

By **Leif Azzopardi** and **David Maxwell**

Python web
development
with Django

Tango with Django

A beginner's guide to web development
with Django 1.9.

REVISED AND
UPDATED FOR
django
version 1.9



Available from www.tangowithdjango.com

How to Tango with Django 1.9

A beginners guide to Python/Django

Leif Azzopardi and David Maxwell

This book is for sale at <http://leanpub.com/tangowithdjango19>

This version was published on 2016-08-04



Leanpub

This is a [Leanpub](#) book. Leanpub empowers authors and publishers with the Lean Publishing process. [Lean Publishing](#) is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

© 2016 Leif Azzopardi and David Maxwell

Tweet This Book!

Please help Leif Azzopardi and David Maxwell by spreading the word about this book on [Twitter!](#)

The suggested tweet for this book is:

I'm now ready to Tango with Django 1.9 @tangowithdjango

The suggested hashtag for this book is #tangowithdjango.

Find out what other people are saying about the book by clicking on this link to search for this hashtag on Twitter:

<https://twitter.com/search?q=#tangowithdjango>

Contents

1. Overview	1
1.1 Why Work with this Book?	1
1.2 What You will Learn	2
1.3 Technologies and Services	3
1.4 Rango: Initial Design and Specification	4
1.5 Summary	10
2. Getting Ready to Tango	12
2.1 Python	12
2.2 The Python Package Manager	13
2.3 Virtual Environments	14
2.4 Integrated Development Environment	14
2.5 Code Repository	15
3. Django Basics	16
3.1 Testing Your Setup	16
3.2 Creating Your Django Project	17
3.3 Creating a Django Application	20
3.4 Creating a View	22
3.5 Mapping URLs	23
3.6 Basic Workflows	25
4. Templates and Static Media	27
4.1 Using Templates	27
4.2 Serving Static Media	32
4.3 Basic Workflow	38
5. Models and Databases	40
5.1 Rango's Requirements	40
5.2 Telling Django about Your Database	41
5.3 Creating Models	42
5.4 Creating and Migrating the Database	44
5.5 Django Models and the Shell	46
5.6 Configuring the Admin Interface	47

CONTENTS

5.7	Creating a Population Script	50
5.8	Workflow: Model Setup	55
6.	Automated Testing	59
6.1	Running Tests	59
6.2	Coverage Testing	63
7.	Deploying Your Project	66
7.1	Creating a PythonAnywhere Account	66
7.2	The PythonAnywhere Web Interface	66
7.3	Creating a Virtual Environment	68
7.4	Setting up Your Web Application	71
7.5	Log Files	75
8.	Final Thoughts	76
9.	Setting up your System	77
9.1	Installing the Software	77
9.2	Virtual Environments	82
9.3	Exercises	84
10.	A Crash Course in UNIX-based Commands	85
10.1	Using the Terminal	85
10.2	Core Commands	89
11.	Virtual Environments	91
12.	A Git Crash Course	93
12.1	Why Use Version Control?	93
12.2	How Git Works	94
12.3	Setting up Git	95
12.4	Basic Commands and Workflow	98
12.5	Recovering from Mistakes	105

1. Overview

The aim of this book is to provide you with a practical guide to web development using *Django 1.9* and *Python*. The book is designed primarily for students, providing a walkthrough of the steps involved in getting your first web applications up and running, as well as deploying them to a web server.

This book seeks to complement the [official Django Tutorials](#) and many of the other excellent tutorials available online. By putting everything together in one place, this book fills in many of the gaps in the official Django documentation providing an example-based design driven approach to learning the Django framework. Furthermore, this book provides an introduction to many of the aspects required to master web application development.

1.1 Why Work with this Book?

This book will save you time. On many occasions we've seen clever students get stuck, spending hours trying to fight with Django and other aspects of web development. More often than not, the problem was usually because a key piece of information was not provided, or something was not made clear. While the occasional blip might set you back 10-15 minutes, sometimes they can take hours to resolve. We've tried to remove as many of these hurdles as possible. This will mean you can get on with developing your application instead of stumbling along.

This book will lower the learning curve. Web application frameworks can save you a lot of hassle and lot of time. Well, that is if you know how to use them in the first place! Often the learning curve is steep. This book tries to get you going - and going fast by explaining how all the pieces fit together.

This book will improve your workflow. Using web application frameworks requires you to pick up and run with a particular design pattern - so you only have to fill in certain pieces in certain places. After working with many students, we heard lots of complaints about using web application frameworks - specifically about how they take control away from them (i.e. inversion of control). To help you, we've created a number of *workflows* to focus your development process so that you can regain that sense of control and build your web application in a disciplined manner.

This book is not designed to be read. Whatever you do, *do not read this book!* It is a hands-on guide to building web applications in Django. Reading is not doing. To increase the value you gain from this experience, go through and develop the application. When you code up the application, *do not just cut and paste the code*. Type it in, think about what it does, then read the explanations we have provided to describe what is going on. If you still do not understand, then check out the Django documentation, go to [Stack Overflow](#) or other helpful websites and fill in this gap in your knowledge. If you think it is worth mentioning, please get in touch with us so that we can improve the book - we already have a number of contributors and we will happily acknowledge your contribution!

1.2 What You will Learn

In this book, we will be taking an example-based approach. The book will show you how to design a web application called *Rango* (see the [Design Brief](#) below). Along the way, we'll show you how to perform the following key tasks.

- **How to setup your development environment** - including how to use the terminal, your virtual environment, the `pip` installer, how to work with Git, and more.
- **Setup a Django project** and create a basic Django application.
- **Configure the Django project** to serve static media and other media files.
- Work with Django's *Model-View-Template* design pattern.
- **Create database models** and use the *object relational mapping (ORM)* functionality provided by Django.
- **Create forms** that can utilise your database models to create dynamically generated web-pages.
- Use the **user authentication** services provided by Django.
- Incorporate **external services** into your Django application.
- Include *Cascading Styling Sheets (CSS)* and *JavaScript* within a web application.
- **Apply CSS** to give your application a professional look and feel.
- Work with **cookies and sessions** with Django.
- Include more advanced functionality like *AJAX* into your application.
- **Deploy your application** to a web server using *PythonAnywhere*.

At the end of each chapter, we have included a number of exercises designed to push you harder and to see if you can apply what you have learned. The later chapters of the book provide a number of open development exercises along with coded solutions and explanations.



Exercises will be clearly delineated like this!

