

This setter method for Date has been promoted to Event.

This setter method for Date has been promoted to Event.

This setter method for Date has been promoted to Event.

These getter methods for Date have been promoted to Event.

```
2019
5
27
```

And if you prefer to use dot operator chaining to call methods on the `Date` value directly, you can do that too:

Get the Event's Date field, then call getter methods on it.

```
fmt.Println(event.Date.Year())
fmt.Println(event.Date.Month())
fmt.Println(event.Date.Day())
```

```
2019
5
27
```

Encapsulating the Event Title field

Because the `Event` struct's `Title` field is exported, we can still access it directly:

main.go

```
// Package, imports omitted
func main() {
        event := calendar.Event{}
        event.Title = "Mom's birthday"
        fmt.Println(event.Title)
}
```

```
Mom's birthday
```

event.go

```
package calendar

type Event struct {
        Title string
        Date
}
```

Exported field → `Title`

This exposes us to the same sort of issues that we had with the `Date` fields, though. For example, there's no limit on the length of the `Title` string:

main.go

```
func main() {
        event := calendar.Event{}
        event.Title = "An extremely long title that is impractical to print"
        fmt.Println(event.Title)
}
```

```
An extremely long title that is impractical to print
```

It seems like a good idea to encapsulate the title field as well, so we can validate new values. Here's an update to the `Event` type that does so. We change the field's name to `title` so it's unexported, then add getter and setter methods. The `RuneCountInString` function from the `unicode/utf8` package is used to ensure there aren't too many runes (characters) in the string.

event.go

```
package calendar

import (
        "errors"
        "unicode/utf8"
)

type Event struct {
        title string
        Date
}

func (e *Event) Title() string {
        return e.title
}
func (e *Event) SetTitle(title string) error {
        if utf8.RuneCountInString(title) > 30 {
                return errors.New("invalid title")
        }
        e.title = title
        return nil
}
```

Add this package for creating error values.

Add this package so we can count the number of runes in a string.

Change to unexported.

Getter method

Must use pointer

Setter method

If the title has more than 30 characters, return an error.

Promoted methods live alongside the outer type's methods

Now that we've added setter and getter methods for the `title` field, our programs can report an error if a title longer than 30 characters is used. An attempt to set a 39-character title causes an error to be returned:


main.go

```
// Package, imports omitted
func main() {
        event := calendar.Event{}
        err := event.SetTitle("An extremely long and impractical title")
        if err != nil {
                log.Fatal(err)
        }
}
```

```
2018/03/23 20:44:17 invalid title
exit status 1
```

The `Event` type's `Title` and `SetTitle` methods live alongside the methods promoted from the embedded `Date` type. Importers of the `calendar` package can treat all the methods as if they belong to the `Event` type, without worrying about which type they're actually defined on.


main.go

```
// Package, imports omitted
func main() {
        event := calendar.Event{}
        err := event.SetTitle("Mom's birthday") <-- Defined on Event itself
        if err != nil {
                log.Fatal(err)
        }
        err = event.SetYear(2019) <-- Promoted from Date
        if err != nil {
                log.Fatal(err)
        }
        err = event.SetMonth(5) <-- Promoted from Date
        if err != nil {
                log.Fatal(err)
        }
        err = event.SetDay(27) <-- Promoted from Date
        if err != nil {
                log.Fatal(err)
        }
        fmt.Println(event.Title()) <-- Defined on Event itself
        fmt.Println(event.Year()) <-- Promoted from Date
        fmt.Println(event.Month()) <-- Promoted from Date
        fmt.Println(event.Day()) <-- Promoted from Date
}
```

```
Mom's birthday
2019
5
27
```

Our calendar package is complete!



Now we can call the `Title` and `SetTitle` methods directly on an `Event`, and call the methods to set a year, month, and day as if they belonged to the `Event`. They're actually defined on `Date`, but we don't have to worry about that. Our work here is done!

Method promotion allows you to easily use one type's methods as if they belonged to another. You can use this to compose types that combine the methods of several other types. This can help you keep your code clean, without sacrificing convenience!



Exercise

We completed the code for the `Coordinates` type in a previous exercise. You won't need to make any updates to it this time; it's just here for reference. On the next page, we're going to embed it in the `Landmark` type (which we also saw back in Chapter 8), so that its methods are promoted to `Landmark`.



coordinates.go

```
package geo

import "errors"

type Coordinates struct {
       latitude  float64
       longitude float64
}

func (c *Coordinates) Latitude() float64 {
       return c.latitude
}
func (c *Coordinates) Longitude() float64 {
       return c.longitude
}

func (c *Coordinates) SetLatitude(latitude float64) error {
       if latitude < -90 || latitude > 90 {
              return errors.New("invalid latitude")
       }
       c.latitude = latitude
       return nil
}

func (c *Coordinates) SetLongitude(longitude float64) error {
       if longitude < -180 || longitude > 180 {
              return errors.New("invalid longitude")
       }
       c.longitude = longitude
       return nil
}
```



Exercise (Continued)

Here's an update to the `Landmark` type. We want its name field to be encapsulated, accessible only by a `Name` getter method and a `SetName` setter method. `SetName` should return an error if its argument is an empty string, or set the name field and return a `nil` error otherwise. `Landmark` should also have an anonymous `Coordinates` field, so that the methods of `Coordinates` are promoted to `Landmark`.

Fill in the blanks to complete the code for the `Landmark` type.



landmark.go

```
package geo

import "errors"

type Landmark struct {
        _____ string
        ___________
}

func (l *Landmark) _____() string {
        return l.name
}

func (l *Landmark) _______(name string) error {
        if name == "" {
                return errors.New("invalid name")
        }
        l.name = name
        return nil
}
```

If the blanks in the code for `Landmark` are completed correctly, the code in the `main` package should run and produce the output shown.



main.go

```
package main
// Imports omitted
func main() {
        location := geo.Landmark{}
        err := location.SetName("The Googleplex")
        if err != nil {
                log.Fatal(err)
        }
        err = location.SetLatitude(37.42)
        if err != nil {
                log.Fatal(err)
        }
        err = location.SetLongitude(-122.08)
        if err != nil {
                log.Fatal(err)
        }
        fmt.Println(location.Name())
        fmt.Println(location.Latitude())
        fmt.Println(location.Longitude())
}
```

Output

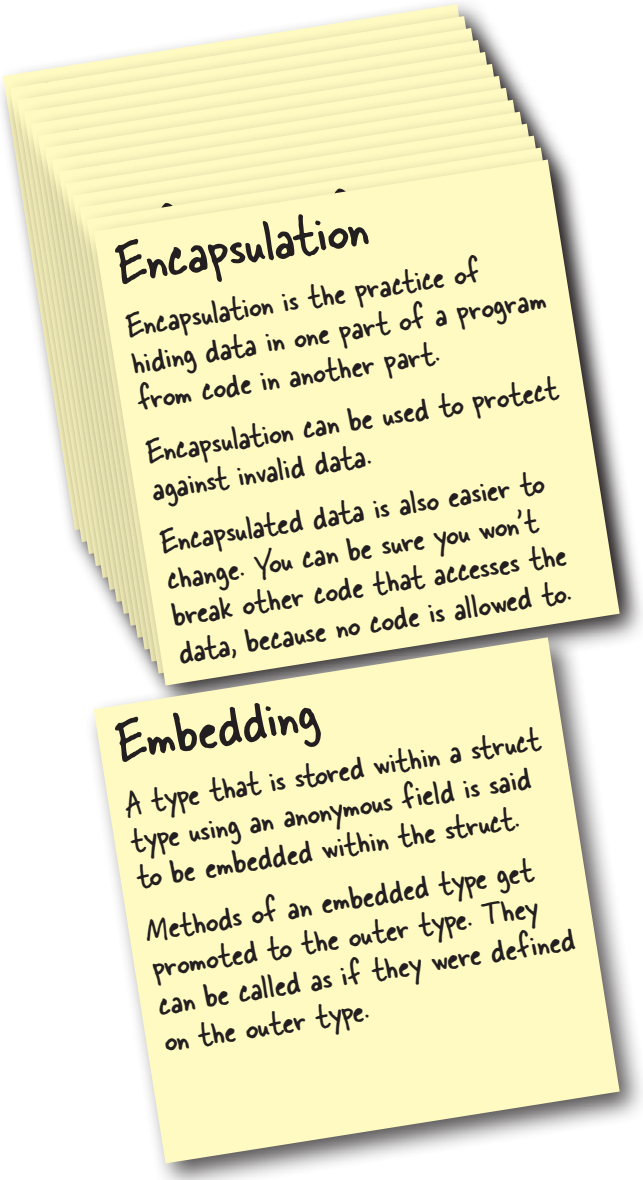
```
The Googleplex
37.42
-122.08
```

Answers on page 320.



Your Go Toolbox

That's it for Chapter 10! You've added encapsulation and embedding to your toolbox.



BULLET POINTS

- In Go, data is encapsulated within packages, using unexported package variables or struct fields.
- Unexported variables, struct fields, functions, methods, and the like can still be accessed by exported functions and methods defined in the same package.
- The practice of ensuring that data is valid before accepting it is known as **data validation**.
- A method that is primarily used to set the value of an encapsulated field is known as a **setter method**. Setter methods often include validation logic, to ensure the new value being provided is valid.
- Since setter methods need to modify their receiver, their receiver parameter should have a pointer type.
- It's conventional for setter method names to be in the form `SetX` where *X* is the name of the field being set.
- A method that is primarily used to get the value of an encapsulated field is known as a **getter method**.
- It's conventional for getter method names to be in the form *X* where *X* is the name of the field being set. Some other programming languages favor the form `GetX` for getter method names, but you should *not* use that form in Go.
- Methods defined on an outer struct type live alongside methods promoted from an embedded type.
- An embedded type's unexported methods don't get promoted to the outer type.



We need to add setter methods to the `Coordinates` type for each of its fields. Fill in the blanks in the *coordinates.go* file below, so that the code in *main.go* will run and produce the output shown.

```
package geo

type Coordinates struct {
        Latitude  float64
        Longitude float64
}

func (c *Coordinates) SetLatitude(latitude float64) {
        c.Latitude = latitude
}

func (c *Coordinates) SetLongitude(longitude float64) {
        c.Longitude = longitude
}
```

Must use a pointer type so we can modify the receiver.

Must use a pointer type so we can modify the receiver.

coordinates.go

```
package main

import (
        "fmt"
        "geo"
)

func main() {
        coordinates := geo.Coordinates{}
        coordinates.SetLatitude(37.42)
        coordinates.SetLongitude(-122.08)
        fmt.Println(coordinates)
}
```

main.go

Output

```
{37.42 -122.08}
```



Your goal with updating this code was to encapsulate the fields of the `Coordinates` type and add validation to its setter methods.

- Update the fields of `Coordinates` so they're unexported.
- Add getter methods for each field.
- Add validation to the setter methods. `SetLatitude` should return an error if the passed-in value is less than `-90` or greater than `90`. `SetLongitude` should return an error if the new value is less than `-180` or greater than `180`.



coordinates.go

```
package geo

import "errors"

type Coordinates struct {
        latitude  float64
        longitude float64
}
func (c *Coordinates) Latitude() float64 {
        return c.latitude
}
func (c *Coordinates) Longitude() float64 {
        return c.longitude
}
func (c *Coordinates) SetLatitude(latitude float64) error {
        if latitude < -90 || latitude > 90 {
                return errors.New("invalid latitude")
        }
        c.latitude = latitude
        return nil
}
func (c *Coordinates) SetLongitude(longitude float64) error {
        if longitude < -180 || longitude > 180 {
                return errors.New("invalid longitude")
        }
        c.longitude = longitude
        return nil
}
```

Fields should be unexported.

Getter method name should be same as field, but capitalized.

Same type as the field

Getter method name should be same as field, but capitalized.

Same type as the field

Need to return error type

Return a new error value.

Return nil if no error.

Need to return error type

Return a new error value.

Return nil if no error.



Exercise Solution (Continued)

Your next task was to update the `main` package code to make use of the revised `Coordinates` type.

- For each call to a setter method, store the `error` return value.
- If the `error` is not `nil`, use the `log.Fatal` function to log the error message and exit.
- If there were no errors setting the fields, call both getter methods to print the field values.

The call to `SetLatitude` below is successful, but we're passing an invalid value to `SetLongitude`, so it logs an error and exits at that point.




```
package main

import (
        "fmt"
        "geo"
        "log"
)

func main() {
        coordinates := geo.Coordinates{}
        err := coordinates.SetLatitude(37.42)
        if err != nil {
                log.Fatal(err)
        }
        err = coordinates.SetLongitude(-1122.08)
        if err != nil {
                log.Fatal(err)
        }
        fmt.Println(coordinates.Latitude())
        fmt.Println(coordinates.Longitude())
}
```

Store the returned error value.

If there was an error, log it and exit.

(An invalid value!)

If there was an error, log it and exit.

Call the getter methods.

Output

```
2018/03/23 20:12:49 invalid longitude
exit status 1
```



Here's an update to the `Landmark` type (which we also saw in Chapter 8). We want its name field to be encapsulated, accessible only by getter and setter methods. The `SetName` method should return an error if its argument is an empty string, or set the name field and return a `nil` error otherwise. `Landmark` should also have an anonymous `Coordinates` field, so that the methods of `Coordinates` are promoted to `Landmark`.

landmark.go

```
package geo

import "errors"

type Landmark struct {
        name string
        Coordinates
}

func (l *Landmark) Name() string {
        return l.name
}

func (l *Landmark) SetName(name string) error {
        if name == "" {
                return errors.New("invalid name")
        }
        l.name = name
        return nil
}
```

Ensure "name" field is unexported so it's encapsulated.

Embed using anonymous field.

Same name as field (but exported)

Same name as field, but with "Set" prefix

main.go

```
package main
// Imports omitted
func main() {
        location := geo.Landmark{}
        err := location.SetName("The Googleplex")
        if err != nil {
                log.Fatal(err)
        }
        err = location.SetLatitude(37.42)
        if err != nil {
                log.Fatal(err)
        }
        err = location.SetLongitude(-122.08)
        if err != nil {
                log.Fatal(err)
        }
        fmt.Println(location.Name())
        fmt.Println(location.Latitude())
        fmt.Println(location.Longitude())
}
```

Create Landmark value.

Defined on Landmark itself

Promoted from Coordinates

Promoted from Coordinates

Defined on Landmark

Promoted from Coordinates

Output

```
The Googleplex
37.42
-122.08
```

11 what can you do?

Interfaces



Sometimes you don't care about the particular type of a value.

You don't care about what it *is*. You just need to know that it will be able to *do* certain things. That you'll be able to call *certain methods* on it. You don't care whether you have a `Pen` or a `Pencil`, you just need something with a `Draw` method. You don't care whether you have a `Car` or a `Boat`, you just need something with a `Steer` method.

That's what Go **interfaces** accomplish. They let you define variables and function parameters that will hold *any* type, as long as that type defines certain methods.

Two different types that have the same methods

Remember audio tape recorders? (We suppose some of you will be too young.) They were great, though. They let you easily record all your favorite songs together on a single tape, even if they were by different artists. Of course, the recorders were usually too bulky to carry around with you. If you wanted to take your tapes on the go, you needed a separate, battery-powered tape player. Those usually didn't have recording capabilities. Ah, but it was so great making custom mixtapes and sharing them with your friends!



We're so overwhelmed with nostalgia that we've created a `gadget` package to help us reminisce. It includes a type that simulates a tape recorder, and another type that simulates a tape player.



The `TapePlayer` type has a `Play` method to simulate playing a song, and a `Stop` method to stop the virtual playback.

```
package gadget

import "fmt"

type TapePlayer struct {
       Batteries string
}
func (t TapePlayer) Play(song string) {
       fmt.Println("Playing", song)
}
func (t TapePlayer) Stop() {
       fmt.Println("Stopped!")
}
```

The `TapeRecorder` type also has `Play` and `Stop` methods, and a `Record` method as well.

```
type TapeRecorder struct {
       Microphones int
}
func (t TapeRecorder) Play(song string) {
       fmt.Println("Playing", song)
}
func (t TapeRecorder) Record() {
       fmt.Println("Recording")
}
func (t TapeRecorder) Stop() {
       fmt.Println("Stopped!")
}
```

Has a Play method just like TapePlayer's

Has a Stop method just like TapePlayer's

A method parameter that can only accept one type

Here's a sample program that uses the `gadget` package. We define a `playList` function that takes a `TapePlayer` value, and a slice of song titles to play on it. The function loops over each title in the slice, and passes it to the `TapePlayer`'s `Play` method. When it's done playing the list, it calls `Stop` on the `TapePlayer`.

Then, in the `main` method, all we have to do is create the `TapePlayer` and the slice of song titles, and pass them to `playList`.

```
package main

import "github.com/headfirstgo/gadget"

func playList(device gadget.TapePlayer, songs []string) {
        for _, song := range songs {
                device.Play(song)
        }
        device.Stop()
}

func main() {
        player := gadget.TapePlayer{}
        mixtape := []string{"Jessie's Girl", "Whip It", "9 to 5"}
        playList(player, mixtape)
}
```

Import our package.

Loop over each song.

Play the current song.

Stop the player once we're done.

Create a TapePlayer.

Create a slice of song titles.

Play the songs using the TapePlayer.

```
Playing Jessie's Girl
Playing Whip It
Playing 9 to 5
Stopped!
```

The `playList` function works great with a `TapePlayer` value. You might hope that it would work with a `TapeRecorder` as well. (After all, a tape recorder is basically just a tape player with an extra record function.) But `playList`'s first parameter has a type of `TapePlayer`. Try to pass it an argument of any other type, and you'll get a compile error:

```
func main() {
        player := gadget.TapeRecorder{}
        mixtape := []string{"Jessie's Girl", "Whip It", "9 to 5"}
        playList(player, mixtape)
}
```

Create a TapeRecorder instead of a TapePlayer.

Pass the TapeRecorder to playList.

Error

```
cannot use player (type gadget.TapeRecorder)
as type gadget.TapePlayer in argument to playList
```

A method parameter that can only accept one type (continued)



That's too bad... All the `playList` function really needs is a value whose type defines `Play` and `Stop` methods. Both `TapePlayer` and `TapeRecorder` have those!

```
func playList(device gadget.TapePlayer, songs []string) {
        for _, song := range songs {
                device.Play(song)
        }
        device.Stop()
}

type TapePlayer struct {
        Batteries string
}
func (t TapePlayer) Play(song string) {
        fmt.Println("Playing", song)
}
func (t TapePlayer) Stop() {
        fmt.Println("Stopped!")
}

type TapeRecorder struct {
        Microphones int
}
func (t TapeRecorder) Play(song string) {
        fmt.Println("Playing", song)
}
func (t TapeRecorder) Record() {
        fmt.Println("Recording")
}
func (t TapeRecorder) Stop() {
        fmt.Println("Stopped!")
}
```

Needs the value to have a Play method with a string parameter

Needs the value to have a Stop method with no parameters

TapePlayer has a Play method with a string parameter.

TapePlayer has a Stop method with no parameters.

TapeRecorder also has a Play method with a string parameter.

TapeRecorder also has a Stop method with no parameters.

In this case, it does seem like the Go language's type safety is getting in our way, rather than helping us. The `TapeRecorder` type defines all the methods that the `playList` function needs, but we're being blocked from using it because `playList` only accepts `TapePlayer` values.

So what can we do? Write a second, nearly identical `playListWithRecorder` function that takes a `TapeRecorder` instead?

Actually, Go offers another way...

Interfaces

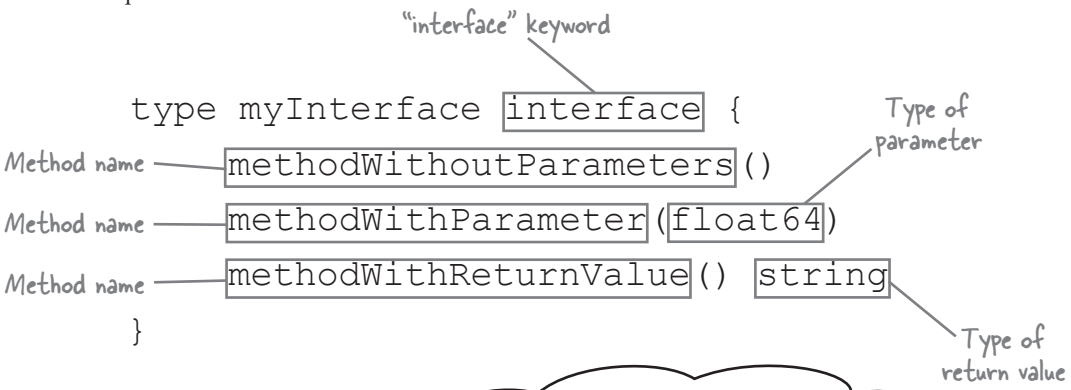
When you install a program on your computer, you usually expect the program to provide you with a way to interact with it. You expect a word processor to give you a place to type text. You expect a backup program to give you a way to select which files to save. You expect a spreadsheet to give you a way to insert columns and rows for data. The set of controls a program provides you so you can interact with it is often called its *interface*.

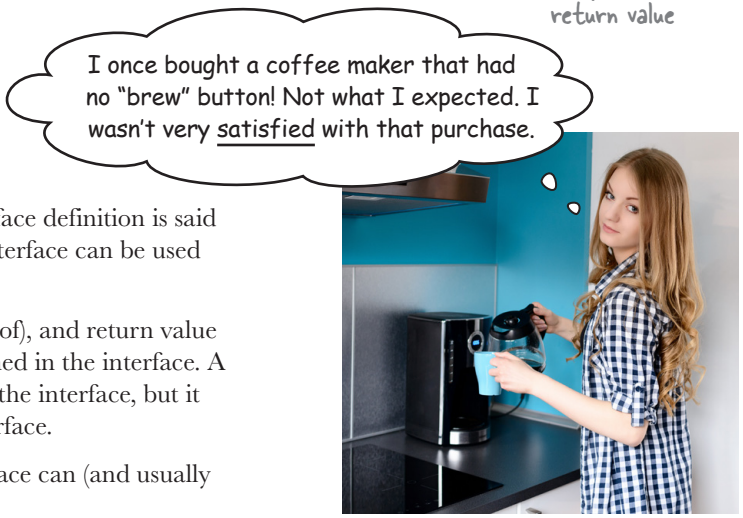
An interface is a set of methods that certain values are expected to have.

Whether you've actually thought about it or not, you probably expect Go values to provide you with a way to interact with them, too. What's the most common way to interact with a Go value? Through its methods.

In Go, an **interface** is defined as a set of methods that certain values are expected to have. You can think of an interface as a set of actions you need a type to be able to perform.

You define an interface type using the `interface` keyword, followed by curly braces containing a list of method names, along with any parameters or return values the methods are expected to have.





Any type that has all the methods listed in an interface definition is said to **satisfy** that interface. A type that satisfies an interface can be used anywhere that interface is called for.

The method names, parameter types (or lack thereof), and return value types (or lack thereof) all need to match those defined in the interface. A type can have methods *in addition* to those listed in the interface, but it mustn't be *missing* any, or it doesn't satisfy that interface.

A type can satisfy multiple interfaces, and an interface can (and usually should) have multiple types that satisfy it.

Defining a type that satisfies an interface

The code below sets up a quick experimental package, named `mypkg`. It defines an interface type named `MyInterface` with three methods. Then it defines a type named `MyType` that satisfies `MyInterface`.

There are three methods required to satisfy `MyInterface`: a `MethodWithoutParameters` method, a `MethodWithParameter` method that takes a `float64` parameter, and a `MethodWithReturnValue` method that returns a `string`.

Then we declare another type, `MyType`. The underlying type of `MyType` doesn't matter in this example; we just used `int`. We define all the methods on `MyType` that it needs to satisfy `MyInterface`, plus one extra method that isn't part of the interface.




```
package mypkg

import "fmt"

type MyInterface interface {
       MethodWithoutParameters()
       MethodWithParameter(float64)
       MethodWithReturnValue() string
}

type MyType int

func (m MyType) MethodWithoutParameters() {
       fmt.Println("MethodWithoutParameters called")
}
func (m MyType) MethodWithParameter(f float64) {
       fmt.Println("MethodWithParameter called with", f)
}
func (m MyType) MethodWithReturnValue() string {
       return "Hi from MethodWithReturnValue"
}
func (my MyType) MethodNotInInterface() {
       fmt.Println("MethodNotInInterface called")
}
```

Declare an interface type.

A type satisfies this interface if it has this method...

...aAnd this method (with a float64 parameter)...

...and this method (with a string return value).

Declare a type. We'll make it satisfy myInterface.

First required method

Second required method (with a float64 parameter)

Third required method (with a string return value)

A type can still satisfy an interface even if it has methods that aren't part of the interface.

Many other languages would require us to explicitly say that `MyType` satisfies `MyInterface`. But in Go, this happens *automatically*. If a type has all the methods declared in an interface, then it can be used anywhere that interface is required, with no further declarations needed.

Defining a type that satisfies an interface (continued)

Here's a quick program that will let us try `mypkg` out.

A variable declared with an interface type can hold any value whose type satisfies that interface. This code declares a `value` variable with `MyInterface` as its type, then creates a `MyType` value and assigns it to `value`. (Which is allowed, because `MyType` satisfies `MyInterface`.) Then we call all the methods on that value that are part of the interface.

```
package main

import (
        "fmt"
        "mypkg"
)

func main() {
        var value mypkg.MyInterface
        value = mypkg.MyType(5)
        value.MethodWithoutParameters()
        value.MethodWithParameter(127.3)
        fmt.Println(value.MethodWithReturnValue())
}
```

Declare a variable using the interface type.

Values of myType satisfy myInterface, so we can assign this value to a variable with a type of myInterface.

We can call any method that's part of myInterface.

```
MethodWithoutParameters called
MethodWithParameter called with 127.3
Hi from MethodWithReturnValue
```

Concrete types, interface types

All the types we've defined in previous chapters have been concrete types. A **concrete type** specifies not only what its values can *do* (what methods you can call on them), but also what they *are*: they specify the underlying type that holds the value's data.

Interface types don't describe what a value *is*: they don't say what its underlying type is, or how its data is stored. They only describe what a value can *do*: what methods it has.

Suppose you need to write down a quick note. In your desk drawer, you have values of several concrete types: `Pen`, `Pencil`, and `Marker`. Each of these concrete types defines a `Write` method, so you don't really care which type you grab. You just want a `WritingInstrument`: an interface type that is satisfied by any concrete type with a `Write` method.



Assign any type that satisfies the interface

When you have a variable with an interface type, it can hold values of any type that satisfies the interface.

Suppose we have `Whistle` and `Horn` types, each of which has a `MakeSound` method. We can create a `NoiseMaker` interface that represents any type with a `MakeSound` method. If we declare a `toy` variable with a type of `NoiseMaker`, we'll be able to assign either `Whistle` or `Horn` values to it. (*Or* any other type that we later declare, as long as it has a `MakeSound` method.)

We can then call the `MakeSound` method on any value assigned to the `toy` variable. Although we don't know exactly what concrete type the value in `toy` *is*, we know what it can *do*: make sounds. If its type didn't have a `MakeSound` method, then it wouldn't satisfy the `NoiseMaker` interface, and we wouldn't have been able to assign it to the variable.

```
package main

import "fmt"

type Whistle string

func (w Whistle) MakeSound() {
       fmt.Println("Tweet!")
}

type Horn string

func (h Horn) MakeSound() {
       fmt.Println("Honk!")
}

type NoiseMaker interface {
       MakeSound()
}

func main() {
       var toy NoiseMaker
       toy = Whistle("Toyco Canary")
       toy.MakeSound()
       toy = Horn("Toyco Blaster")
       toy.MakeSound()
}
```

Has a MakeSound method

Also has a MakeSound method

Represents any type with a MakeSound method

Declare a NoiseMaker variable

Assign a value of a type that satisfies NoiseMaker to the variable

Assign a value of another type that satisfies NoiseMaker to the variable.

```
Tweet!
Honk!
```

You can declare function parameters with interface types as well. (After all, function parameters are really just variables too.) If we declare a `play` function that takes a `NoiseMaker`, for example, then we can pass any value from a type with a `MakeSound` method to `play`:

```
func play(n NoiseMaker) {
       n.MakeSound()
}

func main() {
       play(Whistle("Toyco Canary"))
       play(Horn("Toyco Blaster"))
}
```

```
Tweet!
Honk!
```

You can only call methods defined as part of the interface

Once you assign a value to a variable (or method parameter) with an interface type, you can *only* call methods that are specified by the interface on it.

Suppose we created a `Robot` type, which in addition to a `MakeSound` method, also has a `Walk` method. We add a call to `Walk` in the `play` function, and pass a new `Robot` value to `play`.

But the code doesn't compile, saying that `NoiseMaker` values don't have a `Walk` method.

Why is that? `Robot` values *do* have a `Walk` method; the definition is right there!

But it's *not* a `Robot` value that we're passing to the `play` function; it's a `NoiseMaker`. What if we had passed a `Whistle` or `Horn` to `play` instead? Those don't have `Walk` methods!

When we have a variable of an interface type, the only methods we can be sure it has are the methods that are defined in the interface. And so those are the only methods Go allows you to call. (There *is* a way to get at the value's concrete type, so that you can call more specialized methods. We'll look at that shortly.)

```
package main

import "fmt"

type Whistle string

func (w Whistle) MakeSound() {
      fmt.Println("Tweet!")
}

type Horn string

func (h Horn) MakeSound() {
      fmt.Println("Honk!")
}

type Robot string

func (r Robot) MakeSound() {
      fmt.Println("Beep Boop")
}

func (r Robot) Walk() {
      fmt.Println("Powering legs")
}

type NoiseMaker interface {
      MakeSound()
}

func play(n NoiseMaker) {
      n.MakeSound()
      n.Walk()
}

func main() {
      play(Robot("Botco Ambler"))
}
```

Declare a new Robot type.

Robot satisfies the NoiseMaker interface.

An additional method

OK. Part of NoiseMaker interface.

Not OK! Not part of NoiseMaker!

Error

```
n.Walk undefined
(type NoiseMaker has no
field or method Walk)
```

Note that it *is* just fine to assign a type that *has* other methods to a variable with an interface type. As long as you don't actually call those other methods, everything will work.

```
func play(n NoiseMaker) {
      n.MakeSound()
}

func main() {
      play(Robot("Botco Ambler"))
}
```

Call only methods that are part of the interface.

```
Beep Boop
```



Breaking Stuff is Educational!

Here are a couple concrete types, `Fan` and `CoffeePot`. We also have an `Appliance` interface with a `TurnOn` method. `Fan` and `CoffeePot` both have `TurnOn` methods, so they both satisfy the `Appliance` interface.

That's why, in the `main` function, we're able to define an `Appliance` variable, and assign both `Fan` and `CoffeePot` variables to it.

Make one of the changes below and try to compile the code. Then undo your change and try the next one. See what happens!

```
type Appliance interface {
        TurnOn()
}

type Fan string
func (f Fan) TurnOn() {
        fmt.Println("Spinning")
}

type CoffeePot string
func (c CoffeePot) TurnOn() {
        fmt.Println("Powering up")
}
func (c CoffeePot) Brew() {
        fmt.Println("Heating Up")
}

func main() {
        var device Appliance
        device = Fan("Windco Breeze")
        device.TurnOn()
        device = CoffeePot("LuxBrew")
        device.TurnOn()
}
```

If you do this...	...the code will break because...
Call a method from the concrete type that isn't defined in the interface: `device.Brew()`	When you have a value in a variable with an interface type, you can only call methods defined as part of that interface, regardless of what methods the concrete type had.
Remove the method that satisfies the interface from a type: ~~`func (c CoffeePot) TurnOn() {`~~ ~~`        fmt.Println("Powering up")`~~ ~~`}`~~	If a type doesn't satisfy an interface, you can't assign values of that type to variables that use that interface as their type.

Add a new return value or parameter on the method that satisfies the interface:
`func (f Fan) TurnOn() error {`
`        fmt.Println("Spinning")`
`        return nil`
`}`
	If the number and types of all parameters and return values don't match between a concrete type's method definition and the method definition in the interface, then the concrete type does not satisfy the interface.

Fixing our playList function using an interface

Let's see if we can use an interface to allow our `playList` function to work with the `Play` and `Stop` methods on both of our concrete types: `TapePlayer` and `TapeRecorder`.

```
// TapePlayer type definition here
func (t TapePlayer) Play(song string) {
       fmt.Println("Playing", song)
}
func (t TapePlayer) Stop() {
       fmt.Println("Stopped!")
}
// TapeRecorder type definition here
func (t TapeRecorder) Play(song string) {
       fmt.Println("Playing", song)
}
func (t TapeRecorder) Record() {
       fmt.Println("Recording")
}
func (t TapeRecorder) Stop() {
       fmt.Println("Stopped!")
}
```

In our `main` package, we declare a `Player` interface. (We could define it in the `gadget` package instead, but defining the interface in the same package where we use it gives us more flexibility.) We specify that the interface requires both a `Play` method with a `string` parameter, and a `Stop` method with no parameters. This means that both the `TapePlayer` and `TapeRecorder` types will satisfy the `Player` interface.

We update the `playList` function to take any value that satisfies `Player` instead of `TapePlayer` specifically. We also change the type of the `player` variable from `TapePlayer` to `Player`. This allows us to assign either a `TapePlayer` or a `TapeRecorder` to `player`. We then pass values of both types to `playList`!

```
package main

import "github.com/headfirstgo/gadget"

type Player interface {
       Play(string)
       Stop()
}

func playList(device Player, songs []string) {
       for _, song := range songs {
              device.Play(song)
       }
       device.Stop()
}

func main() {
       mixtape := []string{"Jessie's Girl", "Whip It", "9 to 5"}
       var player Player = gadget.TapePlayer{}
       playList(player, mixtape)
       player = gadget.TapeRecorder{}
       playList(player, mixtape)
}
```

Define an interface type.

Require a Play method with a string parameter.

Also require a Stop method.

Accept any Player, not just a TapePlayer.

Update the variable to hold any Player.

Pass a TapePlayer to playList.

Pass a TapeRecorder to playList.

```
Playing Jessie's Girl
Playing Whip It
Playing 9 to 5
Stopped!
Playing Jessie's Girl
Playing Whip It
Playing 9 to 5
Stopped!
```



Watch it!

If a type declares methods with pointer receivers, then you'll only be able to use pointers to that type when assigning to interface variables.

The `toggle` *method on the* `Switch` *type below has to use a pointer receiver so it can modify the receiver.*

```
package main

import "fmt"

type Switch string
func (s *Switch) toggle() {
       if *s == "on" {
              *s = "off"
       } else {
              *s = "on"
       }
       fmt.Println(*s)
}

type Toggleable interface {
       toggle()
}

func main() {
       s := Switch("off")
       var t Toggleable = s
       t.toggle()
       t.toggle()
}
```

But that results in an error when we assign a `Switch` *value to a variable with the interface type* `Toggleable`*:*

```
Switch does not implement Toggleable
(toggle method has pointer receiver)
```

When Go decides whether a value satisfies an interface, pointer methods aren't included for direct values. But they are included for pointers. So the solution is to assign a pointer to a `Switch` *to the* `Toggleable` *variable, instead of a direct* `Switch` *value:*

```
var t Toggleable = &s
```

Assign a pointer instead.

Make that change, and the code should work correctly.

there are no Dumb Questions

Q: Should interface type names begin with a capital letter or a lowercase letter?

A: The rules for interface type names are the same as the rules for any other type. If the name begins with a lowercase letter, then the interface type will be *unexported* and will not be accessible outside the current package. Sometimes you won't need to use the interface you're declaring from other packages, so making it unexported is fine. But if you *do* want to use it in other packages, you'll need to start the interface type's name with a capital letter, so that it's exported.



Exercise

The code at the right defines `Car` and `Truck` types, each of which have `Accelerate`, `Brake`, and `Steer` methods. Fill in the blanks to add a `Vehicle` interface that includes those three methods, so that the code in the `main` function will compile and produce the output shown.

```
package main

import "fmt"

type Car string
func (c Car) Accelerate() {
       fmt.Println("Speeding up")
}
func (c Car) Brake() {
       fmt.Println("Stopping")
}
func (c Car) Steer(direction string) {
       fmt.Println("Turning", direction)
}

type Truck string
func (t Truck) Accelerate() {
       fmt.Println("Speeding up")
}
func (t Truck) Brake() {
       fmt.Println("Stopping")
}
func (t Truck) Steer(direction string) {
       fmt.Println("Turning", direction)
}
func (t Truck) LoadCargo(cargo string) {
       fmt.Println("Loading", cargo)
}
```

Your code here! →

```
_____
       _____
       _____
       _____
_
```

```
func main() {
       var vehicle Vehicle = Car("Toyoda Yarvic")
       vehicle.Accelerate()
       vehicle.Steer("left")

       vehicle = Truck("Fnord F180")
       vehicle.Brake()
       vehicle.Steer("right")
}
```

```
Speeding up
Turning left
Stopping
Turning right
```

→ Answers on page 348.

Type assertions

We've defined a new `TryOut` function that will let us test the various methods of our `TapePlayer` and `TapeRecorder` types. `TryOut` has a single parameter with the `Player` interface as its type, so that we can pass in either a `TapePlayer` or `TapeRecorder`.

Within `TryOut`, we call the `Play` and `Stop` methods, which are both part of the `Player` interface. We also call the `Record` method, which is *not* part of the `Player` interface, but *is* defined on the `TapeRecorder` type. We're only passing a `TapeRecorder` value to `TryOut` for now, so we should be fine, right?

Unfortunately, no. We saw earlier that if a value of a concrete type is assigned to a variable with an interface type (including function parameters), then you can only call methods on it that are part of that interface, regardless of what other methods the concrete type has. Within the `TryOut` function, we don't have a `TapeRecorder` value (the concrete type), we have a `Player` value (the interface type). And the `Player` interface doesn't have a `Record` method!

```
type Player interface {
        Play(string)
        Stop()
}

func TryOut(player Player) {
        player.Play("Test Track")
        player.Stop()
        player.Record()
}

func main() {
        TryOut(gadget.TapeRecorder{})
}
```

These are fine; they're part of the Player interface.

Not part of Player!

Pass a TapeRecorder (which satisfies Player) to the function.

Error

```
player.Record undefined (type Player
has no field or method Record)
```

We need a way to get the concrete type value (which *does* have a `Record` method) back.

Your first instinct might be to try a type conversion to convert the `Player` value to a `TapeRecorder` value. But type conversions aren't meant for use with interface types, so that generates an error. The error message suggests trying something else:

```
func TryOut(player Player) {
        player.Play("Test Track")
        player.Stop()
        recorder := gadget.TapeRecorder(player)
        recorder.Record()
}
```

A type conversion won't work!

Error

```
cannot convert player (type Player) to type
gadget.TapeRecorder: need type assertion
```

A "type assertion"? What's that?

Type assertions (continued)

When you have a value of a concrete type assigned to a variable with an interface type, a **type assertion** lets you get the concrete type back. It's *kind of* like a type conversion. Its syntax even looks like a cross between a method call and a type conversion. After an interface value, you type a dot, followed by a pair of parentheses with the concrete type. (Or rather, what you're *asserting* the value's concrete type is.)

```
var noiseMaker NoiseMaker = Robot("Botco Ambler")
var robot Robot = noiseMaker.(Robot)
```

Interface value

Asserted type

In plain language, the type assertion above says something like "I know this variable uses the interface type `NoiseMaker`, but I'm pretty sure *this* `NoiseMaker` is actually a `Robot`."

Once you've used a type assertion to get a value of a concrete type back, you can call methods on it that are defined on that type, but aren't part of the interface.

This code assigns a `Robot` to a `NoiseMaker` interface value. We're able to call `MakeSound` on the `NoiseMaker`, because it's part of the interface. But to call the `Walk` method, we need to use a type assertion to get a `Robot` value. Once we have a `Robot` (rather than a `NoiseMaker`), we can call `Walk` on it.

```
type Robot string
func (r Robot) MakeSound() {
       fmt.Println("Beep Boop")
}
func (r Robot) Walk() {
       fmt.Println("Powering legs")
}

type NoiseMaker interface {
       MakeSound()
}
func main() {
       var noiseMaker NoiseMaker = Robot("Botco Ambler")
       noiseMaker.MakeSound()
       var robot Robot = noiseMaker.(Robot)
       robot.Walk()
}
```

Define a variable with an interface type...

...and assign a value of a type that satisfies the interface.

Call a method that's part of the interface.

Convert back to the concrete type using a type assertion.

Call a method that's defined on the concrete type (not the interface).

```
Beep Boop
Powering legs
```

Type assertion failures

Previously, our `TryOut` function wasn't able to call the `Record` method on a `Player` value, because it's not part of the `Player` interface. Let's see if we can get this working using a type assertion.

Just like before, we pass a `TapeRecorder` to `TryOut`, where it gets assigned to a parameter that uses the `Player` interface as its type. We're able to call the `Play` and `Stop` methods on the `Player` value, because those are both part of the `Player` interface.

Then, we use a type assertion to convert the `Player` back to a `TapeRecorder`. And we call `Record` on the `TapeRecorder` value instead.

```
type Player interface {
        Play(string)
        Stop()
}

func TryOut(player Player) {
        player.Play("Test Track")
        player.Stop()
        recorder := player.(gadget.TapeRecorder)
        recorder.Record()
}

func main() {
        TryOut(gadget.TapeRecorder{})
}
```

Store the TapeRecorder value.

Use a type assertion to get a TapeRecorder value.

Call the method that's only defined on the concrete type.

```
Playing Test Track
Stopped!
Recording
```

Everything seems to be working great...with a `TapeRecorder`. But what happens if we try to pass a `TapePlayer` to `TryOut`? How well will that work, considering we have a type assertion that says the parameter to `TryOut` is actually a `TapeRecorder`?

```
func main() {
        TryOut(gadget.TapeRecorder{})
        TryOut(gadget.TapePlayer{})
}
```

Pass a TapePlayer as well...

Everything compiles successfully, but when we try to run it, we get a runtime panic! As you might expect, trying to assert that a `TapePlayer` is actually a `TapeRecorder` did not go well. (It's simply not true, after all.)

Panic!

```
Playing Test Track
Stopped!
Recording
Playing Test Track
Stopped!
panic: interface conversion: main.Player
is gadget.TapePlayer, not gadget.TapeRecorder
```

Avoiding panics when type assertions fail

If a type assertion is used in a context that expects only one return value, and the original type doesn't match the type in the assertion, the program will panic at runtime (*not* when compiling):

```
var player Player = gadget.TapePlayer{}
recorder := player.(gadget.TapeRecorder)
```

Assert that the original type is TapeRecorder, when it's actually TapePlayer...

Panic!

```
panic: interface conversion: main.Player
is gadget.TapePlayer, not gadget.TapeRecorder
```

If type assertions are used in a context where multiple return values are expected, they have a second, optional return value that indicates whether the assertion was successful or not. (And the assertion won't panic if it's unsuccessful.) The second value is a `bool`, and it will be `true` if the value's original type was the asserted type, or `false` if not. You can do whatever you want with this second return value, but by convention it's usually assigned to a variable named `ok`.

This is another place Go follows the "comma ok idiom" that we first saw when accessing maps in Chapter 7.

Here's an update to the above code that assigns the results of the type assertion to a variable for the concrete type's value, and a second `ok` variable. It uses the `ok` value in an `if` statement to determine whether it can safely call `Record` on the concrete value (because the `Player` value had an original type of `TapeRecorder`), or if it should skip doing so (because the `Player` had some other concrete value).

```
var player Player = gadget.TapePlayer{}
recorder, ok := player.(gadget.TapeRecorder)
if ok {
        recorder.Record()
} else {
        fmt.Println("Player was not a TapeRecorder")
}
```

Assign the second return value to a variable.

If the original type was TapeRecorder, call Record on the value.

Otherwise, report that the assertion failed.

```
Player was not a TapeRecorder
```

In this case, the concrete type was `TapePlayer`, not `TapeRecorder`, so the assertion is unsuccessful, and `ok` is `false`. The `if` statement's `else` clause runs, printing `Player was not a TapeRecorder`. A runtime panic is averted.

When using type assertions, if you're not absolutely sure which original type is behind the interface value, then you should use the optional `ok` value to handle cases where it's a different type than you expected, and avoid a runtime panic.

Testing TapePlayers and TapeRecorders using type assertions

Let's see if we can use what we've learned to fix our `TryOut` function for `TapePlayer` and `TapeRecorder` values. Instead of ignoring the second return value from our type assertion, we'll assign it to an `ok` variable. The `ok` variable will be `true` if the type assertion is successful (indicating the `recorder` variable holds a `TapeRecorder` value, ready for us to call `Record` on it), or `false` otherwise (indicating it's *not* safe to call `Record`). We wrap the call to the `Record` method in an `if` statement to ensure it's only called when the type assertion is successful.

```
type Player interface {
        Play(string)
        Stop()
}

func TryOut(player Player) {
        player.Play("Test Track")
        player.Stop()
        recorder, ok := player.(gadget.TapeRecorder)
        if ok {
                recorder.Record()
        }
}

func main() {
        TryOut(gadget.TapeRecorder{})
        TryOut(gadget.TapePlayer{})
}
```

Call the Record method only if the original value was a TapeRecorder.

Assign the second return value to a variable.

```
Playing Test Track
Stopped!
Recording
Playing Test Track
Stopped!
```

TapeRecorder passed in...

...type assertion succeeds, Record called.

TapePlayer passed in...

...type assertion does not succeed, Record not called.

As before, in our `main` function, we first call `TryOut` with a `TapeRecorder` value. `TryOut` takes the `Player` interface value it receives, and calls the `Play` and `Stop` methods on it. The assertion that the `Player` value's concrete type is `TapeRecorder` succeeds, and the `Record` method is called on the resulting `TapeRecorder` value.

Then, we call `TryOut` again with a `TapePlayer`. (This is the call that halted the program previously because the type assertion panicked.) `Play` and `Stop` are called, as before. The type assertion fails, because the `Player` value holds a `TapePlayer` and not a `TapeRecorder`. But because we're capturing the second return value in the `ok` value, the type assertion doesn't panic this time. It just sets `ok` to `false`, which causes the code in our `if` statement not to run, which causes `Record` not to be called. (Which is good, because `TapePlayer` values don't have a `Record` method.)

Thanks to type assertions, we've got our `TryOut` function working with both `TapeRecorder` and `TapePlayer` values!



Pool Puzzle

Updated code from our previous exercise is at the right. We're creating a `TryVehicle` method that calls all the methods from the `Vehicle` interface. Then, it should attempt a type assertion to get a concrete `Truck` value. If successful, it should call `LoadCargo` on the `Truck` value.

Your **job** is to take code snippets from the pool and place them into the blank lines in this code. **Don't** use the same snippet more than once, and you won't need to use all the snippets. Your **goal** is to make a program that will run and produce the output shown.

```
type Truck string
func (t Truck) Accelerate() {
       fmt.Println("Speeding up")
}
func (t Truck) Brake() {
       fmt.Println("Stopping")
}
func (t Truck) Steer(direction string) {
       fmt.Println("Turning", direction)
}
func (t Truck) LoadCargo(cargo string) {
       fmt.Println("Loading", cargo)
}

type Vehicle interface {
       Accelerate()
       Brake()
       Steer(string)
}

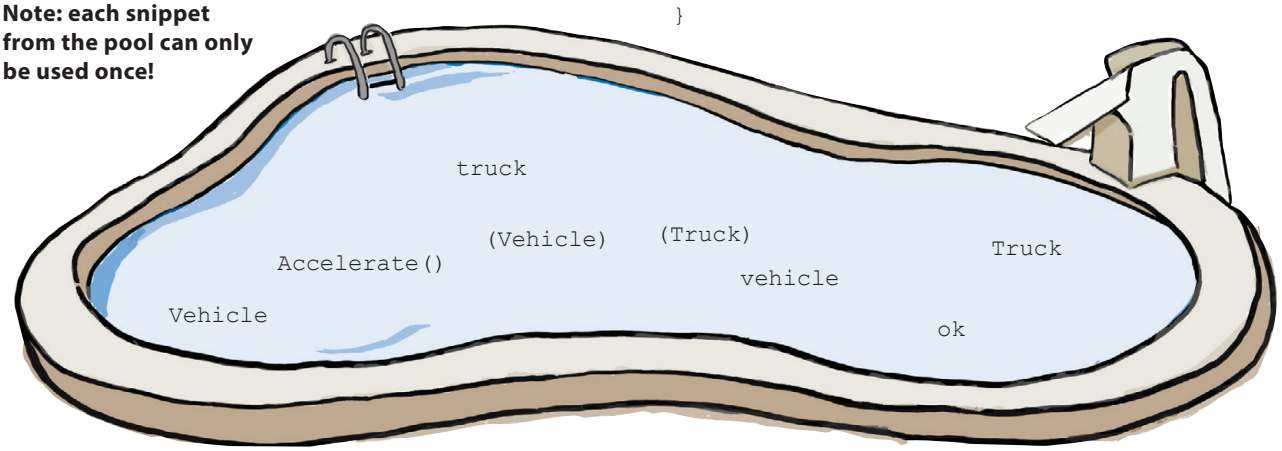
func TryVehicle(vehicle _________) {
       vehicle._________________
       vehicle.Steer("left")
       vehicle.Steer("right")
       vehicle.Brake()
       truck, ___ := vehicle._________
       if ok {
              ______.LoadCargo("test cargo")
       }
}

func main() {
       TryVehicle(Truck("Fnord F180"))
}
```

Output →

```
Speeding up
Turning left
Turning right
Stopping
Loading test cargo
```

Note: each snippet from the pool can only be used once!



Answers on page 348.