In each chapter we have added a number of exercises to test your knowledge and skill.

You will need to complete these exercises as the subsequent chapters are dependent on them.

Don't worry if you get stuck, though. If you want to, you can check out our solutions to all the exercises on our [GitHub repository](#).

1.3 Technologies and Services

Through the course of this book, we will use various technologies and external services including:

- [Python](#)
- [Pip package manager](#)
- [Django](#)
- [Git](#)
- [GitHub](#)
- [HTML](#)
- [CSS](#)
- [JavaScript](#)
- [jQuery](#)
- [Twitter Bootstrap](#)
- [Bing Search API](#) via [Azure Datamarket](#)
- [PythonAnywhere](#)

We've selected these technologies and services as they are either fundamental to web development, and/or enable us to provide examples on how to integrate your web application with CSS toolkits like *Twitter Bootstrap*, external services like those provided by *Microsoft Azure* and deploy your application quickly and easily with *PythonAnywhere*.

1.4 Rango: Initial Design and Specification

As previously mentioned, the focus of this book will be to develop an application called *Rango*. As we develop this application, it will cover the core components that need to be developed when building any web application.

To see a fully-functional version of the application, you can visit the [How to Tango with Django website](#).

Design Brief

Your client would like you to create a website called *Rango* that lets users browse through user-defined categories to access various web pages. In [Spanish](#), the word *rango* is used to mean “*a league ranked by quality*” or “*a position in a social hierarchy*”.

- For the **main page** of the Rango website, your client would like visitors to be able to see:
 - the *five most viewed pages*;
 - the *five most viewed (or rango’ed) categories*; and
 - *some way for visitors to browse or search through categories*.
- When a user views a **category page**, your client would like Rango to display:
 - the *category name, the number of visits, the number of likes*, along with the list of associated pages in that category (showing the page’s title, and linking to its URL); and
 - *some search functionality (via Bing’s Search API) to find other pages that can be linked to this category*.
- For a particular category, the client would like: the *name of the category to be recorded*; the *number of times each category page has been visited*; and how many users have *clicked a “like” button* (i.e. the page gets rango’ed, and voted up the social hierarchy).
- *Each category should be accessible via a readable URL* - for example, `/rango/books-about-django/`.
- *Only registered users will be able to search and add pages to categories*. Visitors to the site should therefore only be able to register for an account.

At first glance, the specified application to develop seems reasonably straightforward. In essence, it is just a list of categories which link to pages, right? However, there are a number of complexities and challenges that need to be addressed. First, let’s try and build up a better picture of what needs to be developed by laying down some high-level designs.



Exercises

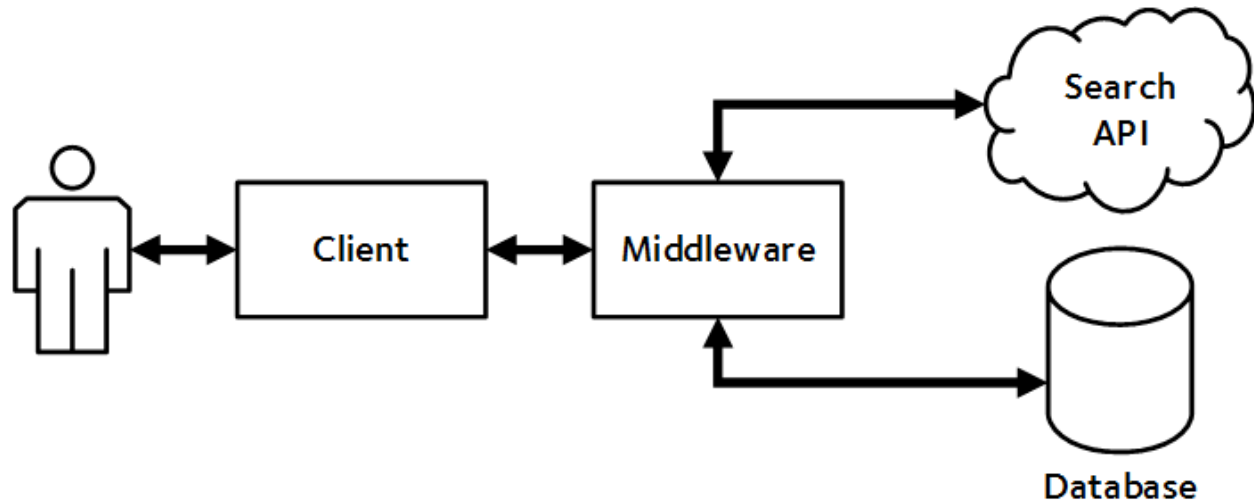
Before going any further, think about these specifications and draw up the following design artefacts.

- An **N-Tier or System Architecture** diagram.
- **Wireframes** of the main and category pages.
- A series of **URL mappings** for the application.
- An ***Entity-Relationship (ER)*** diagram to describe the data model that we'll be implementing.

It's good practice for you to try these exercises out before moving on to examine how we went about creating all of the above.

N-Tier Architecture

The high-level architecture for most web applications is a *3-Tier architecture*. Rango will be a variant on this architecture as it interfaces with an external service.



Overview of the 3-tier system architecture for our Rango application.

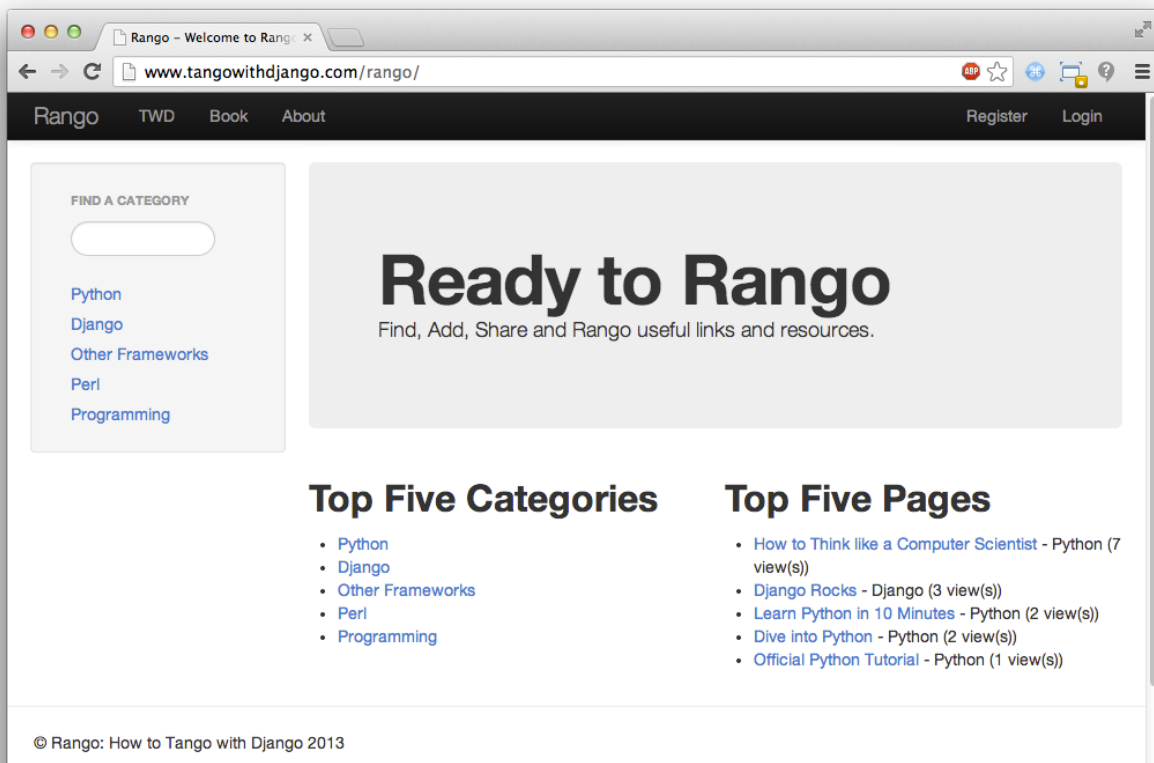
Since we are building a web application with Django, we will use the following technologies for the following tiers.

- The **client** will be a Web browser (such as *Chrome*, *Firefox*, and *Safari*) which will render HTML/CSS pages.
- The **middleware** will be a *Django* application, and will be dispatched through Django's built-in development Web server while we develop.
- The **database** will be the Python-based *SQLite3* Database engine.
- The **search API** will be the *Bing Search API*.

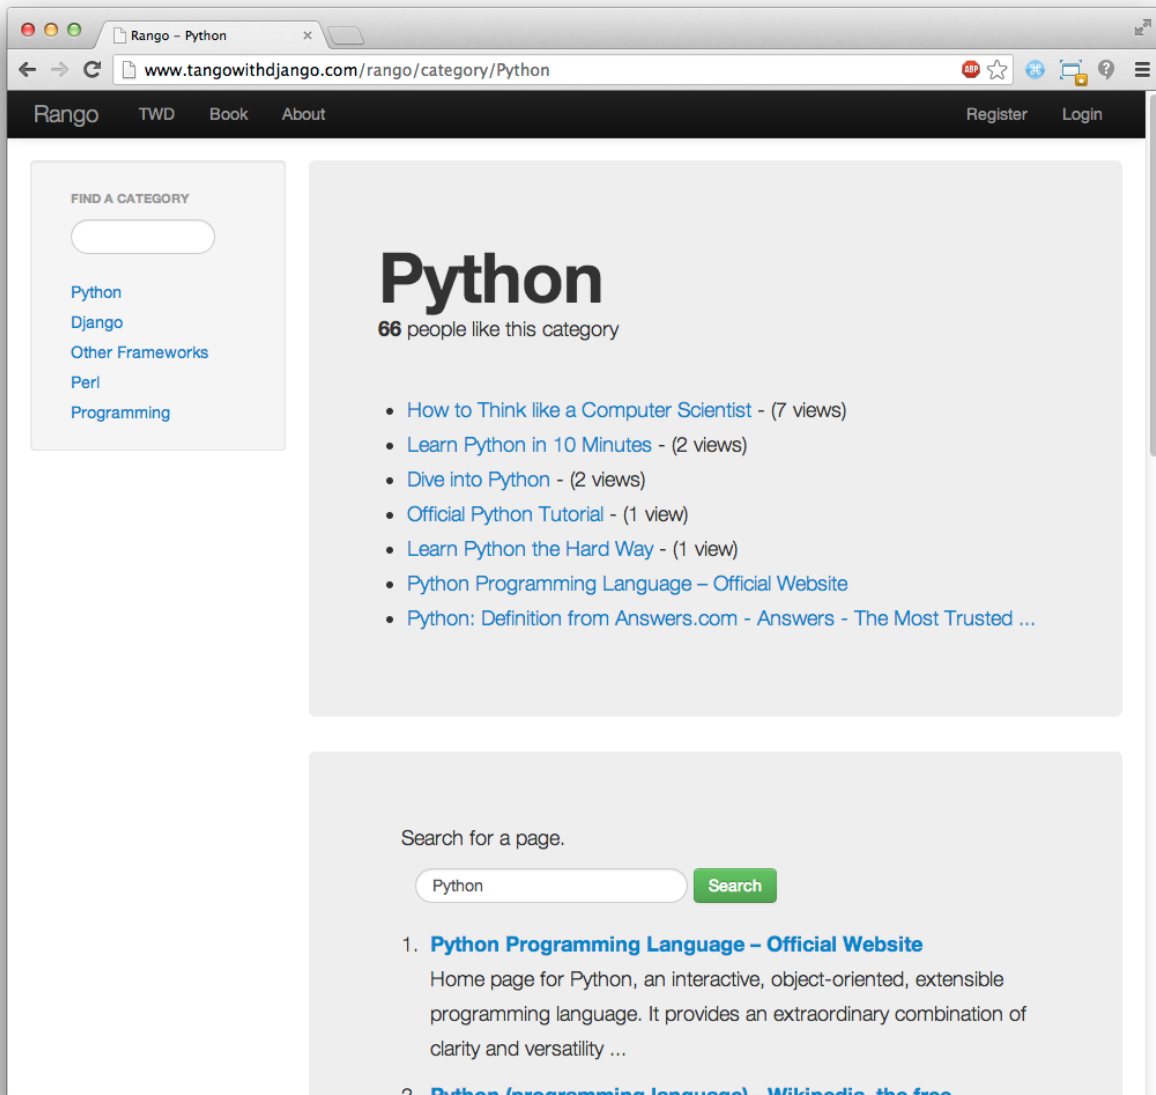
For the most part, this book will focus on developing the middleware. It should however be quite evident from the [system architecture diagram](#) that we will have to interface with all the other components.