The "error" interface

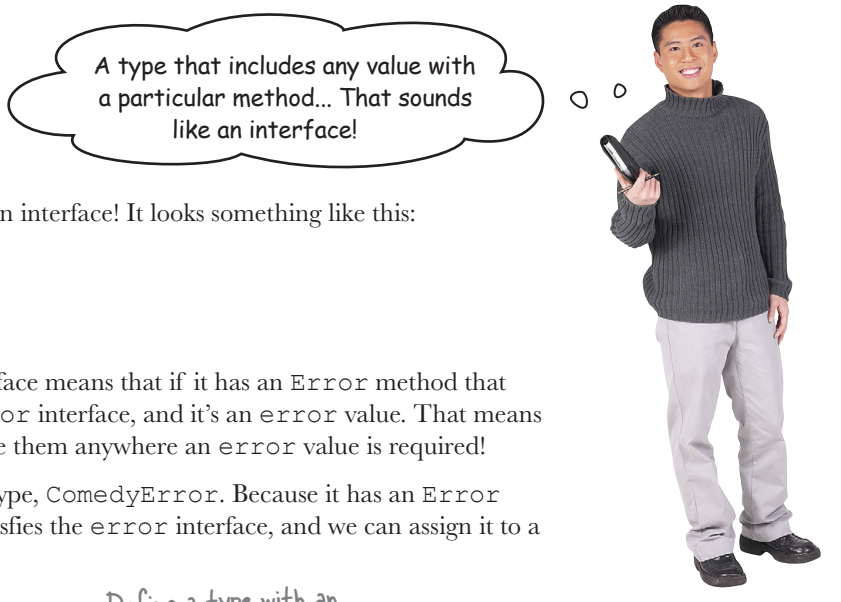
We'd like to wrap up the chapter by looking at a few interfaces that are built into Go. We haven't covered these interfaces explicitly, but you've actually been using them all along.

In Chapter 3, we learned how to create our own `error` values. We said, "An error value is any value with a method named `Error` that returns a string."

Returns an error value
Prints the error message
Also prints the error message

```
err := fmt.Errorf("a height of %0.2f is invalid", -2.33333)
fmt.Println(err.Error())
fmt.Println(err)
```

```
a height of -2.33 is invalid
a height of -2.33 is invalid
```



That's right. The `error` type is just an interface! It looks something like this:

```
type error interface {
      Error() string
}
```

Declaring the `error` type as an interface means that if it has an `Error` method that returns a `string`, it satisfies the `error` interface, and it's an `error` value. That means you can define your own types and use them anywhere an `error` value is required!

For example, here's a simple defined type, `ComedyError`. Because it has an `Error` method that returns a `string`, it satisfies the `error` interface, and we can assign it to a variable with the type `error`.

Define a type with an underlying type of "string".
Satisfy the error interface.
The Error method needs to return a string, so do a type conversion.
Set up a variable with a type of "error".
ComedyError satisfies the error interface, so we can assign a ComedyError to the variable.

```
type ComedyError string
func (c ComedyError) Error() string {
      return string(c)
}

func main() {
      var err error
      err = ComedyError("What's a programmer's favorite beer? Logger!")
      fmt.Println(err)
}
```

```
What's a programmer's favorite beer? Logger!
```

The "error" interface (continued)

If you need an `error` value, but also need to track more information about the error than just an error message string, you can create your own type that satisfies the `error` interface *and* stores the information you want.

Suppose you're writing a program that monitors some equipment to ensure it doesn't overheat. Here's an `OverheatError` type that might be useful. It has an `Error` method, so it satisfies `error`. But more interestingly, it uses `float64` as its underlying type, allowing us to track the degrees over capacity.

```
type OverheatError float64
func (o OverheatError) Error() string {
       return fmt.Sprintf("Overheating by %0.2f degrees!", o)
}
```

Define a type with an underlying type of float64.

Satisfy the error interface.

Use the temperature in the error message.

Here's a `checkTemperature` function that uses `OverheatError`. It takes the system's actual temperature and the temperature that's considered safe as parameters. It specifies that it returns a value of type `error`, not an `OverheatError` specifically, but that's okay because `OverheatError` satisfies the `error` interface. If the `actual` temperature is over the `safe` temperature, `checkTemperature` returns a new `OverheatError` that records the excess.

```
func checkTemperature(actual float64, safe float64) error {
       excess := actual - safe
       if excess > 0 {
              return OverheatError(excess)
       }
       return nil
}

func main() {
       var err error = checkTemperature(121.379, 100.0)
       if err != nil {
              log.Fatal(err)
       }
}
```

Specify that the function returns an ordinary error value.

If the actual temperature is in excess of the safe temperature...

...return an OverheatError that records the excess.

```
2018/04/02 19:27:44 Overheating by 21.38 degrees!
```

there are no Dumb Questions

Q: How is it we've been using the `error` interface type in all these different packages, without importing it? Its name begins with a lowercase letter. Doesn't that mean it's unexported, from whatever package it's declared in? What package is `error` declared in, anyway?

A: The `error` type is a "predeclared identifier," like `int` or `string`. And so, like other predeclared identifiers, it's not part of *any* package. It's part of the "universe block," meaning it's available everywhere, regardless of what package you're in.

Remember how there are `if` and `for` blocks, which are encompassed by function blocks, which are encompassed by package blocks? Well, the universe block encompasses all package blocks. That means you can use anything defined in the universe block from any package, without importing it. And that includes `error` and all other predeclared identifiers.

The Stringer interface

Remember our `Gallons`, `Liters`, and `Milliliters` types, which we created back in Chapter 9 to distinguish between various units for measuring volume? We're discovering that it's not so easy to distinguish between them after all. Twelve gallons is a very different amount than 12 liters or 12 milliliters, but they all look the same when printed. If there are too many decimal places of precision on a value, that looks awkward when printed, too.

```
type Gallons float64
type Liters float64
type Milliliters float64

func main() {
      fmt.Println(Gallons(12.09248342))
      fmt.Println(Liters(12.09248342))
      fmt.Println(Milliliters(12.09248342))
}
```

Create and print a Gallons value.

Create and print a Liters value.

Create and print a Milliliters value.

All three values look identical!

```
12.09248342
12.09248342
12.09248342
```

You can use `Printf` to round the number off and add an abbreviation indicating the unit of measure, but doing that every place you need to use these types would quickly get tedious.

Format the numbers and add abbreviations.

```
fmt.Printf("%0.2f gal\n", Gallons(12.09248342))
fmt.Printf("%0.2f L\n", Liters(12.09248342))
fmt.Printf("%0.2f mL\n", Milliliters(12.09248342))
```

```
12.09 gal
12.09 L
12.09 mL
```

That's why the `fmt` package defines the `fmt.Stringer` interface: to allow any type to decide how it will be displayed when printed. It's easy to set up any type to satisfy `Stringer`; just define a `String()` method that returns a `string`. The interface definition looks like this:

```
type Stringer interface {
      String() string
}
```

Any type is a fmt.Stringer if it has a String method that returns a string.

For example, here we've set up this `CoffeePot` type to satisfy `Stringer`:

```
type CoffeePot string
func (c CoffeePot) String() string {
      return string(c) + " coffee pot"
}

func main() {
      coffeePot := CoffeePot("LuxBrew")
      fmt.Println(coffeePot.String())
}
```

Satisfy the Stringer interface.

Method needs to return a string.

```
LuxBrew coffee pot
```

The Stringer interface (continued)

Many functions in the `fmt` package check whether the values passed to them satisfy the `Stringer` interface, and call their `String` methods if so. This includes the `Print`, `Println`, and `Printf` functions and more. Now that `CoffeePot` satisfies `Stringer`, we can pass `CoffeePot` values directly to these functions, and the return value of the `CoffeePot`'s `String` method will be used in the output:

Create a CoffeePot value.

```
coffeePot := CoffeePot("LuxBrew")
fmt.Print(coffeePot, "\n")
fmt.Println(coffeePot)
fmt.Printf("%s", coffeePot)
```

Pass the CoffeePot to various fmt functions.

```
LuxBrew coffee pot
LuxBrew coffee pot
LuxBrew coffee pot
```

The return value of String is used in the output.

Now for a more serious use of this interface type. Let's make our `Gallons`, `Liters`, and `Milliliters` types satisfy `Stringer`. We'll move our code to format their values to `String` methods associated with each type. We'll call the `Sprintf` function instead of `Printf`, and return the resulting value.

Make Gallons satisfy Stringer.

```
type Gallons float64
func (g Gallons) String() string {
      return fmt.Sprintf("%0.2f gal", g)
}
```

Make Liters satisfy Stringer.

```
type Liters float64
func (l Liters) String() string {
      return fmt.Sprintf("%0.2f L", l)
}
```

Make Milliliters satisfy Stringer.

```
type Milliliters float64
func (m Milliliters) String() string {
      return fmt.Sprintf("%0.2f mL", m)
}
```

```
func main() {
    fmt.Println(Gallons(12.09248342))
    fmt.Println(Liters(12.09248342))
    fmt.Println(Milliliters(12.09248342))
}
```

Pass values of each type to Println.

```
12.09 gal
12.09 L
12.09 mL
```

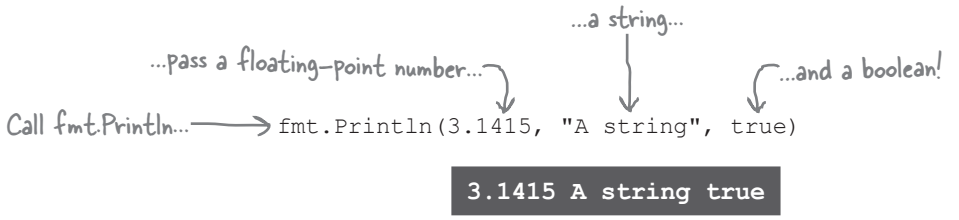
The return values of each type's String method are used in the output.

Now, any time we pass `Gallons`, `Liters`, and `Milliliters` values to `Println` (or most other `fmt` functions), their `String` methods will be called, and the return values used in the output. We've set up a useful default format for printing each of these types!

The empty interface



Something's been bothering me. For most of the functions we've seen so far, we can only call them using values of specific types. But some `fmt` functions like `fmt.Println` can take values of any type! How does that work?



Good question! Let's run **`go doc`** to bring up the documentation for `fmt.Println` and see what type its parameters are declared as...

View documentation for the "fmt" package's "Println" function.

The "..." shows it's a variadic function. But what's this "interface{}" type?

```
File Edit Window Help
$ go doc fmt Println
func Println(a ...interface{}) (n int, err error)
    Println formats using the default formats for its operands and writes to
    standard output. Spaces are always added between operands and a newline...
```

As we saw in Chapter 6, the `...` means that it's a variadic function, meaning it can take any number of parameters. But what's this `interface{}` type?

Remember, an interface declaration specifies the methods that a type is required to have in order to satisfy that interface. For example, our `NoiseMaker` interface is satisfied by any type that has a `MakeSound` method.

```
type NoiseMaker interface {
      MakeSound()
}
```

But what would happen if we declared an interface type that didn't require any methods at all? It would be satisfied by *any* type! It would be satisfied by *all* types!

```
type Anything interface {
}
```

The empty interface (continued)

The type `interface{}` is known as **the empty interface**, and it's used to accept values of *any* type. The empty interface doesn't have any methods that are required to satisfy it, and so *every* type satisfies it.

If you declare a function that accepts a parameter with the empty interface as its type, then you can pass it values of any type as an argument:

Accepts a parameter with the empty interface as its type

```
func AcceptAnything(thing interface{}) {
}

func main() {
    AcceptAnything(3.1415)
    AcceptAnything("A string")
    AcceptAnything(true)
    AcceptAnything(Whistle("Toyco Canary"))
}
```

These are all valid types to pass to our function!

The empty interface doesn't require any methods to satisfy it, and so it's satisfied by all types.

But don't rush out and start using the empty interface for all your function parameters! If you have a value with the empty interface as its type, there's not much you can *do* with it.

Most of the functions in `fmt` accept empty-interface values, so you can pass it on to those:

```
func AcceptAnything(thing interface{}) {
    fmt.Println(thing)
}

func main() {
    AcceptAnything(3.1415)
    AcceptAnything(Whistle("Toyco Canary"))
}
```

```
3.1415
Toyco Canary
```

But don't try calling any methods on an empty-interface value! Remember, if you have a value with an interface type, you can only call methods on it that are part of the interface. And the empty interface doesn't *have* any methods. That means there are *no* methods you can call on a value with the empty interface type!

```
func AcceptAnything(thing interface{}) {
    fmt.Println(thing)
    thing.MakeSound()
}
```

Try to call a method on the empty-interface value...

Error

```
thing.MakeSound undefined (type interface {} is interface with no methods)
```

The empty interface (continued)

To call methods on a value with the empty interface type, you'd need to use a type assertion to get a value of the concrete type back.

```
func AcceptAnything(thing interface{}) {
        fmt.Println(thing)
        whistle, ok := thing.(Whistle)
        if ok {
                whistle.MakeSound()
        }
}

func main() {
        AcceptAnything(3.1415)
        AcceptAnything(Whistle("Toyco Canary"))
}
```

Use a type assertion to get a Whistle.

Call the method on the Whistle.

```
3.1415
Toyco Canary
Tweet!
```

And by that point, you're probably better off writing a function that accepts only that specific concrete type.

```
func AcceptWhistle(whistle Whistle) {
        fmt.Println(whistle)
        whistle.MakeSound()
}
```

Accept a Whistle.

Call the method. No type conversion needed.

So there are limits to the usefulness of the empty interface when defining your own functions. But you'll use the empty interface all the time with the functions in the `fmt` package, and in other places too. The next time you see an `interface{}` parameter in a function's documentation, you'll know exactly what it means!

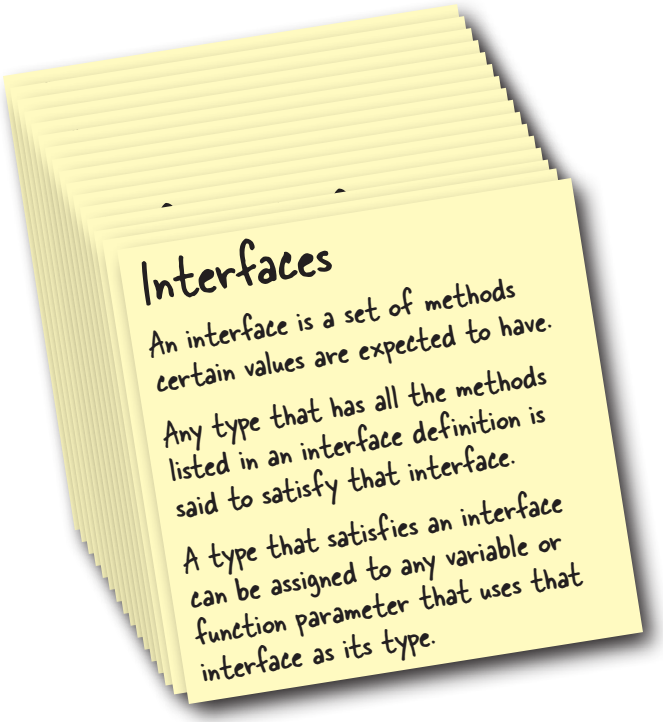
When you're defining variables or function parameters, often you'll know exactly what the value you'll be working with *is*. You'll be able to use a concrete type like `Pen`, `Car`, or `Whistle`. Other times, though, you only care about what the value can *do*. In that case, you're going to want to define an interface type, like `WritingInstrument`, `Vehicle`, or `NoiseMaker`.

You'll define the methods you need to be able to call as part of the interface type. And you'll be able to assign to your variables or call your functions without worrying about the concrete type of your values. If it has the right methods, you'll be able to use it!



Your Go Toolbox

That's it for Chapter 11! You've added interfaces to your toolbox.



BULLET POINTS

- A concrete type specifies not only what its values can *do* (what methods you can call on them), but also what they *are*: they specify the underlying type that holds the value's data.
- An interface type is an abstract type. Interfaces don't describe what a value *is*: they don't say what its underlying type is or how its data is stored. They only describe what a value can *do*: what methods it has.
- An interface definition needs to contain a list of method names, along with any parameters or return values those methods are expected to have.
- To satisfy an interface, a type must have all the methods the interface specifies. Method names, parameter types (or lack thereof), and return value types (or lack thereof) all need to match those defined in the interface.
- A type can have methods in addition to those listed in the interface, but it mustn't be missing any, or it doesn't satisfy that interface.
- A type can satisfy multiple interfaces, and an interface can have multiple types that satisfy it.
- Interface satisfaction is automatic. There is no need to explicitly declare that a concrete type satisfies an interface in Go.
- When you have a variable of an interface type, the only methods you can call on it are those defined in the interface.
- If you've assigned a value of a concrete type to a variable with an interface type, you can use a **type assertion** to get the concrete type value back. Only then can you call methods that are defined on the concrete type (but not the interface).
- Type assertions return a second `bool` value that indicates whether the assertion was successful.
  `car, ok := vehicle.(Car)`



Exercise Solution

```
type Car string
func (c Car) Accelerate() {
   fmt.Println("Speeding up")
}
func (c Car) Brake() {
   fmt.Println("Stopping")
}
func (c Car) Steer(direction string) {
   fmt.Println("Turning", direction)
}

type Truck string
func (t Truck) Accelerate() {
   fmt.Println("Speeding up")
}
func (t Truck) Brake() {
   fmt.Println("Stopping")
}
func (t Truck) Steer(direction string) {
   fmt.Println("Turning", direction)
}
func (t Truck) LoadCargo(cargo string) {
   fmt.Println("Loading", cargo)
}

type Vehicle interface {
   Accelerate()
   Brake()
   Steer(string)
}

func main() {
   var vehicle Vehicle = Car("Toyoda Yarvic")
   vehicle.Accelerate()
   vehicle.Steer("left")

   vehicle = Truck("Fnord F180")
   vehicle.Brake()
   vehicle.Steer("right")
}
```

Don't forget to specify that Steer takes a parameter!

```
Speeding up
Turning left
Stopping
Turning right
```

Pool Puzzle Solution

```
type Truck string
func (t Truck) Accelerate() {
    fmt.Println("Speeding up")
}
func (t Truck) Brake() {
    fmt.Println("Stopping")
}
func (t Truck) Steer(direction string) {
    fmt.Println("Turning", direction)
}
func (t Truck) LoadCargo(cargo string) {
    fmt.Println("Loading", cargo)
}

type Vehicle interface {
    Accelerate()
    Brake()
    Steer(string)
}

func TryVehicle(vehicle Vehicle) {
    vehicle.Accelerate()
    vehicle.Steer("left")
    vehicle.Steer("right")
    vehicle.Brake()
    truck, ok := vehicle.(Truck)
    if ok {
        truck.LoadCargo("test cargo")
    }
}

func main() {
    TryVehicle(Truck("Fnord F180"))
}
```

Was type assertion successful?

Holds a Truck, not (just) a Vehicle, so we can call LoadCargo.

```
Speeding up
Turning left
Turning right
Stopping
Loading test cargo
```

Recovering from Failure



Every program encounters errors. You should plan for them.

Sometimes handling an error can be as simple as reporting it and exiting the program. But other errors may require additional action. You may need to close opened files or network connections, or otherwise clean up, so your program doesn't leave a mess behind. In this chapter, we'll show you how to **defer** cleanup actions so they happen even when there's an error. We'll also show you how to make your program **panic** in those (rare) situations where it's appropriate, and how to **recover** afterward.

Reading numbers from a file, revisited

We've talked about handling errors in Go quite a lot. But the techniques we've shown thus far don't work in every situation. Let's look at one such scenario.

We want to create a program, *sum.go*, that reads `float64` values from a text file, adds them all together, and prints their sum.




```
Shell Edit View Window Help
$ go run sum.go data.txt
Opening data.txt
Closing file
Sum: 9.62
```

In Chapter 6 we created a `GetFloats` function that opened a text file, converted each line of the file to a `float64` value, and returned those values as a slice.

Here, we've moved `GetFloats` to the `main` package and updated it to rely on two new functions, `OpenFile` and `CloseFile`, to open and close the text file.

```
package main

import (
    "bufio"
    "fmt"
    "log"
    "os"
    "strconv"
)

func OpenFile(fileName string) (*os.File, error) {
    fmt.Println("Opening", fileName)
    return os.Open(fileName)
}
func CloseFile(file *os.File) {
    fmt.Println("Closing file")
    file.Close()
}

func GetFloats(fileName string) ([]float64, error) {
    var numbers []float64
    file, err := OpenFile(fileName)
    if err != nil {
        return nil, err
    }
    scanner := bufio.NewScanner(file)
    for scanner.Scan() {
        number, err := strconv.ParseFloat(scanner.Text(), 64)
        if err != nil {
            return nil, err
        }
        numbers = append(numbers, number)
    }
    CloseFile(file)
    if scanner.Err() != nil {
        return nil, scanner.Err()
    }
    return numbers, nil
}
```

We've moved all this code to the "main" package, in the sum.go source file.

Opens the file and returns a pointer to it, along with any error encountered

Closes the file

Instead of calling os.Open directly, call OpenFile.

Instead of calling file.Close directly, call CloseFile.

Reading numbers from a file, revisited (continued)

We want to specify the name of the file we're going to read as a command-line argument. You may recall using the `os.Args` slice in Chapter 6—it's a slice of `string` values containing all the arguments used when the program is run.

So in our `main` function, we get the name of the file to open from the first command-line argument by accessing `os.Args[1]`. (Remember, the `os.Args[0]` element is the name of the program being run; the actual program arguments appear in `os.Args[1]` and later elements.)

We then pass that filename to `GetFloats` to read the file, and get a slice of `float64` values back.

If any errors are encountered along the way, they'll be returned from the `GetFloats` function, and we'll store them in the `err` variable. If `err` is not `nil`, it means there was an error, so we simply log it and exit.

Otherwise, it means the file was read successfully, so we use a `for` loop to add every value in the slice together, and end by printing the total.

Use the first command-line argument as a filename.

Store the slice of numbers read from the file, along with any error.

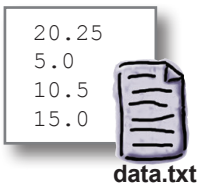
If there was an error, log it and exit.

Add up all the numbers in the slice.

Print the total.

```
func main() {
    numbers, err := GetFloats(os.Args[1])
    if err != nil {
        log.Fatal(err)
    }
    var sum float64 = 0
    for _, number := range numbers {
        sum += number
    }
    fmt.Printf("Sum: %0.2f\n", sum)
}
```

Let's save all this code together in a file named *sum.go*. Then, let's create a plain-text file filled with numbers, one number per line. We'll name it *data.txt* and save it in the same directory as *sum.go*.



We can run the program with `go run sum.go data.txt`. The string `"data.txt"` will be the first argument to the *sum.go* program, so that's the filename that will be passed to `GetFloats`.

We can see when the `OpenFile` and `CloseFile` functions get called, since they both include calls to `fmt.Println`. And at the end of the output, we can see the total of all the numbers in *data.txt*. Looks like everything's working!

Pass data.txt as a command-line argument.

Here's where OpenFile gets called.

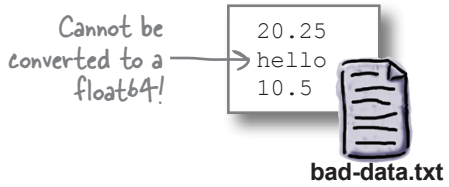
Here's where CloseFile gets called.

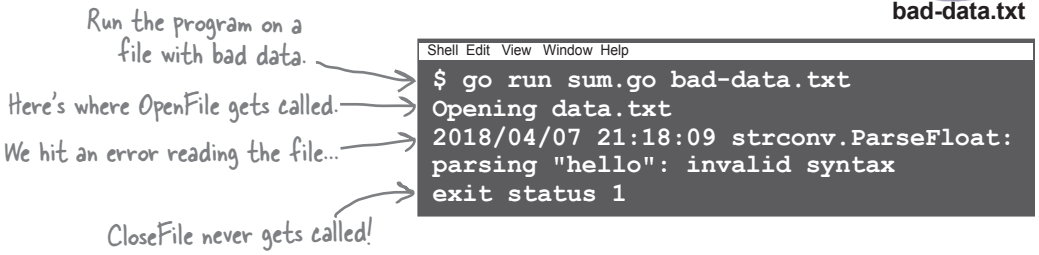
Here's the total of all the numbers in the file.

```
Shell Edit View Window Help
$ go run sum.go data.txt
Opening data.txt
Closing file
Sum: 50.75
```

Any errors will prevent the file from being closed!

If we give the *sum.go* program an improperly formatted file, though, we run into problems. A file with a line that can't be parsed into a `float64` value, for example, results in an error.





Now, that in itself is fine; every program receives invalid data occasionally. But the `GetFloats` function is supposed to call the `CloseFile` function when it's done. We don't see "`Closing file`" in the program output, which would suggest that `CloseFile` isn't getting called!

The problem is that when we call `strconv.ParseFloat` with a string that can't be converted to a `float64`, it returns an error. Our code is set up to return from the `GetFloats` function at that point.

But that return happens *before* the call to `CloseFile`, which means the file never gets closed!

```
func GetFloats(fileName string) ([]float64, error) {
    var numbers []float64
    file, err := OpenFile(fileName)
    if err != nil {
        return nil, err
    }
    scanner := bufio.NewScanner(file)
    for scanner.Scan() {
        number, err := strconv.ParseFloat(scanner.Text(), 64)
        if err != nil {
            return nil, err
        }
        numbers = append(numbers, number)
    }
    CloseFile(file)
    if scanner.Err() != nil {
        return nil, scanner.Err()
    }
    return numbers, nil
}
```

ParseFloat returns an error when it can't convert the text line to a float64...

...which causes GetFloats to return an error...

...which means CloseFile never gets called!

Deferring function calls

Now, failing to close a file may not seem like such a big deal. And for a simple program that just opens a single file, it's probably fine. But each file that's left open continues to consume operating system resources. Over time, multiple files left open can build up and cause a program to fail, or even hamper performance of the entire system. It's really important to get in the habit of ensuring that files are closed when your program is done with them.

But how can we accomplish this? The `GetFloats` function is set up to immediately exit if it encounters an error reading the file, even if `CloseFile` hasn't been called yet!

If you have a function call that you want to ensure is run, *no matter what*, you can use a `defer` statement. You can place the `defer` keyword before any ordinary function or method call, and Go will defer (that is, delay) making the function call until after the current function exits.

Normally, function calls are executed as soon as they're encountered. In this code, the `fmt.Println("Goodbye!")` call runs before the other two `fmt.Println` calls.

```
package main

import "fmt"

func Socialize() {
       fmt.Println("Goodbye!")
       fmt.Println("Hello!")
       fmt.Println("Nice weather, eh?")
}

func main() {
       Socialize()
}
```

```
Goodbye!
Hello!
Nice weather, eh?
```

But if we add the `defer` keyword before the `fmt.Println("Goodbye!")` call, then that call won't be run until all the remaining code in the `Socialize` function runs, and `Socialize` exits.

```
package main

import "fmt"

func Socialize() {
       defer fmt.Println("Goodbye!")
       fmt.Println("Hello!")
       fmt.Println("Nice weather, eh?")
}

func main() {
       Socialize()
}
```

Add the "defer" keyword before the function call.

The first function call is deferred until after Socialize exits!

```
Hello!
Nice weather, eh?
Goodbye!
```

Recovering from errors using deferred function calls



That's cool, but you said something about `defer` being used for function calls that need to happen "no matter what." Mind explaining?

The `defer` keyword ensures a function call takes place even if the calling function exits early, say, by using the `return` keyword.

Below, we've updated our `Socialize` function to return an `error` because we don't feel like talking. `Socialize` will exit before the `fmt.Println("Nice weather, eh?")` call. But because we include a `defer` keyword before the `fmt.Println("Goodbye!")` call, `Socialize` will always be polite enough to print "`Goodbye!`" before ending the conversation.

The "defer" keyword ensures a function call takes place, even if the calling function exits early.

```
package main

import (
        "fmt"
        "log"
)

func Socialize() error {
        defer fmt.Println("Goodbye!")
        fmt.Println("Hello!")
        return fmt.Errorf("I don't want to talk.")
        fmt.Println("Nice weather, eh?")
        return nil
}

func main() {
        err := Socialize()
        if err != nil {
                log.Fatal(err)
        }
}
```

Defer printing "Goodbye!"

Return an error.

This code won't be run!

The deferred function call is still made when Socialize returns.

```
Hello!
Goodbye!
2018/04/08 19:24:48 I don't want to talk.
```

Ensuring files get closed using deferred function calls

Because the `defer` keyword can ensure a function call is made "no matter what," it's usually used for code that needs to be run even in the event of an error. One common example of this is closing files after they've been opened.

And that's exactly what we need in our *sum.go* program's `GetFloats` function. After we call the `OpenFile` function, we need it to call `CloseFile`, even if there's an error parsing the file contents.

We can achieve this by simply moving the call to `CloseFile` immediately after the call to `OpenFile` (and its accompanying error handling code), and placing the `defer` keyword in front of it.

```
func OpenFile(fileName string) (*os.File, error) {
    fmt.Println("Opening", fileName)
    return os.Open(fileName)
}
func CloseFile(file *os.File) {
    fmt.Println("Closing file")
    file.Close()
}

func GetFloats(fileName string) ([]float64, error) {
    var numbers []float64
    file, err := OpenFile(fileName)
    if err != nil {
        return nil, err
    }
    defer CloseFile(file)
    scanner := bufio.NewScanner(file)
    for scanner.Scan() {
        number, err := strconv.ParseFloat(scanner.Text(), 64)
        if err != nil {
            return nil, err
        }
        numbers = append(numbers, number)
    }
    if scanner.Err() != nil {
        return nil, scanner.Err()
    }
    return numbers, nil
}
```

Move this right after the call to OpenFile (and its error handling code).

Add "defer" so it doesn't run until GetFloats exits.

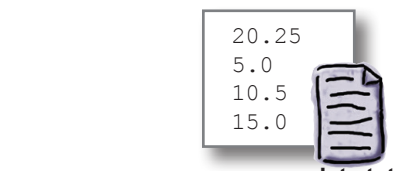
Now, even if an error is returned here, CloseFile will still be called!

CloseFile would be called if an error were returned here, too!

And of course, CloseFile is called if GetFloats completes normally!

Using `defer` ensures `CloseFile` will be called when `GetFloats` exits, whether it completes normally or there's an error parsing the file.

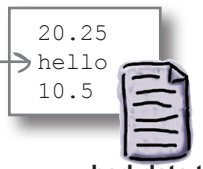
Now, even if *sum.go* is given a file with bad data, it will still close the file before exiting!



Deferred CloseFile call is executed!

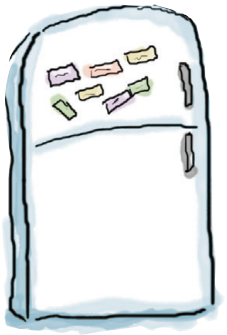
```
Shell Edit View Window Help
$ go run sum.go data.txt
Opening data.txt
Closing file
Sum: 50.75
```

This file has an error.



Deferred CloseFile call is executed!

```
Shell Edit View Window Help
$ go run sum.go bad-data.txt
Opening data.txt
Closing file
2018/04/09 21:30:42 strconv.ParseFloat:
parsing "hello": invalid syntax
exit status 1
```



Code Magnets

This code sets up a `Refrigerator` type that simulates a refrigerator. `Refrigerator` uses a slice of strings as its underlying type; the strings represent the names of foods the refrigerator contains. The type has an `Open` method that simulates opening the door, and a corresponding `Close` method to close it again (we don't want to waste energy, after all). The `FindFood` method calls `Open` to open the door, calls a `find` function we've written to search the underlying slice for a particular food, and then calls `Close` to close the door again.

But there's a problem with `FindFood`. It's set up to return an error value if the food we're searching for isn't found. But when that happens, it's returning before `Close` gets called, leaving the virtual refrigerator door wide open!

(Continued on next page...)

```
func find(item string, slice []string) bool {
        for _, sliceItem := range slice {
                if item == sliceItem {
                        return true
                }
        }
        return false
}

type Refrigerator []string

func (r Refrigerator) Open() {
        fmt.Println("Opening refrigerator")
}
func (r Refrigerator) Close() {
        fmt.Println("Closing refrigerator")
}
func (r Refrigerator) FindFood(food string) error {
        r.Open()
        if find(food, r) {
                fmt.Println("Found", food)
        } else {
                return fmt.Errorf("%s not found", food)
        }
        r.Close()
        return nil
}

func main() {
        fridge := Refrigerator{"Milk", "Pizza", "Salsa"}
        for _, food := range []string{"Milk", "Bananas"} {
                err := fridge.FindFood(food)
                if err != nil {
                        log.Fatal(err)
                }
        }
}
```

Returns true if the string is found in the slice...

...or false if the string is not found.

The Refrigerator type is based on a slice of strings, which will hold the names of the foods the refrigerator contains.

Simulates opening the refrigerator

Simulates closing the refrigerator

If this Refrigerator contains the food we want...

...print that we found it.

Otherwise, return an error.

But if we return an error, this never gets called!

Refrigerator is opened, but never gets closed!

```
Opening refrigerator
Found Milk
Closing refrigerator
Opening refrigerator
2018/04/09 22:12:37 Bananas not found
```

Answers on page 377.

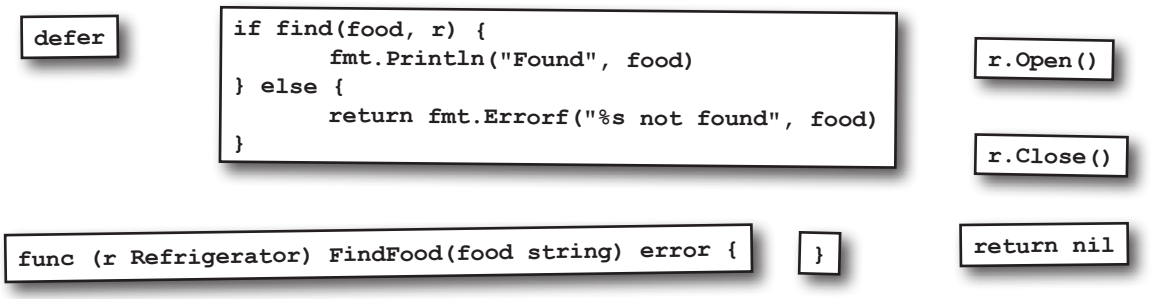
Code Magnets (continued)

Use the magnets below to create an updated version of the `FindFood` method. It should defer the call to the `Close` method, so that it runs when `FindFood` exits (regardless of whether the food was found successfully).

Refrigerator's Close method should be called when food is found.

Close should also be called when food is not found.

```
Opening refrigerator
Found Milk
Closing refrigerator
Opening refrigerator
Closing refrigerator
2018/04/09 22:12:37 Bananas not found
```



there are no Dumb Questions

Q: So I can defer function and method calls... Can I defer other statements too, like `for` loops or variable assignments?

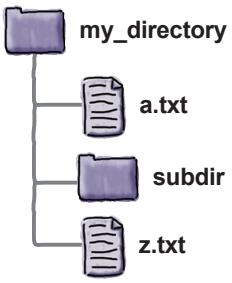
A: No, only function and method calls. You can write a function or method to do whatever you want and then defer a call to that function or method, but the `defer` keyword itself can only be used with a function or method call.

Listing the files in a directory



Go has a couple more features to help you handle errors, and we'll be showing you a program that demonstrates them in a bit. But that program uses a couple new tricks, which we'll need to show you before we dive in. First up, we're going to need to know how to read the contents of a directory.

Try creating a directory, named *my_directory*, that includes two files and a subdirectory, as shown at the right. The program below will list the contents of *my_directory*, indicating the name of each item it contains, and whether it's a file or a subdirectory.



The `io/ioutil` package includes a `ReadDir` function that will let us read the directory contents. You pass `ReadDir` the name of a directory, and it will return a slice of values, one for each file or subdirectory the directory contains (along with any error it encounters).

Each of the slice's values satisfies the `FileInfo` interface, which includes a `Name` method that returns the file's name, and an `IsDir` method that returns `true` if it's a directory.

So our program calls `ReadDir`, passing it the name of *my_directory* as an argument. It then loops over each value in the slice it gets back. If `IsDir` returns `true` for the value, it prints `"Directory:"` and the file's name. Otherwise, it prints `"File:"` and the file's name.




```
package main

import (
        "fmt"
        "io/ioutil"
        "log"
)

func main() {
        files, err := ioutil.ReadDir("my_directory")
        if err != nil {
                log.Fatal(err)
        }

        for _, file := range files {
                if file.IsDir() {
                        fmt.Println("Directory:", file.Name())
                } else {
                        fmt.Println("File:", file.Name())
                }
        }
}
```

Get a slice full of values representing the contents of "my_directory".

For each file in the slice...

If this file is a directory... ...print "Directory" and the filename.

Otherwise, print "File:" and the filename.

Save the above code as *files.go*, in the same directory as *my_directory*. In your terminal, change to that parent directory, and type **`go run files.go`**. The program will run and produce a list of the files and directories *my_directory* contains.

```
Shell Edit View Window Help
$ cd work
$ go run files.go
File: a.txt
Directory: subdir
File: z.txt
```

Listing the files in subdirectories (will be trickier)



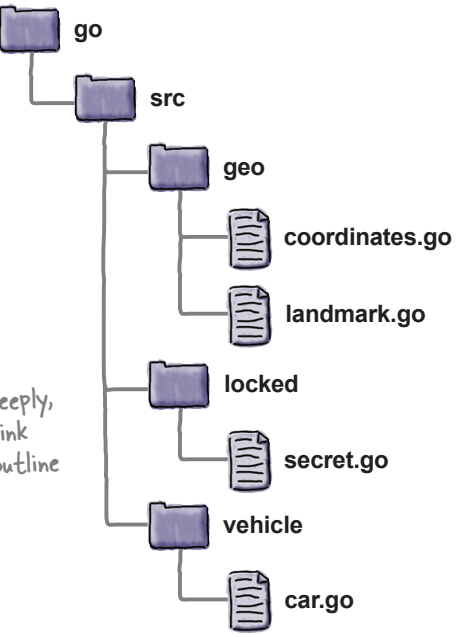
A program that reads the contents of a single directory isn't too complicated. But suppose we wanted to list the contents of something more complicated, like a Go workspace directory. That would contain an entire tree of subdirectories nested within subdirectories, some containing files, some not.

Normally, such a program would be quite complicated. In outline form, the logic would be something like this:

I. Get a list of files in the directory.

- **A.** Get the next file.
- **B.** Is the file a directory?
  - **1.** If yes: get a list of files in the directory.
    - **a.** Get the next file.
    - **b.** Is the file a directory?
      - **01.** If yes: Get a list of the files in the directory...
  - **2.** If no: just print the filename.

This logic is nested so deeply, we can't think of enough outline levels!



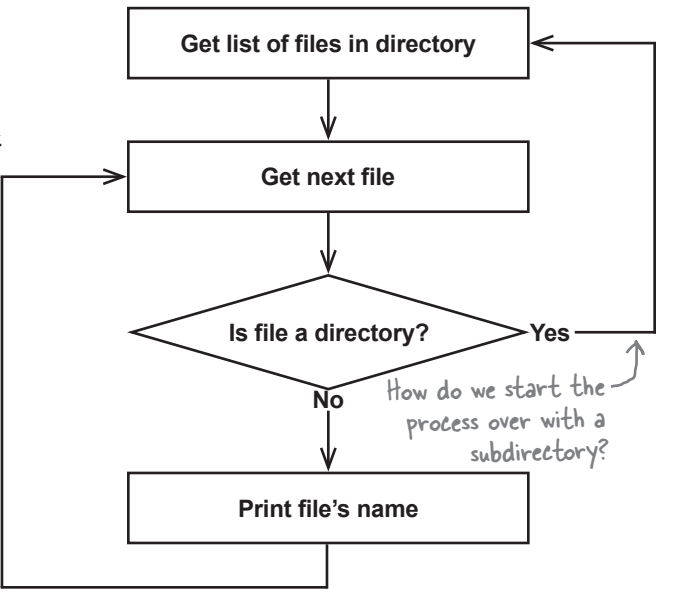
Pretty complicated, right? We'd rather not have to write *that* code!

But what if there were a simpler way? Some logic like this:

I. Get a list of files in the directory.

- **A.** Get the next file.
- **B.** Is the file a directory?
  - **1.** If yes: start over at step **I** with this directory.
  - **2.** If no: just print the filename.

It's not clear how to handle the "Start the logic over with this new directory" part, though. To achieve this, we'll need a new programming concept...



Recursive function calls



That brings us to the second (and last) trick we'll need to show you before we end our detour and get back to handling errors.

Go is one of many programming languages that support **recursion**, which allows a function to call itself.

If you do this carelessly, you'll just wind up with an infinite loop where the function calls itself over and over:

```
package main

import "fmt"

func recurses() {
	fmt.Println("Oh, no, I'm stuck!")
	recurses()
}

func main() {
	recurses()
}
```

The "recurses" function calls itself!

Call "recurses" for the first time.

```
Oh, no, I'm stuck!
Oh, no, I'm stuck!
Oh, no, I'm stuck!
Oh, no, ^Csignal: interrupt
```

Anyone running this program will have to press Ctrl–C to break out of the infinite loop!

But if you make sure that the recursion loop stops itself eventually, recursive functions can actually be useful.

Here's a recursive `count` function that counts from a starting number up to an ending number. (Normally a loop would be more efficient, but this is a simple way of demonstrating how recursion works.)

```
package main

import "fmt"

func count(start int, end int) {
	fmt.Println(start)
	if start < end {
		count(start+1, end)
	}
}

func main() {
	count(1, 3)
}
```

Print the current starting number.

If we haven't reached the ending number...

...the "count" function calls itself, with a starting number 1 higher than before.

Call "count" for the first time, specifying it should count from 1 to 3.

```
1
2
3
```

Recursive function calls (continued)



Here's the sequence the program follows:

1. `main` calls `count` with a `start` parameter of `1` and an `end` of `3`
2. `count` prints the `start` parameter: `1`
3. `start` (1) is less than `end` (3), so `count` calls itself with a `start` of `2` and an `end` of `3`
4. This second invocation of `count` prints its new `start` parameter: `2`
5. `start` (2) is less than `end` (3), so `count` calls itself with a `start` of `3` and an `end` of `3`
6. The third invocation of `count` prints its new `start` parameter: `3`
7. `start` (3) is *not* less than `end` (3), so `count` does *not* call itself again; it just returns
8. The previous two invocations of `count` return as well, and the program ends

If we add calls to `Printf` showing each time `count` is called and each time the function exits, this sequence will be a little more obvious:

```
package main

import "fmt"

func count(start int, end int) {
	fmt.Printf("count(%d, %d) called\n", start, end)
	fmt.Println(start)
	if start < end {
		count(start+1, end)
	}
	fmt.Printf("Returning from count(%d, %d) call\n", start, end)
}

func main() {
	count(1, 3)
}
```

```
count(1, 3) called
1
count(2, 3) called
2
count(3, 3) called
3
Returning from count(3, 3) call
Returning from count(2, 3) call
Returning from count(1, 3) call
```

So that's a simple recursive function. Let's try applying recursion to our *files.go* program, and see if it can help us list the contents of subdirectories...

Recursively listing directory contents

We want our *files.go* program to list the contents of all of the subdirectories in our Go workspace directory. We're hoping to achieve that using recursive logic like this:

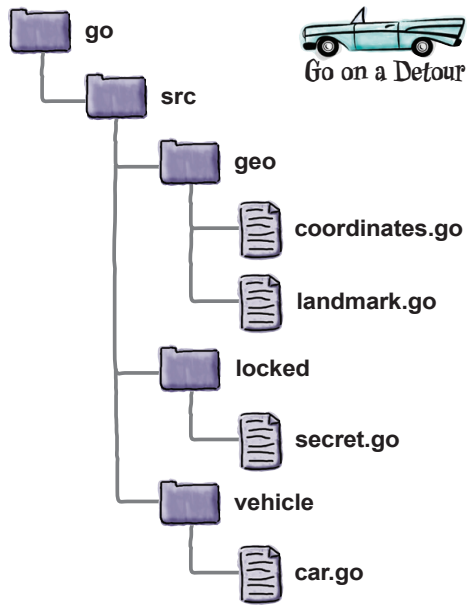
I. Get a list of files in the directory.

  A. Get the next file.

  B. Is the file a directory?

    1. If yes: start over at step **I** with this directory.

    2. If no: just print the filename.



We've removed the code from the `main` function that reads the directory contents; `main` now simply calls a recursive `scanDirectory` function. The `scanDirectory` function takes the path of the directory it should scan, so we pass it the path of the `"go"` subdirectory.

The first thing `scanDirectory` does is print the current path, so we know what directory we're working in. Then it calls `ioutil.ReadDir` on that path, to get the directory contents.

It loops over the slice of `FileInfo` values that `ReadDir` returns, processing each one. It calls `filepath.Join` to join the current directory path and the current filename together with slashes (so `"go"` and `"src"` are joined to become `"go/src"`).

If the current file isn't a directory, `scanDirectory` just prints its full path, and moves on to the next file (if there are any more in the current directory).

But if the current file *is* a directory, the recursion kicks in: `scanDirectory` calls itself with the subdirectory's path. If that subdirectory has any subdirectories, `scanDirectory` will call itself with each of *those* subdirectories, and so on through the whole file tree.

```
package main

import (
    "fmt"
    "io/ioutil"
    "log"
    "path/filepath"
)

// A recursive function that takes the path to scan
// We'll return any errors we encounter.
func scanDirectory(path string) error {
    fmt.Println(path)                          // Print the current directory.
    files, err := ioutil.ReadDir(path)         // Get a slice with the directory's contents.
    if err != nil {
        return err
    }
    for _, file := range files {
        // Join the directory path and filename with a slash.
        filePath := filepath.Join(path, file.Name())
        if file.IsDir() {                      // If this is a subdirectory...
            err := scanDirectory(filePath)     // ...recursively call scanDirectory, this time with the subdirectory's path.
            if err != nil {
                return err
            }
        } else {
            fmt.Println(filePath)              // If this is a regular file, just print its path.
        }
    }
    return nil
}

// Kick the process off by calling scanDirectory on the top directory.
func main() {
    err := scanDirectory("go")
    if err != nil {
        log.Fatal(err)
    }
}
```

Recursively listing directory contents (continued)

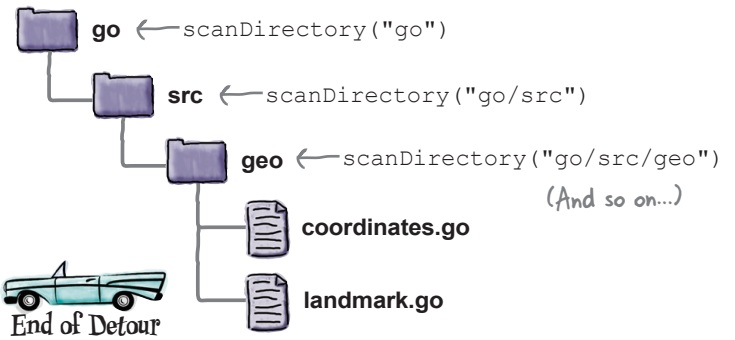


Save the preceding code as *files.go* in the directory that contains your Go workspace (probably your user's home directory). In your terminal, change to that directory, and run the program with **`go run files.go`**.

```
Shell Edit View Window Help
$ cd /Users/jay
$ go run files.go
go
go/src
go/src/geo
go/src/geo/coordinates.go
go/src/geo/landmark.go
go/src/locked
go/src/locked/secret.go
go/src/vehicle
go/src/vehicle/car.go
```

When you see the `scanDirectory` function at work, you'll see the real beauty of recursion. For our sample directory structure, the process goes something like this:

1. `main` calls `scanDirectory` with a path of `"go"`
2. `scanDirectory` prints the path it's passed, `"go"`, indicating the directory it's working in
3. It calls `ioutil.ReadDir` with the `"go"` path
4. There's only one entry in the returned slice: `"src"`
5. Calling `filepath.Join` with the current directory path of `"go"` and a filename of `"src"` gives a new path of `"go/src"`
6. *src* is a subdirectory, so `scanDirectory` is called again, this time with a path of `"go/src"` ← Recursion!
7. `scanDirectory` prints the new path: `"go/src"`
8. It calls `ioutil.ReadDir` with the `"go/src"` path
9. The first entry in the returned slice is `"geo"`
10. Calling `filepath.Join` with the current directory path of `"go/src"` and a filename of `"geo"` gives a new path of `"go/src/geo"`
11. *geo* is a subdirectory, so `scanDirectory` is called again, this time with a path of `"go/src/geo"` ← Recursion!
12. `scanDirectory` prints the new path: `"go/src/geo"`
13. It calls `ioutil.ReadDir` with the `"go/src/geo"` path
14. The first entry in the returned slice is `"coordinates.go"`
15. *coordinates.go* is *not* a directory, so its name is simply printed
16. And so on...



Recursive functions can be tricky to write, and they often consume more computing resources than nonrecursive solutions. But sometimes, recursive functions offer solutions to problems that would be very difficult to solve using other means.

Now that our *files.go* program is set up, we can end our detour. Up next, we'll return to our discussion of Go's error handling features.

Error handling in a recursive function

If `scanDirectory` encounters an error while scanning any subdirectory (for example, if a user doesn't have permission to access that directory), it will return an error. This is expected behavior; the program doesn't have any control over the filesystem, and it's important to report errors when they inevitably occur.

```
Shell Edit View Window Help
$ go run files.go
go
go/src
go/src/geo
go/src/geo/coordinates.go
go/src/geo/landmark.go
go/src/locked
2018/04/09 19:09:21 open
go/src/locked: permission denied
exit status 1
```

But if we add a couple `Printf` statements showing the errors being returned, we'll see that the *way* this error is handled isn't ideal:

```
func scanDirectory(path string) error {
    fmt.Println(path)
    files, err := ioutil.ReadDir(path)
    if err != nil {
        fmt.Printf("Returning error from scanDirectory(\"%s\") call\n", path)
        return err
    }

    for _, file := range files {
        filePath := filepath.Join(path, file.Name())
        if file.IsDir() {
            err := scanDirectory(filePath)
            if err != nil {
                fmt.Printf("Returning error from scanDirectory(\"%s\") call\n", path)
                return err
            }
        } else {
            fmt.Println(filePath)
        }
    }
    return nil
}

func main() {
    err := scanDirectory("go")
    if err != nil {
        log.Fatal(err)
    }
}
```

Print debug info for error in ReadDir call.