Wireframes

Wireframes are a great way to provide clients with some idea of what the application should look like when complete. They save a lot of time, and can vary from hand drawn sketches to exact mockups depending on the tools that you have at your disposal. For our Rango application, we'd like to make the index page of the site look like the [screenshot below](#). Our category page is also [shown below](#).



The index page with the categories bar on the left, also showing the top five pages and top five categories.



The category page showing the pages in the category (along with the number of views). Below, a search for *Python* has been conducted, with the results shown underneath.

Pages and URL Mappings

From the specification, we have already identified two pages that our application will present to the user at different points in time. To access each page we will need to describe URL mappings. Think of a URL mapping as the text a user will have to enter into a browser's address bar to reach the given page. The basic URL mappings for Rango are shown below.

- `/rango/` will point to the main (or index) page.


- `/rango/about/` will point to the about page.
- `/rango/category/<category_name>/` will point to the category page for `<category_name>`, where the category might be:
 - games;
 - python-recipes; or
 - code-and-compilers.

As we build our application, we will probably need to create other URL mappings. However, the ones listed above will get us started and give us an idea of the different pages.

As we progress through the book, we will flesh out how to construct these pages using the Django framework and use its [Model-View-Template](#) design pattern. However, now that we have a gist of the URL mappings and what the pages are going to look like, we need to define the data model that will house the data for our Web application.

Entity-Relationship Diagram

Given the specification, it should be clear that we have at least two entities: a *category* and a *page*. It should also be clear that a *category* can house many *pages*. We can formulate the following ER Diagram to describe this simple data model.



The Entity Relationship Diagram of Rango's two main entities.

Note that this specification is rather vague. A single page could in theory exist in one or more categories. Working with this assumption, we could model the relationship between categories and pages as a [many-to-many relationship](#). This approach however introduces a number of complexities, so we will make the simplifying assumption that *one category contains many pages, but one page is assigned to one category*. This does not preclude that the same page can be assigned to different categories - but the page would have to be entered twice, which is not ideal.



Take Note!

It's good practice to note down any working assumptions you make, just like the one-to-many relationship assumption we assume above. You never know when they may come back to bite you later on! By noting them down, this means you can communicate it with your development team and make sure that the assumption is sensible and that they are happy to proceed under such an assumption.

With this assumption, we then produce a series of tables that describe each entity in more detail. The tables contain information on what fields are contained within each entity. We use Django `ModelField` types to define the type of each field (i.e. `IntegerField`, `CharField`, `URLField` or `ForeignKey`). Note that in Django *primary keys* are implicit such that Django adds an `id` to each Model, but we will talk more about that later in the Models and Database chapter.

Category Model

Field	Type
<code>name</code>	<code>CharField</code>
<code>views</code>	<code>IntegerField</code>
<code>likes</code>	<code>IntegerField</code>

Page Model

Field	Type
<code>category</code>	<code>ForeignKey</code>
<code>title</code>	<code>CharField</code>
<code>url</code>	<code>URLField</code>
<code>views</code>	<code>IntegerField</code>

We will also have model for the `User` so that they can register and login. We have not shown it here, but shall introduce it later in the book when we discuss User Authentication. In the following chapters will we see how to instantiate these models in Django and how to use Django's ORM to connect to the database.

1.5 Summary

These high level design and specifications will serve as a useful reference point when building our Web application. While we will be focusing on using specific technologies, these steps are common to most database driven websites. It's a good idea to become familiar with reading and producing such specifications and designs, which warrants another book on the topic. Here we will be focusing on using Django and the related technologies to implement this specification.



Cut and Paste Coding

As you progress through the tutorial, you'll most likely be tempted to cut and paste the code from the book to your code editor. **However, it is better to type in the code.** We know that this is a hassle, but it will help you to remember the process better and the commands that you will be using later on.

Furthermore, cutting and pasting Python code is asking for trouble. Whitespace can end up being interpreted as spaces, tabs or a mixture of spaces and tabs. This will lead to all sorts of weird errors, and not necessarily indent errors. If you do cut and paste code be wary of this. Pay particular attention to this if you're using Python 3 - inconsistent use of tabs and spaces in your code's indentation will lead to a `TabError`.

Most code editors will show the whitespace and whether it is tabs or spaces. If so, turn it on and save yourself a lot of confusion.

2. Getting Ready to Tango

Before we get down to coding, it's really important that we get our development environment setup so that you can *Tango with Django!* You'll need to ensure that you have all the necessary components installed on your computer. This chapter outlines the five key components that you need to be aware of, setup and use. These are listed below.

- Working with the [terminal](#) or [Command Prompt](#).
- *Python* and your *Python* installation.
- The Python Package Manager *pip* and *virtual environments*.
- Your *Integrated Development Environment (IDE)*, if you choose to use one.
- A *Version Control System (VCS)*, *Git*.

If you already have Python 2.7/3.5 and Django 1.9 installed on your computer, and are familiar with the technologies mentioned, then you can skip straight to the [Django Basics chapter](#). Otherwise, we below provide an overview of the different components and why they are important. We also provide a series of pointers on how to setup the various components.



Your Development Environment

Setting up your development environment is pretty tedious and often frustrating. It's not something that you'd do everyday. Below, we have put together the list of core technologies you need to get started and pointers on how to install them.

From experience, we can also say that it's a good idea when setting your development environment up to note down the steps you took. You'll need them again one day - whether because you have purchased a new computer, or you have been asked to help someone else set their computer up! Taking a note of everything you do will save you time and effort in the future. Don't just think short term!

2.1 Python

To work with Tango with Django, we require you to have installed on your computer a copy of the Python programming language. Any version from the 2.7 family - with a minimum of 2.7.5 - or version 3.4+ will work fine. If you're not sure how to install Python and would like some assistance, have a look at [the chapter dealing with installing components](#).



Not sure how to use Python?

If you haven't used Python before - or you simply wish to brush up on your skills - then we highly recommend that you check out and work through one or more of the following guides:

- [Learn Python in 10 Minutes](#) by Stavros;
- [The Official Python Tutorial](#);
- [Think Python: How to Think like a Computer Scientist](#) by Allen B. Downey; or
- [Learn to Program](#) by Jennifer Campbell and Paul Gries.

2.2 The Python Package Manager

Pip is the python [package manager](#). The package manager allows you install various libraries for the Python programming language to enhance its functionality.