Print debug info for error in recursive scanDirectory call.

If an error occurs in one of the recursive `scanDirectory` calls, that error has to be returned up the entire chain until it reaches the `main` function!

```
Shell Edit View Window Help
$ go run files.go
go
go/src
go/src/geo
go/src/geo/coordinates.go
go/src/geo/landmark.go
go/src/locked
Returning error from scanDirectory("go/src/locked") call
Returning error from scanDirectory("go/src") call
Returning error from scanDirectory("go") call
2018/06/11 11:01:28 open go/src/locked: permission denied
exit status 1
```

Starting a panic

Our `scanDirectory` function is a rare example of a place it might be appropriate for a program to panic at runtime.

We've encountered panics before. We've seen them when accessing invalid indexes in arrays and slices:

```
notes := [7]string{"do", "re", "mi", "fa", "so", "la", "ti"}

for i := 0; i <= len(notes); i++ {
        fmt.Println(i, notes[i])
}
```

The highest value the "i" variable will reach is 7!

Returns the length of the array, 7

Accessing index 7 causes a panic!

```
0 do
1 re
2 mi
3 fa
4 so
5 la
6 ti
panic: runtime error: index out of range

goroutine 1 [running]:
main.main()
        /tmp/sandbox094804331/main.go:11 +0x140
```

We've also seen them when a type assertion fails (if we didn't use the optional `ok` Boolean value):

```
var player Player = gadget.TapePlayer{}
recorder := player.(gadget.TapeRecorder)
```

Assert that the original type is TapeRecorder, when it's actually TapePlayer...

Panic!

```
panic: interface conversion: main.Player
is gadget.TapePlayer, not gadget.TapeRecorder
```

When a program panics, the current function stops running, and the program prints a log message and crashes.

You can cause a panic yourself simply by calling the built-in `panic` function.

```
package main

func main() {
        panic("oh, no, we're going down")
}
```

```
panic: oh, no, we're going down

goroutine 1 [running]:
main.main()
        /tmp/main.go:4 +0x40
```

The `panic` function expects a single argument that satisfies the empty interface (that is, it can be of any type). That argument is converted to a string (if necessary) and printed as part of the panic's log message.

Stack traces

Each function that's called needs to return to the function that called it. To enable this, like other programming languages, Go keeps a **call stack**, a list of the function calls that are active at any given point.

When a program panics, a **stack trace**, or listing of the call stack, is included in the panic output. This can be useful in determining what caused the program to crash.

```
package main

func main() {
        one()
}
func one() {
        two()
}
func two() {
        three()
}
func three() {
        panic("This call stack's too deep for me!")
}
```

This function call gets added to the stack.

Add another call to the stack.

Add a third.

Panic! The stack trace will include all the above calls.

```
panic: This call stack's too deep for me!

goroutine 1 [running]:
main.three()
        /tmp/main.go:13 +0x40
main.two()
        /tmp/main.go:10 +0x20
main.one()
        /tmp/main.go:7 +0x20
main.main()
        /tmp/main.go:4 +0x20
```

The stack trace includes the list of function calls that have been made.

Deferred calls completed before crash

When a program panics, all deferred function calls will still be made. If there's more than one deferred call, they'll be made in the reverse of the order they were deferred in.

The code below defers two calls to `Println` and then panics. The top of the program output shows the two calls being completed before the program crashes.

```
func main() {
    one()
}
func one() {
    defer fmt.Println("deferred in one()")
    two()
}
func two() {
    defer fmt.Println("deferred in two()")
    panic("Let's see what's been deferred!")
}
```

This function call was deferred first, so it'll be made last.

This function call was deferred last, so it'll be made first.

Deferred calls completed before crash

```
deferred in two()
deferred in one()
panic: Let's see what's been deferred!

goroutine 1 [running]:
main.two()
    /tmp/main.go:14 +0xa0
main.one()
    /tmp/main.go:10 +0xa0
main.main()
    /tmp/main.go:6 +0x20
```

Using "panic" with scanDirectory

The `scanDirectory` function at the right has been updated to call `panic` instead of returning an error value. This greatly simplifies the error handling.

First, we remove the `error` return value from the `scanDirectory` declaration. If an `error` value is returned from `ReadDir`, we pass it to `panic` instead. We can remove the error handling code from the recursive call to `scanDirectory`, and the call to `scanDirectory` in `main`, as well.

```
package main

import (
    "fmt"
    "io/ioutil"
    "path/filepath"
)

func scanDirectory(path string) {
    fmt.Println(path)
    files, err := ioutil.ReadDir(path)
    if err != nil {
        panic(err)
    }

    for _, file := range files {
        filePath := filepath.Join(path, file.Name())
        if file.IsDir() {
            scanDirectory(filePath)
        } else {
            fmt.Println(filePath)
        }
    }
}

func main() {
    scanDirectory("go")
}
```

Error return value no longer needed.

Instead of returning the error value, pass it to "panic".

No more need to store or check error return value.

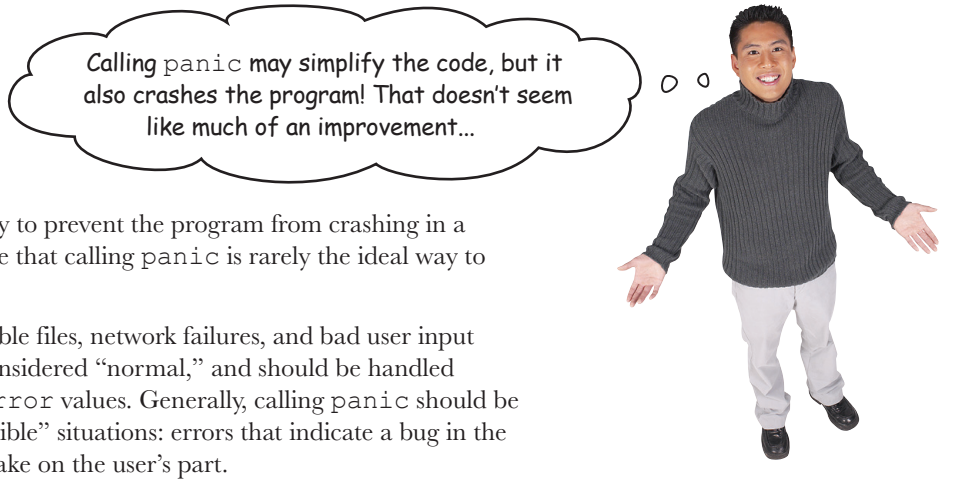
No more need to store or check error return value.

Now, when `scanDirectory` encounters an error reading a directory, it simply panics. All the recursive calls to `scanDirectory` exit.

```
Shell Edit View Window Help
$ go run files.go
go
go/src
go/src/geo
go/src/geo/coordinates.go
go/src/geo/landmark.go
go/src/locked
panic: open go/src/locked: permission denied

goroutine 1 [running]:
main.scanDirectory(0xc420014220, 0xd)
        /Users/jay/files.go:37 +0x29a
main.scanDirectory(0xc420014130, 0x6)
        /Users/jay/files.go:43 +0x1ed
main.scanDirectory(0x10c4148, 0x2)
        /Users/jay/files.go:43 +0x1ed
main.main()
        /Users/jay/files.go:52 +0x36
exit status 2
```

When to panic



We'll show you a way to prevent the program from crashing in a moment. But it's true that calling `panic` is rarely the ideal way to deal with errors.

Things like inaccessible files, network failures, and bad user input should usually be considered "normal," and should be handled gracefully though `error` values. Generally, calling `panic` should be reserved for "impossible" situations: errors that indicate a bug in the program, not a mistake on the user's part.

Here's a program that uses `panic` to indicate a bug. It awards a prize hidden behind one of three virtual doors. The `doorNumber` variable is populated not with user input, but with a random number chosen by the `rand.Intn` function. If `doorNumber` contains any number other than 1, 2, or 3, it's not user error, it's a bug in the program.

So it makes sense to call `panic` if `doorNumber` contains an invalid value. It *should* never happen, and if it does, we want to stop the program before it behaves in unexpected ways.

```
package main

import (
        "fmt"
        "math/rand"
        "time"
)

func awardPrize() {
        doorNumber := rand.Intn(3) + 1
        if doorNumber == 1 {
                fmt.Println("You win a cruise!")
        } else if doorNumber == 2 {
                fmt.Println("You win a car!")
        } else if doorNumber == 3 {
                fmt.Println("You win a goat!")
        } else {
                panic("invalid door number")
        }
}

func main() {
        rand.Seed(time.Now().Unix())
        awardPrize()
}
```

Generate a random integer between 1 and 3.

No other number should be generated, but if it is, panic.

```
You win a cruise!
```



A code sample and its output are shown below, but we've left some blanks in the output. See if you can fill them in.

```
package main

import "fmt"

func snack() {
       defer fmt.Println("Closing refrigerator")
       fmt.Println("Opening refrigerator")
       panic("refrigerator is empty")
}

func main() {
       snack()
}
```

Output:

```
_____________________
_____________________
panic: _____________________

goroutine 1 [running]:
main._______()
       /tmp/main.go:8 +0xe0
main.main()
       /tmp/main.go:12 +0x20
```

Answers on page 378.

The "recover" function

Changing our `scanDirectory` function to use `panic` instead of returning an error greatly simplified the error handling code. But panicking is also causing our program to crash with an ugly stack trace. We'd rather just show users the error message.

Go offers a built-in `recover` function that can stop a program from panicking. We'll need to use it to exit the program gracefully.

When you call `recover` during normal program execution, it just returns `nil` and does nothing else:

```
package main

import "fmt"

func main() {
      fmt.Println(recover())
}
```

If you call "recover" in a program that isn't panicking...

```
<nil>
```

...it does nothing, and returns nil.

If you call `recover` when a program is panicking, it will stop the panic. But when you call `panic` in a function, that function stops executing. So there's no point calling `recover` in the same function as `panic`, because the panic will continue anyway:

```
func freakOut() {
      panic("oh no")
      recover()
}
func main() {
      freakOut()
      fmt.Println("Exiting normally")
}
```

The panic stops the rest of the freakOut function from running...

...so this will never be run!

Program crashes anyway!

```
panic: oh no

goroutine 1 [running]:
main.freakOut()
      /tmp/main.go:4 +0x40
main.main()
      /tmp/main.go:8 +0x20
```

But there *is* a way to call `recover` when a program is panicking... During a panic, any deferred function calls are completed. So you can place a call to `recover` in a separate function, and use `defer` to call that function before the code that panics.

```
func calmDown() {
      recover()
}
func freakOut() {
      defer calmDown()
      panic("oh no")
}
func main() {
      freakOut()
      fmt.Println("Exiting normally")
}
```

Call "recover" in this other function.

Defer a call to the function that recovers.

If the program panics after that, the deferred function call will recover!

Program exits normally.

```
Exiting normally
```

The "recover" function (continued)

Calling `recover` will *not* cause execution to resume at the point of the panic, at least not exactly. The function that panicked will return immediately, and none of the code in that function's block following the panic will be executed. After the function that panicked returns, however, normal execution resumes.

```
func calmDown() {
       recover()
}
func freakOut() {
       defer calmDown()
       panic("oh no")
       fmt.Println("I won't be run!")
}
func main() {
       freakOut()
       fmt.Println("Exiting normally")
}
```

When we recover, freakOut returns at this point.

This code after the panic will never be run!

But this code runs after freakOut returns.

```
Exiting normally
```

The panic value is returned from recover

As we mentioned, when there is no panic, calls to `recover` return `nil`.

If you call "recover" in a program that isn't panicking...

```
func main() {
       fmt.Println(recover())
}
```

```
<nil>
```

...it does nothing, and returns nil.

But when there *is* a panic, `recover` returns whatever value was passed to `panic`. This can be used to gather information about the panic, to aid in recovering or to report errors to the user.

```
func calmDown() {
       fmt.Println(recover())
}
func main() {
       defer calmDown()
       panic("oh no")
}
```

Call "recover" and print the panic value.

This is the value that will be returned from "recover".

```
oh no
```

The panic value is returned from recover (continued)

Back when we introduced the `panic` function, we mentioned the type for its argument is `interface{}`, the empty interface, so that `panic` can accept any value. Likewise, the type for `recover`'s return value is also `interface{}`. You can pass `recover`'s return value to `fmt` functions like `Println` (which accept `interface{}` values), but you won't be able to call methods on it directly.

Here's some code that passes an `error` value to `panic`. But in doing so, the `error` is converted to an `interface{}` value. When the deferred function calls `recover` later, that `interface{}` value is what's returned. So even though the underlying `error` value has an `Error` method, attempting to call `Error` on the `interface{}` value results in a compile error.

```
func calmDown() {
        p := recover()
        fmt.Println(p.Error())
}
func main() {
        defer calmDown()
        err := fmt.Errorf("there's an error")
        panic(err)
}
```

Returns an interface{} value

Even though the underlying "error" value has an Error method, the interface{} value doesn't!

Instead of a string, pass an error value to "panic".

Compile error!

```
p.Error undefined (type interface {} is interface with no methods)
```

To call methods or do anything else with the panic value, you'll need to convert it back to its underlying type using a type assertion.

Here's an update to the above code that takes the return value of `recover` and converts it back to an `error` value. Once that's done, we can safely call the `Error` method.

```
func calmDown() {
        p := recover()
        err, ok := p.(error)
        if ok {
                fmt.Println(err.Error())
        }
}
func main() {
        defer calmDown()
        err := fmt.Errorf("there's an error")
        panic(err)
}
```

Assert that the type of the panic value is "error".

Now that we have an "error" value, we can call the Error method.

```
there's an error
```

Recovering from panics in scanDirectory

When we last left our *files.go* program, adding a call to `panic` in the `scanDirectory` function cleaned up our error handling code, but it also caused the program to crash. We can take everything we've learned so far about `defer`, `panic`, and `recover` and use it to print an error message and exit the program gracefully.

We do this by adding a `reportPanic` function, which we'll call using `defer` in `main`. We do this *before* calling `scanDirectory`, which could potentially panic.

Within `reportPanic`, we call `recover` and store the panic value it returns. If the program is panicking, this will stop the panic.

But when `reportPanic` is called, we *don't* know whether the program is actually panicking or not. The deferred call to `reportPanic` will be made regardless of whether `scanDirectory` calls `panic` or not. So the first thing we do is test whether the panic value returned from `recover` is `nil`. If it is, it means there's no panic, so we return from `reportPanic` without doing anything further.

But if the panic value is *not* `nil`, it means there's a panic, and we need to report it.

Because `scanDirectory` passes an `error` value to `panic`, we use a type assertion to convert the `interface{}` panic value to an `error` value. If that conversion is successful, we print the `error` value.

With these changes in place, instead of an ugly panic log and stack trace, our users will simply see an error message!

```
Shell Edit View Window Help
$ go run files.go
go
go/src
go/src/geo
go/src/geo/coordinates.go
go/src/geo/landmark.go
go/src/locked
open go/src/locked: permission denied
```

```
package main

import (
    "fmt"
    "io/ioutil"
    "path/filepath"
)

func reportPanic() {
    p := recover()
    if p == nil {
        return
    }
    err, ok := p.(error)
    if ok {
        fmt.Println(err)
    }
}

func scanDirectory(path string) {
    fmt.Println(path)
    files, err := ioutil.ReadDir(path)
    if err != nil {
        panic(err)
    }

    for _, file := range files {
        filePath := filepath.Join(path, file.Name())
        if file.IsDir() {
            scanDirectory(filePath)
        } else {
            fmt.Println(filePath)
        }
    }
}

func main() {
    defer reportPanic()
    scanDirectory("go")
}
```

Add this new function.

Call "recover" and store its return value.

If "recover" returned nil, there is no panic...

...so do nothing.

Otherwise, get the underlying "error" value...

...and print it.

Before calling code that might panic, defer a call to our new reportPanic function.

Reinstating a panic

There's one other potential issue with `reportPanic` that we need to address. Right now, it intercepts *any* panic, even ones that didn't originate from `scanDirectory`. And if the panic value can't be converted to an `error` type, `reportPanic` won't print it.

We can test this out by adding another call to `panic` within `main` using a `string` argument:

```
func main() {
        defer reportPanic()
        panic("some other issue")
        scanDirectory("go")
}
```

Introduce a new panic, with a string panic value.

```
Shell Edit View Window Help
$ go run files.go
$
```

No output!

The `reportPanic` function recovers from the new panic, but because the panic value isn't an `error`, `reportPanic` doesn't print it. Our users are left wondering why the program failed!

A common strategy for dealing with unanticipated panics you're not prepared to recover from is to simply renew the panic state. Panicking again is usually appropriate because, after all, this is an unanticipated situation.

The code at right updates `reportPanic` to handle unanticipated panics. If the type assertion to convert the panic value to an `error` succeeds, we simply print it as before. But if it fails, we simply call `panic` again with the same panic value.

Running *files.go* again shows that the fix works: `reportPanic` recovers from our test call to `panic`, but then panics again when the `error` type assertion fails. Now we can remove the call to `panic` in `main`, confident that any other unanticipated panics will be reported!

```
Shell Edit View Window Help
$ go run files.go
panic: some other issue [recovered]
        panic: some other issue

goroutine 1 [running]:
main.reportPanic()
        /Users/jay/files.go:27 +0xd7
panic(0x109ee80, 0x10d1c80)
        /go/.../panic.go:505 +0x229
main.main()
        /Users/jay/files.go:52 +0x55
exit status 2
```

```
func reportPanic() {
    p := recover()
    if p == nil {
        return
    }
    err, ok := p.(error)
    if ok {
        fmt.Println(err)
    } else {
        panic(p)
    }
}
```

If the panic value isn't an error, resume panicking with the same value.

```
func scanDirectory(path string) {
    fmt.Println(path)
    files, err := ioutil.ReadDir(path)
    if err != nil {
        panic(err)
    }
    // Code here omitted
}
```

Don't forget to remove this test panic once you're sure reportPanic works!

```
func main() {
    defer reportPanic()
    panic("some other issue")
    scanDirectory("go")
}
```

there are no Dumb Questions

Q: I've seen other programming languages that have "exceptions." The `panic` and `recover` functions seem to work in a similar way. Can I use them like exceptions?

A: We strongly recommend against it, and so do the Go language maintainers. It can even be said that using `panic` and `recover` is discouraged by the design of the language itself. In a conference keynote in 2012, Rob Pike (one of the creators of Go) described `panic` and `recover` as "intentionally clumsy." That means that when designing Go, its creators didn't try to make `panic` and `recover` easy or pleasant to use, so that they'd be used *less* often.

This is the Go designers' response to one of the major weaknesses of exceptions: they can make program flow much more complex. Instead, Go developers are encouraged to handle errors the exact same way they handle the other parts of their program: with `if` and `return` statements, along with `error` values. Sure, dealing with errors directly within a function can make that function's code a little longer, but that beats not dealing with the errors at all. (The Go creators found many developers using exceptions would just raise an exception and then not properly handle it later.) Dealing with errors directly also makes it immediately obvious how the error is handled—you don't have to go look at a different part of the program to see the error handling code.

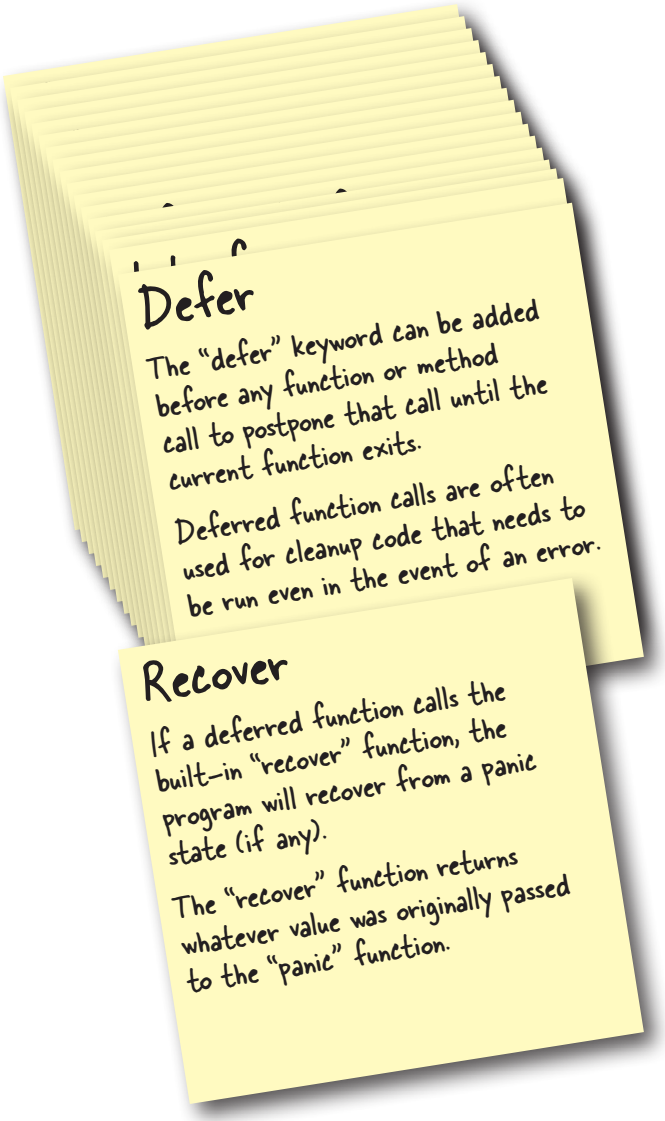
So don't look for an equivalent to exceptions in Go. That feature has been left out, on purpose. It may require a period of adjustment for developers used to using exceptions, but the Go maintainers believe it makes for better software in the end.

You can review a summary of Rob Pike's talk at: https://talks.golang.org/2012/splash.article#TOC_16.



Your Go Toolbox

That's it for Chapter 12! You've added deferred function calls and recovery from panics to your toolbox.



BULLET POINTS

- Returning early from a function with an error value is a good way to indicate an error has occurred, but it can prevent cleanup code later in the function from being run.
- You can use the `defer` keyword to call your cleanup function immediately after the code that requires cleanup. That will set up the cleanup code to run when the current function exits, whether or not there was an error.
- You can call the built-in `panic` function to cause your program to panic.
- Unless the built-in `recover` function is called, a panicking program will crash with a log message.
- You can pass any value as an argument to `panic`. That value will be converted to a string and printed as part of the log message.
- A panic log message includes a stack trace, a list of all active function calls that can be useful for debugging.
- When a program panics, any deferred function calls will still be made, allowing cleanup code to be executed before a crash.
- Deferred functions can also call the built-in `recover` function, which will cause the program to resume normal execution.
- If `recover` is called when there is no panic, it simply returns `nil`.
- If `recover` is called during a panic, it returns the value that was passed to `panic`.
- Most programs should panic only in the event of an unanticipated error. You should think about all possible errors your program might encounter (such as missing files or badly formatted data), and handle those using `error` values instead.

Code Magnets Solution

```
func find(item string, slice []string) bool {
	for _, sliceItem := range slice {
		if item == sliceItem {
			return true
		}
	}
	return false
}

type Refrigerator []string

func (r Refrigerator) Open() {
	fmt.Println("Opening refrigerator")
}
func (r Refrigerator) Close() {
	fmt.Println("Closing refrigerator")
}

func main() {
	fridge := Refrigerator{"Milk", "Pizza", "Salsa"}
	for _, food := range []string{"Milk", "Bananas"} {
		err := fridge.FindFood(food)
		if err != nil {
			log.Fatal(err)
		}
	}
}
```

```
func (r Refrigerator) FindFood(food string) error {
	r.Open()
	defer r.Close()
	if find(food, r) {
		fmt.Println("Found", food)
	} else {
		return fmt.Errorf("%s not found", food)
	}
	return nil
}
```

Close will be called when FindFood exits, regardless of whether there was an error.

Refrigerator's Close method is called when food is found.

Close is also called when food is not found.

```
Opening refrigerator
Found Milk
Closing refrigerator
Opening refrigerator
Closing refrigerator
2018/04/09 22:12:37 Bananas not found
```



A code sample and its output are shown below, but we've left some blanks in the output. See if you can fill them in.

```
package main

import "fmt"

func snack() {
       defer fmt.Println("Closing refrigerator")
       fmt.Println("Opening refrigerator")
       panic("refrigerator is empty")
}

func main() {
       snack()
}
```

This call was deferred, so it's not made until the "snack" function exits (during the panic).

Output:

```
Opening refrigerator
Closing refrigerator
panic: refrigerator is empty

goroutine 1 [running]:
main.snack()
       /tmp/main.go:8 +0xe0
main.main()
       /tmp/main.go:12 +0x20
```

Goroutines and Channels



Working on one thing at a time isn't always the fastest way to finish a task. Some big problems can be broken into smaller tasks. **Goroutines** let your program work on several different tasks at once. Your goroutines can coordinate their work using **channels**, which let them send data to each other *and* synchronize so that one goroutine doesn't get ahead of another. Goroutines let you take full advantage of computers with multiple processors, so that your programs run as fast as possible!

Retrieving web pages



This chapter is going to be about finishing work faster by doing several tasks simultaneously. But first, we need a big task that we can break into little parts. So bear with us for a couple pages while we set the scene...

The smaller a web page is, the faster it loads in visitors' browsers. We need a tool that can measure the sizes of pages, in bytes.

This shouldn't be too difficult, thanks to Go's standard library. The program below uses the `net/http` package to connect to a site and retrieve a web page with just a few function calls.

We pass the URL of the site we want to the `http.Get` function. It will return an `http.Response` object, plus any error it encountered.

The `http.Response` object is a struct with a `Body` field that represents the content of the page. `Body` satisfies the `io` package's `ReadCloser` interface, meaning it has a `Read` method (which lets us read the page data), and a `Close` method that releases the network connection when we're done.

We `defer` a call to `Close`, so the connection gets released after we're done reading from it. Then we pass the response body to the `ioutil` package's `ReadAll` function, which will read its entire contents and return it as a slice of `byte` values.

We haven't covered the `byte` type yet; it's one of Go's basic types (like `float64` or `bool`), and it's used for holding raw data, such as you might read from a file or network connection. A slice of `byte` values won't show us anything meaningful if we print it directly, but if you do a type conversion from a slice of `byte` values to a `string`, you'll get readable text back. (That is, assuming the data represents readable text.) So we end by converting the response body to a `string`, and printing it.

```
package main

import (
        "fmt"
        "io/ioutil"
        "log"
        "net/http"
)

func main() {
        response, err := http.Get("https://example.com")
        if err != nil {
                log.Fatal(err)
        }
        defer response.Body.Close()
        body, err := ioutil.ReadAll(response.Body)
        if err != nil {
                log.Fatal(err)
        }
        fmt.Println(string(body))
}
```

Call http.Get with the URL we want to retrieve.

Release the network connection once the "main" function exits.

Read all the data in the response.

Convert the data to a string and print it.

```
File  Edit  Window  Help
$ go run temp.go
<!doctype html>
<html>
<head>
    <title>Example Domain</title>
    <meta charset="utf-8" />
...
```

The HTML page content

If we save this code to a file and run it with `go run`, it will retrieve the HTML content of the *https://example.com* page, and display it.

Retrieving web pages (continued)



If you want more info on the functions and types used in this program, you can get it via the `go doc` command (which we learned about back in Chapter 4) in your terminal. Try the commands at the right to bring up the documentation. (Or if you prefer, you can look them up in your browser using your favorite search engine.)

```
File  Edit  Window  Help
go doc http Get
go doc http Response
go doc io ReadCloser
go doc ioutil ReadAll
```

Go's documentation will give you more insight into how this program works!

From there, it's not too difficult to convert the program to print the size of multiple pages.

We can move the code that retrieves the page to a separate `responseSize` function, which takes the URL to retrieve as a parameter. We'll print the URL we're retrieving just for debugging purposes. The code to call `http.Get`, read the response, and release the connection will be mostly unchanged. Finally, instead of converting the slice of bytes from the response to a `string`, we simply call `len` to get the slice's length. This gives us the length of the response in bytes, which we print.

We update our `main` function to call `responseSize` with several different URLs. When we run the program, it will print the URLs and page sizes.

```
package main

import (
	"fmt"
	"io/ioutil"
	"log"
	"net/http"
)

func main() {
	responseSize("https://example.com/")
	responseSize("https://golang.org/")
	responseSize("https://golang.org/doc")
}

func responseSize(url string) {
	fmt.Println("Getting", url)
	response, err := http.Get(url)
	if err != nil {
		log.Fatal(err)
	}
	defer response.Body.Close()
	body, err := ioutil.ReadAll(response.Body)
	if err != nil {
		log.Fatal(err)
	}
	fmt.Println(len(body))
}
```

Get the sizes of several pages.

Take the URL as a parameter.

Move the code that gets the page to a separate function.

Print which URL we're retrieving.

Get the given URL.

The size of the slice of bytes is the same as the size of the page.

```
Getting https://example.com/
1270
Getting https://golang.org/
8766
Getting https://golang.org/doc
13078
```

Page URLs and page sizes (in bytes)

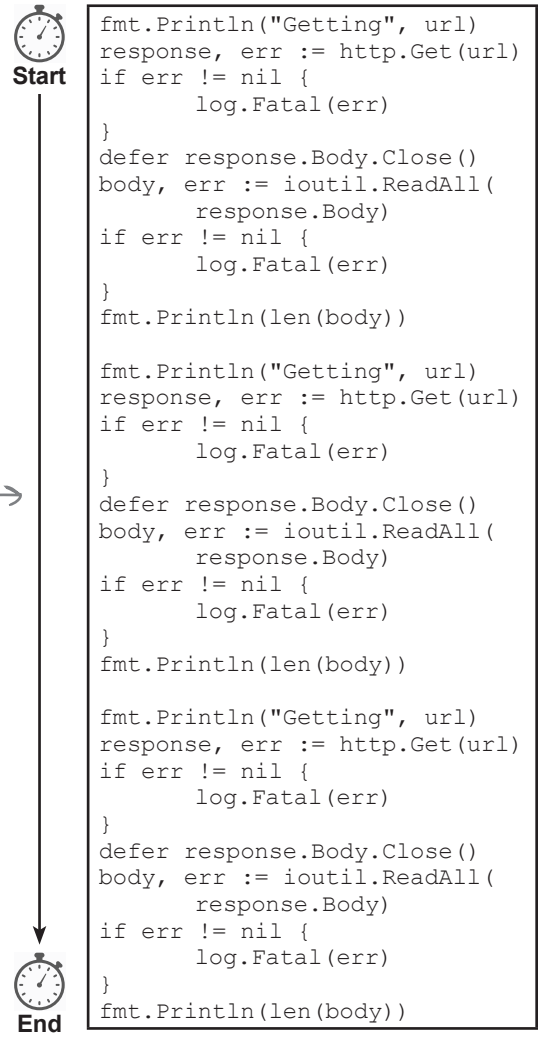


Multitasking

And now we get to the point of this chapter: finding a way to speed programs up by performing multiple tasks at the same time.

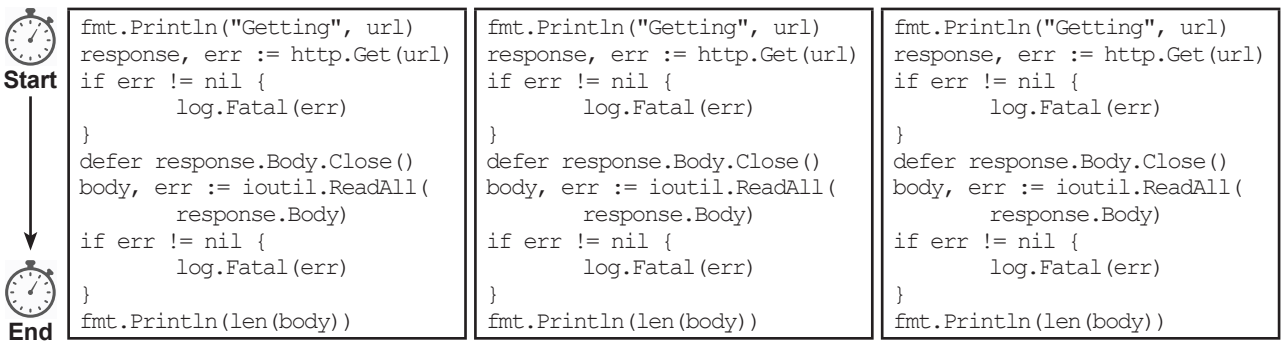
Our program makes several calls to `responseSize`, one at a time. Each call to `responseSize` establishes a network connection to the website, waits for the site to respond, prints the response size, and returns. Only when one call to `responseSize` returns can the next begin. If we had one big long function where the all code was repeated three times, it would take the same amount of time to run as our three calls to `responseSize`.

Three sequential calls to responseSize take this long...



But what if there were a way to run all three calls to `responseSize` at once? The program could complete in as little as a third of the time!

If all the calls to responseSize ran at the same time, the program would complete much faster!



Concurrency using goroutines

When `responseSize` makes the call to `http.Get`, your program has to sit there and wait for the remote website to respond. It's not doing anything useful while it waits.

A different program might have to wait for user input. And another might have to wait while data is read in from a file. There are lots of situations where programs are just sitting around waiting.

Concurrency allows a program to pause one task and work on other tasks. A program waiting for user input might do other processing in the background. A program might update a progress bar while reading from a file. Our `responseSize` program might make other network requests while it waits for the first request to complete.

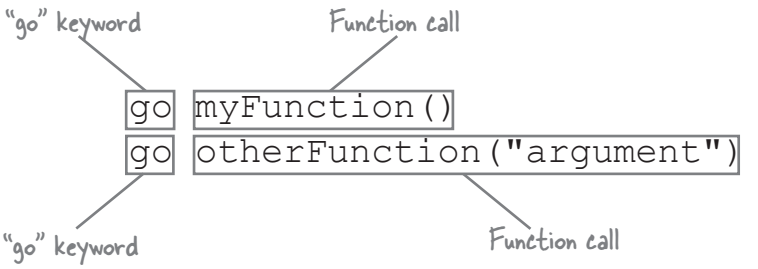
If a program is written to support concurrency, then it may also support **parallelism**: running tasks *simultaneously*. A computer with only one processor can only run one task at a time. But most computers these days have multiple processors (or one processor with multiple cores). Your computer may divide concurrent tasks among different processors to run them at the same time. (It's rare to manage this directly; the operating system usually handles it for you.)

Breaking large tasks into smaller subtasks that can be run concurrently can sometimes mean big speed increases for your programs.

In Go, concurrent tasks are called **goroutines**. Other programming languages have a similar concept called *threads*, but goroutines require less computer memory than threads, and less time to start up and stop, meaning you can run more goroutines at once.

They're also easier to use. To start another goroutine, you use a `go` statement, which is just an ordinary function or method call with the `go` keyword in front of it:

Goroutines allow for concurrency: pausing one task to work on others. And in some situations they allow parallelism: working on multiple tasks simultaneously!



Notice that we say *another* goroutine. The `main` function of every Go program is started using a goroutine, so every Go program runs at least one goroutine. You've been using goroutines all along, without knowing it!

Using goroutines

Here's a program that makes function calls one at a time. The `a` function uses a loop to print the string `"a"` 50 times, and the `b` function prints the string `"b"` 50 times. The `main` function calls `a`, then `b`, and finally prints a message when it exits.

```
package main

import "fmt"

func a() {
	for i := 0; i < 50; i++ {
		fmt.Print("a")
	}
}

func b() {
	for i := 0; i < 50; i++ {
		fmt.Print("b")
	}
}

func main() {
	a()
	b()
	fmt.Println("end main()")
}
```

```
aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaabbbb
```

Using goroutines (continued)

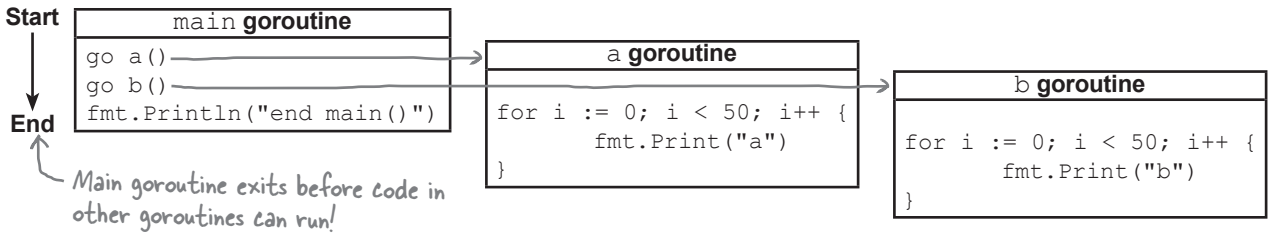
But if we run the program now, the only output we'll see is from the `Println` call at the end of the `main` function—we won't see anything from the `a` or `b` functions!

```
func main() {
        go a()
        go b()
        fmt.Println("end main()")
}
```

Where's the output from the "a" and "b" functions?

```
end main()
```

Here's the problem: Go programs stop running as soon as the `main` goroutine (the goroutine that calls the `main` function) ends, even if other goroutines are still running. Our `main` function completes before the code in the `a` and `b` functions has a chance to run.



We need to keep the `main` goroutine running until the goroutines for the `a` and `b` functions can finish. To do this properly, we're going to need another feature of Go called *channels*, but we won't be covering those until later in the chapter. So for now, we'll just pause the `main` goroutine for a set amount of time so the other goroutines can run.

We'll use a function from the `time` package, called `Sleep`, which pauses the current goroutine for a given amount of time. Calling `time.Sleep(time.Second)` within the `main` function will cause the `main` goroutine to pause for 1 second.

```
func main() {
        go a()
        go b()
        time.Sleep(time.Second)
        fmt.Println("end main()")
}
```

Pause the main goroutine for 1 second.

This gives the other goroutines enough time to run.

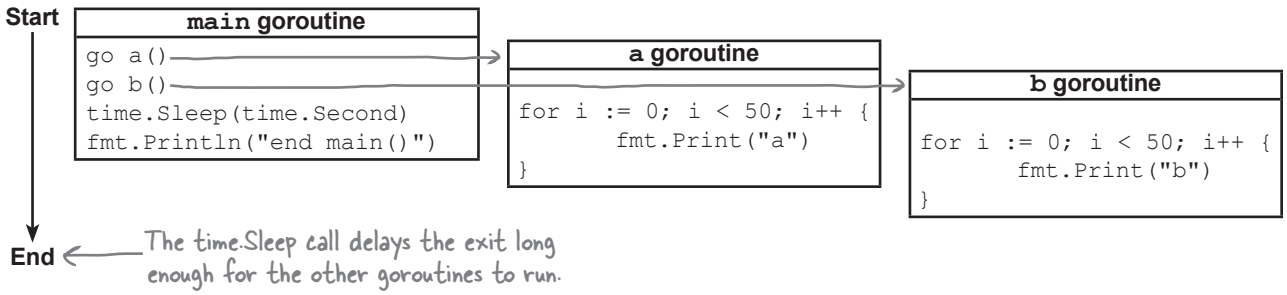
```
aaaaaaaaaaaaaaaaaaaaaabbbbbaaa
aaaaaaaabbbbbbbbbbbaaaaaaaaaa
abbaaaaabbbbbbbbbbbbbbbbbbbbb
bbbbbbbbbbend main()
```

When time.Sleep returns, the main goroutine finishes running.

If we rerun the program, we'll see the output from the `a` and `b` functions again as their goroutines finally get a chance to run. The output of the two will be mixed as the program switches between the two goroutines. (The pattern you get may be different than what's shown here.) When the `main` goroutine wakes back up, it makes its call to `fmt.Println` and exits.

Using goroutines (continued)

The call to `time.Sleep` in the `main` goroutine gives more than enough time for both the `a` and `b` goroutines to finish running.



Using goroutines with our responseSize function

It's pretty easy to adapt our program that prints web page sizes to use goroutines. All we have to do is add the `go` keyword before each of the calls to `responseSize`.

To prevent the `main` goroutine from exiting before the `responseSize` goroutines can finish, we'll also need to add a call to `time.Sleep` in the `main` function.

```
package main

import (
	"fmt"
	"io/ioutil"
	"log"
	"net/http"
	"time"
)

func main() {
	go responseSize("https://example.com/")
	go responseSize("https://golang.org/")
	go responseSize("https://golang.org/doc")
	time.Sleep(5 * time.Second)
}

func responseSize(url string) {
	fmt.Println("Getting", url)
	response, err := http.Get(url)
	if err != nil {
		log.Fatal(err)
	}
	defer response.Body.Close()
	body, err := ioutil.ReadAll(response.Body)
	if err != nil {
		log.Fatal(err)
	}
	fmt.Println(len(body))
}
```

Add the "time" package.

Convert the responseSize calls to Go statements.

Sleep for 5 seconds.

Sleeping for just 1 second may not be enough time for the network requests to complete, though. Calling `time.Sleep(5 * time.Second)` will make the goroutine sleep for 5 seconds. (If you're trying this on a slow or unresponsive network, you may need to increase that time.)

Using goroutines with our responseSize function (continued)

If we run the updated program, we'll see it print the URLs it's retrieving all at once, as the three `responseSize` goroutines start up concurrently.

The three calls to `http.Get` are made concurrently as well; the program doesn't wait until one response comes back before sending out the next request. As a result the three response sizes are printed much sooner using goroutines than they were with the earlier, sequential version of the program. The program still takes 5 seconds to finish, however, as we wait for the call to `time.Sleep` in `main` to complete.

PrintIn calls at the start of responseSize run all at once.

Response sizes are printed as soon as each site responds.

```
Getting https://example.com/
Getting https://golang.org/doc
Getting https://golang.org/
1270
8766
13078
```

We're not exerting any control over the order that calls to `responseSize` are executed in, so if we run the program again, we may see the requests happen in a different order.

Requests for pages may happen in a different order.

```
Getting https://golang.org/doc
Getting https://golang.org/
Getting https://example.com/
1270
8766
13078
```

The program takes 5 seconds to complete even if all the sites respond faster than that, so we're still not getting that great a speed gain from the switch to goroutines. Even worse, 5 seconds may not be *enough* time if the sites take a long time to respond. Sometimes, you may see the program end before all the responses have arrived.

The call to time.Sleep could finish and the program could end before all the sites respond!

```
Getting https://golang.org/doc
Getting https://golang.org/
Getting https://example.com/
1270
```

It's becoming clear that `time.Sleep` is not the ideal way to wait for other goroutines to complete. Once we look at channels in a few pages, we'll have a better alternative.

We don't directly control when goroutines run

We may see the `responseSize` goroutines run in a different order each time the program is run:

```
Getting https://example.com/
Getting https://golang.org/doc
Getting https://golang.org/
```

```
Getting https://golang.org/doc
Getting https://golang.org/
Getting https://example.com/
```

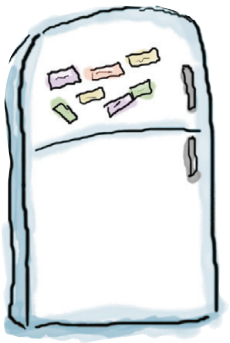
We also had no way of knowing when the previous program would switch between the `a` and `b` goroutines:

```
aaaaaabbbbbbbbbbbbbbb
bbbbbbbaaaaaaaaaaaaaaa
aaaaaaaaaaaaaaaaaaaaab
bbbbbbbbbbbbbbbbbbbbb
bbbbaaaaaaaaend main()
```

```
bbbbbbbbbbbbbbbbbbaaa
aaaabbbbbbbbbbbbaaaaa
aaaaaaaaaaaaaaaaaaaaa
aaaaaaaaaaaaaaaabbbbbb
bbbbbbbbbbbbend main()
```

```
aaaaaaaaaaaaaaaaaaaaa
aaaaaaaaaaaaaaaaaaaaa
aaaaaabbbbbbbbbbbbbbb
bbbbbbbbbbbbbbbbbbbbb
bbbbbbbbbbbbend main()
```

Under normal circumstances, Go makes no guarantees about when it will switch between goroutines, or for how long. This allows goroutines to run in whatever way is most efficient. But if the order your goroutines run in is important to you, you'll need to synchronize them using channels (which we'll look at shortly).



Code Magnets

A program that uses goroutines is scrambled up on the fridge. Can you reconstruct the code snippets to make a working program that will produce output *similar* to the given sample? (It's not possible to predict the order of execution of goroutines, so don't worry, your program's output doesn't need to exactly match the output shown.)

```
(s string)
```

```
repeat
```

```
repeat
```

```
time.Sleep(time.Second)
```

```
()
```

```
("x")
```

```
for i := 0; i < 25; i++ {
        fmt.Print(s)
}
```

```
go
```

```
("y")
```

```
go
```

```
package main

import (
        "fmt"
        "time"
)
```

```
{
```

```
}
```

```
{
```

```
}
```

```
func repeat
```

```
func main
```

One possible output →

```
yyyyyyyyyyyyyxxxxxxxxxxxxyyyyyyyxxxxxxxxxxxxyyyyyxx
```

Answers on page 400.

Go statements can't be used with return values

Switching to goroutines brings up another problem we'll need to solve: we can't use function return values in a `go` statement. Suppose we wanted to change the `responseSize` function to return the page size instead of printing it directly:

```
func main() {
        var size int
        size = go responseSize("https://example.com/")
        fmt.Println(size)
        size = go responseSize("https://golang.org/")
        fmt.Println(size)
        size = go responseSize("https://golang.org/doc")
        fmt.Println(size)
        time.Sleep(5 * time.Second)
}

func responseSize(url string) int {
        fmt.Println("Getting", url)
        response, err := http.Get(url)
        if err != nil {
                log.Fatal(err)
        }
        defer response.Body.Close()
        body, err := ioutil.ReadAll(response.Body)
        if err != nil {
                log.Fatal(err)
        }
        return len(body)
}
```

This code is actually invalid!

Add a return value.

Return the response size instead of printing it.

Compile errors

```
./pagesize.go:13:9: syntax error: unexpected go, expecting expression
./pagesize.go:15:9: syntax error: unexpected go, expecting expression
./pagesize.go:17:9: syntax error: unexpected go, expecting expression
```

We'll get compile errors. The compiler stops you from attempting to get a return value from a function called with a `go` statement.

This is actually a good thing. When you call `responseSize` as part of a `go` statement, you're saying, "Go run `responseSize` in a separate goroutine. I'm going to keep running the instructions in this function." The `responseSize` function isn't going to return a value immediately; it has to wait for the website to respond. But the code in your `main` goroutine would expect a return value immediately, and there wouldn't be one yet!

This is true of any function called in a `go` statement, not just long-running functions like `responseSize`. You can't rely on the return values being ready in time, and so the Go compiler blocks any attempt to use them.

You're saying, "go run this; I'm not going to wait."

```
size = go responseSize("https://example.com/")
fmt.Println(size)
```

But then what is the return value?

Go statements can't be used with return values (continued)

Go won't let you use the return value from a function called with a `go` statement, because there's no guarantee the return value will be ready before we attempt to use it:

```
func greeting() string {
      return "hi"
}

func main() {
      fmt.Println(go greeting())
}
```

Function called as goroutine.

Immediate attempt to use function's return value (which may not be ready yet).

Compile error

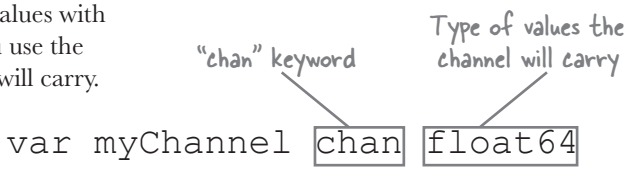
```
syntax error: unexpected go, expecting expression
```

But there *is* a way to communicate between goroutines: **channels**. Not only do channels allow you to send values from one goroutine to another, they ensure the sending goroutine has sent the value before the receiving goroutine attempts to use it.

The only practical way to use a channel is to communicate from one goroutine to another goroutine. So to demonstrate channels, we'll need to be able to do a few things:

- Create a channel.
- Write a function that receives a channel as a parameter. We'll run this function in a separate goroutine, and use it to send values over the channel.
- Receive the sent values in our original goroutine.

Each channel only carries values of a particular type, so you might have one channel for `int` values, and another channel for values with a struct type. To declare a variable that holds a channel, you use the `chan` keyword, followed by the type of values that channel will carry.



To actually create a channel, you need to call the built-in `make` function (the same one you can use to create maps and slices). You pass `make` the type of the channel you want to create (which should be the same as the type of the variable you want to assign it to).

```
var myChannel chan float64
myChannel = make(chan float64)
```

Declare a variable to hold a channel.

Actually create the channel.

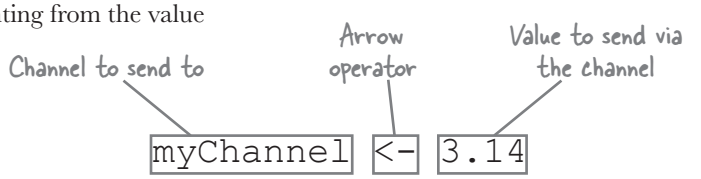
Rather than declare the channel variable separately, in most cases it's easier to just use a short variable declaration:

```
myChannel := make(chan float64)
```

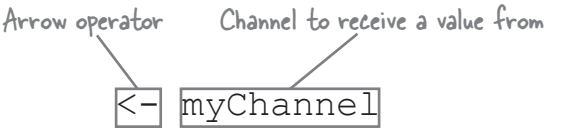
Create a channel and declare a variable at once.

Sending and receiving values with channels

To send a value on a channel, you use the <- operator (that's a less-than symbol followed by a dash). It looks like an arrow pointing from the value you're sending to the channel you're sending it on.



You also use the <- operator to *receive* values from a channel, but the positioning is different: you place the arrow to the *left* of the channel you're receiving from. (It kind of looks like you're pulling a value out of the channel.)



Here's the `greeting` function from the previous page, rewritten to use channels. We've added a `myChannel` parameter to `greeting`, which takes a channel that carries `string` values. Instead of returning a string value, `greeting` now sends a string via `myChannel`.

In the `main` function, we create the channel that we're going to pass to `greeting` using the built-in `make` function. Then we call `greeting` as a new goroutine. Using a separate goroutine is important, because channels should only be used to communicate *between* goroutines. (We'll talk about why in a little bit.) Finally, we receive a value from the channel we passed to `greeting`, and print the string it returns.

```
func greeting(myChannel chan string) {
        myChannel <- "hi"
}

func main() {
        myChannel := make(chan string)
        go greeting(myChannel)
        fmt.Println(<-myChannel)
}
```

Take a channel as a parameter.

Send a value over the channel.

Create a new channel.

Pass the channel to function running in a new goroutine.

Receive a value from the channel.

```
hi
```

We didn't have to pass the value received from the channel straight to `Println`. You can receive from a channel in any context where you need a value. (That is, anywhere you might use a variable or the return value of a function.) So, for example, we could have assigned the received value to a variable first instead:

```
receivedValue := <-myChannel
fmt.Println(receivedValue)
```

We could also have stored the received value in a variable instead.

Synchronizing goroutines with channels

We mentioned that channels also ensure the sending goroutine has sent the value before the receiving channel attempts to use it. Channels do this by **blocking**—by pausing all further operations in the current goroutine. A send operation blocks the sending goroutine until another goroutine executes a receive operation on the same channel. And vice versa: a receive operation blocks the receiving goroutine until another goroutine executes a send operation on the same channel. This behavior allows goroutines to **synchronize** their actions—that is, to coordinate their timing.



Here's a program that creates two channels and passes them to functions in two new goroutines. The `main` goroutine then receives values from those channels and prints them. Unlike our program with the goroutines that printed `"a"` or `"b"` repeatedly, we can predict the output for this program: it will always print `"a"`, then `"d"`, `"b"`, `"e"`, `"c"`, and `"f"` in that order.

```
func abc(channel chan string) {
        channel <- "a"
        channel <- "b"
        channel <- "c"
}

func def(channel chan string) {
        channel <- "d"
        channel <- "e"
        channel <- "f"
}

func main() {
        channel1 := make(chan string)
        channel2 := make(chan string)
        go abc(channel1)
        go def(channel2)
        fmt.Print(<-channel1)
        fmt.Print(<-channel2)
        fmt.Print(<-channel1)
        fmt.Print(<-channel2)
        fmt.Print(<-channel1)
        fmt.Print(<-channel2)
        fmt.Println()
}
```

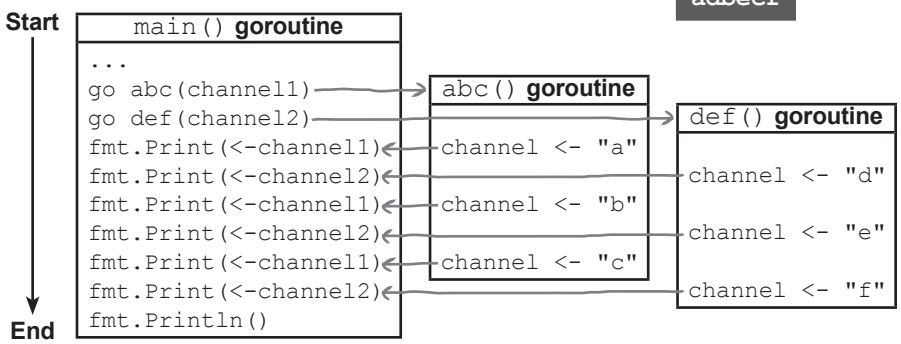
Create two channels.

Pass each channel to a function running in a new goroutine.

Receive and print values from the channels, in order.

```
adbecf
```

We know what the order will be because the `abc` goroutine blocks each time it sends a value to a channel until the `main` goroutine receives from it. The `def` goroutine does the same. The `main` goroutine becomes the orchestrator of the `abc` and `def` goroutines, allowing them to proceed only when it's ready to read the values they're sending.



Observing goroutine synchronization

The `abc` and `def` goroutines send their values over their channels so quickly that it's hard to see what's going on. Here's another program that slows things down so you can see the blocking happen.

We start with a `reportNap` function that causes the current goroutine to sleep for a specified number of seconds. Every second the goroutine is asleep, it will print an announcement that it's still sleeping.

We add a `send` function that will run in a goroutine and send two values to a channel. Before it sends anything, though, it first calls `reportNap` so its goroutine sleeps for 2 seconds.

Name of sleeping goroutine

Time to sleep for

```
func reportNap(name string, delay int) {
        for i := 0; i < delay; i++ {
                fmt.Println(name, "sleeping")
                time.Sleep(1 * time.Second)
        }
        fmt.Println(name, "wakes up!")
}

func send(myChannel chan string) {
        reportNap("sending goroutine", 2)
        fmt.Println("***sending value***")
        myChannel <- "a"
        fmt.Println("***sending value***")
        myChannel <- "b"
}

func main() {
        myChannel := make(chan string)
        go send(myChannel)
        reportNap("receiving goroutine", 5)
        fmt.Println(<-myChannel)
        fmt.Println(<-myChannel)
}
```

Will block on this send while "main" is still asleep

In the `main` goroutine, we create a channel and pass it to `send`. Then we call `reportNap` again so that *this* goroutine sleeps for *5* seconds (3 seconds longer than the `send` goroutine). Finally, we do two receive operations on the channel.

When we run this, we'll see both goroutines sleep for the first 2 seconds. Then the `send` goroutine wakes up and sends its value. But it doesn't do anything further; the send operation blocks the `send` goroutine until the `main` goroutine receives the value.

That doesn't happen right away, because the `main` goroutine still needs to sleep for 3 more seconds. When it wakes up, it receives the value from the channel. Only then is the `send` goroutine unblocked so it can send its second value.

```
receiving goroutine sleeping
sending goroutine sleeping
sending goroutine sleeping
receiving goroutine sleeping
receiving goroutine sleeping
sending goroutine wakes up!
***sending value***
receiving goroutine sleeping
receiving goroutine sleeping
receiving goroutine wakes up!
a
***sending value***
b
```

Both the sending and receiving goroutines are asleep.

The sending goroutine wakes up, and sends a value.

The receiving goroutine is still sleeping.

The receiving goroutine wakes up, and receives a value.

Only then is the sending goroutine unblocked so it can send its second value.



Breaking Stuff is Educational!

Here's the code again for our earliest, simplest demonstration of channels: the `greeting` function, which runs in a goroutine and sends a string value to the `main` goroutine.

Make one of the changes below and try to run the code. Then undo your change and try the next one. See what happens!

```
func greeting(myChannel chan string) {
       myChannel <- "hi"
}

func main() {
       myChannel := make(chan string)
       go greeting(myChannel)
       fmt.Println(<-myChannel)
}
```

If you do this...	...the code will break because...
Send a value to the channel from within the `main` function: `myChannel <- "hi from main"`	You'll get an "`all goroutines are asleep - deadlock!`" error. This happens because the `main` goroutine blocks, waiting for another goroutine to receive from the channel. But the other goroutine doesn't do any receive operations, so the `main` goroutine stays blocked.
Remove the `go` keyword from before the call to `greeting`: ~~`go`~~ `greeting(myChannel)`	This will cause the `greeting` function to run within the `main` goroutine. This also fails with a deadlock error, for the same reason as above: the send operation in `greeting` causes the `main` goroutine to block, but there's no other goroutine to do a receive operation, so it stays blocked.
Delete the line that sends a value to the channel: ~~`myChannel <- "hi"`~~	This also causes a deadlock, but for a different reason: the `main` goroutine tries to *receive* a value, but now there's nothing to *send* a value.
Delete the line that receives a value from the channel: ~~`fmt.Println(<-myChannel)`~~	The send operation in `greeting` causes that goroutine to block. But since there's no receive operation to make the `main` goroutine block as well, `main` completes immediately, and the program ends without producing any output.



Fill in the blanks so that the code below uses values received from two channels to produce the output shown.

```
package main

import "fmt"

func odd(channel chan int) {
       channel __ 1
       channel __ 3
}

func even(channel chan int) {
       channel __ 2
       channel __ 4
}

func main() {
       channelA := _____________
       channelB := _____________
       __ odd(channelA)
       __ even(channelB)
       fmt.Println(_____________)
       fmt.Println(_____________)
       fmt.Println(_____________)
       fmt.Println(_____________)
}
```

Output

```
1
3
2
4
```

Answers on page 400.

Fixing our web page size program with channels

We still have two problems with our program that reports the size of web pages:

- We can't use a return value from the `responseSize` function in a `go` statement.
- Our `main` goroutine was completing before the response sizes were received, so we added a call to `time.Sleep` for 5 seconds. But 5 seconds is too long some times, and too short other times.

```
func main() {
        var size int
        size = go responseSize("https://example.com/")
        fmt.Println(size)
        size = go responseSize("https://golang.org/")
        fmt.Println(size)
        size = go responseSize("https://golang.org/doc")
        fmt.Println(size)
        time.Sleep(5 * time.Second)
}
```

Getting a return value from a go statement is invalid!

The program might exit before all the page sizes are retrieved!

We can use channels to fix both problems at the same time!

First, we remove the `time` package from the `import` statement; we won't be needing `time.Sleep` anymore. Then we update `responseSize` to accept a channel of `int` values. Instead of returning the page size, we'll have `responseSize` send the size via the channel.

```
package main

import (
        "fmt"
        "io/ioutil"
        "log"
        "net/http"
)

func responseSize(url string, channel chan int) {
        fmt.Println("Getting", url)
        response, err := http.Get(url)
        if err != nil {
                log.Fatal(err)
        }
        defer response.Body.Close()
        body, err := ioutil.ReadAll(response.Body)
        if err != nil {
                log.Fatal(err)
        }
        channel <- len(body)
}
```

We won't be using time.Sleep, so remove the "time" package.

We'll pass a channel in to responseSize for it to send page sizes over.

Instead of returning the page size, send it over the channel.

Fixing our web page size program with channels (continued)

In the `main` function, we call `make` to create the channel of `int` values. We update each of the calls to `responseSize` to add the channel as an argument. And finally, we do three receive operations on the channel, one for each value `responseSize` sends.

```
func main() {
        sizes := make(chan int)
        go responseSize("https://example.com/", sizes)
        go responseSize("https://golang.org/", sizes)
        go responseSize("https://golang.org/doc", sizes)
        fmt.Println(<-sizes)
        fmt.Println(<-sizes)
        fmt.Println(<-sizes)
}
```

Make a channel of int values.

Pass the channel with each call to responseSize.

There will be three sends on the channel, so do three receives.

```
Getting https://golang.org/doc
Getting https://example.com/
Getting https://golang.org/
8766
13078
1270
```

If we run this, we'll see that the program completes as rapidly as the websites respond. That time can vary, but in our testing we saw completion times as short as 1 second!

Another improvement we can make is to store the list of URLs we want to retrieve in a slice, and then use loops to call `responseSize`, and to receive values from the channel. This will make our code less repetitive, and will be important if we want to add more URLs later.

We don't need to change `responseSize` at all, just the `main` function. We create a slice of `string` values with the URLs we want. Then we loop over the slice, and call `responseSize` with the current URL and the channel. Finally, we do a second, separate loop that runs once for each URL in the slice, and receives and prints a value from the channel. (It's important to do this in a separate loop. If we received values in the same loop that starts the `responseSize` goroutines, the `main` goroutine would block until the receive completes, and we'd be back to requesting pages one at a time.)

```
func main() {
        sizes := make(chan int)
        urls := []string{"https://example.com/",
                "https://golang.org/", "https://golang.org/doc"}
        for _, url := range urls {
                go responseSize(url, sizes)
        }
        for i := 0; i < len(urls); i++ {
                fmt.Println(<-sizes)
        }
}
```

Move the URLs into a slice.

Call responseSize with each URL.

Receive from the channel once for each send responseSize does.

```
Getting https://golang.org/
Getting https://golang.org/doc
Getting https://example.com/
1270
8766
13078
```

Using loops is much cleaner, but still gets us the same result!

Updating our channel to carry a struct

There's still one issue we need to fix with the `responseSize` function. We have no idea which order the websites will respond in. And because we're not keeping the page URL together with the response size, we have no idea which size belongs to which page!

```
Getting https://golang.org/
Getting https://golang.org/doc
Getting https://example.com/
1270
8766
13078
```

Which response size belongs to which URL?

This won't be difficult to fix, though. Channels can carry composite types like slices, maps, and structs just as easily as they can carry basic types. We can just create a struct type that will store a page URL together with its size, so we can send both over the channel together.

We'll declare a new `Page` type with an underlying `struct` type. `Page` will have a `URL` field that records the page's URL, and a `Size` field for the page's size.

We'll update the channel parameter on `responseSize` to hold the new `Page` type rather than just the `int` page size. We'll have `responseSize` create a new `Page` value with the current URL and the page size, and send that to the channel.

In `main`, we'll update the type the channel holds in the call to `make` as well. When we receive a value from the channel, it will be a `Page` value, so we'll print both its `URL` and `Size` fields.

```
type Page struct {                          <- Declare a struct type with the fields we need.
        URL  string
        Size int
}

func responseSize(url string, channel chan Page) {
        // Omitting identical code...
        channel <- Page{URL: url, Size: len(body)}
}

func main() {
        pages := make(chan Page)            <- Update the type the channel holds.
        urls := []string{"https://example.com/",
                "https://golang.org/", "https://golang.org/doc"}
        for _, url := range urls {
                go responseSize(url, pages)
        }
        for i := 0; i < len(urls); i++ {
                page := <-pages             <- Receive the Page.
                fmt.Printf("%s: %d\n", page.URL, page.Size)
        }
}
```

Channel we pass to responseSize will carry Pages, not ints.

Send back a Page with both the current URL and the page size.

Pass the channel to responseSize.

Print its URL and size together.

```
https://example.com/: 1270
https://golang.org/: 8766
https://golang.org/doc: 13078
```

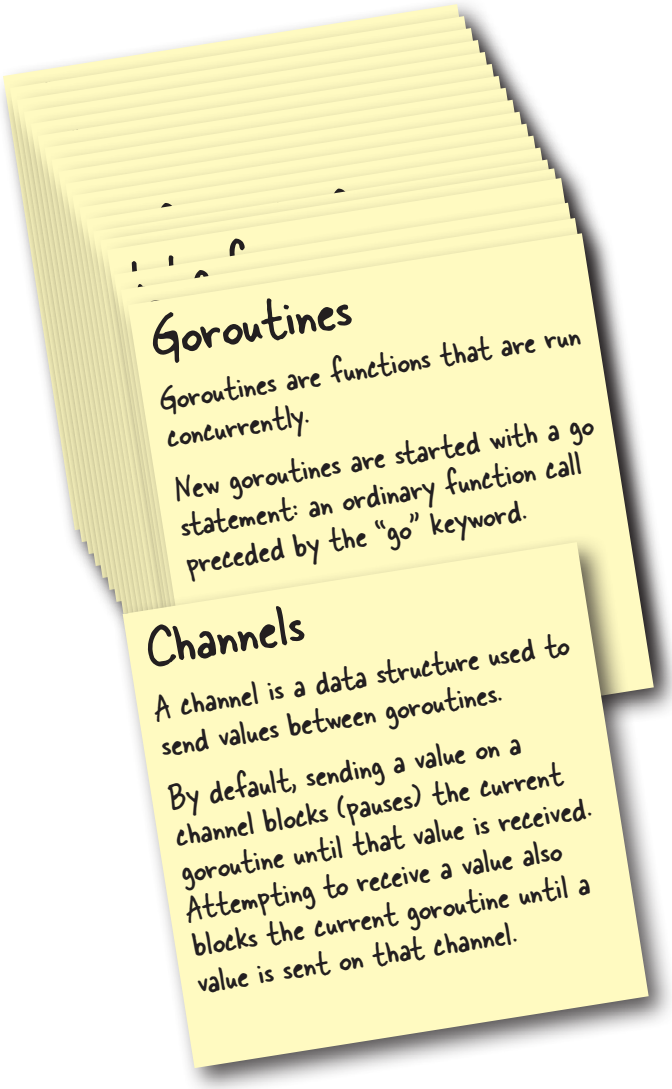
Now the output will pair the page sizes with their URLs. It'll finally be clear again which size belongs to which page.

Before, our program had to request pages one at a time. Goroutines let us start processing the next request while we're waiting for a website to respond. The program completes in as little as one-third of the time!



Your Go Toolbox

That's it for Chapter 13! You've added goroutines and channels to your toolbox.



BULLET POINTS

- All Go programs have at least one goroutine: the one that calls the `main` function when the program starts.
- Go programs end when the `main` goroutine stops, even if other goroutines have not completed their work yet.
- The `time.Sleep` function pauses the current goroutine for a set amount of time.
- Go makes no guarantees about when it will switch between goroutines, or how long it will keep running one goroutine for. This allows the goroutines to run more efficiently, but it means you can't count on operations happening in a particular order.
- Function return values can't be used in a `go` statement, in part because the return value wouldn't be ready when the calling function attempted to use it.
- If you need a value from a goroutine, you'll need to pass it a channel to send the value back on.
- Channels are created by calling the built-in `make` function.
- Each channel only carries values of one particular type; you specify that type when creating the channel.
  `myChannel := make(chan MyType)`
- You send values to channels using the `<-` operator:
  `myChannel <- "a value"`
- The `<-` operator is also used to receive values from a channel:
  `value := <-myChannel`

Code Magnets Solution

```
package main

import (
        "fmt"
        "time"
)
```

```
func repeat (s string) {
        for i := 0; i < 25; i++ {
                fmt.Print(s)
        }
}
```

```
func main () {
        go repeat ("x")
        go repeat ("y")
        time.Sleep(time.Second)
}
```

Run the same function in two different goroutines.

Prevent main goroutine from ending before the others finish.

One possible output

```
yyyyyyyyyyyyyxxxxxxxxxxxxxy
yyyyyyxxxxxxxxxxxxxyyyyyyxx
```



Exercise Solution

```
package main

import "fmt"

func odd(channel chan int) {
        channel <- 1
        channel <- 3
}

func even(channel chan int) {
        channel <- 2
        channel <- 4
}

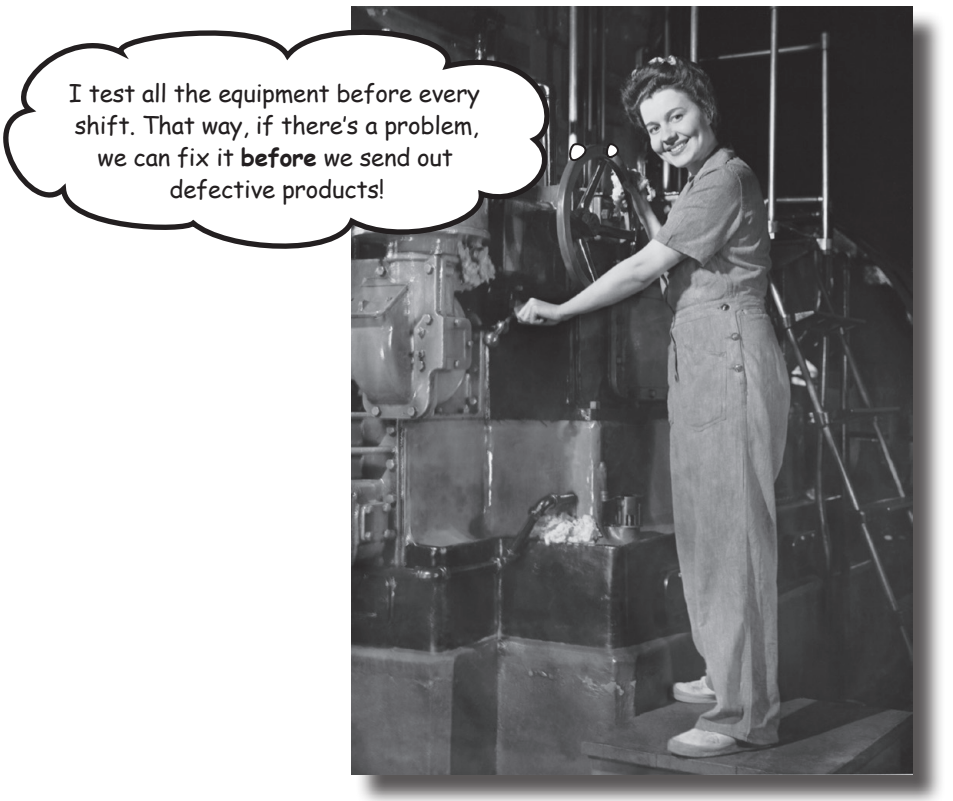
func main() {
        channelA := make(chan int)
        channelB := make(chan int)
        go odd(channelA)
        go even(channelB)
        fmt.Println( <-channelA )
        fmt.Println( <-channelA )
        fmt.Println( <-channelB )
        fmt.Println( <-channelB )
}
```

1
3
2
4

One channel carries values from the "odd" function; the other carries values from "even".

14 code quality assurance

Automated Testing



Are you sure your software is working right now? Really sure?

Before you sent that new version to your users, you presumably tried out the new features to ensure they all worked. But did you try the *old* features to ensure you didn't break any of them? *All* the old features? If that question makes you worry, your program needs **automated testing**. Automated tests ensure your program's components work correctly, even after you change your code. Go's `testing` package and `go test` tool make it easy to write automated tests, using the skills that you've already learned!

Automated tests find your bugs before someone else does

Developer A runs into Developer B at a restaurant they both frequent…

Developer A:

How's the new job going?

Ouch. How did *that* get onto your billing server?

Wow, that long ago… And your tests didn't catch it?

Your automated tests. They didn't fail when the bug got introduced?

What?!

Developer B:

Not so great. I have to head back into the office after dinner. We found a bug that's causing some customers to be billed twice as often as they should be.

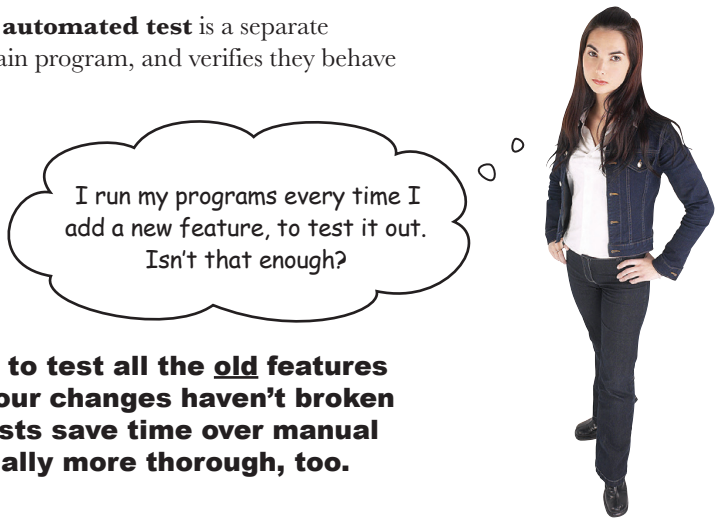
We think it might have gotten introduced a couple of months ago. One of our devs made some changes to the billing code then.

Tests?

Um, we don't have any of those.

Your customers rely on your code. When it fails, it can be disastrous. Your company's reputation is damaged. And *you'll* have to put in overtime fixing the bugs.

That's why automated tests were invented. An **automated test** is a separate program that executes components of your main program, and verifies they behave as expected.



Not unless you're going to test all the old features as well, to make sure your changes haven't broken anything. Automated tests save time over manual testing, and they're usually more thorough, too.

A function we should have had automated tests for

Let's look at an example of a bug that could be caught by automated tests. Here we have a simple package with a function that joins several strings into a single string suitable for use in an English sentence. If there are two items, they'll be joined with the word *and* (as in "apple and orange"). If there are more than two items, commas will be added as appropriate (as in "apple, orange and pear").

One last, great example borrowed from Head First Ruby (which also has a chapter on testing)!



We need this so we can use the strings.Join function.

```
package prose

import "strings"

func JoinWithCommas(phrases []string) string {
        result := strings.Join(phrases[:len(phrases)-1], ", ")
        result += " and "
        result += phrases[len(phrases)-1]
        return result
}
```

Accept a slice of strings to join.

Return the joined string.

Join every phrase except the last with commas.

Insert the word "and" before the last phrase.

Add the last phrase.

The code makes use of the `strings.Join` function, which takes a slice of strings and a string to join them all together with. `Join` returns a single string with all the items from the slice combined, with the joining string separating each entry.

A slice of strings to join

A string to join them together with

```
fmt.Println(strings.Join([]string{"05", "14", "2018"}, "/"))
fmt.Println(strings.Join([]string{"state", "of", "the", "art"}, "-"))
```

```
05/14/2018
state-of-the-art
```

In `JoinWithCommas`, we use the slice operator to gather every phrase in the slice except the last, and pass them to `strings.Join` to join them together in a single string, with a comma and a space between each. Then we add the word *and* (surrounded by spaces), and end the string with the final phrase.

```
[]string{"apple", "orange", "pear", "banana"}
```

```
apple, orange, pear and banana
```

All phrases except the last joined with commas.

Last phrase added following "and".

A function we should have had automated tests for (continued)

Here's a quick program to try our new function. We import our `prose` package and pass a couple slices to `JoinWithCommas`.




```
package main

import (
	"fmt"
	"github.com/headfirstgo/prose"
)

func main() {
	phrases := []string{"my parents", "a rodeo clown"}
	fmt.Println("A photo of", prose.JoinWithCommas(phrases))
	phrases = []string{"my parents", "a rodeo clown", "a prize bull"}
	fmt.Println("A photo of", prose.JoinWithCommas(phrases))
}
```

```
A photo of my parents and a rodeo clown
A photo of my parents, a rodeo clown and a prize bull
```

It works, but there's a small problem with the results. Maybe we're just immature, but we can imagine this leading to jokes that the parents *are* a rodeo clown and a prize bull. And formatting lists in this way could cause other misunderstandings, too.

To resolve any confusion, let's update our package code to place an additional comma before the *and* (as in "apple, orange, and pear"):

```
func JoinWithCommas(phrases []string) string {
	result := strings.Join(phrases[:len(phrases)-1], ", ")
	result += ", and " <-- Add a comma before the "and".
	result += phrases[len(phrases)-1]
	return result
}
```

If we rerun our program, we'll see commas before the *and* in both the resulting strings. Now it should be clear that the parents were in the photo *with* the clown and the bull.

```
A photo of my parents, and a rodeo clown
A photo of my parents, a rodeo clown, and a prize bull
```

There's the new comma!

We've introduced a bug!



Wait! The new code is working correctly with **three** items in the list, but not with **two** items. You've introduced a bug!

Oh, that's true! The function used to return `"my parents and a rodeo clown"` for this list of two items, but an extra comma got included here as well! We were so focused on fixing the list of *three* items that we introduced a bug with lists of *two* items...

A comma doesn't belong here!

```
A photo of my parents, and a rodeo clown
```

If we had automated tests for this function, this problem could have been avoided.

An automated test runs your code with a particular set of inputs and looks for a particular result. As long as your code's output matches the expected value, the test will "pass."

But suppose that you accidentally introduced a bug in your code (like we did with the extra comma). Your code's output would no longer match the expected value, and the test would "fail." You'd know about the bug immediately.



Pass.

For `[]string{"apple", "orange", "pear"}`, `JoinWithCommas` should return `"apple, orange, and pear"`.

Fail!

For `[]string{"apple", "orange"}`, `JoinWithCommas` should return `"apple and orange"`.

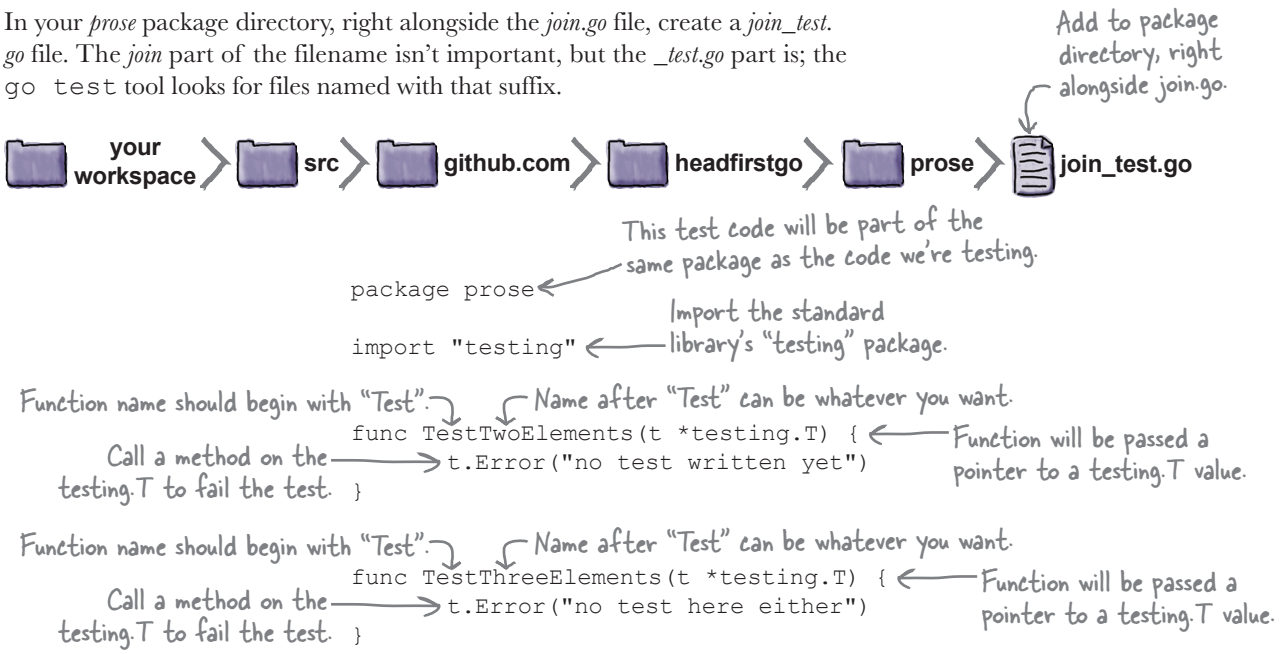
Having automated tests is like having your code inspected for bugs automatically every time you make a change!

Writing tests

Go includes a `testing` package that you can use to write automated tests for your code, and a `go test` command that you can use to run those tests.

Let's start by writing a simple test. We won't test anything practical at first, we're just going to show you how tests work. Then we'll actually use tests to help us fix our `JoinWithCommas` function.

In your *prose* package directory, right alongside the *join.go* file, create a *join_test.go* file. The *join* part of the filename isn't important, but the *_test.go* part is; the `go test` tool looks for files named with that suffix.



The code within the test file consists of ordinary Go functions, but it needs to follow certain conventions in order to work with the `go test` tool:

- You're not required to make your tests part of the same package as the code you're testing, but if you want to access unexported types or functions from the package, you'll need to.
- Tests are required to use a type from the `testing` package, so you'll need to import that package at the top of each test file.
- Test function names should begin with `Test`. (The rest of the name can be whatever you want, but it should begin with a capital letter.)
- Test functions should accept a single parameter: a pointer to a `testing.T` value.
- You can report that a test has failed by calling methods (such as `Error`) on the `testing.T` value. Most methods accept a string with a message explaining the reason the test failed.

Running tests with the "go test" command

To run tests, you use the `go test` command. The command takes the import paths of one or more packages, just like `go install` or `go doc`. It will find all files in those package directories whose names end in *_test.go*, and run every function contained in those files whose name starts with `Test`.

Let's run the tests we just added to our `prose` package. In your terminal, run this command:

```
go test github.com/headfirstgo/prose
```

The test functions will run and print their results.



Because both test functions make a call to the `Error` method on the `testing.T` value passed to them, both tests fail. The name of each failing test function is printed, as well as the line containing the call to `Error`, and the failure message that was given.

At the bottom of the output is the status for the entire `prose` package. If any test within the package fails (as ours did), a status of "`FAIL`" will be printed for the package as a whole.

If we remove the calls to the `Error` method within the tests...

```
func TestTwoElements(t *testing.T) {
}
```
← Remove call to t.Error.

```
func TestThreeElements(t *testing.T) {
}
```
← Remove call to t.Error.

...then we'll be able to rerun the same `go test` command and the tests will pass. Since every test is passing, `go test` will only print a status of "`ok`" for the entire `prose` package.



Testing our actual return values

We can make our tests pass, and we can make them fail. Now let's try writing some tests that will actually help us troubleshoot our `JoinWithCommas` function.

We'll update `TestTwoElements` to show the return value we *expect* from the `JoinWithCommas` function when it's called with a two-element slice. We'll do the same for `TestThreeElements` with a three-element slice. We'll run the tests, and confirm that `TestTwoElements` is currently failing and `TestThreeElements` is passing.

Once our tests are set up the way we want, we'll alter the `JoinWithCommas` function to make the all the tests pass. At that point, we'll know our code is fixed!

In `TestTwoElements`, we'll pass a slice with two elements, `[]string{"apple", "orange"}`, to `JoinWithCommas`. If the result doesn't equal `"apple and orange"`, we'll fail the test. Likewise, in `TestThreeElements`, we'll pass a slice with three elements, `[]string{"apple", "orange", "pear"}`. If the result doesn't equal `"apple, orange, and pear"`, we'll fail the test.

```
func TestTwoElements(t *testing.T) {
        list := []string{"apple", "orange"}
        if JoinWithCommas(list) != "apple and orange" {
                t.Error("didn't match expected value")
        }
}

func TestThreeElements(t *testing.T) {
        list := []string{"apple", "orange", "pear"}
        if JoinWithCommas(list) != "apple, orange, and pear" {
                t.Error("didn't match expected value")
        }
}
```

Pass a list with two elements.

If JoinWithCommas doesn't return the expected string...

...fail the test.

Pass a list with three elements.

If JoinWithCommas doesn't return the expected string...

...fail the test.

If we rerun the tests, the `TestThreeElements` test will pass, but the `TestTwoElements` test will fail.

Only the TestTwoElements test fails.

```
File  Edit  Window  Help
$ go test github.com/headfirstgo/prose
--- FAIL: TestTwoElements (0.00s)
        lists_test.go:13: didn't match expected value
FAIL
FAIL    github.com/headfirstgo/prose    0.006s
```

Testing our actual return values (continued)

This is a *good* thing; it matches what we expected to see based on the output of our `join` program. It means that we'll be able to rely on our tests as an indicator of whether `JoinWithCommas` is working as it should be!

Pass. ✓ For `[]string{"apple", "orange", "pear"}`, `JoinWithCommas` should return `"apple, orange, and pear"`.

Fail! ✗ For `[]string{"apple", "orange"}`, `JoinWithCommas` should return `"apple and orange"`.

Incorrect → `A photo of my parents, and a rodeo clown`
Correct → `A photo of my parents, a rodeo clown, and a prize bull`



Fill in the blanks in the test code below.

your workspace > src > arithmetic > math.go

```
package arithmetic

func Add(a float64, b float64) float64 {
        return a + b
}
func Subtract(a float64, b float64) float64 {
        return a - b
}
```

your workspace > src > arithmetic > math_test.go

```
package _____

import _____

func _____Add(t _____) {
        if _____(1, 2) != 3 {
                _____("1 + 2 did not equal 3")
        }
}

func _____Subtract(t _____) {
        if _____(8, 4) != 4 {
                _____("8 - 4 did not equal 4")
        }
}
```

Answers on page 423.

More detailed test failure messages with the "Errorf" method

Our test failure message isn't very helpful in diagnosing the problem right now. We know there was some value that was expected, and we know the return value from `JoinWithCommas` was different than that, but we don't know what those values were.

```
--- FAIL: TestTwoElements (0.00s)
        lists_test.go:13: didn't match expected value
FAIL
FAIL    github.com/headfirstgo/prose    0.006s
```

What was the expected value? What did we get?

A test function's `testing.T` parameter also has an `Errorf` method you can call. Unlike `Error`, `Errorf` takes a string with formatting verbs, just like the `fmt.Printf` and `fmt.Sprintf` functions. You can use `Errorf` to include additional information in your test's failure messages, such as the arguments you passed to a function, the return value you got, and the value you were expecting.

Here's an update to our tests that uses `Errorf` to generate more detailed failure messages. So that we don't have to repeat strings within each test, we add a `want` variable (as in "the value we *want*") to hold the return value we expect `JoinWithCommas` to return. We also add a `got` variable (as in "the value we actually *got*") to hold the actual return value. If `got` isn't equal to `want`, we'll call `Errorf` and have it generate an error message that includes the slice we passed to `JoinWithCommas` (we use a format verb of `%#v` so the slice is printed the same way it would appear in Go code), the return value we got, and the return value we wanted.

```
func TestTwoElements(t *testing.T) {
    list := []string{"apple", "orange"}
    want := "apple and orange"
    got := JoinWithCommas(list)
    if got != want {
        t.Errorf("JoinWithCommas(%#v) = \"%s\", want \"%s\"", list, got, want)
    }
}
```

The return value we want

The return value we actually got

Display the slice passed to JoinWithCommas, in debug format.

Include the return value we got for this slice.

Include the return value we wanted for this slice.

```
func TestThreeElements(t *testing.T) {
    list := []string{"apple", "orange", "pear"}
    want := "apple, orange, and pear"
    got := JoinWithCommas(list)
    if got != want {
        t.Errorf("JoinWithCommas(%#v) = \"%s\", want \"%s\"", list, got, want)
    }
}
```

The return value we want

The return value we actually got

Display the slice passed to JoinWithCommas, in debug format.

Include the return value we got for this slice.

Include the return value we wanted for this slice.

If we rerun the tests, we'll see exactly what the failure was.

```
--- FAIL: TestTwoElements (0.00s)
    lists_test.go:15: JoinWithCommas([]string{"apple", "orange"}) =
                      "apple, and orange", want "apple and orange"
FAIL
FAIL    github.com/headfirstgo/prose    0.006s
```

Test "helper" functions

You aren't limited to only having test functions in your *_test.go* files. You can reduce repeated code in your tests by moving it to other "helper" functions within your test file. The `go test` command only uses functions whose names begin with `Test`, so as long as you name your functions anything else, you'll be fine.

There's a fairly cumbersome call to `t.Errorf` that's duplicated between our `TestTwoElements` and `TestThreeElements` functions (with the possibility for more duplication as we add more tests). One solution might be to move the string generation out to a separate `errorString` function the tests can call.

We'll have `errorString` accept the slice that's passed to `JoinWithCommas`, the `got` value, and the `want` value. Then, instead of calling `Errorf` on a `testing.T` value, we'll have `errorString` call `fmt.Sprintf` to generate an (identical) error string for us to return. The test itself can then call `Error` with the returned string to indicate a test failure. This code is slightly cleaner, but still gets us the same output.

```
import (
    "fmt"
    "testing"
)

func TestTwoElements(t *testing.T) {
    list := []string{"apple", "orange"}
    want := "apple and orange"
    got := JoinWithCommas(list)
    if got != want {
        t.Error(errorString(list, got, want))
    }
}

func TestThreeElements(t *testing.T) {
    list := []string{"apple", "orange", "pear"}
    want := "apple, orange, and pear"
    got := JoinWithCommas(list)
    if got != want {
        t.Error(errorString(list, got, want))
    }
}

func errorString(list []string, got string, want string) string {
    return fmt.Sprintf("JoinWithCommas(%#v) = \"%s\", want \"%s\"", list, got, want)
}
```

Need "fmt" so we can call fmt.Sprintf

Instead of calling t.Errorf, call our new helper function.

Instead of calling t.Errorf, call our new helper function.

This function name doesn't begin with "Test", so it's not treated as a test.

```
--- FAIL: TestTwoElements (0.00s)
        lists_test.go:18: JoinWithCommas([]string{"apple", "orange"}) =
                          "apple, and orange", want "apple and orange"
FAIL
FAIL    github.com/headfirstgo/prose    0.006s
```

Same output

Getting the tests to pass

Now that our tests are set up with useful failure messages, it's time to look at using them to fix our main code.

We have two tests for our `JoinWithCommas` function. The test that passes a slice with three items passes, but the test that passes a slice with two items fails.

This is because `JoinWithCommas` currently includes a comma even when returning a list of just two items.

Pass. For `[]string{"apple", "orange", "pear"}`, `JoinWithCommas` should return `"apple, orange, and pear"`.

Fail! For `[]string{"apple", "orange"}`, `JoinWithCommas` should return `"apple and orange"`.

A comma doesn't belong here!

```
A photo of my parents, and a rodeo clown
```

Let's modify `JoinWithCommas` to fix this. If there are just two elements in the slice of strings, we'll simply join them together with `" and "`, then return the resulting string. Otherwise, we'll follow the same logic we always have.

If the slice has just two items, just join them together with "and".

```
func JoinWithCommas(phrases []string) string {
        if len(phrases) == 2 {
                return phrases[0] + " and " + phrases[1]
        } else {
                result := strings.Join(phrases[:len(phrases)-1], ", ")
                result += ", and "
                result += phrases[len(phrases)-1]
                return result
        }
}
```

Otherwise, use the same code we always have.

We've updated our code, but is it working correctly? Our tests can tell us immediately! If we rerun our tests now, `TestTwoElements` will pass, meaning all tests are passing.

All tests pass!

```
File  Edit  Window  Help
$ go test github.com/headfirstgo/prose
ok      github.com/headfirstgo/prose    0.006s
```

Pass For `[]string{"apple", "orange", "pear"}`, `JoinWithCommas` should return `"apple, orange, and pear"`.

Now passing! For `[]string{"apple", "orange"}`, `JoinWithCommas` should return `"apple and orange"`.

Getting the tests to pass (continued)

We can say with certainty that `JoinWithCommas` works with a slice of two strings now, because the corresponding unit test now passes. And we don't need to worry about whether it still works correctly with slices of three strings; we have a unit test assuring us that's fine, too.

This is reflected in the output of our `join` program, too. If we rerun it now, we'll see that both slices are formatted correctly!

```
func main() {
        phrases := []string{"my parents", "a rodeo clown"}
        fmt.Println("A photo of", prose.JoinWithCommas(phrases))
        phrases = []string{"my parents", "a rodeo clown", "a prize bull"}
        fmt.Println("A photo of", prose.JoinWithCommas(phrases))
}
```

No extra comma when there are two items

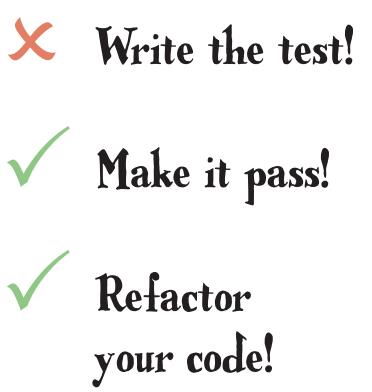
Still works with three items →

```
A photo of my parents and a rodeo clown
A photo of my parents, a rodeo clown, and a prize bull
```

Test-driven development

Once you have some experience with unit testing, you'll probably fall into a cycle known as *test-driven development*:

1. **Write the test:** You write a test for the feature you *want*, even though it doesn't exist yet. Then you run the test to ensure that it *fails*.
2. **Make it pass:** You implement the feature in your main code. Don't worry about whether the code you're writing is sloppy or inefficient; your only goal is to get it working. Then you run the test to ensure that it *passes*.
3. **Refactor your code:** Now, you're free to *refactor* the code, to change and improve it, however you please. You've watched the test *fail*, so you know it will fail again if your app code breaks. You've watched the test *pass*, so you know it will continue passing as long as your code is working correctly.



This freedom to *change* your code without worrying about it breaking is the real reason you want unit tests. Anytime you see a way to make your code shorter or easier to read, you won't hesitate to do it. When you're finished, you can simply run your tests again, and you'll be confident that everything is still working.

Another bug to fix

It's possible that `JoinWithCommas` could be called with a slice containing only a single phrase. But it doesn't behave very well in that case, treating that one item as if it appeared at the end of a list:

```
phrases = []string{"my parents"}
fmt.Println("A photo of", prose.JoinWithCommas(phrases))
```

Our function treats a single item as if it were at the end of a list!

```
A photo of , and my parents
```

What *should* `JoinWithCommas` return in this case? If we have a list of one item, we don't really need commas, the word *and*, or anything at all. We could simply return a string with that one item.

```
A photo of my parents
```

A list of one item should look like this.

Let's express this as a new test in *join_test.go*. We'll add a new test function called `TestOneElement` alongside the existing `TestTwoElements` and `TestThreeElements` tests. Our new test will look just like the others, but we'll pass a slice with just one string to `JoinWithCommas`, and expect a return value with that one string.

```
func TestOneElement(t *testing.T) {
      list := []string{"apple"}
      want := "apple"
      got := JoinWithCommas(list)
      if got != want {
            t.Error(errorString(list, got, want))
      }
}
```

Pass a slice with just one string.

Expect the return value to consist of just that one string.

```
--- FAIL: TestOneElement (0.00s)
        lists_test.go:13: JoinWithCommas([]string{"apple"}) =
                          ", and apple", want "apple"
FAIL
FAIL    github.com/headfirstgo/prose    0.006s
```

As you might expect knowing that there's a bug in our code, the test fails, showing that `JoinWithCommas` returned `", and apple"` rather than just `"apple"`.

Another bug to fix (continued)

Updating `JoinWithCommas` to fix our broken test is pretty simple. We test whether the given slice contains only one string, and if so, we simply return that string.

```
func JoinWithCommas(phrases []string) string {
        if len(phrases) == 1 {
                return phrases[0]
        } else if len(phrases) == 2 {
                return phrases[0] + " and " + phrases[1]
        } else {
                result := strings.Join(phrases[:len(phrases)-1], ", ")
                result += ", and "
                result += phrases[len(phrases)-1]
                return result
        }
}
```

With our code fixed, if we rerun the test, we'll see that everything's passing.

All tests pass!

```
File  Edit  Window  Help
$ go test github.com/headfirstgo/prose
ok      github.com/headfirstgo/prose    0.006s
```

And when we use `JoinWithCommas` in our code, it will behave as it should.

```
phrases = []string{"my parents"}
fmt.Println("A photo of", prose.JoinWithCommas(phrases))
```

Now it's working correctly!

```
A photo of my parents
```

there are no Dumb Questions

Q: Isn't all this test code going to make my program bigger and slower?

A: Don't worry! Just as the `go test` command has been set up to only work with files whose names end in *_test.go*, the various other commands in the `go` tool (such as `go build` and `go install`) have been set up to *ignore* files whose names end in *_test.go*. The `go` tool can compile your program code into an executable file, but it will ignore your test code, even when it's saved in the same package directory.

Code Magnets

Oops! We've created a `compare` package with a `Larger` function that is supposed to return the larger of two integers passed into it. But we got the comparison wrong, and `Larger` is returning the *smaller* integer instead!

We've started writing tests to help diagnose the problem. Can you reconstruct the code snippets to make working tests that will produce the output shown? You'll need to create a helper function that returns a string with the test failure message, and then add two calls to that helper function within the tests.

your workspace > src > compare > larger.go

```
package compare

func Larger(a int, b int) int {
        if a < b {
                return a
        } else {
                return b
        }
}
```

Oops! This comparison is backward!

your workspace > src > compare > larger_test.go

```
package compare

import (
        "fmt"
        "testing"
)

func TestFirstLarger(t *testing.T) {
        want := 2
        got := Larger(2, 1)
        if got != want {
                t.Error(
        }
}

func TestSecondLarger(t *testing.T) {
        want := 8
        got := Larger(4, 8)
        if got != want {
                t.Error(
        }
}
```

Call your helper function here.

Call your helper function here.

Define your helper function here.

```
"Larger(%d, %d) = %d, want %d",

(4, 8, got, want)        func
(2, 1, got, want)        string
fmt.Sprintf(   )         return
(   )   {   }            want int
)
errorString    a int,
errorString    b int,
errorString    got int,
a, b, got, want
)
```

```
File Edit Window Help
$ go test compare
--- FAIL: TestFirstLarger (0.00s)
        larger_test.go:12:
        Larger(2, 1) = 1, want 2
--- FAIL: TestSecondLarger (0.00s)
        larger_test.go:20:
        Larger(4, 8) = 4, want 8
FAIL
FAIL    compare 0.007s
```

Answers on page 424.

Running specific sets of tests

Sometimes you'll want to run only a few specific tests, rather than your whole collection. The `go test` command provides a couple of command-line flags that help you do this. A **flag** is an argument, usually a dash (-) followed by one or more letters, that you provide to a command-line program to change the program's behavior.

The first flag that's worth remembering for the `go test` command is the `-v` flag, which stands for "verbose." If you add it to any `go test` command, it will list the name and status of each test function it runs. Normally passing tests are omitted to keep the output "quiet," but in verbose mode, `go test` will list even passing tests.

Add the "-v" flag to the command.

Name and status of each test will be listed.

```
File Edit Window Help
$ go test github.com/headfirstgo/prose -v
=== RUN   TestOneElement
--- PASS: TestOneElement (0.00s)
=== RUN   TestTwoElements
--- PASS: TestTwoElements (0.00s)
=== RUN   TestThreeElements
--- PASS: TestThreeElements (0.00s)
PASS
ok      github.com/headfirstgo/prose    0.007s
```

Once you have the name of one or more tests (either from the `go test -v` output or from looking them up in your test code files), you can add the `-run` option to limit the set of tests that are run. Following `-run`, you specify part or all of a function name, and only test functions whose name matches what you specify will be run.

If we add `-run Two` to our `go run` command, only test functions with `Two` in their name will be matched. In our case, that means only `TestTwoElements` will be run. (You can use `-run` with or without the `-v` flag, but we find that adding `-v` helps avoid confusion about which tests are running.)

Run tests with "Two" in their names.

```
File Edit Window Help
$ go test github.com/headfirstgo/prose -v -run Two
=== RUN   TestTwoElements
--- PASS: TestTwoElements (0.00s)
PASS
ok      github.com/headfirstgo/prose    0.007s
```

If we add `-run Elements` instead, both `TestTwoElements` and `TestThreeElements` will be run. (But not `TestOneElement`, because it doesn't have an `s` at the end of its name.)