A package manager, whether for Python, your [operating system](#) or [some other environment](#), is a software tool that automates the process of installing, upgrading, configuring and removing *packages* - that is, a package of software which you can use on your computer. This is opposed to downloading, installing and maintaining software manually. Maintaining Python packages can be difficult - with other packages are *dependencies* of the package you are attempting to install, to maintaining system paths - pip should handle this all for you.

Try and run pip with the command `$ pip`. If the command is not found, you'll need to install pip itself - check out the [system setup chapter](#) for more information. You should also ensure that the following packages are installed on your system. Run the following commands to install Django and [pillow](#) (an image manipulation library for Python).

```
$ pip install -U django==1.9.5
$ pip install pillow
```



Problems Installing pillow?

When installing Pillow, you may receive an error stating that the installation failed due to a lack of JPEG support. This error is shown as the following:

```
ValueError: jpeg is required unless explicitly disabled using
--disable-jpeg, aborting
```

If you receive this error, try installing Pillow *without* JPEG support enabled, with the following command.

```
pip install pillow --global-option="build_ext"
--global-option="--disable-jpeg"
```

While you obviously will have a lack of support for handling JPEG images, Pillow should then install without problem. Getting Pillow installed is enough for you to get started with this tutorial. For further information, check out the [Pillow documentation](#).

2.3 Virtual Environments

We're almost all set to go! However, before we continue, it's worth pointing out that while this setup is fine to begin with, there are some drawbacks. What if you had another Python application that requires a different version to run, or you wanted to switch to the new version of Django, but still wanted to maintain your Django 1.9 project?

The solution to this is to use [virtual environments](#). Virtual environments allow multiple installations of Python and their relevant packages to exist in harmony. This is the generally accepted approach to configuring a Python setup nowadays.

Setting up a virtual environment is not necessarily but it is highly recommended. The [virtual environment chapter](#) details how to setup, create and use virtual environments.

2.4 Integrated Development Environment

While not absolutely necessary, a good Python-based IDE can be very helpful to you during the development process. Several exist, with perhaps [PyCharm](#) by JetBrains and [PyDev](#) (a plugin of the [Eclipse IDE](#)) standing out as popular choices. The [Python Wiki](#) provides an up-to-date list of Python IDEs.

Research which one is right for you, and be aware that some may require you to purchase a licence. Ideally, you'll want to select an IDE that supports integration with Django.

We use PyCharm as it supports virtual environments and Django integration - though you will have to configure the IDE accordingly. We don't cover that here - although JetBrains do provide a [guide on setting PyCharm up](#). PyCharm is also available for students - and JetBrains provides a graduate discount, too.

2.5 Code Repository

We should also point out that when you develop code, you should always house your code within a version-controlled repository such as [SVN](#) or [GIT](#). We won't be explaining this right now, so that we can get stuck into developing an application in Django. We have however written a [chapter providing a crash course on GIT](#) for your reference which you can refer to later on. **We highly recommend that you set up a Git repository for your own projects. Doing so will could save you from disaster, as we can no doubt tell you.**



Exercises

To get comfortable with your environment, try out the following exercises.

- Install Python 2.7.5+/3.4+ and Pip.
- Play around with your CLI and create a directory called `code`, which we use to create our projects in.
- Setup your Virtual Environment (optional)
- Install the Django and Pillow packages
- Setup an account on a Git Repository site like: GitHub, BitBucket, etc if you haven't already done so.
- Download and setup an Integrated Development Environment like [PyCharm](#)

As previously stated, we've made the code for the book and application available on our [GitHub repository](#).

- If you spot any errors or problem, please let us know by making a change request on GitHub.
- If you have any problems with the exercises, you can check out the repository to see how we completed them.

3. Django Basics

Let's get started with Django! In this chapter, we'll be giving you an overview of the creation process. You'll be setting up a new project and a new Web application. By the end of this chapter, you will have a simple Django powered website up and running!

3.1 Testing Your Setup

Let's start by checking that your Python and Django installations are correct for this tutorial. To do this, open a new terminal window and issue the following command, which tells you what Python version you have:

```
$ python --version
```

The response should be something like 2.7.11 or 3.5.1, but any 2.7.5+ or 3.4+ versions of Python should work fine. If you need to upgrade or install Python go to the chapter on [setting up your system](#).

If you are using a virtual environment, then ensure that you have activated it - if you don't remember how go back to our chapter on [virtual environments](#).

After verifying your Python installation, check your Django installation. In your terminal window, run the Python interpreter by issuing the following command.

```
$ python
Python 2.7.10 (default, Jul 14 2015, 19:46:27)
[GCC 4.2.1 Compatible Apple LLVM 6.0 (clang-600.0.39)] on darwin
Type "help", "copyright", "credits" or "license" for more information.
>>>
```

At the prompt, enter the following commands:

```
>>> import django
>>> django.get_version()
'1.9.5'
>>> exit()
```

All going well you should see the correct version of Django, and then can use `exit()` to leave the Python interpreter. If `import django` fails to import, then check that you are in your virtual environment, and check what packages are installed with `pip list` at the terminal window.

If you have problems with installing the packages or have a different version installed, go to [System Setup](#) chapter or consult the [Django Documentation on Installing Django](#).

Prompts

In this book, there's two things you should look out for when we include code snippets.

Snippets beginning with a dollar sign (\$) indicates that the remainder of the following line is a terminal or Command Prompt command.

Whenever you see `>>>`, the following is a command that should be entered into the interactive Python interpreter. This is launched by issuing `$ python`. See what we did there? You can also exit the Python interpreter by entering `>>> quit()`.

3.2 Creating Your Django Project

To create a new Django Project, go to your workspace directory, and issue the following command:

```
$ django-admin.py startproject tango_with_django_project
```

If you don't have a workspace directory, then create one, so that you can house your Django projects and other code projects within this directory. We will refer to your workspace directory in the code as `<workspace>`. You will have to substitute in the path to your workspace directory, for example: `/Users/leifos/Code/` or `/Users/maxwelled90/Workspace/`.

Running Windows?

On Windows, you may have to use the full path to the `django-admin.py` script, for example:

```
python c:\python27\scripts\django-admin.py  
startproject tango_with_django_project
```

as suggested on [StackOverflow](#).

This command will invoke the `django-admin.py` script, which will set up a new Django project called `tango_with_django_project` for you. Typically, we append `_project` to the end of our Django project directories so we know exactly what they contain - but the naming convention is entirely up to you.