Run tests with "Elements" in their names.

```
File Edit Window Help
$ go test github.com/headfirstgo/prose -v -run Elements
=== RUN   TestTwoElements
--- PASS: TestTwoElements (0.00s)
=== RUN   TestThreeElements
--- PASS: TestThreeElements (0.00s)
PASS
ok      github.com/headfirstgo/prose    0.007s
```

Table-driven tests

There's quite a bit of duplicated code between our three test functions. Really, the only things that vary between tests are the slice we pass to `JoinWithCommas`, and the string we expect it to return.

```
func TestOneElement(t *testing.T) {
       list := []string{"apple"}
       want := "apple"
       got := JoinWithCommas(list)
       if got != want {
              t.Error(errorString(list, got, want))
       }
}

func TestTwoElements(t *testing.T) {
       list := []string{"apple", "orange"}
       want := "apple and orange"
       got := JoinWithCommas(list)
       if got != want {
              t.Error(errorString(list, got, want))
       }
}

func TestThreeElements(t *testing.T) {
       list := []string{"apple", "orange", "pear"}
       want := "apple, orange, and pear"
       got := JoinWithCommas(list)
       if got != want {
              t.Error(errorString(list, got, want))
       }
}

func errorString(list []string, got string, want string) string {
       return fmt.Sprintf("JoinWithCommas(%#v) = \"%s\", want \"%s\"", list, got, want)
}
```

Duplicated code

Instead of maintaining separate test functions, we can build a "table" of input data and the corresponding output we expect, then use a single test function to check each entry in the table.

There's no standard format for the table, but one common solution is to define a new type, specifically for use in your tests, that holds the input and expected output for each test. Here's a `testData` type we might use, which has a `list` field to hold the slice of strings we'll pass to `JoinWithCommas`, and a `want` field to hold the corresponding string we expect it to return.

```
type testData struct {
       list []string
       want string
}
```

The slice we'll pass to JoinWithCommas.

The string we expect JoinWithCommas to return for the above slice.

Table-driven tests (continued)

We can define the `testData` type right in the *lists_test.go* file where it will be used.

Our three test functions can be merged into a single `TestJoinWithCommas` function. At the top, we set up a `tests` slice, and move the values for the `list` and `want` variables from the old `TestOneElement`, `TestTwoElements`, and `TestThreeElements` into `testData` values within the `tests` slice.

We then loop through each `testData` value in the slice. We pass the `list` slice to `JoinWithCommas`, and store the string it returns in a `got` variable. If `got` isn't equal to the string in the `testData` value's `want` field, we call `Errorf` and use it to format a test failure message, just like we did in the `errorString` helper function. (And since that makes the `errorString` function redundant, we can delete it.)

```
import "testing"

type testData struct {
    list []string
    want string
}

func TestJoinWithCommas(t *testing.T) {
    tests := []testData{
        testData{list: []string{"apple"}, want: "apple"},
        testData{list: []string{"apple", "orange"}, want: "apple and orange"},
        testData{list: []string{"apple", "orange", "pear"}, want: "apple, orange, and pear"},
    }
    for _, test := range tests {
        got := JoinWithCommas(test.list)
        if got != test.want {
            t.Errorf("JoinWithCommas(%#v) = \"%s\", want \"%s\"", test.list, got, test.want)
        }
    }
}
```

We can define our testData type right within the test file.

This single function will replace our three old functions.

The data from TestOneElement

The data from TestTwoElements

The data from TestThreeElements

Create a slice of testData values.

Process each testData value in the slice.

Pass the slice to JoinWithCommas.

If the return value we got doesn't equal the value we want...

Format an error string and fail the test.

This updated code is much shorter and less repetitive, but the tests in the table pass just like they did when they were separate test functions!

```
File Edit Window Help
$ go test github.com/headfirstgo/prose
ok      github.com/headfirstgo/prose    0.006s
```

Fixing panicking code using a test

The best thing about table-driven tests, though, is that it's easy to add new tests when you need them. Suppose we weren't sure how `JoinWithCommas` would behave when it's passed an empty slice. To find out, we simply add a new `testData` struct in the `tests` slice. We'll specify that if an empty slice is passed to `JoinWithCommas`, an empty string should be returned:

```
func TestJoinWithCommas(t *testing.T) {
    tests := []testData{
        testData{list: []string{}, want: ""},
        testData{list: []string{"apple"}, want: "apple"},
        testData{list: []string{"apple", "orange"}, want: "apple and orange"},
        testData{list: []string{"apple", "orange", "pear"}, want: "apple, orange, and pear"},
    }
    // Additional code omitted...
}
```

Add a new testData value that will pass an empty slice to JoinWithCommas.

It looks like we were right to be worried. If we run the test, it panics with a stack trace:

```
--- FAIL: TestJoinWithCommas (0.00s)
panic: runtime error: slice bounds out of range [recovered]
        panic: runtime error: slice bounds out of range

goroutine 5 [running]:
testing.tRunner.func1(0xc4200a20f0)
        /usr/go/1.10/libexec/src/testing/testing.go:742 +0x29d
panic(0x110a480, 0x11d6fd0)
        /usr/go/1.10/libexec/src/runtime/panic.go:505 +0x229
github.com/headfirstgo/prose.JoinWithCommas(0x11fa400, 0x0, 0x0, 0x10afead, 0x11ae270)
        /Users/jay/go/src/github.com/headfirstgo/prose/lists.go:11 +0x1bf
github.com/headfirstgo/prose.TestJoinWithCommas(0xc4200a20f0)
        /Users/jay/go/src/github.com/headfirstgo/prose/lists_test.go:20 +0x250
...
FAIL    github.com/headfirstgo/prose    0.009s
```

Apparently some code tried to access an index that's out of bounds for a slice (it tried to access an element that doesn't exist).

```
panic: runtime error: slice bounds out of range
```

Looking at the stack trace, we see the panic occurred at line 11 of the *lists.go* file, within the `JoinWithCommas` function:

Error occurred in the JoinWithCommas function.

```
github.com/headfirstgo/prose.JoinWithCommas(0x11fa400, 0x0, 0x0, 0x10afead, 0x11ae270)
        /Users/jay/go/src/github.com/headfirstgo/prose/lists.go:11 +0x1bf
```

Error is at line 11 of the lists.go file.

Fixing panicking code using a test (continued)

So the panic occurs at line 11 of the *lists.go* file... That's where we access all the elements in the slice except the last, and join them together with commas. But since the `phrases` slice we're passing in is empty, there *are* no elements to access.

```
func JoinWithCommas(phrases []string) string {
        if len(phrases) == 1 {
                return phrases[0]
        } else if len(phrases) == 2 {
                return phrases[0] + " and " + phrases[1]
        } else {
                result := strings.Join(phrases[:len(phrases)-1], ", ")
                result += ", and "
                result += phrases[len(phrases)-1]
                return result
        }
}
```

Panic occurs here, when we try to access elements from an empty slice.

If the `phrases` slice is empty, we really shouldn't be attempting to access *any* elements from it. There's nothing to join, so all we have to do is return an empty string. Let's add another clause to the `if` statement that returns an empty string when `len(phrases)` is `0`.

```
func JoinWithCommas(phrases []string) string {
        if len(phrases) == 0 {
                return ""
        } else if len(phrases) == 1 {
                return phrases[0]
        } else if len(phrases) == 2 {
                return phrases[0] + " and " + phrases[1]
        } else {
                result := strings.Join(phrases[:len(phrases)-1], ", ")
                result += ", and "
                result += phrases[len(phrases)-1]
                return result
        }
}
```

If the slice is empty, just return an empty string.

After that, if we run the tests again, everything passes, even the test that calls `JoinWithCommas` with an empty slice!

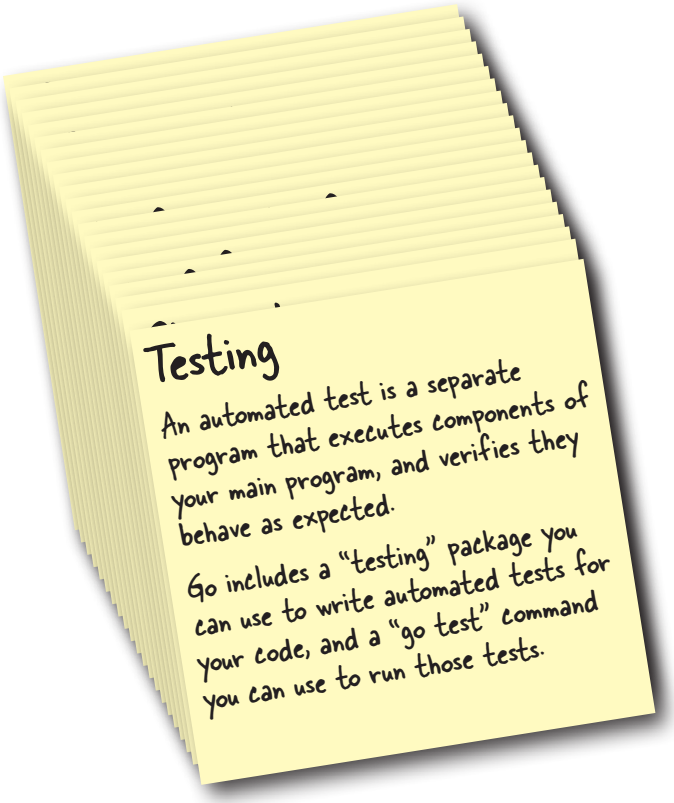
```
File Edit Window Help
$ go test github.com/headfirstgo/prose
ok      github.com/headfirstgo/prose    0.006s
```

Maybe you can imagine further changes and improvements you'd like to make to `JoinWithCommas`. Go ahead! You can do so without fear of breaking anything. If you run your tests after each change, you'll know for certain whether everything is working as it should be. (And if it's not, you'll have a clear indicator of what you need to fix!)



Your Go Toolbox

That's it for Chapter 14! You've added testing to your toolbox.



BULLET POINTS

- An automated test runs your code with a particular set of inputs, and looks for a particular result. If the code's output matches the expected value, the test will "pass"; otherwise, it will "fail."
- The `go test` tool is used to run tests. It looks for files within a specified package whose names end in *_test.go*.
- You're not required to make your tests part of the same package as the code you're testing, but doing so will allow you to access unexported types or functions from that package.
- Tests are required to use a type from the `testing` package, so you'll need to import that package at the top of each test file.
- A *_test.go* file can contain one or more test functions, whose names begin with `Test`. The rest of the name can be whatever you want.
- Test functions must accept a single parameter: a pointer to a `testing.T` value.
- Your test code can make ordinary calls to the functions and methods in your package, then check that the return values match the expected values. If they don't, the test should fail.
- You can report that a test has failed by calling methods (such as `Error`) on the `testing.T` value. Most methods accept a string with a message explaining the reason the test failed.
- The `Errorf` method works similarly to `Error`, but it accepts a formatting string just like the `fmt.Printf` function.
- Functions within a *_test.go* file whose names do not begin with `Test` are not run by `go test`. They can be used by tests as "helper" functions.
- **Table-driven tests** are tests that process "tables" of inputs and expected outputs. They pass each set of input to the code being tested, and check that the code's output matches the expected values.



math.go

```
package arithmetic

func Add(a float64, b float64) float64 {
       return a + b
}

func Subtract(a float64, b float64) float64 {
       return a - b
}
```

math_test.go

Same package as code being tested

```
package arithmetic

import "testing"

func TestAdd(t *testing.T) {
    if Add(1, 2) != 3 {
        t.Error("1 + 2 did not equal 3")
    }
}

func TestSubtract(t *testing.T) {
    if Subtract(8, 4) != 4 {
        t.Error("8 - 4 did not equal 4")
    }
}
```

Must import this package for its testing.T type

Test functions must receive a *testing.T.

Call the code being tested. If the return value isn't as expected, fail the test.

Test functions must receive a *testing.T.

Call the code being tested. If the return value isn't as expected, fail the test.

Code Magnets Solution




```
package compare

func Larger(a int, b int) int {
        if a < b {
                return a
        } else {
                return b
        }
}
```

Backward!

```
File Edit Window Help
$ go test compare
--- FAIL: TestFirstLarger (0.00s)
        larger_test.go:12: Larger(2, 1) = 1, want 2
--- FAIL: TestSecondLarger (0.00s)
        larger_test.go:20: Larger(4, 8) = 4, want 8
FAIL
FAIL    compare 0.007s
```




```
package compare

import (
        "fmt"
        "testing"
)

func TestFirstLarger(t *testing.T) {
        want := 2
        got := Larger(2, 1)
        if got != want {
                t.Error( errorString (2, 1, got, want) )
        }
}

func TestSecondLarger(t *testing.T) {
        want := 8
        got := Larger(4, 8)
        if got != want {
                t.Error( errorString (4, 8, got, want) )
        }
}
```

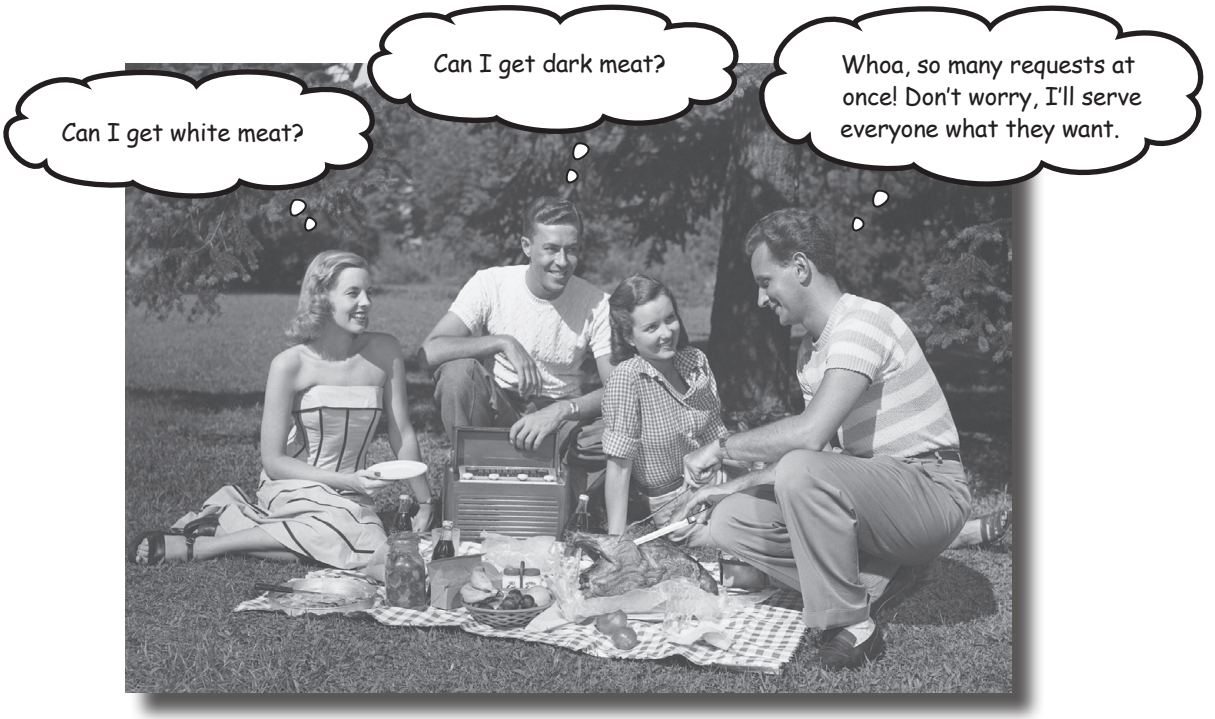
Call the helper function, and use its return value as the test failure message.

Call the helper function, and use its return value as the test failure message.

```
func errorString ( a int, b int, got int, want int ) string {
        return fmt.Sprintf( "Larger(%d, %d) = %d, want %d", a, b, got, want )
}
```

15 responding to requests

Web Apps



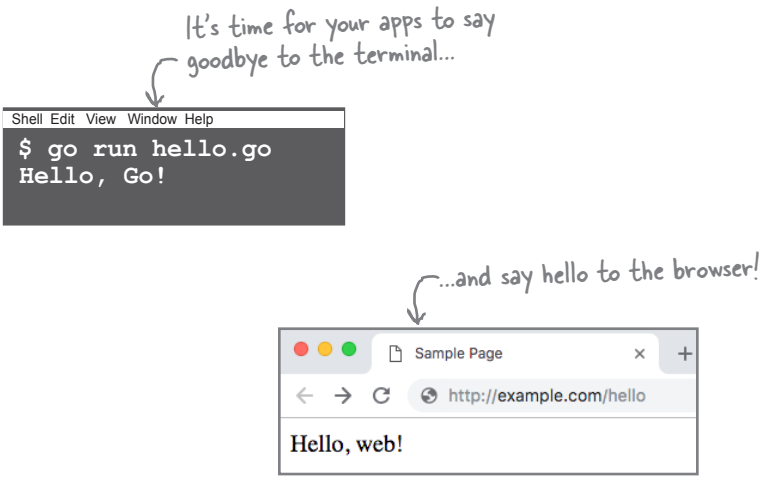
This is the 21st century. Users want web apps. Go's got you covered there, too! The Go standard library includes packages to help you host your own web applications and make them accessible from any web browser. So we're going to spend the final two chapters of the book showing you how to build web apps.

The first thing your web app needs is the ability to respond when a browser sends it a request. In this chapter, we'll learn to use the `net/http` package to do just that.

Writing web apps in Go

An app that runs in your terminal is great—for your own use. But ordinary users have been spoiled by the internet and the World Wide Web. They don't want to learn to use a terminal so they can use your app. They don't even want to install your app. They want it to be ready to use the moment they click a link in their browser.

But don't worry! Go can help you write apps for the web, too.



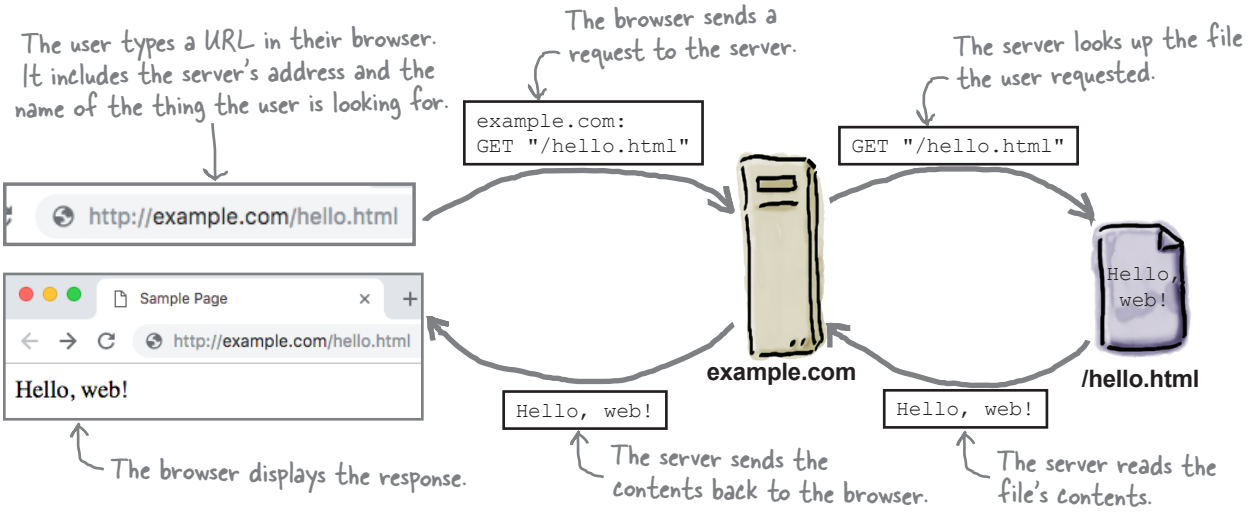
We won't lead you on—writing a web app is not a small task. This is going to require all of the skills you've learned so far, plus a few new ones. But Go has some excellent packages available that will make the process easier!

This includes the `net/http` package. HTTP stands for "**H**yper**T**ext **T**ransfer **P**rotocol," and it's used for communication by web browsers and web servers. With `net/http`, you'll be able to create your very own web apps using Go!

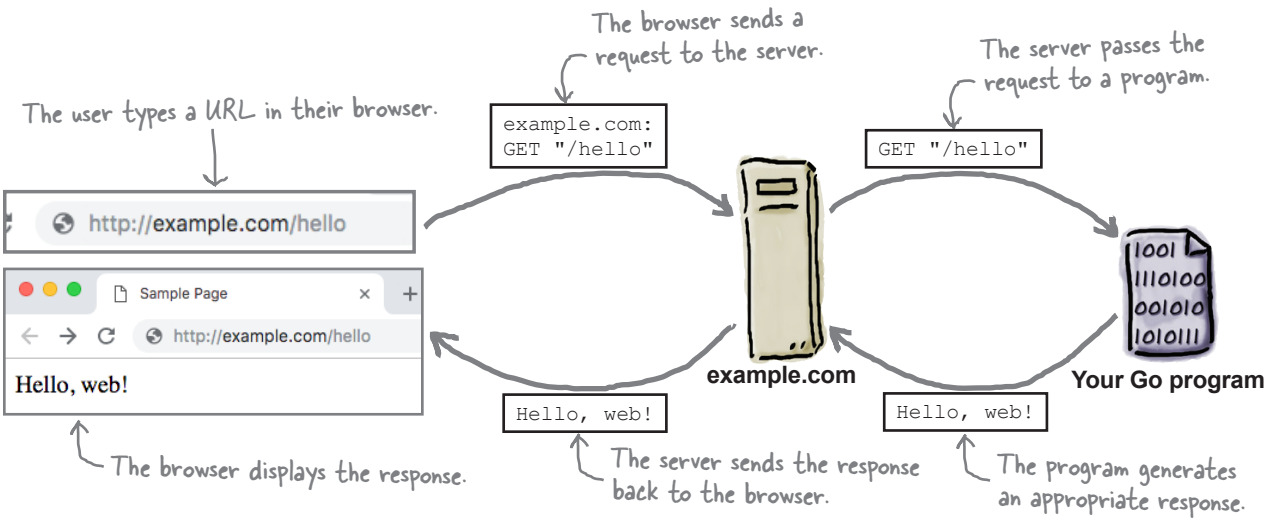
Browsers, requests, servers, and responses

When you type a URL into your browser, you're actually sending a *request* for a web page. That request goes to a *server*. A server's job is to get the appropriate page and send it back to the browser in a *response*.

In the early days of the web, the server usually read the contents of an HTML file on the server's hard drive and sent that HTML back to the browser.



But today, it's much more common for the server to communicate with a *program* to fulfill the request, instead of reading from a file. This program can be written in pretty much any language you want, including Go!



A simple web app

Handling a request from a browser is a lot of work. Fortunately, we don't have to do it all ourselves. Back in Chapter 13, we used the `net/http` package to make requests to servers. The `net/http` package also includes a small web server, so it's also able to *respond* to requests. All *we* have to do is write the code that fills those responses with data.

Here's a program that uses `net/http` to serve simple responses to the browser. Although the program is short, there's a lot going on here, some of it new. We'll run the program first, then go back and explain it piece by piece.

```
package main

import (
       "log"
       "net/http"
)

func viewHandler(writer http.ResponseWriter, request *http.Request) {
       message := []byte("Hello, web!")
       _, err := writer.Write(message)
       if err != nil {
              log.Fatal(err)
       }
}

func main() {
       http.HandleFunc("/hello", viewHandler)
       err := http.ListenAndServe("localhost:8080", nil)
       log.Fatal(err)
}
```

Import the "net/http" package.

A value for updating the response that will be sent to the browser

A value representing the request from the browser

Add "Hello, web!" to the response.

If we receive a request for a URL ending in "/hello"...

...then call the viewHandler function to generate a response.

Listen for browser requests, and respond to them.

Save the above code to a file of your choosing, and run it from your terminal using **`go run`**:

Run the server.

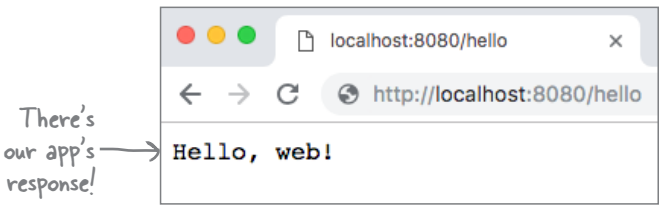
```
File  Edit  Window  Help
$ go run hello.go
```

We're running our own web app! Now we just need to connect a web browser to it and test it out. Open your browser and type this URL into the address bar. (If the URL looks a little strange to you, don't worry; we'll explain what it means in a moment.)

```
http://localhost:8080/hello
```

The browser will send a request to the app, which will respond with "`Hello, web!`". We've just sent our first response to the browser!

The app will keep listening for requests until we stop it. When you're done with the page, press Ctrl-C in your terminal to signal the program to exit.



Your computer is talking to itself

When we launched our little web app, it started its very own web server, right there on your computer.



Because the app is running *on* your computer (and not somewhere out on the internet), we use the special hostname `localhost` in the URL. This tells your browser that it needs to establish a connection *from* your computer *to* that same computer.

```
http://localhost:8080/hello
       Host    Port
```

We also need to specify a port as part of the URL. (A *port* is a numbered network communication channel that an application can listen for messages on.) In our code, we specified that the server should listen on port 8080, so we include that in the URL, following the hostname.

Here's the port number.

```
http.ListenAndServe("localhost:8080", nil)
```

there are no Dumb Questions

Q: I got an error saying the browser was unable to connect!

A: Your server might not actually be running. Look for error messages in your terminal. Also check the hostname and port number in your browser, in case you mistyped them.

Q: Why do I have to specify a port number in the URL? I don't have to do that with other websites!

A: Most web servers listen for HTTP requests on port 80, because that's the port that web browsers make HTTP requests to by default. But on many operating systems, you need special permissions to run a service that listens on port 80, for security reasons. That's why we set up our server to listen on port 8080 instead.

Q: My browser just displays the message, "404 page not found."

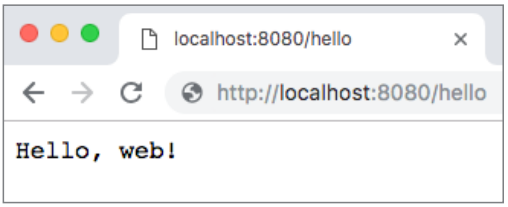
A: That's a response from the server, which is good, but it also means the resource you requested wasn't found. Check that your URL ends in */hello*, and ensure you haven't made a typo in the server program code.

Q: When I tried to run my app, I got an error saying "listen tcp 127.0.0.1:8080: bind: address already in use"!

A: Your program is trying to listen on the same port as another program (which your OS won't allow). Have you run the server program more than once? If so, did you press Ctrl-C in the terminal to stop it when you were done? Be sure to stop the old server before running a new one.

Our simple web app, explained

Now let's take a closer look at the parts of our little web app.



In the `main` function, we call `http.HandleFunc` with the string `"/hello"`, and the `viewHandler` function. (Go supports *first-class functions*, which allow you to pass functions to other functions. We'll talk more about those shortly.) This tells the app to call `viewHandler` whenever a request for a URL ending in */hello* is received.

Then, we call `http.ListenAndServe`, which starts up the web server. We pass it the string `"localhost:8080"`, which will cause it to accept requests only from your own machine on port 8080. (When you're ready to open apps up to requests from other computers, you can use a string of `"0.0.0.0:8080"` instead. You can also change the port number to something other than 8080, if you want.) The `nil` value in the second argument just means that requests will be handled using functions set up via `HandleFunc`.

(Later, if you want to learn about alternate ways to handle requests, look up the documentation for the "ListenAndServe" function, the "Handler" interface, and the "ServeMux" type from the "http" package.)

We call `ListenAndServe` *after* `HandleFunc` because `ListenAndServe` will run forever, unless it encounters an error. If it does, it will return that error, which we log before the program exits. If there are no errors, though, this program will just continue running until we interrupt it by pressing Ctrl-C in the terminal.

If we receive a request for a URL ending in "/hello"...

...then call the viewHandler function to generate a response.

```
func main() {
        http.HandleFunc("/hello", viewHandler)
        err := http.ListenAndServe("localhost:8080", nil)
        log.Fatal(err)
}
```

Listen for browser requests, and respond to them.

Compared to `main`, there's nothing very surprising in the `viewHandler` function. The server passes `viewHandler` an `http.ResponseWriter`, which is used for writing data to the browser response, and a pointer to an `http.Request` value, which represents the browser's request. (We don't use the `Request` value in this program, but handler functions still have to accept one.)

A value for updating the response that will be sent to the browser

A value representing the request from the browser

```
func viewHandler(writer http.ResponseWriter, request *http.Request) {
        ...
}
```

Our simple web app, explained (continued)

Within `viewHandler`, we add data to the response by calling the `Write` method on the `ResponseWriter`. `Write` doesn't accept strings, but it does accept a slice of `byte` values, so we convert our `"Hello, web!"` string to a `[]byte`, then pass it to `Write`.

```
message := []byte("Hello, web!")
_, err := writer.Write(message)
```

Convert "Hello, web!" to a slice of bytes.

Add "Hello, web!" to the response.

You might remember `byte` values from Chapter 13. The `ioutil.Readall` function returned a slice of `byte` values when called on a response retrieved via the `http.Get` function.

We haven't covered the `byte` type yet; it's one of Go's basic types (like `float64` or `bool`), and it's used for holding raw data, such as you might read from a file or network connection. A slice of `byte` values won't show us anything meaningful if we print it directly, but if you do a type conversion from a slice of `byte` values to a `string`, you'll get readable text back. (That is, assuming the data represents readable text.) So we end by converting the response body to a `string`, and printing it.

```
func main() {
        response, err := http.Get("https://example.com")
        if err != nil {
                log.Fatal(err)
        }
        defer response.Body.Close()
        body, err := ioutil.ReadAll(response.Body)
        if err != nil {
                log.Fatal(err)
        }
        fmt.Println(string(body))
}
```

Close the network connection once the "main" function exits.

Read all the data in the response.

Convert the data to a string, and print it.

As we saw in Chapter 13, a `[]byte` can be converted to a `string`:

```
fmt.Println(string([]byte{72, 101, 108, 108, 111}))
```

```
Hello
```

And as you've just seen in this simple web app, a `string` can be converted to a `[]byte`.

```
fmt.Println([]byte("Hello"))
```

```
[72 101 108 108 111]
```

The `ResponseWriter`'s `Write` method returns the number of bytes successfully written, and any error encountered. We can't do anything useful with the number of bytes written, so we ignore that. But if there's an error, we log it and exit the program.

```
_, err := writer.Write(message)
if err != nil {
        log.Fatal(err)
}
```

Resource paths

When we entered a URL in our browser to access our web app, we made sure it ended in */hello*. But why did we need to?

```
http://localhost:8080/hello
```

A server usually has lots of different resources that it can send to a browser, including HTML pages, images, and more.



The part of a URL following the host address and port is the resource *path*. It tells the server which of its many resources you want to act on. The `net/http` server pulls the path off the end of the URL, and uses it in handling the request.

`http://localhost:8080/hello`
Path

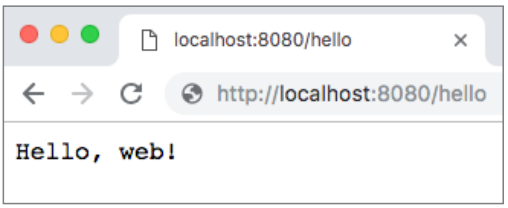
When we called `http.HandleFunc` in our web app, we passed it the string `"/hello"`, and the `viewHandler` function. The string is used as a request resource path to look for. From then on, any time a request with a path of `/hello` is received, the app will call the `viewHandler` function. The `viewHandler` function is then responsible for generating a response that's appropriate for the request it received.

If we receive a request for a URL ending in "/hello"...

...then call the viewHandler function to generate a response.

```
http.HandleFunc("/hello", viewHandler)
```

In this case, that means responding with the text "Hello, web!"



Your app can't just respond "Hello, web!" to every request it receives, though. Most apps will need to respond to different request paths in different ways.

One way to accomplish this is by calling `HandleFunc` once for each path you want to handle, and provide a different function to handle each path. Your app will then be able to respond to requests for any of those paths.

Responding differently for different resource paths

Here's an update to our app that provides greetings in three different languages. We call `HandleFunc` three different times. Requests with a `"/hello"` path cause the `englishHandler` function to be called, requests for `"/salut"` are handled by the `frenchHandler` function, and requests for `"/namaste"` are handled by `hindiHandler`. Each of these handler functions passes its `ResponseWriter` and a string to a new `write` function, which writes the string to the response.

```
package main

import (
        "log"
        "net/http"
)

func write(writer http.ResponseWriter, message string) {
        _, err := writer.Write([]byte(message))
        if err != nil {
                log.Fatal(err)
        }
}

func englishHandler(writer http.ResponseWriter, request *http.Request) {
        write(writer, "Hello, web!")
}
func frenchHandler(writer http.ResponseWriter, request *http.Request) {
        write(writer, "Salut web!")
}
func hindiHandler(writer http.ResponseWriter, request *http.Request) {
        write(writer, "Namaste, web!")
}

func main() {
        http.HandleFunc("/hello", englishHandler)
        http.HandleFunc("/salut", frenchHandler)
        http.HandleFunc("/namaste", hindiHandler)
        err := http.ListenAndServe("localhost:8080", nil)
        log.Fatal(err)
}
```

The ResponseWriter from the handler function

The message to add to the response

Convert the string to a slice of bytes, as before, and write it to the response.

Write this string to the response.

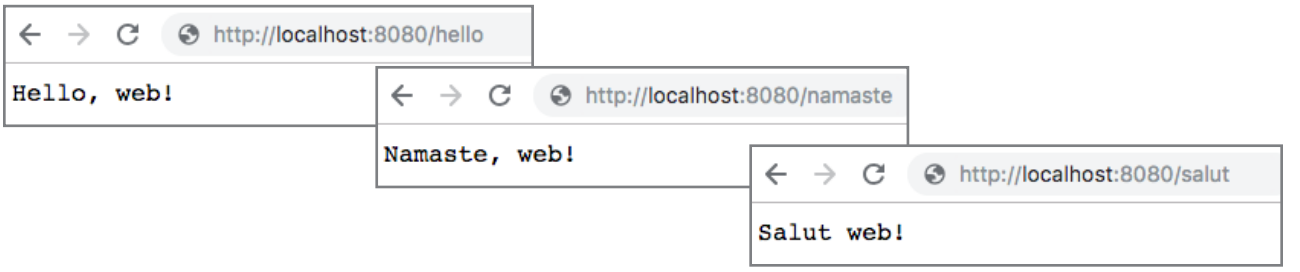
Write this string to the response.

Write this string to the response.

For requests with a path of "/hello", call englishHandler.

For requests with a path of "/salut", call frenchHandler.

For requests with a path of "/namaste", call hindiHandler.





Code for a simple web app is below, followed by several possible responses. Next to each response, write the URL you'd need to type in your browser to generate that response.

```
package main

import (
       "log"
       "net/http"
)

func write(writer http.ResponseWriter, message string) {
       _, err := writer.Write([]byte(message))
       if err != nil {
              log.Fatal(err)
       }
}

func d(writer http.ResponseWriter, request *http.Request) {
       write(writer, "z")
}
func e(writer http.ResponseWriter, request *http.Request) {
       write(writer, "x")
}
func f(writer http.ResponseWriter, request *http.Request) {
       write(writer, "y")
}

func main() {
       http.HandleFunc("/a", f)
       http.HandleFunc("/b", d)
       http.HandleFunc("/c", e)
       err := http.ListenAndServe("localhost:4567", nil)
       log.Fatal(err)
}
```

Response:	URL to generate response:
x	.....
y	.....
z	.....

→ Answers on page 442.

First-class functions

When we call `http.HandleFunc` with handler functions, we're not calling the handler function and passing its result to `HandleFunc`. We are passing the *function itself* to `HandleFunc`. That function is stored to be called later when a matching request path is received.

```
func main() {
        http.HandleFunc("/hello", englishHandler)
        http.HandleFunc("/salut", frenchHandler)
        http.HandleFunc("/namaste", hindiHandler)
        err := http.ListenAndServe("localhost:8080", nil)
        log.Fatal(err)
}
```

Pass the englishHandler function to HandleFunc.

Pass the frenchHandler function to HandleFunc.

Pass the hindiHandler function to HandleFunc.

The Go language supports **first-class functions**; that is, functions in Go are treated as "first-class citizens."

In a programming language with first-class functions, functions can be assigned to variables, and then called from those variables.

The code below first defines a `sayHi` function. In our `main` function, we declare a `myFunction` variable with a type of `func()`, meaning the variable can hold a function.

Then we assign the `sayHi` function itself to `myFunction`. Notice that we don't put any parentheses—we don't write `sayHi()`—because doing so would *call* `sayHi`. We type only the function name, like this:

```
myFunction = sayHi
```

This causes the `sayHi` function itself to be assigned to the `myFunction` variable.

But on the next line, we *do* include parentheses following the `myFunction` variable name, like this:

```
myFunction()
```

This causes the function stored inside the `myFunction` variable to be called.

Declare a function normally.

```
func sayHi() {
        fmt.Println("Hi")
}

func main() {
        var myFunction func()
        myFunction = sayHi
        myFunction()
}
```

Declare a variable with a type of "func()". This variable can hold a function.

Assign the sayHi function to the variable.

Call the function stored in the variable.

```
Hi
```

Passing functions to other functions

Programming languages with first-class functions also allow you to pass functions as arguments to other functions. This code defines simple `sayHi` and `sayBye` functions. It also defines a `twice` function that takes another function as a parameter named `theFunction`. The `twice` function then calls whatever function is stored in `theFunction` twice.

In `main`, we call `twice` and pass the `sayHi` function as an argument, which causes `sayHi` to be run twice. Then we call `twice` with the `sayBye` function, which causes `sayBye` to be run twice.

```
func sayHi() {
        fmt.Println("Hi")
}
func sayBye() {
        fmt.Println("Bye")
}

func twice(theFunction func()) {
        theFunction()
        theFunction()
}

func main() {
        twice(sayHi)
        twice(sayBye)
}
```

The "twice" function accepts another function as a parameter.

Call the passed-in function.

Call the passed-in function (again).

Pass the "sayHi" function to the "twice" function.

Pass the "sayBye" function to the "twice" function.

```
Hi
Hi
Bye
Bye
```

Functions as types

We can't just use any function as an argument when calling any other function, though. If we tried to pass the `sayHi` function as an argument to `http.HandleFunc`, we'd get a compile error:

```
func sayHi() {
        fmt.Println("Hi")
}

func main() {
        http.HandleFunc("/hello", sayHi)
        err := http.ListenAndServe("localhost:8080", nil)
        log.Fatal(err)
}
```

Attempt to set up sayHi as an HTTP request handler function.

Compile error

```
cannot use sayHi (type func()) as type func(http.ResponseWriter, *http.Request)
in argument to http.HandleFunc
```

Functions as types (continued)

A function's parameters and return value are part of its type. A variable that holds a function needs to specify what parameters and return values that function should have. That variable can only hold functions whose number and types of parameters and return values match the specified type.

This code defines a `greeterFunction` variable with a type of `func()`: it holds a function that accepts no parameters and returns no values. Then we define a `mathFunction` variable with a type of `func(int, int) float64`: it holds a function that accepts two integer parameters and returns a `float64` value.

The code also defines `sayHi` and `divide` functions. If we assign `sayHi` to the `greeterFunction` variable and `divide` to the `mathFunction` variable, everything compiles and runs fine:

```
func sayHi() {
        fmt.Println("Hi")
}
func divide(a int, b int) float64 {
        return float64(a) / float64(b)
}

func main() {
        var greeterFunction func()
        var mathFunction func(int, int) float64
        greeterFunction = sayHi
        mathFunction = divide
        greeterFunction()
        fmt.Println(mathFunction(5, 2))
}
```

This variable will hold a function with no parameters and no return value.

This variable will hold a function with two int parameters and a float64 return value.

Assign the "sayHi" function to the greeterFunction variable.

Assign the "divide" function to the mathFunction variable.

```
Hi
2.5
```

But if we try to reverse the two, we'll get compile errors again:

```
        greeterFunction = divide
        mathFunction = sayHi
```

Compile errors

```
cannot use divide (type func(int, int) float64) as type func() in assignment
cannot use sayHi (type func()) as type func(int, int) float64 in assignment
```

The `divide` function accepts two `int` parameters and returns a `float64` value, so it can't be stored in the `greeterFunction` variable (which expects a function with no parameters and no return value). And the `sayHi` function accepts no parameters and returns no value, so it can't be stored in the `mathFunction` variable (which expects a function with two `int` parameters and a `float64` return value).

Functions as types (continued)

Functions that accept a function as a parameter also need to specify the parameters and return types the passed-in function should have.

Here's a `doMath` function with a `passedFunction` parameter. The passed-in function needs to accept two `int` parameters, and return one `float64` value.

We also define `divide` and `multiply` functions, both of which accept two `int` parameters and return one `float64`. Either `divide` or `multiply` can be passed to `doMath` successfully.

```
func doMath(passedFunction func(int, int) float64) {
        result := passedFunction(10, 2)
        fmt.Println(result)
}
func divide(a int, b int) float64 {
        return float64(a) / float64(b)
}
func multiply(a int, b int) float64 {
        return float64(a * b)
}

func main() {
        doMath(divide)
        doMath(multiply)
}
```

The doMath function accepts another function as a parameter. The passed-in function must accept two integers and return a float64.

Call the passed-in function.

Print the passed-in function's return value.

A function that can be passed into doMath

Another function that can be passed into doMath

Pass the "divide" function to doMath.

Pass the "multiply" function to doMath.

```
5
20
```

A function that doesn't match the specified type can't be passed to `doMath`.

```
func main() {
        doMath(sayHi)
}
```

The sayHi function doesn't have any parameters or a return value.

Compile error

```
cannot use sayHi (type func()) as type func(int, int) float64 in argument to doMath
```

And that's why we get compile errors if we pass the wrong function to `http.HandleFunc`. `HandleFunc` expects to be passed a function that takes a `ResponseWriter` and a pointer to a `Request` as parameters. Pass anything else, and you'll get a compile error.

And really, that's a good thing. A function that can't analyze a request and write a response probably isn't going to be able to handle browser requests. If you try to pass a function with the wrong type, Go will alert you to the problem before your program even compiles.

```
http.HandleFunc("/hello", sayHi)
```

Compile error

```
cannot use sayHi (type func()) as type func(http.ResponseWriter, *http.Request)
in argument to http.HandleFunc
```



Pool Puzzle

Your **job** is to take code snippets from the pool and place them into the blank lines in this code. **Don't** use the same snippet more than once, and you won't need to use all the snippets. Your **goal** is to make a program that will run and produce the output shown.

Output

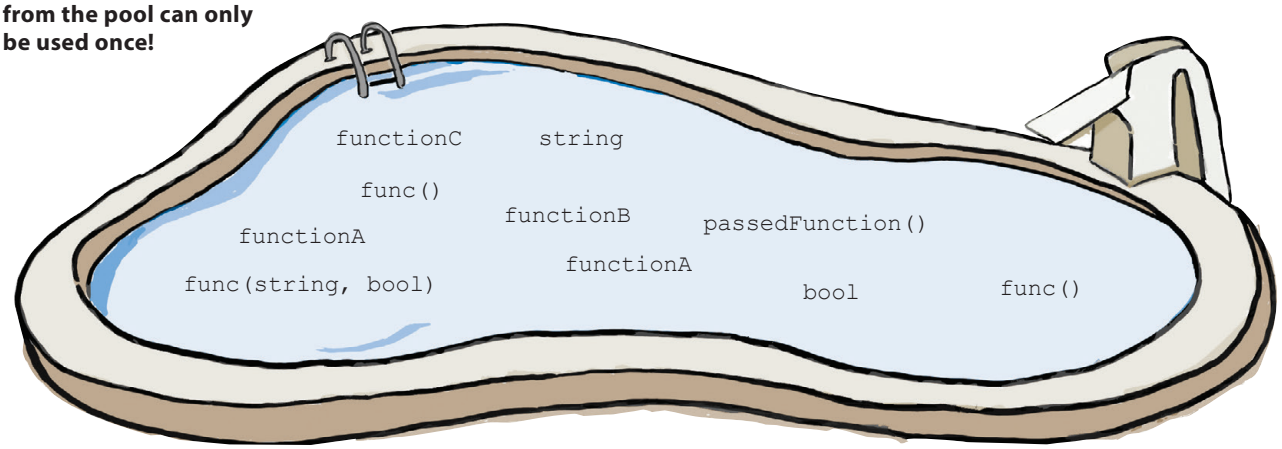
```
function called
function called
function called
function called
This sentence is false
function called
Returning from function
```

```
func callFunction(passedFunction _________) {
      passedFunction()
}
func callTwice(passedFunction _________) {
      passedFunction()
      passedFunction()
}
func callWithArguments(passedFunction ___________________) {
      passedFunction("This sentence is", false)
}
func printReturnValue(passedFunction func() string) {
      fmt.Println(_____________________)
}

func functionA() {
      fmt.Println("function called")
}
func functionB() _________ {
      fmt.Println("function called")
      return "Returning from function"
}
func functionC(a string, b bool) {
      fmt.Println("function called")
      fmt.Println(a, b)
}

func main() {
      callFunction(____________)
      callTwice(____________)
      callWithArguments(functionC)
      printReturnValue(functionB)
}
```

Note: each snippet from the pool can only be used once!



Answers on page 443.

What's next

Now you know how to receive a request from a browser and send a response. The trickiest part is done!

```
package main

import (
        "log"
        "net/http"
)

func viewHandler(writer http.ResponseWriter, request *http.Request) {
        message := []byte("Hello, web!")
        _, err := writer.Write(message)
        if err != nil {
                log.Fatal(err)
        }
}

func main() {
        http.HandleFunc("/hello", viewHandler)
        err := http.ListenAndServe("localhost:8080", nil)
        log.Fatal(err)
}
```

A value for updating the response that will be sent to the browser

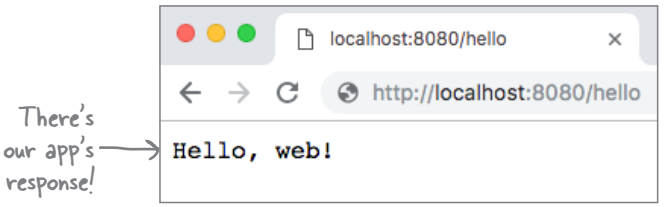
A value representing the request from the browser

Add "Hello, web!" to the response.

If we receive a request for a URL ending in "/hello"...

...then call the viewHandler function to generate a response.

Listen for browser requests and respond to them.



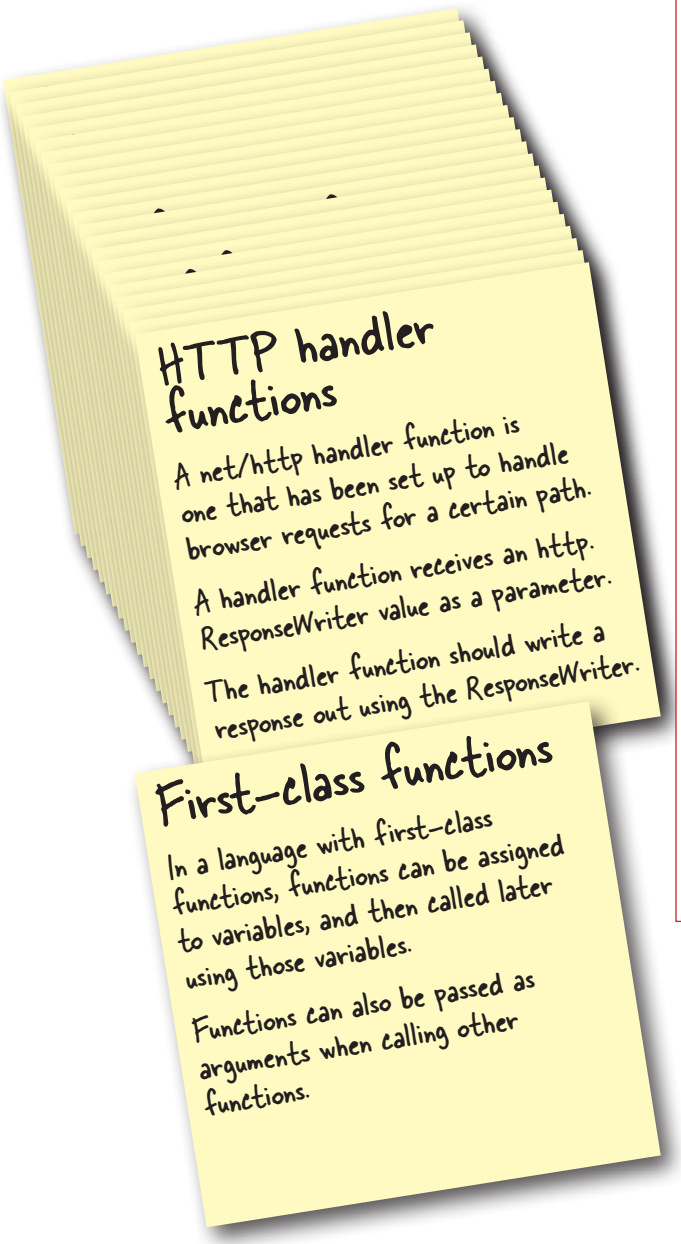
In the final chapter, we'll use this knowledge to build a more complex app.

So far, all our responses have used plain text. We're going to learn to use HTML to give the page more structure. And we'll learn to use the `html/template` package to insert data into our HTML before sending it back to the browser. See you there!



Your Go Toolbox

That's it for Chapter 15! You've added HTTP handler functions and first-class functions to your toolbox.



BULLET POINTS

- The `net/http` package's `ListenAndServe` function runs a web server on a port you specify.
- The `localhost` hostname handles connections from your computer back to itself.
- Each HTTP request includes a resource path, which specifies which of a server's many resources the browser is requesting.
- The `HandleFunc` function takes a path string, and a function that will handle requests for that path.
- You can call `HandleFunc` repeatedly to set up different handler functions for different paths.
- Handler functions must accept an `http.ResponseWriter` value and a pointer to an `http.Request` value as parameters.
- If you call the `Write` method on an `http.ResponseWriter` with a slice of bytes, that data will be added to the response sent to the browser.
- Variables that can hold a function have a function type.
- A function type includes the number and type of parameters that the function accepts (or lack thereof), and the number and type of values that the function returns (or lack thereof).
- If `myVar` holds a function, you can call that function by putting parentheses (containing any arguments the function might require) after the variable name.



Code for a simple web app is below, followed by several possible responses. Next to each response, write the URL you'd need to type in your browser to generate that response.

```
package main

import (
       "log"
       "net/http"
)

func write(writer http.ResponseWriter, message string) {
       _, err := writer.Write([]byte(message))
       if err != nil {
              log.Fatal(err)
       }
}

func d(writer http.ResponseWriter, request *http.Request) {
       write(writer, "z")
}
func e(writer http.ResponseWriter, request *http.Request) {
       write(writer, "x")
}
func f(writer http.ResponseWriter, request *http.Request) {
       write(writer, "y")
}

func main() {
       http.HandleFunc("/a", f)
       http.HandleFunc("/b", d)
       http.HandleFunc("/c", e)
       err := http.ListenAndServe("localhost:4567", nil)
       log.Fatal(err)
}
```

Notice that we specified a different port! Sneaky, huh?

Response:	URL to generate response:
x	http://localhost:4567/c
y	http://localhost:4567/a
z	http://localhost:4567/b

Pool Puzzle Solution

```
func callFunction(passedFunction func() ) {
      passedFunction()
}
func callTwice(passedFunction func() ) {
      passedFunction()
      passedFunction()
}
func callWithArguments(passedFunction func(string, bool) ) {
      passedFunction("This sentence is", false)
}
func printReturnValue(passedFunction func() string) {
      fmt.Println( passedFunction() )
}

func functionA() {
      fmt.Println("function called")
}
func functionB() string {
      fmt.Println("function called")
      return "Returning from function"
}
func functionC(a string, b bool) {
      fmt.Println("function called")
      fmt.Println(a, b)
}

func main() {
      callFunction( functionA )
      callTwice( functionA )
      callWithArguments(functionC)
      printReturnValue(functionB)
}
```

We can tell from the callFunction body that the passed function accepts no parameters.

We can tell from the callTwice body that the passed function accepts no parameters.

We can tell from the callWithArguments body that the passed-in function must accept these parameter types.

Call the passed-in function and print its return value.

If it's going to be passed to printReturnValue, functionB needs to return a string.

Only functionA has the right set of parameters (and the right output).

```
function called
function called
function called
function called
This sentence is false
function called
Returning from function
```

16 a pattern to follow

HTML Templates



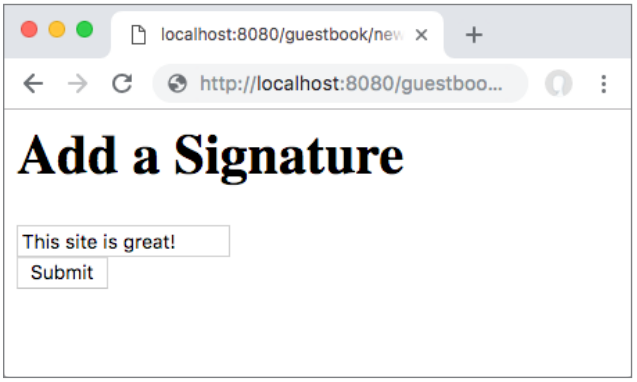
Your web app needs to respond with HTML, not plain text.

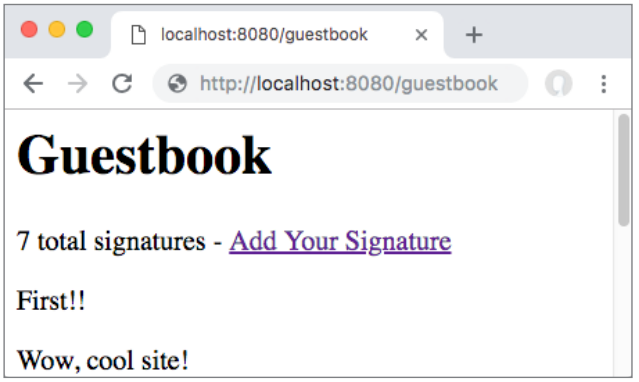
Plain text is fine for emails and social media posts. But your pages need to be formatted. They need headings and paragraphs. They need forms where your users can submit data to your app. To do any of that, you need HTML code.

And eventually, you'll need to insert data into that HTML code. That's why Go offers the `html/template` package, a powerful way to include data in your app's HTML responses. Templates are key to building bigger, better web apps, and in this final chapter, we'll show you how to use them!

A guestbook app

Let's put everything we've learned in Chapter 15 to use. We're going to build a simple guestbook app for a website. Your visitors will be able to enter messages in a form, which will be saved to a file. They'll also be able to view a list of all the previous signatures.





There's a lot left to cover before we can get this app working, but don't worry—we'll be breaking this process down into little steps. Let's take a look at what will be involved...

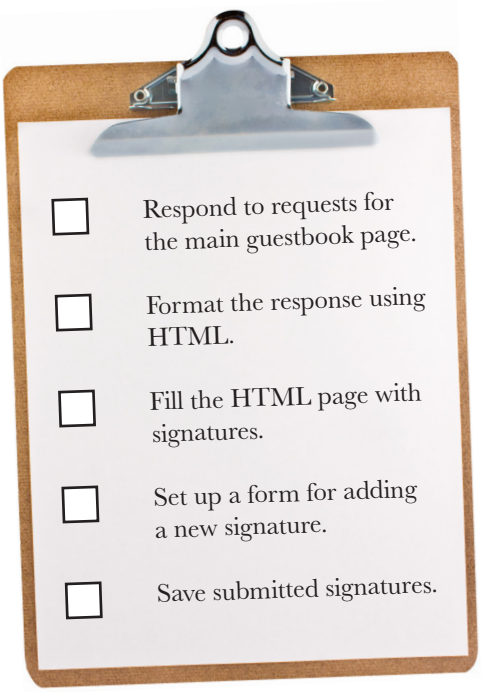
We'll need to set up our app and get it to respond to requests for the main guestbook page. This part won't be too difficult; we've already covered everything we need to know in the previous chapter.

Then we need to include HTML in our response. We'll be creating a simple page using just a few HTML tags, which we'll store in a file. Then we'll load the HTML code in from the file and use that in our app's response.

We'll need to take the signatures that our visitors have entered, and incorporate them into the HTML. We'll show you how to do this, using the `html/template` package.

Then we'll need to create a separate page with a form for adding a signature. We can do this fairly easily using HTML.

Lastly, when a user submits the form, we'll need to save the form contents as a new signature. We'll save it to a text file along with all the other submitted signatures so we can load it back in later.



Functions to handle a request and check errors

Our first task will be to display the main guestbook page. With all the practice we've had writing sample web apps, this shouldn't be too difficult. In our `main` function, we'll call `http.HandleFunc` and set up the app to call a function named `viewHandler` for any request with a path of `"/guestbook"`. Then we'll call `http.ListenAndServe` to start the server.

For now, the `viewHandler` function will look just like the handler functions in our previous examples. It accepts an `http.ResponseWriter` and a pointer to an `http.Request`, just like previous handlers. We'll convert a string for the response to a `[]byte`, and use the `Write` method on the `ResponseWriter` to add it to the response.

The `check` function is the only part of this code that's really new. We're going to have a lot of potential `error` return values in this web app, and we don't want to repeat code to check and report them everywhere. So we'll pass each error to our new `check` function. If the `error` is nil, `check` does nothing, but otherwise it logs the error and exits the program.



guestbook.go

```
package main

import (
        "log"
        "net/http"
)

func check(err error) {
        if err != nil {
                log.Fatal(err)
        }
}

func viewHandler(writer http.ResponseWriter, request *http.Request) {
        placeholder := []byte("signature list goes here")
        _, err := writer.Write(placeholder)
        check(err)
}

func main() {
        http.HandleFunc("/guestbook", viewHandler)
        err := http.ListenAndServe("localhost:8080", nil)
        log.Fatal(err)
}
```

Move our code for reporting errors to this function.

As always, handler functions will be passed a ResponseWriter...

...and also a pointer to a Request value.

We convert a string to a slice of bytes...

...and add it to the response via the Write method.

We then call "check" to report an error (if any).

We set viewHandler up to be called for any request with a path of "/guestbook".

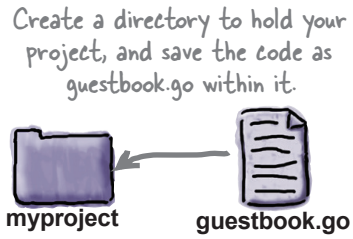
As usual, we set up the server to listen on port 8080.

This error will never be nil, so we don't call "check" on it.

Calling `Write` on the `ResponseWriter` may or may not return an error, so we pass the `error` return value to `check`. Notice that we *don't* pass the `error` return value from `http.ListenAndServe` to `check`, though. That's because `ListenAndServe` always returns an error. (If there is no error, `ListenAndServe` never returns.) Since we know this error will never be `nil`, we just immediately call `log.Fatal` on it.

Setting up a project directory and trying the app

We'll be creating several files for this project, so you might want to take a moment and create a new directory to hold them all. (It doesn't have to be within your Go workspace directory.) Save the preceding code within this directory, in a file named *guestbook.go*.



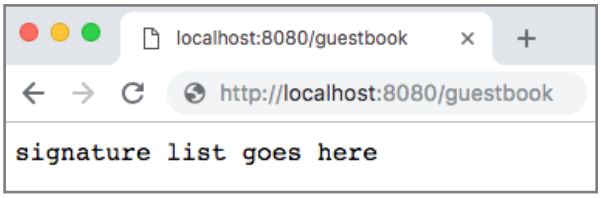
Let's try running it. In your terminal, change to the directory where *guestbook.go* is saved and run it using **`go run`**.



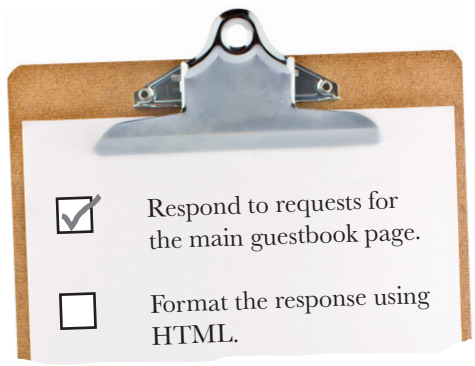
Then visit this URL in your browser:

http://localhost:8080/guestbook

It's the same as the URLs for our previous apps, except for the */guestbook* path on the end. Your browser will make a request to the app, which will respond with our placeholder text:



Our app is now responding to requests. Our first task is complete!



We're just responding using plain text, though. Up next, we're going to format our response using HTML.

Making a signature list in HTML

So far, we've just been sending snippets of text to the browser. We need actual HTML, so that we can apply formatting to the page. HTML uses tags to apply formatting to text.

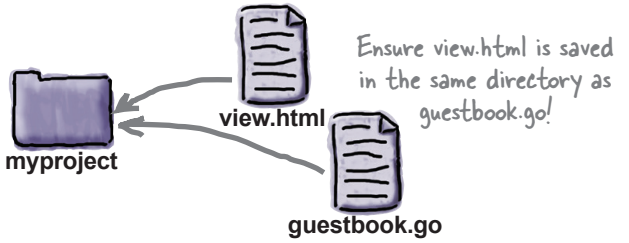
Don't worry if you haven't written HTML before; we'll be covering the basics as we go!

Save the HTML code below in the same directory as *guestbook.go*, in a file named *view.html*.

```
<h1>Guestbook</h1>
<div>
    X total signatures -
    <a href="/guestbook/new">Add Your Signature</a>
</div>
<div>
    <p>Signatures</p>
    <p>go</p>
    <p>here</p>
</div>
```

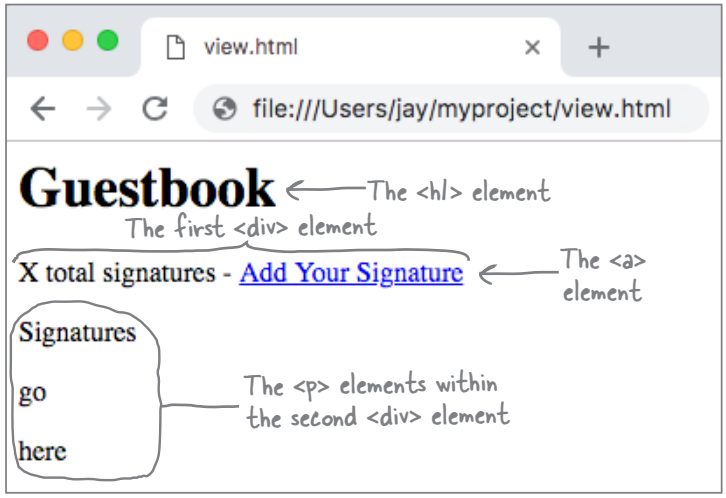

Here are the HTML elements used in this file:

- `<h1>`: A level-one heading. Usually shown in large, bold text.
- `<div>`: A division element. Not directly visible on its own, but it's used for dividing the page into sections.
- `<p>`: A paragraph of text. We'll be treating each signature as a separate paragraph.
- `<a>`: Stands for "anchor." Creates a link.



Now, let's try viewing the HTML in a browser. Launch your favorite web browser, choose "Open File…" from the menu, and open the HTML file you just saved.

Notice how the elements on the page correspond with the HTML code. Each element has a opening tag (`<h1>`, `<div>`, `<p>`, etc.), and a corresponding closing tag (`</h1>`, `</div>`, `</p>`, etc.). Any text between the opening and closing tags is used as the element's content on the page. It's also possible for elements to contain other elements (as the `<div>` elements on this page do).



You can click on the link if you want, but it will only produce a "Page not found" error right now. Before we can fix that, we'll need to figure out how to serve this HTML via our web app...

Making our app respond with HTML

Our HTML works when we load it directly into our browser from the *view.html* file, but we need to serve it via the app. Let's update our *guestbook.go* code to respond with the HTML we've created.

Go provides a package that will load the HTML in from the file *and* insert signatures into it for us: the `html/template` package. For now, we'll just load the contents of *view.html* in as is; inserting signatures will be our next step.

We'll need to update the `import` statement to add the `html/template` package. The only other changes we'll need to make are within the `viewHandler` function. We'll call the `template.ParseFiles` function and pass it the name of the file we want to load: `"view.html"`. This will use the contents of *view.html* to create a `Template` value. `ParseFiles` will return a pointer to this `Template`, and possibly an `error` value, which we pass to our `check` function.

To get output from the `Template` value, we call its `Execute` method with two arguments... We pass our `ResponseWriter` value as the place to write the output. The second value is the data we want to insert into the template, but since we're not inserting anything right now, we just pass `nil`.

```
// Code omitted...
import (
        "html/template"
        "log"
        "net/http"
)

func check(err error) {
        // Code omitted...
}

func viewHandler(writer http.ResponseWriter, request *http.Request) {
        html, err := template.ParseFiles("view.html")
        check(err)
        err = html.Execute(writer, nil)
        check(err)
}
// Code omitted...
```

Import the "html/template" package.

Use the contents of view.html to create a new Template.

Report any errors.

Write the template content to the ResponseWriter.

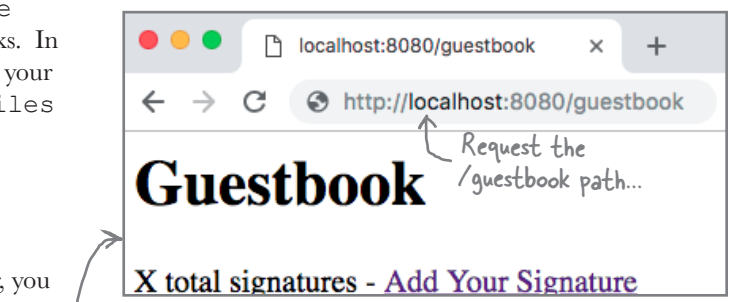
Report any errors.

We'll be learning more about the `html/template` package shortly, but for now let's just see if this works. In your terminal, run *guestbook.go*. (Make sure you're in your project directory when you do this, or the `ParseFiles` function won't be able to find *view.html*.)

In your browser, go back to the URL:

http://localhost:8080/guestbook

Instead of the "signature list goes here" placeholder, you should see the HTML from *view.html*.



The "text/template" package

Our app is responding with our HTML code. That's two tasks complete!

Right now, though, we're just showing a placeholder list of signatures that we hardcoded. Our next task will be to use the `html/template` package to insert a list of signatures into the HTML, one that will be updated when the list changes.

- [x] Respond to requests for the main guestbook page.
- [x] Format the response using HTML.
- [ ] Fill the HTML page with signatures.

The `html/template` package is based on the `text/template` package. You work with the two packages in almost exactly the same way, but `html/template` has some extra security features needed for working with HTML. Let's learn how to use the `text/template` package first, and then later we'll take what we've learned and apply it to the `html/template` package.

The program below uses `text/template` to parse and print a template string. It prints its output to the terminal, so you won't need your web browser to try it.

In `main`, we call the `text/template` package's `New` function, which returns a pointer to a new `Template` value. Then we call the `Parse` method on the `Template`, and pass it the string `"Here's my template!\n"`. `Parse` uses its string argument as the template's text, unlike `ParseFiles`, which loads the template text in from files. `Parse` returns the template and an `error` value. We store the template in the `tmpl` variable, and pass the `error` to a `check` function (identical to the one in *guestbook.go*) to report any non-`nil` errors.

Then we call the `Execute` method on the `Template` value in `tmpl`, just like we did in *guestbook.go*. Instead of an `http.ResponseWriter`, though, we pass `os.Stdout` as the place to write the output. This causes the `"Here's my template!\n"` template string to be displayed as output when the program is run.

```
package main

import (
        "log"
        "os"
        "text/template"
)

func check(err error) {
        if err != nil {
                log.Fatal(err)
        }
}

func main() {
        text := "Here's my template!\n"
        tmpl, err := template.New("test").Parse(text)
        check(err)
        err = tmpl.Execute(os.Stdout, nil)
        check(err)
}
```

We need this package so we can access os.Stdout.

Import text/template instead of html/template.

Identical to our previous "check" function

The template text

Create a new Template value based on the text.

Write out the template text.

Instead of an HTTP response, write the template to the terminal.

```
Here's my template!
```

Using the io.Writer interface with a template's Execute method



So what exactly is this `os.Stdout` value? And how can an `http.ResponseWriter` and `os.Stdout` both be valid values to pass to a `Template`'s `Execute` method?

```
func viewHandler(writer http.ResponseWriter, request *http.Request) {
        html, err := template.ParseFiles("view.html")
        check(err)
        err = html.Execute(writer, nil)
        // ...
}
```

Write the template content to the ResponseWriter.

```
text := "Here's my template!\n"
tmpl, err := template.New("test").Parse(text)
check(err)
err = tmpl.Execute(os.Stdout, nil)
check(err)
```

Write the template content to the terminal.

The `os.Stdout` value is part of the `os` package. `Stdout` stands for "standard output." It acts like a file, but any data written to it is output to the terminal instead of being saved to disk. (Functions like `fmt.Println`, `fmt.Printf`, and so on write data to `os.Stdout` behind the scenes.)

How can `http.ResponseWriter` and `os.Stdout` both be valid arguments for `Template.Execute`? Let's bring up its documentation and see...

File Edit Window Help

```
$ go doc text/template Template.Execute
func (t *Template) Execute(wr io.Writer, data interface{}) error
    Execute applies a parsed template to the specified data object, and writes
    the output to wr. If an error occurs executing the template or writing its
    ...
```

Hmm, this says the first argument to `Execute` should be an `io.Writer`. What's that? Let's check the documentation for the `io` package:

File Edit Window Help

```
$ go doc io Writer
type Writer interface {
        Write(p []byte) (n int, err error)
}
    Writer is the interface that wraps the basic Write method.
    ...
```

It looks like `io.Writer` is an interface! It's satisfied by any type with a `Write` method that accepts a slice of `byte` values, and returns an `int` with the number of bytes written and an `error` value.

ResponseWriters and os.Stdout both satisfy io.Writer

We've already seen that `http.ResponseWriter` values have a `Write` method. We've used `Write` in several earlier examples:

```
func viewHandler(writer http.ResponseWriter, request *http.Request) {
        placeholder := []byte("signature list goes here")
        _, err := writer.Write(placeholder)
        check(err)
}
```

We convert a string to a slice of bytes...

...and add it to the response via the Write method.

It turns out the `os.Stdout` value has a `Write` method, too! If you pass it a slice of `byte` values, that data will be written to the terminal:

Write data to the terminal.

```
func main() {
        _, err := os.Stdout.Write([]byte("hello"))
        check(err)
}
```

```
hello
```

That means both `http.ResponseWriter` values and `os.Stdout` satisfy the `io.Writer` interface, and can be passed to a `Template` value's `Execute` method. `Execute` will write out the template by calling the `Write` method on whatever value is passed to it.

If you pass in an `http.ResponseWriter`, it means the template will be written to the HTTP response. And if you pass in `os.Stdout`, it means the template will be written to the output in the terminal:

```
func main() {
        tmpl, err := template.New("test").Parse("Here's my template!\n")
        check(err)
        err = tmpl.Execute(os.Stdout, nil)
        check(err)
}
```

Write out the template text.

Write the template to the terminal.

```
Here's my template!
```

Inserting data into templates using actions

The second parameter to a `Template` value's `Execute` method allows you to pass in data to insert in the template. Its type is the empty interface, meaning you can pass in a value of any type you want.

```
File Edit Window Help
$ go doc text/template Template.Execute
func (t *Template) Execute(wr io.Writer, data interface{}) error
    Execute applies a parsed template to the specified data object, and writes
    the output to wr. If an error occurs executing the template or writing its
    ...
```

So far, our templates haven't provided any places to insert data, so we've just been passing `nil` for the data value:

This template doesn't provide any places to insert data.

```
func main() {
        tmpl, err := template.New("test").Parse("Here's my template!\n")
        check(err)
        err = tmpl.Execute(os.Stdout, nil)
        check(err)
}
```

Just pass "nil" for the data to insert.

```
Here's my template!
```

To insert data in a template, you add **actions** to the template text. Actions are denoted with double curly braces, `{{ }}`. Inside the double braces, you specify data you want to insert or an operation you want the template to perform. Whenever the template encounters an action, it will evaluate its contents, and insert the result into the template text in place of the action.

Within an action, you can reference the data value that was passed to the `Execute` method with a single period, called "dot."

This code sets up a template with a single action. It then calls `Execute` on the template several times, with a different data value each time. `Execute` replaces the action with the data value before writing the result to `os.Stdout`.

An action that inserts the data value

```
func main() {
        templateText := "Template start\nAction: {{.}}\nTemplate end\n"
        tmpl, err := template.New("test").Parse(templateText)
        check(err)
        err = tmpl.Execute(os.Stdout, "ABC")
        check(err)
        err = tmpl.Execute(os.Stdout, 42)
        check(err)
        err = tmpl.Execute(os.Stdout, true)
        check(err)
}
```

Execute the same template with different data values.

Values are inserted in the template in place of the action.

```
Template start
Action: ABC
Template end
Template start
Action: 42
Template end
Template start
Action: true
Template end
```

Inserting data into templates using actions (continued)

There are lots of other things you can do with template actions, too. Let's set up an `executeTemplate` function that will let us experiment with them more easily. It will take a template string that we'll pass to `Parse` to create a new template, and a data value that we'll pass to `Execute` on that template. As before, each template will be written to `os.Stdout`.

We'll create a template based on this string.

We'll forward this data value to the Execute method on the template.

```
func executeTemplate(text string, data interface{}) {
       tmpl, err := template.New("test").Parse(text)
       check(err)
       err = tmpl.Execute(os.Stdout, data)
       check(err)
}
```

Parse the given text to create a template.

Use the given data value in the template actions.

As we mentioned, you can use a single period to refer to "dot," the current value within the data the template is working with. Although the value of dot can change in various contexts within the template, initially it refers to the value that was passed to `Execute`.

```
func main() {
       executeTemplate("Dot is: {{.}}!\n", "ABC")
       executeTemplate("Dot is: {{.}}!\n", 123.5)
}
```

```
Dot is: ABC!
Dot is: 123.5!
```

Making parts of a template optional with "if" actions

A section of a template between an `{{if}}` action and its corresponding `{{end}}` marker will be included only if a condition is true. Here we execute the same template text twice, once when dot is `true` and once when it's `false`. Thanks to the `{{if}}` action, the "Dot is true!" text is only included in the output when dot is `true`.

This portion of the template will appear only if the dot value is true.

```
executeTemplate("start {{if .}}Dot is true!{{end}} finish\n", true)
executeTemplate("start {{if .}}Dot is true!{{end}} finish\n", false)
```

```
start Dot is true! finish
start  finish
```

Repeating parts of a template with "range" actions

A section of a template between a `{{range}}` action and its corresponding `{{end}}` marker will be repeated for each value collected in an array, slice, map, or channel. Any actions within that section will also be repeated.

Within the repeated section, the value of dot will be set to the current element from the collection, allowing you to include each element in the output or do other processing with it.

This template includes a `{{range}}` action that will output each element in a slice. Before and after the loop, the value of dot will be the slice itself. But *within* the loop, dot refers to the current element of the slice. You'll see this reflected in the output.

This portion of the template will be repeated once for each element in the slice.

```
templateText := "Before loop: {{.}}\n{{range .}}In loop: {{.}}\n{{end}}After loop: {{.}}\n"
```

Before the loop, dot contains the entire slice.

In the loop, dot contains the current value from the slice.

After the loop, dot contains the entire slice again.

```
executeTemplate(templateText, []string{"do", "re", "mi"})
```

Pass a slice as the data value.

```
Before loop: [do re mi]
In loop: do
In loop: re
In loop: mi
After loop: [do re mi]
```

This template works with a slice of `float64` values, which it will display as a list of prices.

This portion of the template will be repeated once for each element in the slice.

```
templateText = "Prices:\n{{range .}}${{.}}\n{{end}}"
executeTemplate(templateText, []float64{1.25, 0.99, 27})
```

```
Prices:
$1.25
$0.99
$27
```

If the value provided to the `{{range}}` action is empty or `nil`, the loop won't be run at all:

```
templateText = "Prices:\n{{range .}}${{.}}\n{{end}}"
executeTemplate(templateText, []float64{})
executeTemplate(templateText, nil)
```

Pass in an empty slice.

Pass in nil.

```
Prices:
Prices:
```

Looped section isn't included.

Looped section isn't included.

Inserting struct fields into a template with actions

Simple types usually can't hold the variety of information needed to fill in a template, though. It's more common to use struct types when executing a template.

If the value in dot is a struct, then an action with dot followed by a field name will insert that field's value in the template. Here we create a `Part` struct type, then set up a template that will output a `Part` value's `Name` and `Count` fields:

```
type Part struct {
       Name  string
       Count int
}
templateText := "Name: {{.Name}}\nCount: {{.Count}}\n"
executeTemplate(templateText, Part{Name: "Fuses", Count: 5})
executeTemplate(templateText, Part{Name: "Cables", Count: 2})
```

Insert the value of the Part's Name field.

Insert the value of the Part's Count field.

```
Name: Fuses
Count: 5
Name: Cables
Count: 2
```

Finally, below we declare a `Subscriber` struct type and a template that prints them. The template will output the `Name` field regardless, but it uses an `{{if}}` action to output the `Rate` field only if the `Active` field is set to `true`.

```
type Subscriber struct {
       Name   string
       Rate   float64
       Active bool
}
templateText = "Name: {{.Name}}\n{{if .Active}}Rate: ${{.Rate}}\n{{end}}"
subscriber := Subscriber{Name: "Aman Singh", Rate: 4.99, Active: true}
executeTemplate(templateText, subscriber)
subscriber = Subscriber{Name: "Joy Carr", Rate: 5.99, Active: false}
executeTemplate(templateText, subscriber)
```

This portion of the template will be output only if the Subscriber's Active field value is true.

Rate section is omitted for an inactive Subscriber.

```
Name: Aman Singh
Rate: $4.99
Name: Joy Carr
```

There's a lot more you can do with templates, and we don't have space to cover it all here. To learn more, look up the documentation for the `text/template` package:

```
File Edit Window Help
$ go doc text/template
package template // import "text/template"

Package template implements data-driven templates for generating textual
output.

To generate HTML output, see package html/template, which has the same
interface as this package but automatically secures HTML output against
certain attacks.
...
```

Reading a slice of signatures in from a file

Now that we know how to insert data into a template, we're almost ready to insert signatures into the guestbook page. But first, we're going to need signatures that we can insert.

In your project directory, save a few lines of text to a plain-text file named *signatures.txt*. These are going to serve as our "signatures" for now.

Save these placeholder "signatures" to a file in your project directory.

```
First signature
Second signature
Third signature
```

signatures.txt

Now we need the ability to load these signatures into our app. In *guestbook.go*, add a new `getStrings` function. This function will work a lot like the `datafile.GetStrings` function we wrote back in Chapter 7, reading a file and appending each line to a slice of strings, which it then returns.

But there are a couple differences. First, the new `getStrings` will rely on our `check` function to report errors rather than returning them.

Second, if the file doesn't exist, `getStrings` will just return `nil` in place of the slice of strings, rather than reporting an error. It does this by passing any `error` value it gets from `os.Open` to the `os.IsNotExist` function, which will return `true` if the error indicates that the file doesn't exist.

```
import (
        "bufio"
        "fmt"
        "html/template"
        "log"
        "net/http"
        "os"
)

// Code omitted...

func getStrings(fileName string) []string {
        var lines []string
        file, err := os.Open(fileName)
        if os.IsNotExist(err) {
                return nil
        }
        check(err)
        defer file.Close()
        scanner := bufio.NewScanner(file)
        for scanner.Scan() {
                lines = append(lines, scanner.Text())
        }
        check(scanner.Err())
        return lines
}

// Code omitted...
```

Used by getStrings

We'll use this within viewHandler in a moment.

Used by getStrings

Open the file.

If an error is returned saying the file doesn't exist...

...return nil instead of the slice of strings.

For any other kind of error, report it and exit.

After the function exits, ensure the file is closed.

Report any scanning error and exit.

Reading a slice of signatures in from a file (continued)

We'll also make a small change to the `viewHandler` function, adding a call to `getStrings` and a temporary `fmt.Printf` call to show us what was loaded from the file.

```
func viewHandler(writer http.ResponseWriter, request *http.Request) {
        signatures := getStrings("signatures.txt")
        fmt.Printf("%#v\n", signatures)
        html, err := template.ParseFiles("view.html")
        check(err)
        err = html.Execute(writer, nil)
        check(err)
}
```

Add a call to getStrings.

Display the loaded signatures.

Let's try the `getStrings` function out. In your terminal, change to your project directory, and run *guestbook.go*. Visit *http://localhost:8080/guestbook* in your browser, so that the `viewHandler` function is called. It will call `getStrings`, which will load and return a slice with the contents of *signatures.txt*.

Loading the page will cause the slice of signatures to be displayed.

```
File  Edit  Window  Help
$ cd myproject
$ go run guestbook.go
[]string{"First signature", "Second signature", "Third signature"}
```

there are no Dumb Questions

Q: What happens if the *signatures.txt* file doesn't exist, and `getStrings` returns `nil`? Won't that cause problems rendering the template?

A: There's no need to worry. Just as we've already seen with the `append` function, other functions in Go are generally set up to treat `nil` slices and maps as if they were empty. For example, the `len` function simply returns `0` if it's passed a `nil` slice:

Since no slice is assigned, mySlice's value will be nil.

```
var mySlice []string
fmt.Printf("%#v, %d\n", mySlice, len(mySlice))
```

```
[]string(nil), 0
```

But "len" returns 0, as if we'd passed in an empty slice!

And template actions treat `nil` slices and maps as if they were empty, too. As we learned, for example, the `{{range}}` action simply skips outputting its contents if it's given a `nil` value. So having `getStrings` return `nil` instead of a slice will be fine; if no signatures are loaded from the file, the template will just skip outputting any signatures.

A struct to hold the signatures and signature count

Now, we could just pass this slice of signatures to our HTML template's `Execute` method, and have the signatures inserted into the template. But we also want our main guestbook page to show the *number* of signatures we've received, along with the signatures themselves.

We only get to pass one value to the template's `Execute` method, though. So we'll need to create a struct type that will hold both the total number of signatures as well as the slice with the signatures themselves.



Near the top of the *guestbook.go* file, add a new declaration for a new `Guestbook` struct type. It should have two fields: a `SignatureCount` field to hold the number of signatures, and a `Signatures` field to hold the slice with the signatures themselves.

```
type Guestbook struct {
        SignatureCount int
        Signatures     []string
}
```

Near the top of guestbook.go, define this new type.

Now we need to update `viewHandler` to create a new `Guestbook` struct and pass it to the template. First, we won't be needing the `fmt.Printf` call that displays the contents of the `signatures` slice anymore, so remove that. (You'll also need to remove `"fmt"` from the `import` section.) Then, create a new `Guestbook` value. Set its `SignatureCount` field to the length of the `signatures` slice, and set its `Signatures` field to the `signatures` slice itself. Finally, we need to actually pass the data into the template. So change the data value being passed as the second argument to the `Execute` method from `nil` to our new `Guestbook` value.

```
func viewHandler(writer http.ResponseWriter, request *http.Request) {
        signatures := getStrings("signatures.txt")
        html, err := template.ParseFiles("view.html")
        check(err)
        guestbook := Guestbook{
                SignatureCount: len(signatures),
                Signatures:     signatures,
        }
        err = html.Execute(writer, guestbook)
        check(err)
}
```

Create a new Guestbook struct.

Set its SignatureCount field to the length of the signatures slice.

Set its Signatures field to the signatures slice itself.

Pass the struct to the Template's Execute method.

Updating our template to include our signatures

Now let's update the template text in *view.html* to display the list of signatures.

We're passing the `Guestbook` struct into the template's `Execute` method, so within the template, dot represents that `Guestbook` struct. In the first `div` element, replace the `X` placeholder in `X total signatures` with an action that inserts the `Guestbook`'s `SignatureCount` field: `{{.SignatureCount}}`.

The second `div` element holds a series of `p` (paragraph) elements, one for each signature. Use a `range` action to loop over each signature in the `Signatures` slice: `{{range .Signatures}}`. (Don't forget the corresponding `{{end}}` marker before the end of the `div` element.) Within the `range` action, include a `p` HTML element with an action that outputs dot nested inside it: `<p>{{.}}</p>`. Remember that dot gets set to each element of a slice in turn, so this will cause a `p` element to be output for each signature in the slice, with its content set to that signature's text.

```
<h1>Guestbook</h1>

<div>
        {{.SignatureCount}} total signatures -
        <a href="/guestbook/new">Add Your Signature</a>
</div>

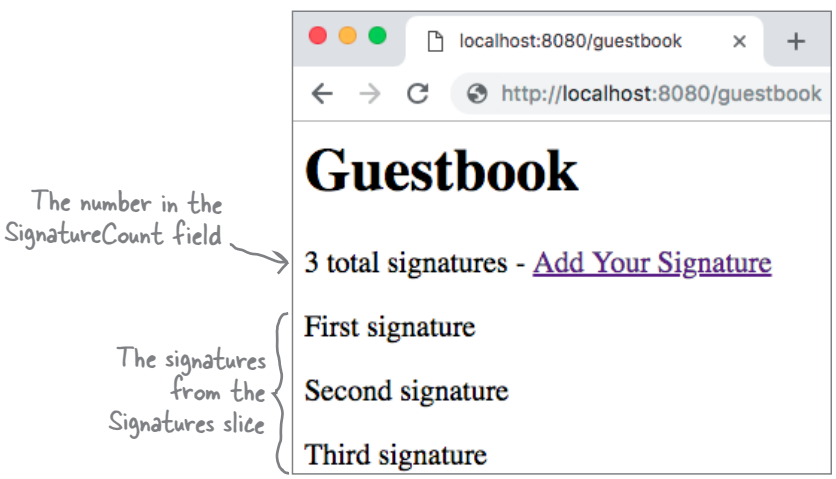
<div>
        {{range .Signatures}}
                <p>{{.}}</p>
        {{end}}
</div>
```

Insert the number of signatures from the Guestbook struct.

Repeat for each string in the Signatures slice.

Insert a <p> element containing the current signature.

Finally, we can test out our template with our data included! Restart the *guestbook.go* app, and visit *http://localhost:8080/guestbook* in your browser again. The response should show your template. The total number of signatures should be at the top, and each signature should appear within its own <p> element!



there are no Dumb Questions

Q: You mentioned the `html/template` package has some "security features." What are they?

A: The `text/template` package inserts values into a template as is, no matter what they contain. But that means that visitors could add HTML code as a "signature," and it would be treated as part of the page's HTML.

You can try this yourself. In *guestbook.go*, change the `html/template` import to `text/template`. (You won't need to change any other code, because the names of all the functions in the two packages are identical.) Then, add the following as a new line in your *signatures.txt* file:

```
<script>alert("hi!");</script>
```

This is an HTML tag containing JavaScript code. If you try running the app and reload the signatures page, you'll see an annoying alert pop up, because the `text/template` package included this code in the page as is.

Now go back to *guestbook.go*, change the import back to `html/template`, and restart the app. If you reload the page, instead of an alert pop up, you'll see text that looks just like the above script tag in the page.

But that's because the `html/template` package automatically "escaped" the HTML, replacing the characters that cause it to be treated as HTML with code that causes it to appear in the page's text instead (where it's harmless). Here's what actually gets inserted into the response:

```
&lt;script&gt;alert(&#34;hi!&#34;);&lt;/script&gt;
```

Inserting script tags like this is just one of many ways unscrupulous users can insert malicious code into your web pages. The `html/template` package makes it easy to protect against this and many other attacks!



Below is a program that loads an HTML template in from a file, and outputs it to the terminal. Fill in the blanks in the *bill.html* file so that the program will run and produce the output shown.

bill.go

```
type Invoice struct {
    Name     string
    Paid     bool
    Charges  []float64
    Total    float64
}

func main() {
    html, err := template.ParseFiles("bill.html")
    check(err)
    bill := Invoice{
       Name:    "Mary Gibbs",
       Paid:    true,
       Charges: []float64{23.19, 1.13, 42.79},
       Total:   67.11,
    }
    err = html.Execute(os.Stdout, bill)
    check(err)
}
```

bill.html

```
<h1>Invoice</h1>

<p>Name: _____</p>

{{if _____}}
<p>Paid - Thank you!</p>
_____

<h1>Fees</h1>

{{range .Charges}}
<p>$_____</p>
{{end}}

<p>Total: $_____</p>
```

Output

```
<h1>Invoice</h1>

<p>Name: Mary Gibbs</p>

<p>Paid - Thank you!</p>

<h1>Fees</h1>

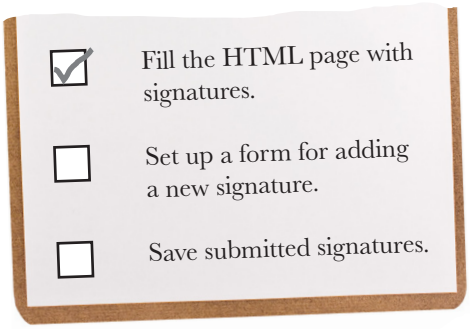
<p>$23.19</p>
<p>$1.13</p>
<p>$42.79</p>

<p>Total: $67.11</p>
```

Answers on page 478.

Letting users add data with HTML forms

That's another task complete. We're getting close: only two tasks left to go!



Up next, we need to allow visitors to add their own signature. We'll need to create an HTML *form* where they can type a signature in. A form usually provides one or more fields that a user can enter data into, and a submit button that allows them to send the data to the server.

In your project directory, create a file called *new.html* with the HTML code below. There are some tags here that we haven't seen before:

- **`<form>`**: This element encloses all the other form components.
- **`<input>` with a `type` attribute of `"text"`**: A text field where the user can enter a string. Its `name` attribute will be used to label the field's value in the data sent to the server (kind of like a map key).
- **`<input>` with a `type` attribute of `"submit"`**: Creates a button that the user can click to submit the form's data.

```
<h1>Add a Signature</h1>

<form>
  <div><input type="text" name="signature"></div>
  <div><input type="submit"></div>
</form>
```

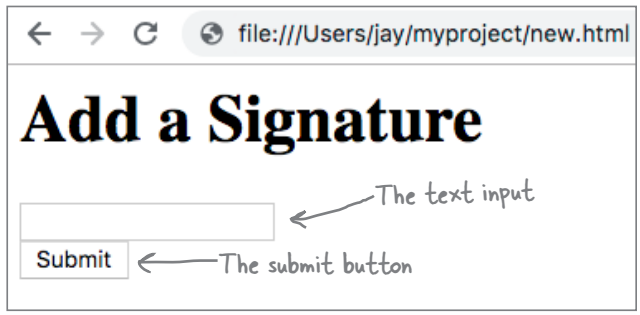
The "signature" text input

A button that submits the form data



new.html

If we were to load this HTML in the browser, it would look like this:



Responding with the HTML form



We already have an "Add Your Signature" link in *view.html* that points to a path of */guestbook/new*. Clicking on this link will take you to a new path on the same server, so it's just like typing in this URL:

http://localhost:8080/guestbook/new

But visiting this path right now just responds with the error "404 page not found." We'll need to set up the app to respond with the form in *new.html* when users click the link.

In *guestbook.go*, add a `newHandler` function. It will look much like the early versions of our `viewHandler` function. Just like `viewHandler`, `newHandler` should take an `http.ResponseWriter` and a pointer to an `http.Request` as parameters. It should call `template.ParseFiles` on the *new.html* file. And then it should call `Execute` on the resulting template, so that the contents of *new.html* get written to the HTTP response. We won't be inserting any data into this template, so we pass `nil` as the data value for the call to `Execute`.

Then we need to ensure that the `newHandler` function is called when the "Add Your Signature" link is clicked. In the `main` function, add another call to `http.HandleFunc`, and set up `newHandler` as the handler function for requests with a path of */guestbook/new*.

Add another handler function, with the same parameters as viewHandler.

Load the contents of new.html as the text of a template.

Write the template to the response (there's no need to insert any data in it).

Set the newHandler function up to handle requests with a path of "/guestbook/new".

```
// Code omitted...

func newHandler(writer http.ResponseWriter, request *http.Request) {
        html, err := template.ParseFiles("new.html")
        check(err)
        err = html.Execute(writer, nil)
        check(err)
}

// Code omitted...

func main() {
        http.HandleFunc("/guestbook", viewHandler)
        http.HandleFunc("/guestbook/new", newHandler)
        err := http.ListenAndServe("localhost:8080", nil)
        log.Fatal(err)
}
```

If we save the above code and restart *guestbook.go*, then click the "Add Your Signature" link, we'll be taken to the */guestbook/new* path. The `newHandler` function will be called, which will load our form HTML from *new.html* and include it in the response.



Form submission requests



We've completed yet another task. Just one to go!

When someone visits the */guestbook/new* path, either by entering it directly or by clicking a link, our form for entering a signature is displayed. But if you fill in that form and click Submit, nothing useful happens.



The browser will just make another request for the */guestbook/new* path. The content of the `"signature"` form field will be added as an ugly-looking parameter on the end of the URL. And because our `newHandler` function doesn't know how to do anything useful with the form data, it will simply be discarded.



Our app can respond to requests to display the form, but there's no way for the form to submit its data back to the app. We'll need to fix this before we can save visitors' signatures.

Path and HTTP method for form submissions

Submitting a form actually requires *two* requests to the server: one to *get* the form, and a second to *send* the user's entries back to the server. Let's update the form's HTML to specify where and how this second request should be sent.

Edit *new.html*, and add two new HTML attributes to the `form` element. The first attribute, `action`, will specify the path to use for the submission request. Instead of letting the path default back to */guestbook/new*, we'll specify a new path: */guestbook/create*.

We'll also need a second attribute, named `method`, which should have a value of `"POST"`.

```
<h1>Add a Signature</h1>

<form action="/guestbook/create" method="POST">
  <div><input type="text" name="signature"></div>
  <div><input type="submit"></div>
</form>
```

Submit the form data to "/guestbook/create".

Submit as a POST request, rather than GET.

new.html

That `method` attribute requires a little explanation... HTTP defines several *methods* that a request can use. These aren't the same as methods on a Go value, but the meaning is similar. GET and POST are among the most common methods:

- **GET**: Used when your browser needs to *get* something from the server, usually because you entered a URL or clicked a link. This could be an HTML page, an image, or some other resource.
- **POST**: Used when your browser needs to *add* some data to the server, usually because you submitted a form with new data.

We're adding new data to the server: a new guestbook signature. So it seems like we should submit the data using a POST request.

Forms are submitted using GET requests by default, though. This is why we needed to add a `method` attribute with a value of `"POST"` to the `form` element.

Now, if we reload the */guestbook/new* page and resubmit the form, the request will use a path of */guestbook/create* instead. We'll get a "404 page not found" error, but that's because we haven't set up a handler for the */guestbook/create* path yet.

We'll also see that the form data is no longer added onto the end of the URL. This is because the form is being submitted using a POST request.



Reload and resubmit the form...



There's no handler for this path yet, but that's OK.

No more ugly parameter on the end of the URL

Getting values of form fields from the request

Now that we're submitting the form using a POST request, the form data is embedded in the request itself, rather than being appended to the request path as a parameter.

Let's address that "404 page not found" error we get when form data is submitted to the */guestbook/create* path. When we do, we'll also see how to access the form data from the POST request.

As usual, we'll do this by adding a request handler function. In the `main` function of *guestbook.go*, call `http.HandleFunc`, and assign requests with a path of `"/guestbook/create"` to a new `createHandler` function.

Then add a definition for the `createHandler` function itself. It should accept an `http.ResponseWriter` and a pointer to an `http.Request`, just like the other handler functions.

Unlike the other handler functions, though, `createHandler` is meant to work with form data. That data can be accessed through the `http.Request` pointer that gets passed to the handler function. (That's right, after ignoring `http.Request` values all this time, we finally get to use one!)

For now, let's just take a look at the data the request contains. Call the `FormValue` method on the `http.Request`, and pass it the string `"signature"`. This will return a string with the value of the `"signature"` form field. Store it in a variable named `signature`.

Let's write the field value to the response so we can see it in the browser. Call the `Write` method on the `http.ResponseWriter`, and pass `signature` to it (but convert it to a slice of bytes first, of course). As always, `Write` will return a number of bytes written and an `error` value. We'll ignore the number of bytes by assigning it to `_`, and call `check` on the `error`.

Define another request handler function, with the same parameters as the others.

```
func createHandler(writer http.ResponseWriter, request *http.Request) {
        signature := request.FormValue("signature")
        _, err := writer.Write([]byte(signature))
        check(err)
}

func main() {
        http.HandleFunc("/guestbook", viewHandler)
        http.HandleFunc("/guestbook/new", newHandler)
        http.HandleFunc("/guestbook/create", createHandler)
        err := http.ListenAndServe("localhost:8080", nil)
        log.Fatal(err)
}
```

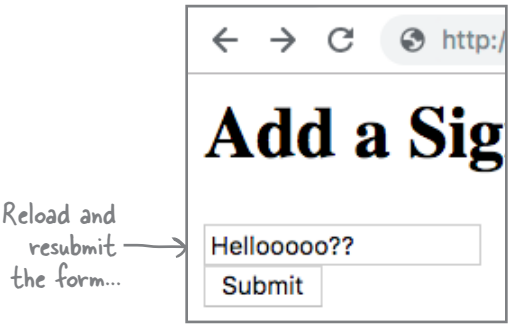
Get the value of the "signature" form field.

Write the field value to the response.

Call createHandler for requests with a path of "/guestbook/create".

Getting values of form fields from the request (continued)

Let's see if our form submissions are getting through to the `createHandler` function. Restart *guestbook.go*, visit the */guestbook/new* page, and submit the form again.



You'll be taken to the */guestbook/create* path, and instead of a "404 page not found" error, the app will respond with the value you entered in the `"signature"` field!



If you want, you can click your browser's back button to return to the */guestbook/new* page, and try different submissions. Whatever you enter will be echoed to the browser.

Setting up a handler for HTML form submissions was a big step. We're getting close!

Saving the form data

Our `createHandler` function is receiving the request with the form data, and is able to retrieve the guestbook signature from it. Now all we need to do is add that signature to our *signatures.txt* file. We'll handle that within the `createHandler` function itself.

First, we'll get rid of the call to the `Write` method on the `ResponseWriter`; we only needed that to confirm we could access the signature form field.

Now, let's add the code below. The `os.OpenFile` function is called in a slightly unusual way, and the details aren't directly relevant to writing a web app, so we won't describe it fully here. (See Appendix A if you want more info.) For now, all you need to know is that this code does three basic things:

1. It opens the *signatures.txt* file, creating it if it doesn't exist.
2. It adds a line of text to the end of the file.
3. It closes the file.

```
import (
    // ...
    "fmt"
    // ...
)

// Code omitted...

func createHandler(writer http.ResponseWriter, request *http.Request) {
    signature := request.FormValue("signature")
    options := os.O_WRONLY | os.O_APPEND | os.O_CREATE
    file, err := os.OpenFile("signatures.txt", options, os.FileMode(0600))
    check(err)
    _, err = fmt.Fprintln(file, signature)
    check(err)
    err = file.Close()
    check(err)
}
```

Reimport the "fmt" package.

(See Appendix A for a full description of os.OpenFile.)

Options for opening the file

Open the file.

Write a signature to the file, on a new line.

Close the file.

The `fmt.Fprintln` function adds a line of text to a file. It takes the file to write to and the string to write (no need to convert to a `[]byte`) as arguments. Just like the `Write` methods we saw earlier in this chapter, `Fprintln` returns the number of bytes successfully written to the file (which we ignore), and any error encountered (which we pass to the `check` function).