You'll now notice within your workspace is a directory set to the name of your new project, `tango_with_django_project`. Within this newly created directory, you should see two items:

- another directory with the same name as your project, `tango_with_django_project`; and
- a Python script called `manage.py`.

For the purposes of this tutorial, we call this nested directory called `tango_with_django_project` the *project configuration directory*. Within this directory, you will find four Python scripts. We will discuss these scripts in detail later on, but for now you should see:

- `__init__.py`, a blank Python script whose presence indicates to the Python interpreter that the directory is a Python package;
- `settings.py`, the place to store all of your Django project's settings;
- `urls.py`, a Python script to store URL patterns for your project; and
- `wsgi.py`, a Python script used to help run your development server and deploy your project to a production environment.

In the project directory, you will see there is a file called `manage.py`. We will be calling this script time and time again as we develop our project. It provides you with a series of commands you can run to maintain your Django project. For example, `manage.py` allows you to run the built-in Django development server, test your application and run various database commands. We will be using the script for virtually every Django command we want to run.



The Django Admin Script

For Further Information on Django admin script, see the Django documentation for more details about the [Admin and Manage scripts](#).

If you run `python manage.py help` you can see the list of commands available for you to run.

You can try using the `manage.py` script now, by issuing the following command.

```
$ python manage.py runserver
```

Executing this command will launch Python, and instruct Django to initiate its lightweight development server. You should see the output in your terminal window similar to the example shown below:

```
$ python manage.py runserver
```

```
Performing system checks...
```

```
System check identified no issues (0 silenced).
```

```
You have unapplied migrations; your app may  
not work properly until they are applied.
```

```
Run 'python manage.py migrate' to apply them.
```

```
April 10, 2016 - 11:07:24
```

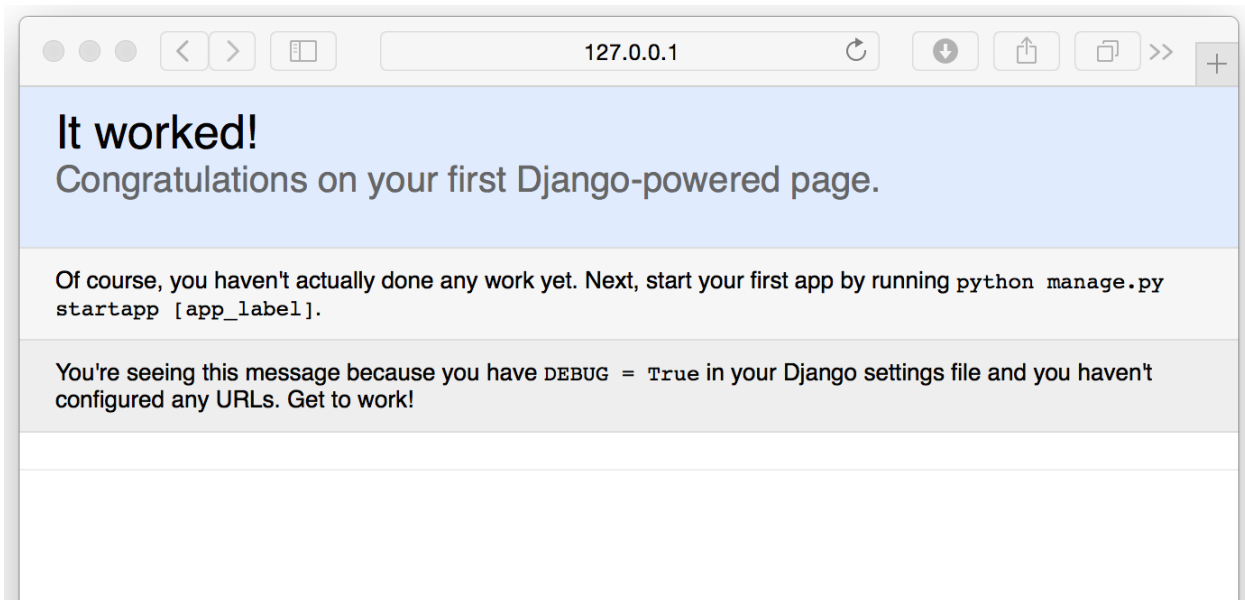
```
Django version 1.9.5, using settings 'tango_with_django_project.settings'
```

```
Starting development server at http://127.0.0.1:8000/
```

```
Quit the server with CONTROL-C.
```

In the output you can see a number of things. First, there are no issues that stop the application from running. Second, however, you will notice that a warning is raised, i.e. unapplied migrations. We will talk about this in more detail when we setup our database, but for now we can ignore it. Third, and most importantly, you can see that a URL has been specified: `http://127.0.0.1:8000/`, which is the address of the Django development webserver.

Now open up your Web browser and enter the URL `http://127.0.0.1:8000/`. You should see a webpage similar to the one shown in below.



A screenshot of the initial Django page you will see when running the development server for the first time.

You can stop the development server at anytime by pushing `CTRL + C` in your terminal or Command Prompt window. If you wish to run the development server on a different port, or allow users from other machines to access it, you can do so by supplying optional arguments. Consider the following command:

```
$ python manage.py runserver <your_machines_ip_address>:5555
```

Executing this command will force the development server to respond to incoming requests on TCP port 5555. You will need to replace `<your_machines_ip_address>` with your computer's IP address or `127.0.0.1`.



Don't know your IP Address?

If you use `0.0.0.0`, Django figures out what your IP address is. Go ahead and try:

```
python manage.py runserver 0.0.0.0:5555
```

When setting ports, it is unlikely that you will be able to use TCP port 80 or 8080 as these are traditionally reserved for HTTP traffic. Also, any port below 1024 is considered to be [privileged](#) by your operating system.

While you won't be using the lightweight development server to deploy your application, it's nice to be able to demo your application on another machine in your network. Running the server with your machine's IP address will enable others to enter in `http://<your_machines_ip_address>:<port>/` and view your Web application. Of course, this will depend on how your network is configured. There may be proxy servers or firewalls in the way which would need to be configured before this would work. Check with the administrator of the network you are using if you can't view the development server remotely.

3.3 Creating a Django Application

A Django project is a collection of *configurations* and *applications* that together make up a given Web application or website. One of the intended outcomes of using this approach is to promote good software engineering practices. By developing a series of small applications, the idea is that you can theoretically drop an existing application into a different Django project and have it working with minimal effort.