Finally, we call the `Close` method on the file. You might notice that we did *not* use the `defer` keyword. This is because we're writing to the file, rather than reading from it. Calling `Close` on a file you're writing to can result in errors that we need to handle, and we can't readily do that if we use `defer`. So, we simply call `Close` as part of the regular program flow and then pass its return value to `check`.

Saving the form data (continued)

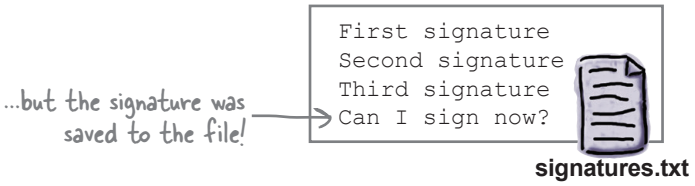
Save the previous code and restart *guestbook.go*. Fill in and submit the form on the */guestbook/go* page.



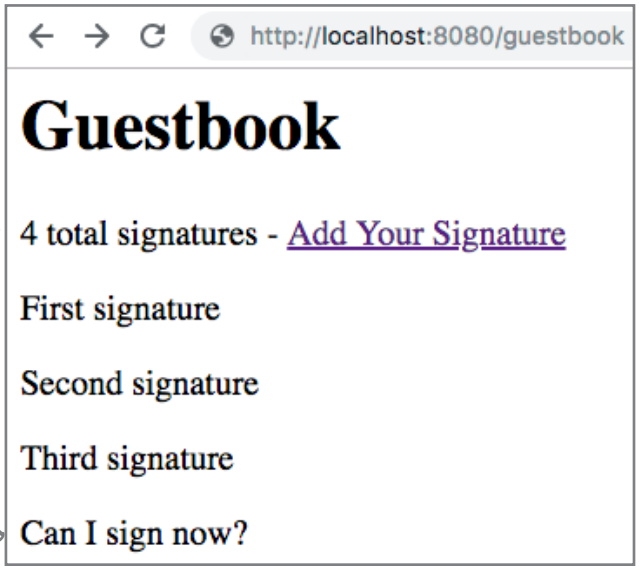
Your browser will load the */guestbook/create* path, which shows as a totally blank page now (because `createHandler` is no longer writing anything to the `http.ResponseWriter`).



But if you look at the contents of the *signatures.txt* file, you'll see a new signature saved at the end!



And if you visit the list of signatures at */guestbook*, you'll see the signature count has increased by one, and the new signature appears in the list!



(By the way, when you created the signatures.txt file, if you didn't press Enter after the final line, your new signature will appear squashed onto the end of the previous signature. That's OK! You can edit signatures.txt to fix it, and all future signatures will be saved on separate lines.)

HTTP redirects

We have our `createHandler` function saving new signatures. There's just one more thing we need to take care of. When a user submits the form, their browser loads the */guestbook/create* path, which shows a blank page.



There's nothing useful to show at the */guestbook/create* path anyway; it's just there to accept requests to add a new signature. Instead, let's have the browser load the */guestbook* path, so the user can see their new signature in the guestbook.

At the end of the `createHandler` function, we'll add a call to `http.Redirect`, which sends a response to the browser directing it to load a different resource than the one it requested. `Redirect` takes an `http.ResponseWriter` and a `*http.Request` as its first two arguments, so we'll just give it the values from the `writer` and `request` parameters to `createHandler`. Then `Redirect` needs a string with a path to redirect the browser to; we'll redirect to `"/guestbook"`.

The last argument to `Redirect` needs to be a status code to give the browser. Every HTTP response needs to include a status code. Our responses so far have had their codes set automatically for us: successful responses had a code of 200 ("OK"), and requests for nonexistent pages had a code of 404 ("Not found"). We need to specify a code for `Redirect`, though, so we'll use the constant `http.StatusFound`, which will cause the redirect response to have a status of 302 ("Found").

```
func createHandler(writer http.ResponseWriter, request *http.Request) {
        signature := request.FormValue("signature")
        options := os.O_WRONLY | os.O_APPEND | os.O_CREATE
        file, err := os.OpenFile("signatures.txt", options, os.FileMode(0600))
        check(err)
        _, err = fmt.Fprintln(file, signature)
        check(err)
        err = file.Close()
        check(err)
        http.Redirect(writer, request, "/guestbook", http.StatusFound)
}
```

The path to redirect to

A response code indicating the request was successful

We need to pass Redirect the ResponseWriter...

...and also the original request.

Now that we've added the call to `Redirect`, submitting the signature form should work something like this:

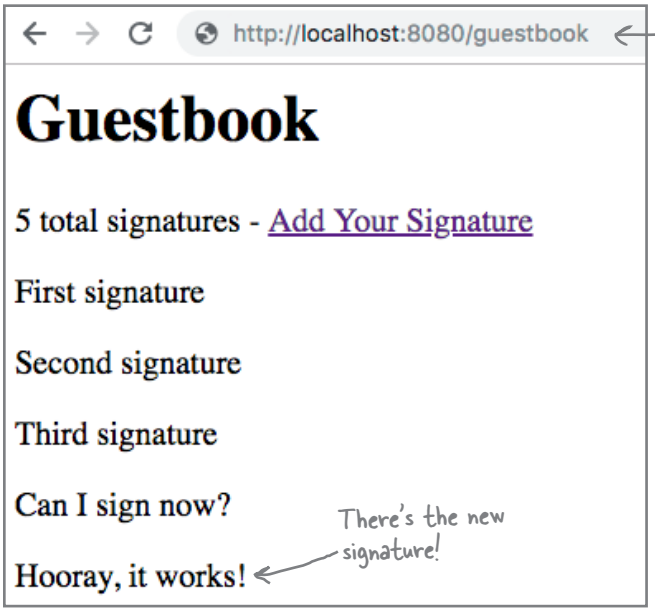
1. The browser submits an HTTP POST request to the */guestbook/create* path.
2. The app responds with a redirect to */guestbook*.
3. The browser sends a GET request for the */guestbook* path.

Let's try it all out!

Let's see if the redirect works! Restart *guestbook.go*, and visit the */guestbook/new* path. Fill in the form and submit it.

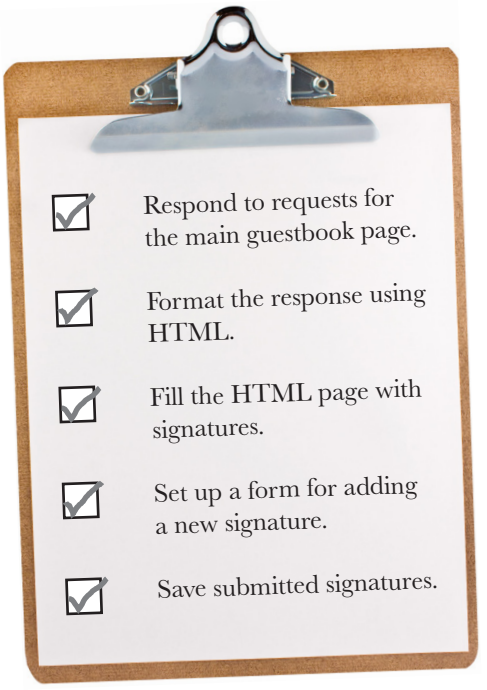
The app will save the form contents to *signatures.txt*, then immediately redirect the browser to the */guestbook* path. When the browser requests */guestbook*, the app will load the updated *signatures.txt* file, and the user will see their new signature in the list!





Our app is saving signatures submitted from the form and displaying them along with all the others. All our features are complete.

It took quite a few components to make it all work, but you now have a usable web app!



Our complete app code

The code for our app has gotten so long, we've only been able to look at it in bits and pieces. Let's take one more moment to look at all the code in one place!

The *guestbook.go* file makes up the bulk of the code for the app. (In an app intended for wide use, we might have split some of this code into multiple packages and source files within our Go workspace directory, and you can do that yourself if you want.) We've gone through and added comments documenting the `Guestbook` type and each of the functions.




```
package main

import (
    "bufio"
    "fmt"
    "html/template"
    "log"
    "net/http"
    "os"
)

// Guestbook is a struct used in rendering view.html.
type Guestbook struct {
    SignatureCount int
    Signatures     []string
}

// check calls log.Fatal on any non-nil error.
func check(err error) {
    if err != nil {
       log.Fatal(err)
    }
}

// viewHandler reads guestbook signatures and displays them together
// with a count of all signatures.
func viewHandler(writer http.ResponseWriter, request *http.Request) {
    signatures := getStrings("signatures.txt")
    html, err := template.ParseFiles("view.html")
    check(err)
    guestbook := Guestbook{
       SignatureCount: len(signatures),
       Signatures:     signatures,
    }
    err = html.Execute(writer, guestbook)
    check(err)
}
```

We can only pass a single value to a Template's Render method, so this struct will hold all the data we need.

This will hold the total number of signatures.

This will hold the signatures themselves.

We'll call this when we need to check an error value returned from a function or method.

Most of the time the value will be nil, but if not...

...output the error and exit the program.

Like all HTTP handler functions, this needs to accept an http.ResponseWriter and a *http.Request.

Read signatures from a file.

Create a template based on the contents of view.html.

Store the number of signatures.

Store the signatures themselves.

Insert the Guestbook struct data into the template and write the result out to the ResponseWriter.



guestbook.go (continued)

```
// newHandler displays a form to enter a signature.
func newHandler(writer http.ResponseWriter, request *http.Request) {
    html, err := template.ParseFiles("new.html")
    check(err)
    err = html.Execute(writer, nil)
    check(err)
}

// createHandler takes a POST request with a signature to add, and
// appends it to the signatures file.
func createHandler(writer http.ResponseWriter, request *http.Request) {
    signature := request.FormValue("signature")
    options := os.O_WRONLY | os.O_APPEND | os.O_CREATE
    file, err := os.OpenFile("signatures.txt", options, os.FileMode(0600))
    check(err)
    _, err = fmt.Fprintln(file, signature)
    check(err)
    err = file.Close()
    check(err)
    http.Redirect(writer, request, "/guestbook", http.StatusFound)
}

// getStrings returns a slice of strings read from fileName, one
// string per line.
func getStrings(fileName string) []string {
    var lines []string
    file, err := os.Open(fileName)
    if os.IsNotExist(err) {
       return nil
    }
    check(err)
    defer file.Close()
    scanner := bufio.NewScanner(file)
    for scanner.Scan() {
       lines = append(lines, scanner.Text())
    }
    check(scanner.Err())
    return lines
}

func main() {
    http.HandleFunc("/guestbook", viewHandler)
    http.HandleFunc("/guestbook/new", newHandler)
    http.HandleFunc("/guestbook/create", createHandler)
    err := http.ListenAndServe("localhost:8080", nil)
    log.Fatal(err)
}
```

Load the HTML form in from a template.

Write the template to the ResponseWriter (there is no data to insert).

Get the value of the "signature" form field.

Open the file for writing. If it exists, append to it. If not, create it.

Add the form field contents to the file.

Close the file.

Redirect the browser to the main guestbook page.

Each line of the file will be appended to this slice as a string.

Open the file.

If we get an error indicating the file doesn't exist...

...return nil instead of a slice.

All other errors should be checked and reported normally.

Create a scanner for the file contents.

For each line of the file...

...append its text to the slice.

Report any errors encountered while scanning.

Return the slice of strings.

Requests to view the signature list will be handled by the viewHandler function.

Requests to get the HTML form will be handled by newHandler.

Requests to submit the form will be handled by createHandler.

Loop forever, passing HTTP requests to the appropriate function for handling.

The *view.html* file provides the HTML template for the list of signatures. Template actions provide places to insert the number of signatures, as well as the entire signature list.

```
<h1>Guestbook</h1>

<div>
    {{.SignatureCount}} total signatures -
    <a href="/guestbook/new">Add Your Signature</a>
</div>

<div>
    {{range .Signatures}}
      <p>{{.}}</p>
    {{end}}
</div>
```

A level-one heading for the top of the page

The "dot" value is the Guestbook struct. Insert its SignatureCount field here.

A link to the path that presents the HTML form

Get the slice from the Guestbook struct's Signatures field, and repeat for each string it contains.

This gets repeated for each slice element. "Dot" gets set to the current signature string. Insert an HTML paragraph element containing the signature.



The *new.html* file simply holds the HTML form for new signatures. No data will be inserted into it, so no template actions are present.

```
<h1>Add a Signature</h1>
<form action="/guestbook/create" method="POST">
  <div><input type="text" name="signature"></div>
  <div><input type="submit"></div>
</form>
```

A level-one heading for the top of the page

An HTML form

Submissions will go to the "/guestbook/create" path.

Submissions will use the POST method.

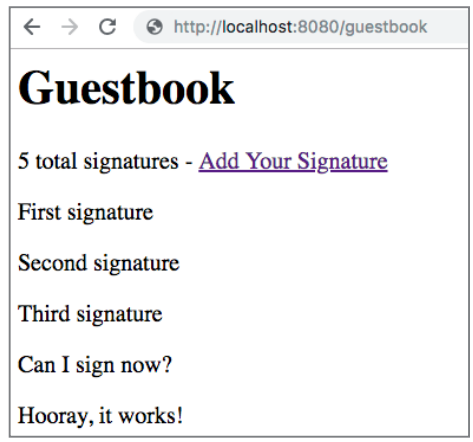
A button to submit the form

A text field whose data can be accessed under the name "signature"



And that's it—a complete web app that can store user-submitted signatures and retrieve them again later!

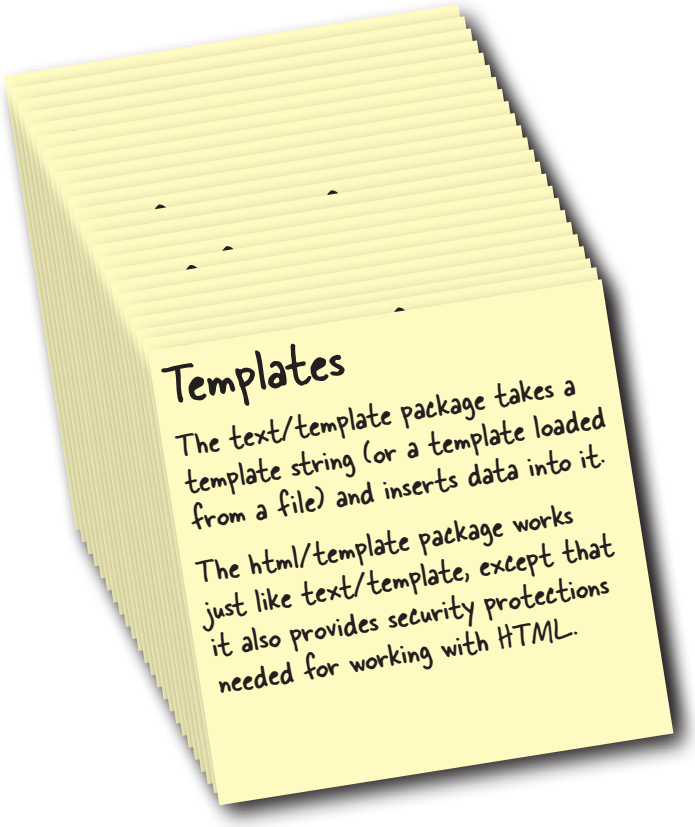
Writing web apps can be complex, but the `net/http` and `html/template` packages leverage the power of Go to make the whole process simpler for you!





Your Go Toolbox

That's it for Chapter 16! You've added templates to your toolbox.



BULLET POINTS

- A template string contains text that will be output verbatim. Within this text, you can insert various **actions** containing simple code that will be evaluated. Actions can be used to insert data into the template text.
- A `Template` value's `Execute` method takes a value that satisfies the `io.Writer` interface, and a data value that can be accessed within actions in the template.
- Template actions can reference the data value passed to `Execute` with `{{.}}`, referred to as "dot." The value of dot can change within various contexts in the template.
- A section of a template between an `{{if}}` action and its corresponding `{{end}}` marker will be included only if a certain condition is true.
- A section of a template between a `{{range}}` action and its corresponding `{{end}}` marker will be repeated for each value within an array, slice, map, or channel. Any actions within that section will also be repeated.
- Within a `{{range}}` section, the value of dot will be updated to refer to the current element of the collection being processed.
- If dot refers to a struct value, the value of fields in that struct can be inserted with `{{.FieldName}}`.
- HTTP GET requests are commonly used when a browser needs to get data from the server.
- HTTP POST requests are used when a browser needs to submit new data to the server.
- Form data from a request can be accessed using an `http.Request` value's `FormValue` method.
- The `http.Redirect` function can be used to direct the browser to request a different path.



Below is a program that loads an HTML template in from a file, and outputs it to the terminal. Fill in the blanks in the *bill.html* file so that the program will run and produce the output shown.

bill.go

```
type Invoice struct {
    Name     string
    Paid     bool
    Charges  []float64
    Total    float64
}

func main() {
    html, err := template.ParseFiles("bill.html")
    check(err)
    bill := Invoice{
       Name:    "Mary Gibbs",
       Paid:    true,
       Charges: []float64{23.19, 1.13, 42.79},
       Total:   67.11,
    }
    err = html.Execute(os.Stdout, bill)
    check(err)
}
```

bill.html

```
<h1>Invoice</h1>

<p>Name: {{.Name}} </p>

{{if .Paid }}
<p>Paid - Thank you!</p>
{{end}}

<h1>Fees</h1>

{{range .Charges}}
<p>$ {{.}} </p>
{{end}}

<p>Total: $ {{.Total}} </p>
```

Is the Invoice's Paid field set to true?

The end of the "if" action

Output a <p> element for each item in the Charges slice.

Output

```
<h1>Invoice</h1>

<p>Name: Mary Gibbs</p>

<p>Paid - Thank you!</p>

<h1>Fees</h1>

<p>$23.19</p>
<p>$1.13</p>
<p>$42.79</p>

<p>Total: $67.11</p>
```

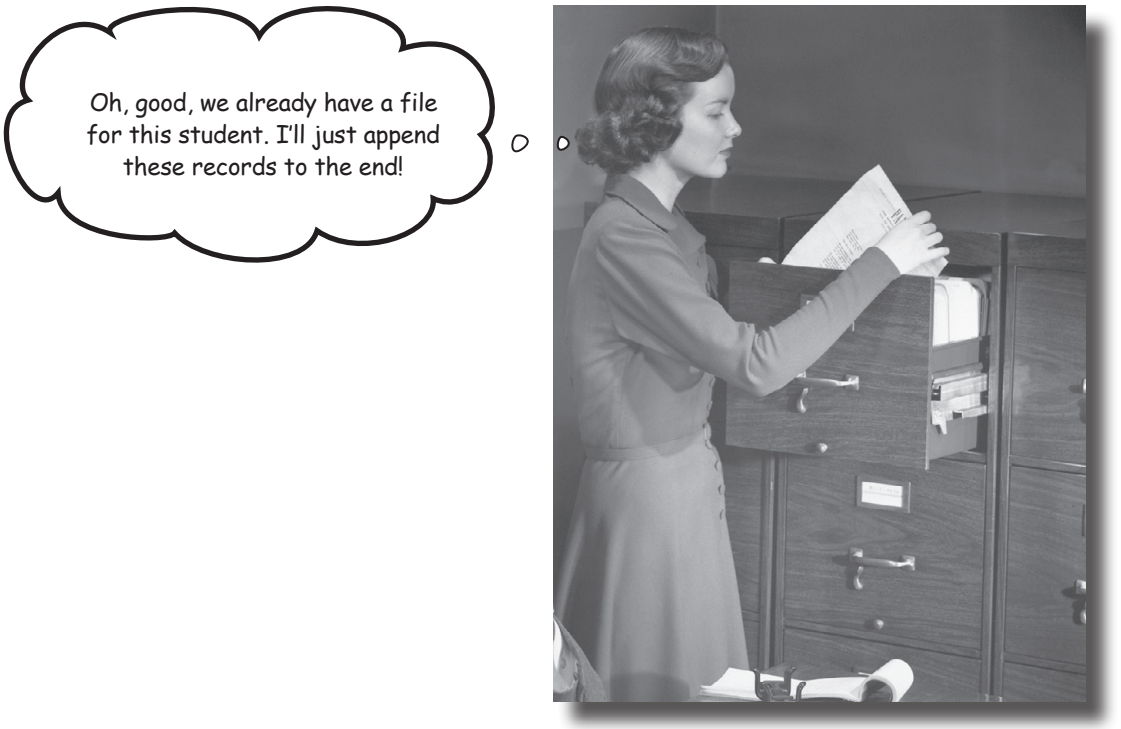


Congratulations!

You made it to the end.

Of course, there's still two appendixes.
And the index.
And then there's the website...
There's no escape, really.

Appendix A: Opening Files



Some programs need to write data to files, not just read data.

Throughout the book, when we've wanted to work with files, you had to create them in your text editor for your programs to read. But some programs *generate* data, and when they do, they need to be able to *write* data to a file.

We used the `os.OpenFile` function to open a file for writing earlier in the book. But we didn't have space then to fully explore how it worked. In this appendix, we'll show you everything you need to know in order to use `os.OpenFile` effectively!

Understanding os.OpenFile

In Chapter 16, we had to use the `os.OpenFile` function to open a file for writing, which required some rather strange-looking code:

```
options := os.O_WRONLY | os.O_APPEND | os.O_CREATE
file, err := os.OpenFile("signatures.txt", options, os.FileMode(0600))
```

Options for opening the file

Open the file.

Back then, we were focused on writing a web app, so we didn't want to take too much time out to fully explain `os.OpenFile`. But you'll almost certainly need to use this function again in your Go-writing career, so we added this appendix to take a closer look at it.

When you're trying to figure out how a function works, it's always good to start with its documentation. In your terminal, run **`go doc os OpenFile`** (or search for the `"os"` package documentation in your browser).

```
File Edit Window Help
$ go doc os OpenFile
func OpenFile(name string, flag int, perm FileMode) (*File, error)
    OpenFile is the generalized open call; most users will use Open or Create
    instead. It opens the named file with specified flag (O_RDONLY etc.) and
    ...
```

Its arguments are a `string` filename, an `int` "flag," and an `os.FileMode` "perm." It's pretty clear that the filename is just the name of the file we want to open. Let's figure out what this "flag" means first, then come back to the `os.FileMode`.

To help keep our code samples in this appendix short, assume that all our programs include a `check` function, just like the one we showed you in Chapter 16. It accepts an `error` value, checks whether it's `nil`, and if not, reports the error and exits the program.

```
func check(err error) {
        if err != nil {
                log.Fatal(err)
        }
}
```

Assume all the programs we're about to show you include this "check" function.

Passing flag constants to os.OpenFile

The description mentions that one possible value for the flag is `os.O_RDONLY`. Let's look that up and see what it means...

```
File Edit Window Help
$ go doc os O_RDONLY
const (
        // Exactly one of O_RDONLY, O_WRONLY, or O_RDWR must be specified.
        O_RDONLY int = syscall.O_RDONLY // open the file read-only.
        O_WRONLY int = syscall.O_WRONLY // open the file write-only.
        O_RDWR   int = syscall.O_RDWR   // open the file read-write.
        // The remaining values may be or'ed in to control behavior.
        O_APPEND int = syscall.O_APPEND // append data to the file when writing.
        O_CREATE int = syscall.O_CREAT  // create a new file if none exists.
        ...
)
    Flags to OpenFile wrapping those of the underlying system. Not all flags may
    be implemented on a given system.
```

From the documentation, it looks like `os.O_RDONLY` is one of several `int` constants intended for passing to the `os.OpenFile` function, which change the function's behavior.

Let's try calling `os.OpenFile` with some of these constants, and see what happens.

First, we'll need a file to work with. Create a plain-text file with a single line of text. Save it in any directory you want, with the name *aardvark.txt*.



Then, in the same directory, create a Go program that includes the `check` function from the previous page, and the following `main` function. In `main`, we call `os.OpenFile` with the `os.O_RDONLY` constant as the second argument. (Ignore the third argument for now; we'll talk about that later.) Then we create a `bufio.Scanner` and use it to print the contents of the file.

```
func main() {
        file, err := os.OpenFile("aardvark.txt", os.O_RDONLY, os.FileMode(0600))
        check(err)
        defer file.Close()
        scanner := bufio.NewScanner(file)
        for scanner.Scan() {
                fmt.Println(scanner.Text())
        }
        check(scanner.Err())
}
```

Open the file for reading.

Print each line of the file.

In your terminal, change to the directory where you saved the *aardvark.txt* file and your program, and use **`go run`** to run the program. It will open *aardvark.txt* and print out its contents.

```
File Edit Window Help
$ cd work
$ go run openfile.go
Aardvarks are...
```

Passing flag constants to os.OpenFile (continued)

Now let's try writing to the file instead. Update your `main` function with the code below. (You'll also need to remove unused packages from the `import` statement.) This time, we'll pass the `os.O_WRONLY` constant to `os.OpenFile`, so that it opens the file for writing. Then we'll call the `Write` method on the file with a slice of bytes to write to the file.

```
func main() {
        file, err := os.OpenFile("aardvark.txt", os.O_WRONLY, os.FileMode(0600))
        check(err)
        _, err = file.Write([]byte("amazing!\n"))
        check(err)
        err = file.Close()
        check(err)
}
```

Open the file for writing.

Write data to the file.

If we run the program, it will produce no output, but it will update the *aardvark.txt* file. But if we open *aardvark.txt*, we'll see that instead of appending the text to the end, the program overwrote part of the file!

The program inserted the new text at the start of the file, overwriting the data that had been there!



That's not how we wanted the program to work. What can we do?

Well, the `os` package has some other constants that might help. This includes an `os.O_APPEND` flag that should cause the program to append data to the file instead of overwriting it.

```
File Edit Window Help
$ go doc os O_RDONLY
...
        // The remaining values may be or'ed in to control behavior.
        O_APPEND int = syscall.O_APPEND // append data to the file when writing.
        O_CREATE int = syscall.O_CREAT  // create a new file if none exists.
        ...
```

But you can't just pass `os.O_APPEND` to `os.OpenFile` by itself; you'll get an error if you try.

Attempt to append to the existing file.

```
file, err := os.OpenFile("aardvark.txt", os.O_APPEND, os.FileMode(0600))
```

Runtime error!

```
write aardvark.txt:
bad file descriptor
```

The documentation says something about how `os.O_APPEND` and `os.O_CREATE` "may be or'ed in." This is referring to the *binary OR* operator. We'll need to take a few pages to explain how that works...

Binary notation

At the lowest level, computers have to represent information using simple switches, which can be either on or off. If one switch were used to represent a number, you could only represent the values `0` (switch "off") or `1` (switch "on"). Computer scientists call this a *bit*.

If you combine multiple bits, you can represent larger numbers. This is the idea behind *binary* notation. In everyday life, we have the most experience with decimal notation, which uses the digits 0 through 9. But binary notation uses only the digits 0 and 1 to represent numbers.

(If you'd like to know more, just type "binary" into your favorite web search engine.)

You can view the binary representation of various numbers (the bits the numbers are composed of) using `fmt.Printf` with the `%b` formatting verb:

Print the number in decimal notation.

Print the number in binary notation.

```
fmt.Printf("%3d: %08b\n", 0, 0)
fmt.Printf("%3d: %08b\n", 1, 1)
fmt.Printf("%3d: %08b\n", 2, 2)
fmt.Printf("%3d: %08b\n", 3, 3)
fmt.Printf("%3d: %08b\n", 4, 4)
fmt.Printf("%3d: %08b\n", 5, 5)
fmt.Printf("%3d: %08b\n", 6, 6)
fmt.Printf("%3d: %08b\n", 7, 7)
fmt.Printf("%3d: %08b\n", 8, 8)
fmt.Printf("%3d: %08b\n", 16, 16)
fmt.Printf("%3d: %08b\n", 32, 32)
fmt.Printf("%3d: %08b\n", 64, 64)
fmt.Printf("%3d: %08b\n", 128, 128)
```

```
  0: 00000000
  1: 00000001
  2: 00000010
  3: 00000011
  4: 00000100
  5: 00000101
  6: 00000110
  7: 00000111
  8: 00001000
 16: 00010000
 32: 00100000
 64: 01000000
128: 10000000
```

Bitwise operators

We've seen operators like `+`, `-`, `*`, and `/` that allow you to do math operations on entire numbers. But Go also has **bitwise operators**, which allow you to manipulate the individual bits a number is composed of. Two of the most common ones are the `&` bitwise AND operator, and the `|` bitwise OR operator.

Operator	Name
`&`	Bitwise AND
`|`	Bitwise OR

The bitwise AND operator

We've seen the `&&` operator. It's a Boolean operator that gives a `true` value only if both the values to its left *and* its right are `true`:

```
fmt.Printf("false && false == %t\n", false && false)
fmt.Printf("true  && false == %t\n", true  && false)
fmt.Printf("true  && true  == %t\n", true  && true)
```

```
false && false == false
true  && false == false
true  && true  == true
```

The `&` operator (with just one ampersand), however, is a *bitwise* operator. It sets a bit to `1` only if the corresponding bit in the value to its left *and* the bit in the value to its right are both `1`. For the numbers `0` and `1`, which require only one bit to represent, this is fairly straightforward:

```
fmt.Printf("%b & %b == %b\n", 0, 0, 0&0)
fmt.Printf("%b & %b == %b\n", 0, 1, 0&1)
fmt.Printf("%b & %b == %b\n", 1, 1, 1&1)
```

```
0 & 0 == 0
0 & 1 == 0
1 & 1 == 1
```

Neither bit is 1.
Only one bit is 1.
Both bits are 1.

For larger numbers, however, it can seem like nonsense!

```
fmt.Println(170 & 15)
fmt.Println( 10 &  7)
fmt.Println(100 & 45)
```

```
10
2
36
```

What do THESE results mean?

It's only when you look at the values of individual bits that bitwise operations make sense. The `&` operator only sets a bit to `1` in the result if the bit in the same place in the left number *and* the bit in the same place in the right number are both `1`.

```
fmt.Printf("%02b\n", 1)
fmt.Printf("%02b\n", 3)
fmt.Printf("%02b\n", 1&3)
```

Second bit is 0. First bit is 1. → 01
Second bit is 1. First bit is 1. → 11
Second bit in result is 0. First bit in result is 1. → 01

```
fmt.Printf("%02b\n", 2)
fmt.Printf("%02b\n", 3)
fmt.Printf("%02b\n", 2&3)
```

Second bit is 1. First bit is 0. → 10
Second bit is 1. First bit is 1. → 11
Second bit in result is 1. First bit in result is 0. → 10

This is true for numbers of any size. The bits of the two values the `&` operator is used on determine the bits at the same places in the resulting value.

```
fmt.Printf("%08b\n", 170)
fmt.Printf("%08b\n", 15)
fmt.Printf("%08b\n", 170&15)
```

```
10101010
00001111
00001010
```

If the bit at a given place in the first number is a 1...
...and the bit at the same place in the second number is a 1...
...then the bit at the same place in the result will be a 1.

The bitwise OR operator

We've also seen the `||` operator. It's a Boolean operator that gives a `true` value if the value to its left *or* the value to its right is `true`.

```
fmt.Printf("false || false == %t\n", false || false)
fmt.Printf("true  || false == %t\n", true  || false)
fmt.Printf("true  || true  == %t\n", true  || true)
```

```
false || false == false
true  || false == true
true  || true  == true
```

The `|` operator sets a bit to `1` in the result if the corresponding bit in the value to its left *or* the bit in the value to its right has a value of `1`.

```
fmt.Printf("%b | %b == %b\n", 0, 0, 0|0)
fmt.Printf("%b | %b == %b\n", 0, 1, 0|1)
fmt.Printf("%b | %b == %b\n", 1, 1, 1|1)
```

```
0 | 0 == 0
0 | 1 == 1
1 | 1 == 1
```

Neither bit is 1.
Only one bit is 1.
Both bits are 1.

Just as with bitwise AND, the bitwise OR operator looks at the bits at a given position in the two values it's operating on to decide the value of the bit at the same position in the result.

```
fmt.Printf("%02b\n", 1)
fmt.Printf("%02b\n", 0)
fmt.Printf("%02b\n", 1|0)
```

Second bit is 0.
Second bit is 0.
Second bit in result is 0.

```
01
00
01
```

First bit is 1.
First bit is 0.
First bit in result is 1.

```
fmt.Printf("%02b\n", 2)
fmt.Printf("%02b\n", 0)
fmt.Printf("%02b\n", 2|0)
```

Second bit is 1.
Second bit is 0.
Second bit in result is 1.

```
10
00
10
```

First bit is 0.
First bit is 0.
First bit in result is 0.

This is true for numbers of any size. The bits of the two values the `|` operator is used on determine the bits at the same places in the resulting value.

```
fmt.Printf("%08b\n", 170)
fmt.Printf("%08b\n", 15)
fmt.Printf("%08b\n", 170|15)
```

```
10101010
00001111
10101111
```

If the bit at a given place in the first number is a 1...
...or the bit at the same place in the second number is a 1...
...then the bit at the same place in the result will be a 1.

Using bitwise OR on the "os" package constants



Well, that was certainly...nerdy. I don't see how any of this is going to help me use the `os.O_APPEND` and `os.O_CREATE` constants!

We showed you all this because you'll need to use the bitwise OR operator to combine the constant values together!

When the documentation says that the `os.O_APPEND` and `os.O_CREATE` values "may be or'ed in" with the `os.O_RDONLY`, `os.O_WRONLY`, or `os.O_RDWR` values, it means that you should use the bitwise OR operator on them.

Behind the scenes, these constants are all just `int` values:

```
fmt.Println(os.O_RDONLY, os.O_WRONLY, os.O_RDWR, os.O_CREATE, os.O_APPEND)
```

```
0 1 2 64 1024
```

If we look at the binary representation of these values, we'll see that just one bit is set to `1` for each, and all the other bits are `0`:

```
fmt.Printf("%016b\n", os.O_RDONLY)
fmt.Printf("%016b\n", os.O_WRONLY)
fmt.Printf("%016b\n", os.O_RDWR)
fmt.Printf("%016b\n", os.O_CREATE)
fmt.Printf("%016b\n", os.O_APPEND)
```

```
0000000000000000
0000000000000001
0000000000000010
0000000001000000
0000010000000000
```

That means we can combine the values with the bitwise OR operator, and none of the bits will interfere with each other:

```
fmt.Printf("%016b\n", os.O_WRONLY|os.O_CREATE)
fmt.Printf("%016b\n", os.O_WRONLY|os.O_CREATE|os.O_APPEND)
```

```
0000000001000001
0000010001000001
```

The `os.OpenFile` function can check whether the first bit is a `1` to determine whether the file should be write-only. If the seventh bit is a `1`, `OpenFile` will know to create the file if it doesn't exist. And if the 11th bit is a `1`, `OpenFile` will append to the file.



Watch it!

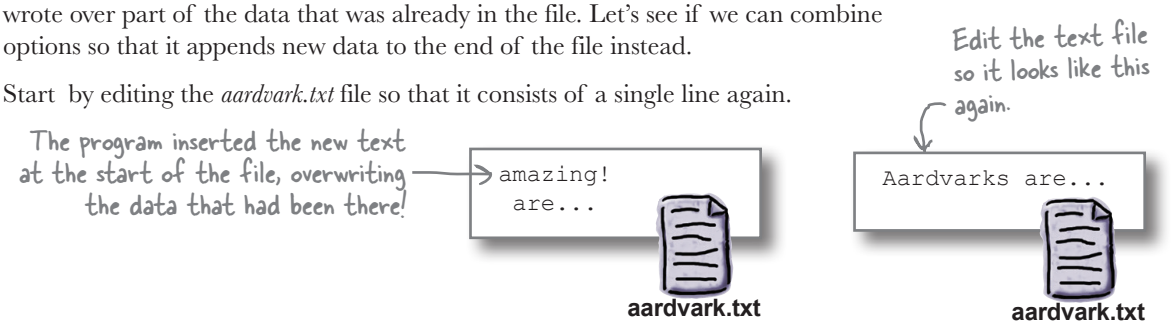
Only use the constant names in your code, never their `int` values!

If you use values like `1` and `1024` in your code in place of the constants, it might work in the short term. But if the Go maintainers ever modified the constants' values, your code would break. Make sure to use the constant names like `os.O_WRONLY` and `os.O_APPEND`, and you'll be safe.

Using bitwise OR to fix our os.OpenFile options

Previously, when we passed only the `os.O_WRONLY` option to `os.OpenFile`, it wrote over part of the data that was already in the file. Let's see if we can combine options so that it appends new data to the end of the file instead.

Start by editing the *aardvark.txt* file so that it consists of a single line again.



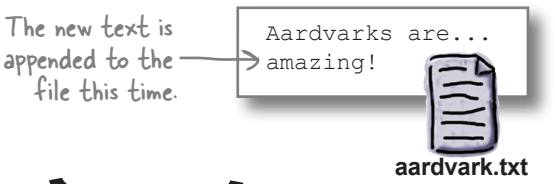
Next, update our program to use the bitwise OR operator to combine the `os.O_WRONLY` and `os.O_APPEND` constant values into a single value. Pass the result to `os.OpenFile`.

```
func main() {
        options := os.O_WRONLY | os.O_APPEND
        file, err := os.OpenFile("aardvark.txt", options, os.FileMode(0600))
        check(err)
        _, err = file.Write([]byte("amazing!\n"))
        check(err)
        err = file.Close()
        check(err)
}
```

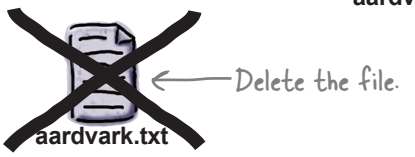
Use bitwise OR to combine the two values.

Pass the result to os.OpenFile.

Run the program again and take another look at the file's contents. You should see the new line of text appended at the end.



Let's also try using the `os.O_CREATE` option, which causes `os.OpenFile` to create the specified file if it doesn't exist. Start by deleting the *aardvark.txt* file.

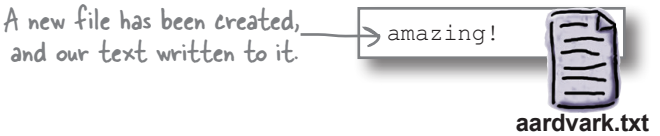


Now update the program to add `os.O_CREATE` to the options being passed to `os.OpenFile`.

```
        options := os.O_WRONLY | os.O_APPEND | os.O_CREATE
        file, err := os.OpenFile("aardvark.txt", options, os.FileMode(0600))
        // ...
```

Use bitwise OR to add the os.O_CREATE value.

When we run the program, it will create a new *aardvark.txt* file and then write the data to it.



Unix-style file permissions

We've been focusing on the second argument to `os.OpenFile`, which controls reading, writing, creating, and appending files. Up until now, we've been ignoring the third argument, which controls the file's *permissions*: which users will be permitted to read from and write to the file after your program creates it.

This argument controls "permissions" on new files.

```
file, err := os.OpenFile("aardvark.txt", options, os.FileMode(0600))
```

```
File Edit Window Help
$ go doc os OpenFile
func OpenFile(name string, flag int, perm FileMode) (*File, error)
    OpenFile is the generalized open call; most users will use Open or Create
    instead. It opens the named file with specified flag (O_RDONLY etc.) and
    ...
```

When developers talk about file permissions, they usually mean permissions as they're implemented on Unix-like systems like macOS and Linux. Under Unix, there are three major permissions a user can have on a file:

Abbreviation	Permission
`r`	The user can **r**ead the file's contents.
`w`	The user can **w**rite the file's contents.
`x`	The user can e**x**ecute the file. (This is only appropriate for files that contain program code.)

If a user doesn't have read permissions on a file, for example, any program they run that tries to access the file's contents will get an error from the operating system:

```
File Edit Window Help
$ cat locked.txt
cat: locked.txt: Permission denied
```

If a user doesn't have execute permissions on a file, they won't be able to execute any code it contains. (Files that don't contain executable code should *not* be marked executable, because attempting to run them could produce unpredictable results.)

```
File Edit Window Help
$ ./hello
-bash: ./hello: Permission denied
```



Watch it!

The permissions argument is ignored on Windows.

Windows doesn't treat file permissions in the same way as Unix-like systems, so files will be created with default permissions on Windows no matter what you do. But that same program will not ignore the permissions argument when it runs on Unix-like machines. It's important to be familiar with how permissions work, and if possible, to test your program on the various operating systems you want it to run on.

Representing permissions with the os.FileMode type

Go's `os` package uses the `FileMode` type to represent file permissions. If a file doesn't already exist, the `FileMode` you pass to `os.OpenFile` determines what permissions the file will be created with, and therefore what kinds of access users will have to it.

`FileMode` values have a `String` method, so if you pass a `FileMode` to functions in the `fmt` package like `fmt.Println`, you'll get a special string representation of the value. That string shows the permissions the `FileMode` represents, in a format similar to the one you might see in the Unix `ls` utility.




```
fmt.Println(os.FileMode(0700))
```

Each file has three sets of permissions, affecting three different classes of users. The first set of permissions applies only to the user that owns the file. (By default, your user account is the owner of any files you create.) The second set of permissions is for the group of users that the file is assigned to. And the third set applies to other users on the system that are neither the file owner nor part of the file's assigned group.

(Look for "Unix file permissions" in a search engine if you'd like more info.)

```
fmt.Println(os.FileMode(0700))
fmt.Println(os.FileMode(0070))
fmt.Println(os.FileMode(0007))
```

```
-rwx------
----rwx---
-------rwx
```

The file's owner will have full permissions.

Users in the file's group will have full permissions.

All other users on the system will have full permissions.

`FileMode` has an underlying type of `uint32`, which stands for "32-bit unsigned integer." It's a basic type that we haven't talked about previously. Because it's unsigned, it can't hold any negative numbers, but it can hold larger numbers within its 32 bits of memory than it would otherwise be able to.

Because `FileMode` is based on `uint32`, you can use a type conversion to convert (almost) any non-negative integer to a `FileMode` value. The results may be a little hard to understand, though:

```
fmt.Println(os.FileMode(17))
fmt.Println(os.FileMode(249))
fmt.Println(os.FileMode(1000))
```

```
-----w---x
--wxrwx--x
-rwxr-x---
```

Garbled permissions that give too much access in some areas, not enough in others.

Octal notation

Instead, it's easier to specify integers for conversion to `FileMode` values using **octal notation**. We've seen decimal notation, which uses 10 digits: 0 through 9. We've seen binary notation, which uses just two digits: 0 and 1. Octal notation uses eight digits: 0 through 7.

You can view the octal representation of various numbers using `fmt.Printf` with the `%0` formatting verb:

```
for i := 0; i <= 19; i++ {
        fmt.Printf("%3d: %04o\n", i, i)
}
```

Print the number in decimal notation.

Print the number in octal notation.

```
  0: 0000
  1: 0001
  2: 0002
  3: 0003
  4: 0004
  5: 0005
  6: 0006
  7: 0007
  8: 0010
  9: 0011
 10: 0012
 11: 0013
 12: 0014
 13: 0015
 14: 0016
 15: 0017
 16: 0020
 17: 0021
 18: 0022
 19: 0023
```

Octal goes up to 7 in the first position...

...then the first position resets to 0, and the second position is incremented to 1.

Up to 7 in the first position again...

...then the first position resets to 0, and the second position is incremented to 2. And so on.

Unlike with binary notation, Go lets you write numbers using octal notation in your program code. Any series of digits preceded by a `0` will be treated as an octal number.

This can be confusing if you're not prepared for it. Decimal `10` is not at all the same as octal `010`, and decimal `100` isn't at all like octal `0100`!

```
fmt.Printf("Decimal   1: %3d Octal   01: %2d\n",   1,   01)
fmt.Printf("Decimal  10: %3d Octal  010: %2d\n",  10,  010)
fmt.Printf("Decimal 100: %3d Octal 0100: %2d\n", 100, 0100)
```

```
Decimal   1:   1 Octal   01:  1
Decimal  10:  10 Octal  010:  8
Decimal 100: 100 Octal 0100: 64
```

Only the digits 0 through 7 are valid in octal numbers. If you include an 8 or a 9, you'll get a compile error.

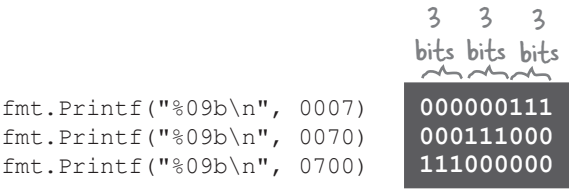
```
fmt.Println(089)
```

```
illegal octal number
```

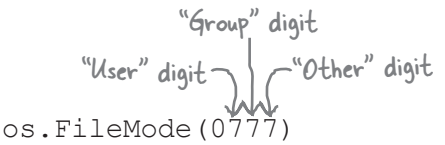
Compile error

Converting octal values to FileMode values

So why use this (arguably strange) octal notation for file permissions? Because each digit of an octal number can be represented using just 3 bits of memory:



Three bits is also the exact amount of data needed to store the permissions for one user class ("user," "group," or "other"). Any combination of permissions you need for a user class can be represented using one octal digit!



Notice the similarity between the binary representation of the octal numbers below and the `FileMode` conversion for the same number. If a bit in the binary representation is `1`, then the corresponding permission is enabled.

Print the binary representation of an octal number.

Print the string for the FileMode conversion of the same number.

If a bit is 1, then the corresponding permission is enabled.

```
fmt.Printf("%09b %s\n", 0000, os.FileMode(0000))
fmt.Printf("%09b %s\n", 0111, os.FileMode(0111))
fmt.Printf("%09b %s\n", 0222, os.FileMode(0222))
fmt.Printf("%09b %s\n", 0333, os.FileMode(0333))
fmt.Printf("%09b %s\n", 0444, os.FileMode(0444))
fmt.Printf("%09b %s\n", 0555, os.FileMode(0555))
fmt.Printf("%09b %s\n", 0666, os.FileMode(0666))
fmt.Printf("%09b %s\n", 0777, os.FileMode(0777))
```

```
000000000 ----------
001001001 ---x--x--x
010010010 --w--w--w-
011011011 --wx-wx-wx
100100100 -r--r--r--
101101101 -r-xr-xr-x
110110110 -rw-rw-rw-
111111111 -rwxrwxrwx
```

For this reason, the Unix `chmod` utility (short for "change mode") has used octal digits to set file permissions for decades now.

```
File Edit Window Help
$ chmod 0000 allow_nothing.txt
$ chmod 0100 execute_only.sh
$ chmod 0200 write_only.txt
$ chmod 0300 execute_write.sh
$ chmod 0400 read_only.txt
$ chmod 0500 read_execute.sh
$ chmod 0600 read_write.txt
$ chmod 0700 read_write_execute.sh
$ chmod 0124 user_execute_group_write_other_read.sh
$ chmod 0777 all_read_write_execute.sh
```

Go's support for octal notation allows you to follow the same convention in your code!

Octal digit	Permission
0	no permissions
1	execute
2	write
3	write, execute
4	read
5	read, execute
6	read, write
7	read, write, execute

Calls to os.OpenFile, explained

Now that we understand both bitwise operators and octal notation, we can finally understand just what calls to `os.OpenFile` do!

This code, for example, will append new data to an existing logfile. The user that owns the file will be able to read from and write to the file. All other users will only be able to read from it.

```
options := os.O_WRONLY | os.O_APPEND
file, err := os.OpenFile("log.txt", options, os.FileMode(0644))
```

Open the file for writing, appending new data to the end.

File will be readable and writable by its owner, and just readable by everyone else.

And this code will create a file if it doesn't exist, then append data to it. The resulting file will be readable and writable by its owner, but no other user will have access to it.

```
options := os.O_WRONLY | os.O_APPEND | os.O_CREATE
file, err := os.OpenFile("log.txt", options, os.FileMode(0600))
```

If the file doesn't exist, create it. Open the file for writing, appending new data to the end.

File will be readable and writable by its owner. No one else will have access.



Watch it!

If the os.Open or os.Create functions will do what you need, use those instead.

The `os.Open` *function can only open files for reading. But if that's all you need, you may find it simpler to use than* `os.OpenFile`*. Likewise, the* `os.Create` *function can only create files that are readable and writable by any user. But if that's all you need, you should consider using it instead of* `os.OpenFile`*. Sometimes less powerful functions result in more readable code.*

there are no Dumb Questions

Q: Octal notation and bitwise operators are a pain! Why is it done this way?

A: To save computer memory! These conventions for handling files have their roots in Unix, which was developed when RAM and disk space were both smaller and more expensive. But even now, when a hard disk can contain millions of files, packing file permissions into a few bits instead of several bytes can save a lot of space (and make your system run faster). Trust us, the effort is worth it!

Q: What's that extra dash at the front of a `FileMode` string?

A: A dash in that position indicates that a file is just an ordinary file, but it can show several other values. For example, if the `FileMode` value represents a directory, it will be a `d` instead.

Get stats on a file or directory. You can look this up in the docs!

Print the FileMode info for the directory.

```
fileInfo, err := os.Stat("my_directory")
if err != nil {
      log.Fatal(err)
}
fmt.Println(fileInfo.Mode())
```

```
drwxr-xr-x
```

Appendix B: Leftovers



We've covered a lot of ground, and you're almost finished with this book. We'll miss you, but before we let you go, we wouldn't feel right about sending you out into the world without a *little* more preparation. We've saved six important topics for this appendix.

#1 Initialization statements for "if"

Here we have a `saveString` function that returns a single `error` value (or `nil` if there was no error). In our `main` function, we might store that return value in an `err` variable before handling it:

```
func saveString(fileName string, str string) error {
        err := ioutil.WriteFile(fileName, []byte(str), 0600)
        return err
}
```

(You can learn more about the WriteFile function with "go doc io/ioutil WriteFile".)

Call saveString and store the return value.

Report any error.

```
func main() {
        err := saveString("hindi.txt", "Namaste")
        if err != nil {
                log.Fatal(err)
        }
}
```

Now suppose we added another call to `saveString` in `main` that also uses an `err` variable. We have to remember to make the first use of `err` a short variable declaration, and change later uses to assignments. Otherwise, we'll get a compile error for attempting to redeclare a variable.

This code also uses a variable named "err".

If we forget to convert the original code from a short declaration to an assignment...