A Django application exists to perform a particular task. You need to create specific applications that are responsible for providing your site with particular kinds of functionality. For example, we could imagine that a project might consist of several applications including a polling app, a registration app, and a specific content related app. In another project, we may wish to re-use the polling and registration apps, and so can include them in other projects. We will talk about this later. For now we are going to create the application for the *Rango* app.

To do this, from within your Django project directory (e.g. `<workspace>/tango_with_django_-project`), run the following command.

```
$ python manage.py startapp rango
```

The `startapp` command creates a new directory within your project's root. Unsurprisingly, this directory is called `rango` - and contained within it are another five Python scripts:

- another `__init__.py`, serving the exact same purpose as discussed previously;
- `admin.py`, where you can register your models so that you can benefit from some Django machinery which creates an admin interface for you;
- `apps.py`, that provides a place for any application specific configuration;
- `models.py`, a place to store your application's data models - where you specify the entities and relationships between data;
- `tests.py`, where you can store a series of functions to test your application's code;
- `views.py`, where you can store a series of functions that handle requests and return responses; and
- `migrations` directory, which stores database specific information related to your models.

`views.py` and `models.py` are the two files you will use for any given application, and form part of the main architectural design pattern employed by Django, i.e. the *Model-View-Template* pattern. You can check out [the official Django documentation](#) to see how models, views and templates relate to each other in more detail.

Before you can get started with creating your own models and views, you must first tell your Django project about your new application's existence. To do this, you need to modify the `settings.py` file, contained within your project's configuration directory. Open the file and find the `INSTALLED_APPS` tuple. Add the `rango` application to the end of the tuple, which should then look like the following example.

```
INSTALLED_APPS = [  
    'django.contrib.admin',  
    'django.contrib.auth',  
    'django.contrib.contenttypes',  
    'django.contrib.sessions',  
    'django.contrib.messages',  
    'django.contrib.staticfiles',  
    'rango',  
]
```

Verify that Django picked up your new application by running the development server again. If you can start the server without errors, your application was picked up and you will be ready to proceed to the next step.



startapp Magic

When creating a new app with the `python manage.py startapp` command, Django may add the new app's name to your `settings.py` `INSTALLED_APPS` list automatically for you. It's nevertheless good practice to check everything is setup correctly before you proceed.

3.4 Creating a View

With our Rango application created, let's now create a simple view. For our first view, let's just send some simple text back to the client - we won't concern ourselves about using models or templates just yet.

In your favourite IDE, open the file `views.py`, located within your newly created `rango` application directory. Remove the comment `# Create your views here.` so that you now have a blank file.

You can now add in the following code.

```
from django.http import HttpResponseRedirect

def index(request):
    return HttpResponseRedirect("Rango says hey there partner!")
```

Breaking down the three lines of code, we observe the following points about creating this simple view.

- We first import the `HttpResponse` object from the `django.http` module.
- Each view exists within the `views.py` file as a series of individual functions. In this instance, we only created one view - called `index`.
- Each view takes in at least one argument - a `HttpRequest` object, which also lives in the `django.http` module. Convention dictates that this is named `request`, but you can rename this to whatever you want if you so desire.
- Each view must return a `HttpResponse` object. A simple `HttpResponse` object takes a string parameter representing the content of the page we wish to send to the client requesting the view.

With the view created, you're only part of the way to allowing a user to access it. For a user to see your view, you must map a [Uniform Resource Locator \(URL\)](#) to the view.

To create an initial mapping, open `urls.py` located in your project directory and add the following lines of code to the `urlpatterns`:

```
from rango import views

urlpatterns = [
    url(r'^$', views.index, name='index'),
    url(r'^admin/', admin.site.urls),
]
```

This maps the basic URL to the `index` view in the `rango` application. Run the development server (e.g. `python manage.py runserver`) and visit `http://127.0.0.1:8000` or whatever address your development server is running on. You'll then see the rendered output of the `index` view.

3.5 Mapping URLs

Rather than directly mapping URLs from the project to the application, we can make our application more modular (and thus re-usable) by changing how we route the incoming URL to a view. To do this, we first need to modify the project's `urls.py` and have it point to the application to handle any specific `rango` application requests. Then, we need to specify how `rango` deals with such requests.

First, open the project's `urls.py` file which is located inside your project configuration directory. As a relative path from your workspace directory, this would be the file `<workspace>/tango_with_django_project/tango_with_django_project/urls.py`. Update the `urlpatterns` list as shown in the example below.

```
from django.conf.urls import url
from django.contrib import admin
from django.conf.urls import include
from rango import views

urlpatterns = [
    url(r'^$', views.index, name='index'),
    url(r'^rango/', include('rango.urls')),
    # above maps any URLs starting
    # with rango/ to be handled by
    # the rango application
    url(r'^admin/', admin.site.urls),
]
```

You will see that the `urlpatterns` is a Python list, which is expected by the Django framework. The added mapping looks for URL strings that match the patterns `^rango/`. When a match is made the remainder of the url string is then passed onto and handled by `rango.urls` through the use of the `include()` function from within `django.conf.urls`.

Think of this as a chain that processes the URL string - as illustrated in the [URL chain figure](#). In this chain, the domain is stripped out and the remainder of the URL string (`rango/`) is passed on to `tango_with_django` project, where it finds a match and strips away `rango/`, leaving an empty string to be passed on to the application `rango`.

Consequently, we need to create a new file called `urls.py` in the `rango` application directory, to handle the remaining URL string (and map the empty string to the `index` view):

```
from django.conf.urls import url
from rango import views


urlpatterns = [
    url(r'^$', views.index, name='index'),
]
```

This code imports the relevant Django machinery for URL mappings and the `views` module from `rango`. This allows us to call the function `url` and point to the `index` view for the mapping in `urlpatterns`.


The URL mapping we have created calls Django's `url()` function, where the first parameter is the regular expression `^$`, which matches to an empty string. Any URL string supplied by the user that matches this pattern means that the view `views.index()` would be invoked by Django. You might be thinking that matching a blank URL is pretty pointless - what use would it serve? Remember that when the URL pattern matching takes place, only a portion of the original URL string is considered. This is because our project will first process the original URL string (i.e. `rango/`) and strip away the `rango/` part, passing on an empty string to the `rango` application to handle.