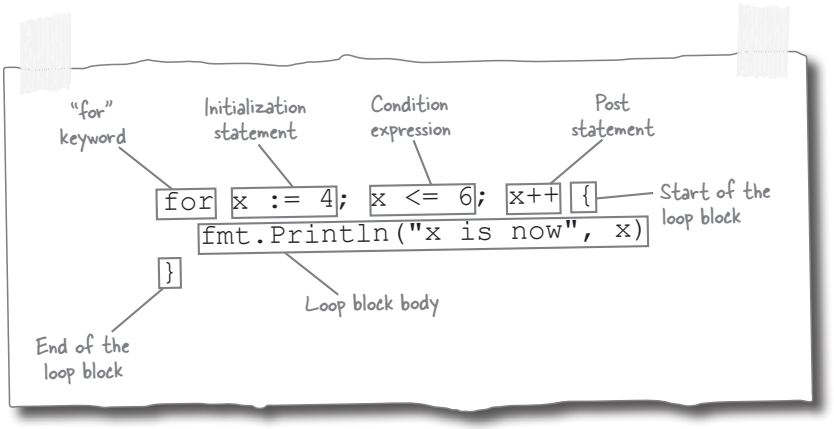
```
func main() {
        err := saveString("english.txt", "Hello")
        if err != nil {
                log.Fatal(err)
        }
        err := saveString("hindi.txt", "Namaste")
        if err != nil {
                log.Fatal(err)
        }
}
```

Compile error!

```
no new variables on left side of :=
```

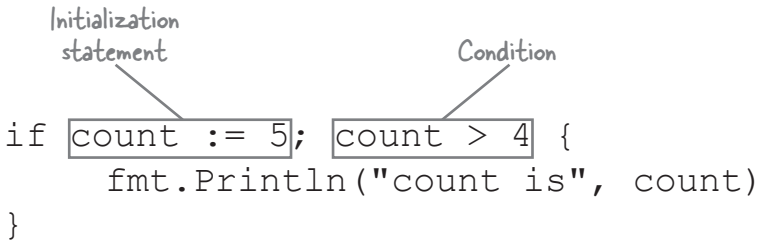
But really, we're only using the `err` variable within the `if` statement and its block. What if there was a way to limit the scope of the variable, so that we could treat each occurrence as a separate variable?

Remember when we first covered `for` loops, back in Chapter 2? We said they could include an initialization statement, where you initialize variables. Those variables were only in scope within the `for` loop's block.

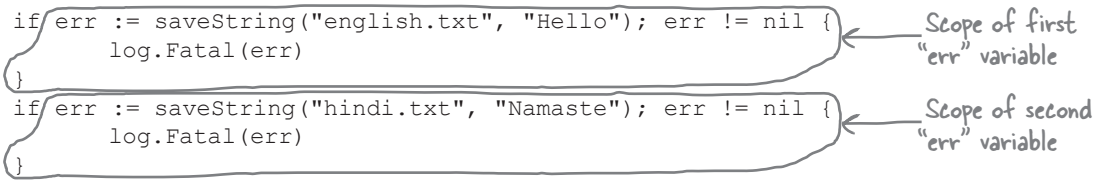


#1 Initialization statements for "if" (continued)

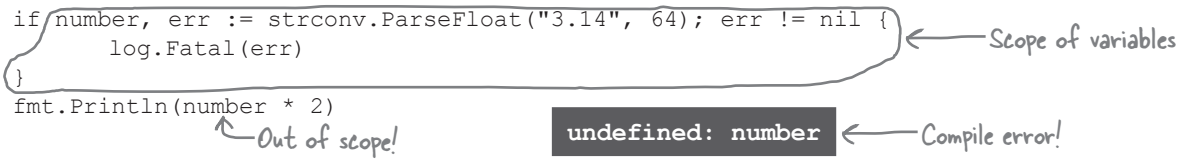
Similar to `for` loops, Go allows you to add an initialization statement before the condition in `if` statements. Initialization statements are usually used to initialize one or more variables for use within the `if` block.



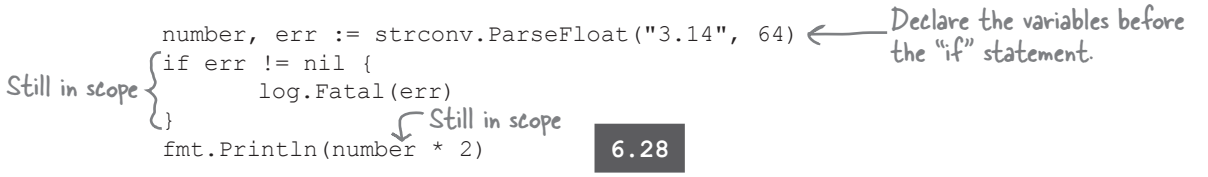
The scope of variables declared within an initialization statement is limited to that `if` statement's conditional expression and its block. If we rewrite our previous sample to use `if` initialization statements, the scope of each `err` variable will be limited to the `if` statement conditional and block, meaning we'll have two completely separate `err` variables. We won't have to worry about which one is defined first.



This limitation on scope cuts both ways. If a function has multiple return values, and you need one of them *inside* the `if` statement and one *outside*, you probably won't be able to call it in an `if` initialization statement. If you try, you'll find the value you need outside the `if` block is out of scope.



Instead, you'll need to call the function prior to the `if` statement, as normal, so that its return values are in scope both inside *and* outside the `if` statement:



#2 The switch statement

When you need to take one of several actions based on the value of an expression, it can lead to a mess of `if` statements and `else` clauses. The `switch` statement is a more efficient way to express these choices.

You write the `switch` keyword, followed by a condition expression. Then you add several `case` expressions, each with a possible value the condition expression could have. The first `case` whose value matches the condition expression is selected, and the code it contains is run. The other `case` expressions are ignored. You can also provide a `default` statement which will be run if no `case` matches.

Here's a reimplementation of a code sample that we wrote with `if` and `else` statements in Chapter 12. This version requires significantly less code. For our `switch` condition, we select a random number from 1 to 3. We provide `case` expressions for each of those values, each of which prints a different message. To alert us to the theoretically impossible situation where no `case` matches, we also provide a `default` statement that panics.

```
import (
        "fmt"
        "math/rand"
        "time"
)

func awardPrize() {
        switch rand.Intn(3) + 1 {
        case 1:
                fmt.Println("You win a cruise!")
        case 2:
                fmt.Println("You win a car!")
        case 3:
                fmt.Println("You win a goat!")
        default:
                panic("invalid door number")
        }
}

func main() {
        rand.Seed(time.Now().Unix())
        awardPrize()
}
```

The condition expression

If the result is 1...

...then print this message.

If the result is 2...

...then print this message.

If the result is 3...

...then print this message.

If the result is none of the above...

...then panic, because it means something's wrong with our code.

```
You win a goat!
```

there are no Dumb Questions

Q: I've seen other languages where you have to provide a "break" statement at the end of each `case`, or it will run the next `case`'s code as well. Does Go not require this?

A: Developers have a history of forgetting the "break" statement in other languages, resulting in bugs. To help avoid this, Go automatically exits the `switch` at the end of a `case`'s code.

There's a `fallthrough` keyword you can use in a case, if you *do* want the next case's code to run as well.

#3 More basic types

Go has additional basic types that we haven't had space to talk about. You probably won't have reason to use these in your own projects, but you'll encounter them in some libraries, so it's best to be aware they exist.

Types	Description

`int8`
`int16`
`int32`
`int64`
	These hold integers, just like `int`, but they're a specific size in memory (the number in the type name specifies the size in bits). Fewer bits consume less RAM or other storage; more bits mean larger numbers can be stored. You should use `int` unless you have a specific reason to use one of these; it's more efficient.

`uint`	This is just like `int`, but it holds only *unsigned* integers; it can't hold negative numbers. This means you can fit larger numbers into the same amount of memory, as long as you're certain the values will never be negative.

`uint8`
`uint16`
`uint32`
`uint64`
	These also hold unsigned integers, but like the `int` variants, they consume a specific number of bits in memory.

`float32`	The `float64` type holds floating-point numbers and consumes 64 bits of memory. This is its smaller 32-bit cousin. (There are no 8-bit or 16-bit variants for floating-point numbers.)

#4 More about runes

We introduced runes very briefly back in Chapter 1, and we haven't talked about them since. But we don't want to end the book without going into a little more detail about them...

Back in the days before modern operating systems, most computing was done using the unaccented English alphabet, with its 26 letters (in upper- and lowercase). There were so few of them, a character could be represented by a single byte (with 1 bit to spare). A standard called ASCII was used to ensure the same byte value was converted to the same letter on different systems.

But of course, the English alphabet isn't the only writing system in the world; there are many others, some with thousands of different characters. The Unicode standard is an attempt to create one set of *4-byte* values that can represent every character in every one of these different writing systems (and many other characters besides).

Go uses values of the `rune` type to represent Unicode values. Usually, one rune represents one character. (There are exceptions, but those are beyond the scope of this book.)

#4 More about runes (continued)

Go uses UTF-8, a standard that represents Unicode characters using 1 to 4 bytes each. Characters from the old ASCII set can still be represented using a single byte; other characters may require anywhere from 2 to 4 bytes.

Here are two strings, one with letters from the English alphabet, and one with letters from the Russian alphabet.

```
asciiString := "ABCDE"
utf8String := "БГДЖИ"
```

These characters are all from the ASCII character set, so they take up 1 byte each.

These Unicode characters take up 2 bytes each.

Generally, you don't need to worry about the details of how characters are stored. That is, *until* you try to convert strings to their component bytes and back. If we try to call the `len` function with our two strings, for example, we get very different results:

```
fmt.Println(len(asciiString))
fmt.Println(len(utf8String))
```

```
5
10
```

This string takes up 5 bytes.

This string takes up 10 bytes.

When you pass a string to the `len` function, it returns the length in *bytes*, not *runes*. The English alphabet string fits into 5 bytes—each rune requires just 1 byte because it's from the old ASCII character set. But the Russian alphabet string takes 10 bytes—each rune requires 2 bytes to store.

If you want the length of a string in *characters*, you should instead use the `unicode/utf8` package's `RuneCountInString` function. This function will return the correct number of characters, regardless of the number of bytes used to store each one.

```
fmt.Println(utf8.RuneCountInString(asciiString))
fmt.Println(utf8.RuneCountInString(utf8String))
```

```
5
5
```

This string has five runes.

This string also holds five runes.

Working with partial strings safely means converting the string to runes, not bytes.

#4 More about runes (continued)

Previously in the book, we've had to convert strings to slices of bytes so we could write them to an HTTP response or to the terminal. This works fine, as long as you make sure to write *all* the bytes in the resulting slice. But if you try to work with just *part* of the bytes, you're asking for trouble.

Here's some code that attempts to strip the first three characters from the previous strings. We convert each string to a slice of bytes, then use the slice operator to gather everything from the fourth element to the end of the slice. Then we convert the partial byte slices back to strings and print them.

```
asciiBytes := []byte(asciiString)
utf8Bytes := []byte(utf8String)
asciiBytesPartial := asciiBytes[3:]
utf8BytesPartial := utf8Bytes[3:]
fmt.Println(string(asciiBytesPartial))
fmt.Println(string(utf8BytesPartial))
```

Convert the strings to slices of bytes.

Omit the first 3 bytes from each slice.

First 3 bytes removed, which removes the first three characters.

DE
□ДЖИ

First 3 bytes removed, which removes the first rune, and 1 byte from the second rune!

This works fine with the English alphabet characters, which each take up 1 byte. But the Russian characters each take *2 bytes*. Cutting off the first 3 bytes of that string omits only the first character, and "half" of the second, resulting in an unprintable character.

Go supports converting from strings to slices of `rune` values, and from slices of runes back to strings. To work with partial strings, you should convert them to a slice of `rune` values rather than a slice of `byte` values. That way, you won't accidentally grab just part of the bytes for a rune.

Here's an update to the previous code that converts the strings to slices of runes instead of slices of bytes. Our slice operators now omit the first three *runes* from each slice, rather than the first *3 bytes*. When we convert the partial slices to strings and print them, we get only the last two (complete) characters from each.

```
asciiRunes := []rune(asciiString)
utf8Runes := []rune(utf8String)
asciiRunesPartial := asciiRunes[3:]
utf8RunesPartial := utf8Runes[3:]
fmt.Println(string(asciiRunesPartial))
fmt.Println(string(utf8RunesPartial))
```

Convert the strings to slices of runes.

Omit the first three runes from each slice.

Convert this slice of runes to a string.

DE
ЖИ

First three runes removed

First three runes removed

#4 More about runes (continued)

You'll encounter similar problems if you try to use a slice of bytes to process each character of a string. Processing 1 byte at a time will work as long as your strings are all characters from the ASCII set. But as soon as a character comes along that requires 2 or more bytes, you'll find yourself working with just part of the bytes for a rune again.

This code uses a `for ... range` loop to print the English alphabet characters, 1 byte per character. Then it tries to do the same with the Russian alphabet characters, 1 byte per character—which fails because each of these characters requires 2 bytes.

```
for index, currentByte := range asciiBytes {
      fmt.Printf("%d: %s\n", index, string(currentByte))
}
for index, currentByte := range utf8Bytes {
      fmt.Printf("%d: %s\n", index, string(currentByte))
}
```

Process each byte in the slice.

Convert the byte to a string and print it.

Process each byte in the slice.

Convert the byte to a string and print it.

Results in printable characters for the ASCII characters...

```
0: A
1: B
2: C
3: D
4: E
0: Đ
1:
2: Đ
3:
4: Đ
5:
6: Đ
7:
8: Đ
9:
```

...but unprintable characters for the Unicode characters!

Go allows you to use a `for...range` loop on a string, which will process a *rune* at a time, not a *byte* at a time. This is a much safer approach. The first variable you provide will be assigned the current byte index (not the rune index) within the string. The second variable will be assigned the current rune.

Here's an update to the above code that uses a `for...range` loop to process the strings themselves, not their byte representations. You can see from the indexes in the output that 1 byte at a time is being processed for the English characters, but *2 bytes* at a time are being processed for the Russian characters.

```
for position, currentRune := range asciiString {
      fmt.Printf("%d: %s\n", position, string(currentRune))
}
for position, currentRune := range utf8String {
      fmt.Printf("%d: %s\n", position, string(currentRune))
}
```

Process each rune in the string.

Convert the rune to a string and print it.

Process each rune in the string.

Convert the rune to a string and print it.

All characters are printable.

```
0: A
1: B
2: C
3: D
4: E
0: Б
2: Г
4: Д
6: Ж
8: И
```

Go's runes make it easy to work with partial strings and not have to worry about whether they contain Unicode characters or not. Just remember, anytime you want to work with just part of a string, convert it to runes, not bytes!

#5 Buffered channels

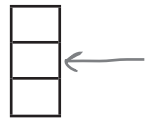
There are two kinds of Go channels: *unbuffered* and *buffered*.

All the channels we've shown you so far have been unbuffered. When a goroutine sends a value on an unbuffered channel, it immediately blocks until another goroutine receives the value. Buffered channels, on the other hand, can hold a certain number of values before causing the sending goroutine to block. Under the right circumstances, this can improve a program's performance.

When creating a channel, you can make a buffered channel by passing a second argument to `make` with the number of values the channel should be able to hold in its buffer.

```
channel := make(chan string, 3)
```

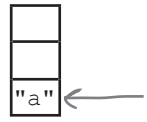
This argument specifies the size of the channel's buffer.



A buffered channel that can hold three values.

When a goroutine sends a value via the channel, that value is added to the buffer. Instead of blocking, the sending goroutine continues running.

```
channel <- "a"
```

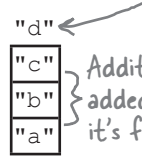


Sent value is added to the buffer.

The sending goroutine can continue sending values on the channel until the buffer is full; only then will an additional send operation cause the goroutine to block.

```
channel <- "b"
channel <- "c"
channel <- "d"
```

Sending a value when the buffer is full causes the sending goroutine to block.

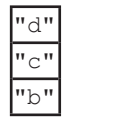


Additional sent values are added to the buffer until it's full.

When another goroutine receives a value from the channel, it pulls the earliest-added value from the buffer.

```
fmt.Println(<-channel)
```

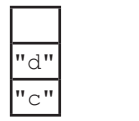
"a"



Additional receive operations will continue to empty the buffer, while additional sends will fill the buffer back up.

```
fmt.Println(<-channel)
```

"b"



#5 Buffered channels (continued)

Let's try running a program with an unbuffered channel, and then change it to a buffered channel so you can see the difference. Below, we define a `sendLetters` function to run as a goroutine. It sends four values to a channel, sleeping 1 second before each value. In `main`, we create an unbuffered channel and pass it to `sendLetters`. Then we put the `main` goroutine to sleep for 5 seconds.

```
func sendLetters(channel chan string) {
    time.Sleep(1 * time.Second)
    channel <- "a"
    time.Sleep(1 * time.Second)
    channel <- "b"
    time.Sleep(1 * time.Second)
    channel <- "c"
    time.Sleep(1 * time.Second)
    channel <- "d"
}

func main() {
    fmt.Println(time.Now())
    channel := make(chan string)
    go sendLetters(channel)
    time.Sleep(5 * time.Second)
    fmt.Println(<-channel, time.Now())
    fmt.Println(<-channel, time.Now())
    fmt.Println(<-channel, time.Now())
    fmt.Println(<-channel, time.Now())
    fmt.Println(time.Now())
}
```

Accepts a channel as a parameter

Send four values, sleeping 1 second before each.

Print the time the program started.

Make an unbuffered channel, just like we've been doing all along.

Launch sendLetters in a new goroutine.

Make the main goroutine sleep 5 seconds.

Receive and print four values along with the current time.

Print the time the program ends.

Here's when the program started.

The first value is already waiting to be received when the main goroutine wakes up.

But the sendLetters goroutine was blocked until the first value was received, so now we have to wait while later values are sent.

```
2018-07-21 11:36:20.676155577 -0700 MST m=+0.000255509
a 2018-07-21 11:36:25.677846276 -0700 MST m=+5.001810208
b 2018-07-21 11:36:26.677931968 -0700 MST m=+5.001895900
c 2018-07-21 11:36:27.679233609 -0700 MST m=+6.003129541
d 2018-07-21 11:36:28.680125059 -0700 MST m=+7.004020991
2018-07-21 11:36:28.680236070 -0700 MST m=+7.004132001
```

The program takes 8 seconds to complete.

When the `main` goroutine wakes up, it receives four values from the channel. But the `sendLetters` goroutine was blocked, waiting for `main` to receive the first value. So the `main` goroutine has to wait 1 second between each remaining value while the `sendLetters` goroutine catches back up.

#5 Buffered channels (continued)

We can speed our program up a bit simply by adding a single-value buffer to the channel.

All we have to do is add a second argument when calling `make`. Interactions with the channel are otherwise identical, so we don't have to make any other changes to the code.

Now, when `sendLetters` sends its first value to the channel, it doesn't block until the `main` goroutine receives it. The sent value goes in the channel's buffer instead. It's only when the second value is sent (and none have yet been received) that the channel's buffer is filled and the `sendLetters` goroutine blocks. Adding a one-value buffer to the channel shaves 1 second off the program's run time.

```
func main() {
        channel := make(chan string, 1)
        // Remaining code unchanged
}
```

Make a buffered channel that can hold one value before blocking.

The first item sent goes on the buffered channel's queue.

After that the queue is full, so the next send causes the sendLetters goroutine to block.

```
2018-07-21 15:29:10.709656836 -0700 MST m=+0.000318261
a 2018-07-21 15:29:15.710058943 -0700 MST m=+5.000584368
b 2018-07-21 15:29:15.710105511 -0700 MST m=+5.000630936
c 2018-07-21 15:29:16.712044927 -0700 MST m=+6.002502352
d 2018-07-21 15:29:17.716495 -0700 MST m=+7.006883143
2018-07-21 15:29:17.716615312 -0700 MST m=+7.007004737
```

The program takes only 7 seconds to complete.

Increasing the buffer size to 3 allows the `sendLetters` goroutine to send three values without blocking. It blocks on the final send, but this is after all of its 1-second `Sleep` calls have completed. So when the `main` goroutine wakes up after 5 seconds, it immediately receives the three values waiting in the buffered channel, as well as the value that caused `sendLetters` to block.

```
channel := make(chan string, 3)
```

Make a buffered channel that can hold three values before blocking.

These three values were waiting in the channel buffer.

This value caused the sendLetters goroutine to block, but only after it was done sleeping.

```
2018-07-21 17:02:20.062202682 -0700 MST m=+0.000341112
a 2018-07-21 17:02:25.066350665 -0700 MST m=+5.004353095
b 2018-07-21 17:02:25.066574585 -0700 MST m=+5.004577015
c 2018-07-21 17:02:25.066583453 -0700 MST m=+5.004585883
d 2018-07-21 17:02:25.066588589 -0700 MST m=+5.004591019
2018-07-21 17:02:25.066593481 -0700 MST m=+5.004595911
```

The program takes only 5 seconds to complete.

This allows the program to complete in only 5 seconds!

#6 Further reading

This is the end of the book. But it's just the beginning of your journey as a Go programmer. We want to recommend a few resources that will help you along the road.

The Head First Go Website
https://headfirstgo.com/

The official website for this book. Here you can download all our code samples, practice with additional exercises, and learn about new topics, all written in the same easy-to-read, incredibly witty prose!

A Tour of Go
https://tour.golang.org

This is an interactive tutorial on Go's basic features. It covers much the same material as this book, but includes some additional details. Examples in the Tour can be edited and run right from your browser (just like in the Go Playground).

Effective Go
https://golang.org/doc/effective_go.html

A guide maintained by the Go team on how to write idiomatic Go code (that is, code that follows community conventions).

The Go Blog
https://blog.golang.org

The official Go blog. Offers helpful articles on using Go and announcements of new Go versions and features.

Package Documentation
https://golang.org/pkg/

Documentation on all the standard packages. These are the same docs available through the `go doc` command, but all the libraries are in one convenient list for browsing. The `encoding/json`, `image`, and `io/ioutil` packages might be interesting places to start.

The Go Programming Language
https://www.gopl.io/

This book is the only resource on this page that isn't free, but it's worth it. It's well known and widely used.

There are two kinds of technical books out there: tutorial books (like the one you're holding) and reference books (like *The Go Programming Language*). And it's a great reference: it covers all the topics we didn't have room for in this book. If you're going to continue using Go, this is a must-read.

Index

Symbols

++ statement 62
+= assignment operator 62, 102
+ arithmetic operator 13
<= comparison operator 13
<- arrow operator 391
< comparison operator 13
== comparison operator 13, 212
-= assignment operator 62
*= assignment operator 102
= assignment statement 16–17, 51, 54
!= comparison operator 13
:= (short variable declaration). *See* short variable declaration (: =)
>= comparison operator 13
> comparison operator 13
|| logical operator 40, 487
| OR bitwise operator 485, 487
%0 formatting verb 492
& address-of operator 103, 243
& AND bitwise operator 485–486
- arithmetic operator 13
/ arithmetic operator 13
* (asterisk). *See* asterisk (*)
\ (backslash) 11
/* */ (block comments) 34
// (comments) 34, 141–142
{} (curly braces. *See* curly braces {}
- (dash) 417, 494
. (dot operator). *See* dot operator (.)
... (ellipsis) 195
%% formatting verb 82
! logical operator 40
&& logical operator 40, 486
() (parentheses). *See* parentheses ()
% (percent sign) 82–83
; (semicolon) 5
‘ (single quotation mark) 12
[] (square brackets). *See* square brackets []
-- statement 62

A

a tag (HTML) 449
action attribute (form tag) 467
actions
  about 454
  inserting data into templates with 454–455
  inserting struct fields into templates with 457
  making parts of templates optional with 455
  repeating parts of templates with 456
address-of operator (&) 103, 243
ampersand (&) 103
AND bitwise operator (&) 485–486
anonymous struct fields 258–259, 276, 308
append built-in function 44, 184–185, 187–189, 459
Args variable (os) 192–193, 351
arguments, command-line 191–192, 194, 417
arguments to functions
  compile errors 14
  parentheses and 5, 7, 87
  passing pointers 107
  structs and 246
  variadic functions and 195–200
arithmetic operators 13, 269
array literals 153

arrays
about 150–151, 153, 171, 232
accessing elements within loops 155–156
average of numbers in 161
fmt package support 154
for loops and 155, 157
for ... range loops and 157–158
indexes for 151, 155–157
len function and 156, 161
panics and 155
reading text files into 166–171
slices and 176, 180–183, 193
sum of numbers in 159–160
zero values in 152
arrow operator (<-) 391
assignment operators 62, 102
assignment statement (=) 16–17, 51, 54
asterisk (*)
accessing struct fields through pointers 244–245
as arithmetic operator 13
as pointer type prefix 104–105, 243
Atoi function (strconv) 59, 95
automated testing
about 401–405
detailed test failure messages 410
fixing bugs 414–415
fixing panicking code using tests 420–421
getting tests to pass 412–413
running specific sets of tests 417
running tests 407
table-driven testing 418–419
test-driven development 413
test helper functions 411
testing return values 408–409
writing tests 406
average of numbers in arrays 161
averages, calculating using variadic functions 198

B

%b formatting verb 485
backslash (\) 11
binary notation 485, 492
bin workspace subdirectory 117
bitwise operators 485–489, 494
blank identifier (_)
error return value and 37
for ... range loops and 158
testing for map values with 216
block comments (/**/) 34
blocks of code
about 39, 49
nesting 49–51
universe block 341
variable scope 50–51, 90
bool (Boolean value)
about 12–13, 15
conditionals and 39–40
formatting verbs and 82
zero values for 17
boolean operators 40, 487
break keyword 68
buffered channels 503–505
bufio package
about 163
Close method 165
Err method 165
NewReader function 35
Reader type 35
ReadString method 35–39, 47
Scan method 165
Scanner type 165
byte type 380, 431

C

calling functions
about 7–8
deferring 353–355, 357, 366
parameters and 87, 100
recursively 360–363
type errors 14
calling methods 32–33, 275, 329–330
call stacks 366
camel case 21, 86
case-sensitivity 6–8, 21, 86
cd command 160, 191–192

chan keyword 390
channels
  about 379, 390
  buffered 503–505
  deadlock errors 394
  return values and 390–391
  struct types and 390, 398
  synchronizing goroutines with 392–393
  usage example 396–397
Close method (bufio) 165
Close method (io) 380
closing files 165, 352
columns, aligning in tables 83
command-line arguments
  about 191
  go test command 417
  os.Args and 192
  updating programs to use 194
comments (//) 34, 71, 141–142
comparison operators 13, 212, 269
comparisons, math operations and 13, 22–23, 60
compiling Go code 5, 26, 45, 130
concrete types 327, 329–330, 334–338
concurrency using goroutines 379, 383–388
conditional statements
  about 31, 39–40
  blocks of code in 49
  Boolean values and 12
  error handling 39, 42
  if/else statements 39, 46, 69
  if statements 39–40, 50, 496–497
condition expressions in loops 61, 63
constants 126–127
const keyword 126
continue keyword 68
conversions
  avoiding errors with 22–23
  between strings and bytes 431
  between types using functions 271–272
  between types using methods 283–284
  between types with same underlying type 267–268
  strings to numbers 46–48, 59
Create function (os) 494
Ctrl+C keyboard shortcut 62, 428
curly braces {}
  actions and 454
  arrays and 153
  blocks of code and 39, 49
  interfaces and 325
  maps and 214
  slices and 177
  structs and 233, 236, 251

D

%d formatting verb 82
dash (-) 417, 494
date validation 291, 296–298
decimals 13, 82, 84
declaring constants 127
declaring functions 86–87
declaring variables
  about 16–17
  compile errors 20, 22, 36
  re-assignment using short declarations 54
  scope and 50–51
  shadowing names 44–45
  short declarations 19–20, 37, 51, 54
defer keyword 353–355, 357, 370, 373
defined types
  about 236–237, 290
  converting using functions 271–272
  converting using methods 283–284
  naming rules for 241, 249–250
  operators and 269
  underlying basic types 267–268
  using to store data 238
  using with functions 239–240
delete built-in function 218
directories
  listing contents recursively 362–363
  listing files in 358–359
  nesting 118, 128–129
  workspace 117–118, 128
div tag (HTML) 449

documentation
creating for packages 141–142
reading for packages 139–140
viewing in web browsers 143–144
dot operator (.)
about 33, 454
accessing functions 8, 234
accessing struct fields 234–235, 243
chaining together 256–257, 309, 311
dot value for templates 454–457, 461, 476
double quotation marks escape sequence (\") 11
downloading packages 137

E

editor (online) 3
ellipsis (...) 195
embedding
about 289
struct types and 258–259, 276, 308, 310
empty interface 344–346
empty strings 17, 152
encapsulation
about 289, 305
exported fields and 312
struct types and 289
unexported fields and 309
{{end}} action marker 455–456, 461
Err method (bufio) 165
error built-in type 45, 340–341
Errorf function (fmt) 96
Errorf method (testing) 410–411
error handling
about 99
conditionals 39, 42
in general 6
recursive functions 364–368
return values 36–38, 95, 99–100
slices 190
type assertions 336–337, 365
Error method 96, 340, 406–407
errors package 96, 299
escape sequences 11–12, 462
Execute method (html/template) 450
Execute method (text/template) 451–455
execute (x) file permission 490
executing Go code
about 27, 119
halting code execution 38, 62, 428
.exe file extension 130
export command 132
exporting from packages
constants 126–127
defined types 241, 249–250
functions 21, 118
methods 277, 301–303, 310–311
struct types 249–250, 298

F

%f formatting verb 82–83
Fatal function (log) 38, 447
file management
closing files 165, 352
file permissions 490–493
listing files in directories 358–359
opening files 165, 380, 481–494
reading names from text files 207–209
reading numbers from files 350–351
reading slices in from files 458–459
reading text files 163–165
reading text files into arrays 166–171
reading text files using slices and append function 187–189
FileMode type (os) 482, 491–494
File type (os) 165, 458, 470, 494
first-class functions 435–436
flags 417
float64 type (floating-point number)
about 15
conversions with 22–23, 47
decimal points and 13, 84
formatting verbs and 82–83
Floor function (math) 8–9, 14

fmt package
about 116
Errorf function 96
error values and 96
Fprintln function 470
Printf function. *See* Printf function (fmt)
Print function 35, 343
Println function 5–9, 154, 343–344
Sprintf function 81, 96, 154, 410
Stringer interface 342–343
for loops
about 61–63
array elements and 155, 157
maps and 223
scope and 63
Format button 3, 5, 27
formatting output. *See* fmt package
forms. *See* HTML forms
form tag (HTML) 464, 467
for ... range loops 157–158, 221–224
forward slashes (//) 34
Fprintln function (fmt) 470
functions
about 4, 85
arguments to. *See* arguments to functions
blocks of code in 49
calling. *See* calling functions
calling recursively 360–361
case sensitivity 6
converting types using 271–272
declaring 86
deferring calling 353–355, 357, 366
defined types and 239–240
dot operator and 8, 234
first-class 435
fixing name conflicts using methods 274
from other packages 8
handler 435–436, 447, 465, 468
interfaces and 331
methods and 32, 277, 331
modifying structs using 242–243
naming rules for 21, 86, 118, 121
nil return value for 39, 47, 95
overloading 272
parameters in 87, 102, 328
passing to other functions 436
pointers and 107–108, 243
recursive 360–363
return values. *See* return values
scope of 50, 90
shadowing names 44–45
slices and 176
type definitions and 237, 239–240
as types 436–438
unexported 21, 121, 302–303, 305
variadic 195–200, 344

G

Get function (net/http) 380–381, 387
GET method (HTTP) 467, 472
getter methods 304–305
GitHub website 133
go build command 26–27, 130
go directory 117–118
go doc command 139–143, 381
godoc tool 144–145
.go file extension 117, 119
go fmt command 5, 26–27, 40, 234
go get command 137
go install command 130, 160
go keyword 383
golang.org site 143–144
GOPATH environment variable 131–132
Go Playground site 3, 25, 27
Go programming language
about 1–2
additional reading 506
case sensitivity 6
compiling Go code 5, 26, 45, 130
executing code 27, 119
halting code execution 38, 62, 428
installing 25
online editor 3
tool support 27
website for 3, 25
goroutines
about 379, 383–387

channels and 390–393, 503–505
deadlock errors 394
return values and 389–390
runtime control and 388
usage example 396–397
go run command 27, 119
go test command 401, 406–407, 411, 417
go version command 25, 27
Griesemer, Robert 2

H

h1 tag (HTML) 449
halting code execution 38, 62, 428
HandleFunc function (net/http) 430, 432–433, 447
handler functions 435–436, 447, 465, 468
Handler interface (net/http) 430
HTML documentation 143–144
HTML elements 449
HTML forms
form submission requests 466–469
getting field values from requests 468–469
responding with 465
saving data 470–471
users adding data with 464
html/template package
about 445, 451, 462
Execute method 450
ParseFiles function 450
Template type 450
HTML templates
about 445
executing with struct types 457
html/template package 445, 450–451, 462
inserting data into 454–455
inserting struct fields into 457
making parts optional 455
repeating parts of 456
text/template package 451–455, 462
updating 461
HTTP methods 467–468, 472
http package. *See* net/http package
hyphen (-) 417, 494

I

{{if}} action 455, 457
if/else statements 39, 46, 69
if statements 39–40, 50, 496–497
import paths
errors 121
go doc command and 139–141
nested directories and 128–129
package names vs. 56–57, 135–137, 139–141
import statement
about 4, 8, 119–120
errors 6
import paths vs. package names 56–57, 135–137, 139–141
nested directories and import paths 128–129
indexes
for arrays 151, 155–157
for maps 213
for slices 180–181
infinite loops 62–63
initialization statements 61, 63, 496–497
input tag (HTML) 464
installing Go 25
installing packages 137
interface keyword 325
interfaces
about 321, 325
empty 344–346
errors in satisfying 330
error type as 340–341
functions and 331
methods and 325, 329–330, 335–336
return values and 325
types satisfying 325–328
interface types
about 327
assigning values with 328–330, 334
defining 325–327, 331
empty interface 344–346
naming rules for 332
type assertions and 334–335
Intn function (math/rand) 56–58

int type (integer)
about 13, 15
conversions with 22–23, 59
zero values for 17, 152
io/ioutil package
ReadAll function 380
ReadDir function 358, 362–363
io package
Close method 380
ReadCloser interface 380
Read method 380
Writer interface 452–453
IsNotExist function (os) 458

J

Join function (strings) 403

K

keyboard, user input via 35–36, 42, 46, 59
key/value pairs (maps) 212–214, 218, 221–224

L

labeling data. *See* maps
len built-in function
arrays and 156, 161
slices and 177, 194
Linux operating systems 132, 490
ListenAndServe function (net/http) 430, 447
lists. *See* arrays; slices
literals
array 153
defined types and 269
map 214
slice 177, 183
string 11–12
struct 251
localhost hostname 429
logical operators 40, 487
log package
error handling and 96
Fatal function 38, 447
loops
about 31, 61–62, 66–67
accessing array elements within 155
breaking out of 69–70
condition expression 61, 63
controlling flow of 68–69
for 61–63, 155, 223
for ... range 157–158, 221–224
infinite 62–63
initialization statements 61, 63
mistakes in initialization 64
post statements 61–63
scope and 63
ls utility (Unix) 491

M

Mac operating systems 132, 490
main function
about 4, 120
goroutines and 383–386
main package 4–6, 118–121
make built-in function
channels and 390
maps and 213
slices and 176, 180, 183
map literals 214, 251
maps
about 205, 212–213, 232–233
assigned values for 216–217
counting names using 219–224
for loops 223
for ... range loops 221–224
indexes for 213
key/value pairs for 212–214, 218, 221–224
nil value for 215
random order processing 223
slices and 212
zero values for 215–217
math operations and comparisons 13, 22–23, 60
math package 8–9, 14, 116
math/rand package 56–58
method attribute (form tag) 467
method promotion 310–311, 313–314

methods
about 32
calling 32–33, 275, 329–330
converting types using 283–284
deferring 353–355, 357
defining 32, 275–276
dot operator and 234
fixing function name conflicts 274
functions and 32, 277, 331
getter 304
interfaces and 325, 329–330, 345–346
naming rules for 86, 277
nil value for 39
receiver parameters. *See* receiver parameters (methods)
return values from 36–38
setter 292–294, 296–297, 301, 304
types and 32, 276, 322–324
unexported 277, 302–303, 305
multitasking 382–387

N

\n (newline character) 11, 46–47, 82
name attribute (input tag) 464
names
counting using maps 219–224
counting with slices 209–211
reading from text files 207–209
naming rules
defined types 241, 249–250
functions 21, 86, 118, 121
getter methods 304
interface types 332
methods 86, 277
packages 123
shadowing names 44–45
struct types 249–250, 298, 301
test function names 406
types 21, 241, 249–250, 298, 332
variables 21, 241
nesting
blocks of code 49–51
directories 118, 128–129
net/http package
about 426
Get function 380–381, 387
HandleFunc function 430, 432–433, 447
Handler interface 430
ListenAndServe function 430, 447
Redirect function 472
Request type 430, 447
Response type 380
ResponseWriter interface 430–431, 433, 447, 452–453
ServeMux type 430
Write method 431, 453
New function (errors) 96
New function (text/template) 451
newline character (\n) 11, 46–47, 82
NewReader function (bufio) 35
NewReplacer function (strings) 33
nil value
for functions 39, 47, 95
for maps 215
for methods 39
for slices 186, 190
Now function (time) 32, 58
numbers and numeric types
converting from strings 46–48, 59
defining 13
getting average of numbers in arrays 161
getting sum of numbers in arrays 159–160
reading numbers from files 350–351
rounding numbers 83
zero values for 17, 152

O

O_APPEND flag (os) 484, 488–489
O_CREATE flag (os) 484, 488–489
octal notation 492–494
OpenFile function (os) 470, 481–484, 489, 494
Open function (os) 165, 458, 494
opening files 165, 380
Open method (io) 380
operations and comparisons (math) 13, 22–23, 60
operators
address of 103
arithmetic 13, 269
assignment 62, 102

bitwise 485, 489, 494
boolean 40, 487
comparison 13, 212, 269
defined types and 269
logical 40, 487
OR bitwise operator (|) 485, 487
O_RDONLY flag (os) 483
O_RDWR flag (os) 488
os package
Args variable 192–193, 351
bitwise operators and 488–489
Create function 494
FileMode type 482, 491–494
File type 165, 458, 470, 494
IsNotExist function 458
O_APPEND flag 484, 488
O_CREATE flag 484, 488
OpenFile function 470, 481–484, 489, 494
Open function 165, 458, 494
O_RDONLY flag 483, 488
O_RDWR flag 488
O_WRONLY flag 484, 488
Stdout file descriptor 451–455
String method 491
Write method 453
output, formatting 81–85
overloading functions 272
O_WRONLY flag (os) 484, 489

P

p tag (HTML) 449
package clause 4–6, 118–121
package names
import paths vs. 56–57, 135–137, 139–141
naming conventions 123
shadowing names 44–45
packages. *See also* specific packages
about 4, 113, 116
accessing contents of 123
accessing unexported fields 302–303, 305
creating 118, 166–167
defined type names and 241, 249–250
documenting 141–142
dot operator and 234
downloading 137
exporting from. *See* exporting from packages
functions from other 8
installing 137
moving shared code to 124–125
moving struct types to different 248, 299–300
nested directories and import paths 128–129
publishing 133–136
reading documentation on 139–140
scope of 50
type definitions and 237–238
workspace directory and 117
padding with spaces 83
panic built-in function 365–368, 370–375
panicking programs and messages
about 170–171
avoiding 155
deferred calls and 366
fixing code using tests 420–421
panic function vs. error values 367–368
recover function and 370–373
reinstating panics 374
stack traces and 366
starting a panic 365
parameters
functions and 87, 102, 328
methods and 275–277, 323–324
parentheses ()
accessing struct fields through pointers 244–245
concrete types and 335
defining method parameters 277
in function calls 5, 7, 87
if statement and 40
import statement and 8
ParseFiles function (html/template) 450
ParseFiles function (text/template) 451
ParseFloat function (strconv) 47–48, 167, 352
Parse method (text/template) 451, 455
pass-by-value languages 102, 242
passing functions to other functions 436
PATH environment variable 130
percent sign (%) 82–83
permissions, file 490–493
Pike, Rob 2

pkg subdirectory 117
Playground site 3, 25, 27
pointers and pointer types
  about 103–104
  accessing struct fields through 244–245
  functions and 107–108, 243
  getting/changing values at pointers 104–105, 243
  passing large structs using 246
  receiver parameters 279–282, 293–294, 304, 332
  test functions and 406
ports 429
POST method (HTTP) 467–468, 472
post statements 61–63
Printf function (fmt)
  file permissions and 491–493
  formatting output 81–85, 154, 410, 485–488
  returning error values 96
  Stringer interface and 342–343
Print function (fmt) 35, 343
Println function (fmt) 5–9, 154, 343–344
publishing packages 133–136

Q

quotation marks 11–12

R

random number generation 56–58
{{range}} action 456, 459, 461
ReadAll function (io/ioutil) 380
ReadCloser interface (io) 380
ReadDir function (io/ioutil) 358, 362–363
Reader type (bufio) 35
reading documentation on packages 139–140
reading text files
  about 163–165
  into arrays 166–171
  reading names 207–209
  reading numbers 350–351
  reading slices 458–459
  using slices and append function 187–189
read (r) file permission 490
ReadString method (bufio) 35–39, 47
receiver parameters (methods)
  about 276–277
  pointers and 279–282, 293–294, 304, 332
  setter methods and 293–294
recover built-in function 370–373, 375
recursive functions
  about 360–361
  error handling 364–368
  listing directory contents 362–363
Redirect function (net/http) 472
reflect package 15, 22, 104
repetitive code 80, 88
Replace method (Replacer) 33
Replacer type (strings) 33
request/response process for web apps 427–431
Request type (net/http) 430, 447
resource paths 432–433
Response type (net/http) 380
ResponseWriter interface (net/http) 430–431, 433, 447, 452–453
return statement 91–92, 94
return values
  about 9, 91–93
  channels and 390–391
  compile errors 14, 92, 94, 100, 389
  error handling 36–38, 95, 99–100
  error values 96
  go statements and 389–390
  interfaces and 325
  methods and 277
  multiple 36–38, 97–98
  testing 408–409
rounding numbers 83
runes 12, 499–502
-run option (go test) 417

S

%s formatting verb 82
saving form data 470–471
Scan method (bufio) 165

Scanner type (bufio) 165
scope
  blocks of code and 50–51, 90
  of constants 127
  functions and 50, 90
  loops and 63
  shadowing names and 44–45
  type definitions and 237
script tag (HTML) 462
Seed function (math/rand) 58
semicolon (;) 5
ServeMux type (net/http) 430
set command 132
setter methods
  about 292, 304
  adding 294
  adding validation to 296–298
  encapsulation and 305
  exported 301–303
  pointer receiver parameters and 293–294
shadowing names 44–45
short variable declaration (:=)
  about 19–20, 54
  array literals and 153
  assignment statements and 51, 54
  blank identifier and 37
  for channels 390
  for maps 213
  for slices 176
  for structs 251
  for types 267
  using type conversions 267
single quotation mark (') 12
Sleep function (time) 386–387, 396
slice literals 177, 183
slice operator 180–181, 183, 193
slices
  about 175–177, 182, 232–233
  append function and 184–185, 187–189
  arrays and 176, 180–183, 193
  counting names with 209–211
  error handling 190
  functions and 176
  indexes for 180–181
  len function and 177, 194
  maps and 212
  nil value for 186, 190
  passing to variadic functions 199–200
  reading in from files 458–459
  zero values for 186
spaces, padding with 83
special characters 11
Sprintf function (fmt) 81, 96, 154, 410
square brackets []
  arrays and 151
  maps and 213
  slices and 176
src subdirectory 117–118
stack traces 366
start indexes (slices) 180–181
statements
  assignment 37, 51
  conditional. *See* conditional statements
  deferring function/method calls 353–355, 357, 366
  errors 6
  initialization 61, 63, 496–497
  in loops 61–63
  semicolons and 5
static types 14
Stdout file descriptor (os) 451–455
stop indexes (slices) 180–181
strconv package
  Atoi function 59, 95
  ParseFloat function 47–48, 167, 352
  reading documentation on 139–140
Stringer interface (fmt) 342–343
string literals 11–12
String method (os) 491
strings package
  about 116
  Join function 403
  NewReplacer function 33
  Replacer type 33
  Title function 8–9, 14
  TrimSpace function 47–48

string type (strings)
about 11–12, 15
converting between bytes 431
converting to numbers 46–48, 59
formatting verbs and 82
operator support 269
zero values for 17, 152
struct keyword 233, 236
struct literals 233, 251, 302
structs. *See also* arrays; maps; slices
about 233, 251
accessing fields through pointers 244–245
accessing fields with dot operator 234–235, 243
anonymous fields 258–259, 276, 308
embedding 258–260
making fields unexported 301
modifying using functions 242–243
moving types to different packages 248, 299–300
passing using pointers 246
setting up within other structs 255–260
storing subscriber data in 235
zero value for 251
struct types
about 292
adding fields to 255
assigning values to fields 235
channels and 390, 398
creating 236–238, 253–254, 290, 460
declaring 233
embedding 258–259, 276, 308, 310
encapsulation and 289
executing templates with 457
functions and 239–240
inserting fields into templates with actions 457
invalid data 291
moving to different packages 248, 299–300
naming rules for 249–250, 298, 301
in struct literals 251
subdirectories 359
subscriber data, storing in structs 235
sum of numbers in arrays 159–160
switch statement 498
synchronizing goroutines with channels 392–393

T

%t formatting verb 82
%T formatting verb 82
\t (tab character) 11, 47
T type (testing) 406, 410–411
table-driven testing 418–419
tables, aligning columns in 83
template package. *See* html/template package
templates. *See* HTML templates
Template type (html/template) 450
Template type (text/template) 451–454
test-driven development 413
testing
about 401–405
Errorf method test failure messages 410
fixing bugs 414–415
fixing panicking code using tests 420–421
getting tests to pass 412–413
return values 408–409
running specific sets of tests 417
table-driven 418
test-driven development 413
using type assertions 338
writing tests 406
testing package
about 401, 406
Errorf method 410–411
T type 406, 410–411
text files
reading 163–165
reading into arrays 166–171
reading names from 207–209
reading slices in from 458–459
reading using slices and append function 187–189
text/template package
about 451, 462
Execute method 451–455
New function 451
ParseFiles function 451
Parse method 451, 455
Template type 451–454
Thompson, Ken 2

threads 383
time package
  Now function 32, 58
  Sleep function 386–387, 396
  Time type 32, 58
  Year method 32
Time type (time) 32, 58
title case 9, 14
Title function (strings) 8–9, 14
TrimSpace function (strings) 47–48
type assertions
  about 334–335
  error handling and 336–337, 365
  panicking and 372
  using 338
type attribute (input tag) 464
type keyword 236
TypeOf function (reflect) 15, 22, 104
types. *See also* specific types
  about 13–15, 380, 499
  arrays and 152
  compile errors 20, 92, 94
  constants and 126
  converting using functions 271–272
  converting using methods 283–284
  converting values 22–23
  creating 236–237
  declaring function parameters 87
  declaring variables 16–17, 19–20
  determining for arguments 14
  functions as 436–438
  math operations and comparisons 22–23
  methods and 32, 276, 322–324
  mixing values of different 232–233
  moving to different packages 248, 299–300
  naming rules for 21, 241, 249–250, 298, 332
  operators and 269
  pointer 104–105, 107–108
  satisfying interfaces 325–328
  shadowing names 44–45
  static 14
  underlying basic types 267–268
  using to store data 238
  using with functions 239–240

U

uint32 type 491, 499
unexported fields
  accessing 301–303, 305
  encapsulation and 309
  functions and 21, 121, 302–303, 305
  methods and 277, 302–303, 305
  promotion and 309–311
  struct types and 301
  types and 21, 241, 249–250, 298
  variables and 21, 241, 302–303, 305
Unicode character code 12
universe block 341
Unix operating systems 490
user input
  via HTML forms 464
  via keyboard 35–36, 42, 46, 59
UTF-8 standard 500

V

-v flag (go test) 417
%v formatting verb 82
%#v formatting verb 82–83, 154, 186
validation, date 291, 296–298
value widths, formatting for Printf 81, 83–84
variables
  about 16
  assigning values to 16–17, 37, 51, 54
  compile errors 20, 22, 36
  declaring 16–17, 37, 51, 54
  defined types and 268
  first-class functions and 435–436
  naming rules for 21, 241
  pointers and 103–105, 107–108
  scope of 50–51, 90
  shadowing names 44–45
  unexported 21, 241, 302–303, 305
  zero values in 17, 152
variadic functions
  about 195, 344
  examples of 196–197

passing slices to 199–200
using to calculate averages 198
var keyword 16
verbs, formatting for Printf 81–83, 485

W

web apps
first-class functions 435–436
functions as types 436–438
handling requests and checking errors 447
HTML forms and 465–471
passing functions to other functions 436
request/response process 427–431
resource paths 432–433
responding with HTML code 450
simple example 428–431
writing 426
web pages, retrieving 380–381, 396–397
website (Go) 3, 25
whitespace characters 47
Windows operating systems 132, 490
workspace directory 117–118, 128
Write method (net/http) 431, 453
Write method (os) 453
Writer interface (io) 452–453
write (w) file permission 490
writing
tests 406
web apps 426

Y

Year method (time) 32

Z

zero values
for arrays 152
for maps 215–217
for numbers 17, 152
for slices 186
for strings 17, 152
for structs 251
for variables 17, 152

This isn't goodbye

Bring your brain over to *headfirstgo.com*



O'REILLY®

There's much more where this came from.

Experience books, videos, live online training courses, and more from O'Reilly and our 200+ partners—all in one place.

Learn more at oreilly.com/online-learning

©2019 O'Reilly Media, Inc. O'Reilly is a registered trademark of O'Reilly Media, Inc. | 175