The next parameter passed to the `url()` function is the `index` view, which will handle the incoming requests, followed by the optional parameter, `name` that is set to a string `'index'`. By naming our URL mappings we can employ *reverse URL matching* later on. That is we can reference the URL mapping by name rather than by the URL. Later we will explain how to use this when creating templates. But do check out [the Official Django documentation on this topic](#) for more information.

Now, restart the Django development server and visit `http://127.0.0.1:8000/rango/`. If all went well, you should see the text `Rango says hey there partner!`. It should look just like the screenshot shown below.



An illustration of a URL, represented as a chain, showing how different parts of the URL following the domain are the responsibility of different `url.py` files.



A screenshot of a Web browser displaying our first Django powered webpage. Hello, Rango!

Within each application, you will create a number of URL mappings. The initial mapping is quite simple, but as we progress through the tutorial we will create more sophisticated, parameterised URL mappings.

It's also important to have a good understanding of how URLs are handled in Django. So, if you are still a bit confused or would like to know more, check out the [official Django documentation on URLs](#) for further details and further examples.



Note on Regular Expressions

Django URL patterns use [regular expressions](#) to perform the matching. It is worthwhile familiarising yourself on how to use regular expressions in Python. The official Python documentation contains a [useful guide on regular expressions](#), while [regexcheatsheet.com](#) provides a [neat summary of regular expressions](#).

If you are using version control, now is a good time to commit the changes you have made to your workspace. Refer to the [chapter providing a crash course on Git](#) if you can't remember the commands and steps involved in doing this.

3.6 Basic Workflows

What you've just learnt in this chapter can be succinctly summarised into a list of actions. Here, we provide these lists for the two distinct tasks you have performed. You can use this section for a quick reference if you need to remind yourself about particular actions later on.

Creating a new Django Project

1. To create the project run, `python django-admin.py startproject <name>`, where `<name>` is the name of the project you wish to create.

Creating a new Django application

1. To create a new application, run `$ python manage.py startapp <appname>`, where `<appname>` is the name of the application you wish to create.

2. Tell your Django project about the new application by adding it to the `INSTALLED_APPS` tuple in your project's `settings.py` file.
3. In your project `urls.py` file, add a mapping to the application.
4. In your application's directory, create a `urls.py` file to direct incoming URL strings to views.
5. In your application's `view.py`, create the required views ensuring that they return a `HttpResponse` object.



Exercises

Now that you have got Django and your new app up and running, give the following exercises a go to reinforce what you've learnt. Getting to this stage is a significant landmark in working with Django. Creating views and mapping URLs to views is the first step towards developing more complex and usable Web applications.

- Revise the procedure and make sure you follow how the URLs are mapped to views.
- Now create a new view called `about` which returns the following: Rango says here is the about page.
- Now map the this view to `/rango/about/`. For this step, you'll only need to edit the `urls.py` of the `rango` application.
- Revise the `HttpResponse` in the `index` view to include a link to the about page.
- In the `HttpResponse` in the `about` view include a link back to the main page.
- If you haven't done so already, it is a good point to go off an complete part one of the official [Django Tutorial](#).



Hints

If you're struggling to get the exercises done, the following hints will hopefully provide you with some inspiration on how to progress.

- Your `index` view should be updated to include a link to the about view. Keep it simple for now - something like `Rango says: Hello world!
 About` will suffice. We'll be going back later to improve the presentation of these pages.
- The regular expression to match `about/` is `r'^about/'` - this will be handy when thinking about your URL pattern.
- The HTML to link back to the index page is `Index`. The link uses the same structure as the link to the about page shown above.

4. Templates and Static Media

In this chapter, we'll be introducing the Django template engine, as well as showing how to serve *static media* files which can be integrated within your app's webpages.

4.1 Using Templates

Up until this point, you have plugged a few things together to create a Django powered webpage. This is coupled to a view, which is in turn coupled with a series of URL mappings. Here we will delve into how to combine *templates* into the mix.

The layout from page to page within a website is often the same. Whether you see a common header or footer on a website's pages, the [repetition of page layouts](#) aids users with navigation, promotes organisation of the website and reinforces a sense of continuity. [Django provides templates](#) to make it easier for developers to achieve this design goal, as well as separating application logic (code within your views) from presentational concerns (look and feel of your app). In this chapter, you'll create a basic template which will be used to create a HTML page. This template will then be dispatched via a Django view. In the [chapter concerning databases and models](#), we will take this a step further by using templates in conjunction with models to dispatch dynamically generated data.



Summary: What is a Template?

In the world of Django, think of a *template* as the scaffolding that is required to build a complete HTML webpage. A template contains the *static parts* of a webpage (that is, parts that never change), complete with special syntax (or *template tags*) which can be overridden and replaced with *dynamic content* that your Django app's views can replace to produce a final HTML response.

Configuring the Templates Directory

To get templates up and running with your Django app, you'll need to create a directory in which template files are stored.

In your Django project's directory (e.g. `<workspace>/tango_with_django_project/`), create a new directory called `templates`. Within the new `templates` directory, create another directory called `rango`. This means that the path `<workspace>/tango_with_django_project/templates/rango/` will be the location in which we will store templates associated with our `rango` application.



Keep your Templates Organised

It's good practice to separate out your templates into subdirectories for each app you have. This is why we've created a `rango` directory within our `templates` directory. If you package your app up to distribute to other developers, it'll be much easier to know which templates belong to which app!

To tell the Django project where templates will be stored, open your project's `settings.py` file. Next, locate the `TEMPLATES` data structure. By default, when you create a new Django 1.9 project, it will look like the following.

```
TEMPLATES = [
    {
        'BACKEND': 'django.template.backends.django.DjangoTemplates',
        'DIRS': [],
        'APP_DIRS': True,
        'OPTIONS': {
            'context_processors': [
                'django.template.context_processors.debug',
                'django.template.context_processors.request',
                'django.contrib.auth.context_processors.auth',
                'django.contrib.messages.context_processors.messages',
            ],
        },
    },
]
```

What we need to do to tell Django where our templates are stored is modify the `DIRS` list, and is set to an empty list by default. Change the dictionary key/value pair to look like the following.

```
'DIRS': ['<workspace>/tango_with_django_project/templates']
```

Note that you are *required to use absolute paths* to locate the templates directory. If you are collaborating with team members or working on different computers, then this will become a problem. You'll have different usernames and different drive structures, meaning the paths to the `<workspace>` directory will be different. One solution would be to the path for each different configuration, for example:

```
'DIRS': [ '/Users/leifos/templates',  
         '/Users/maxwelld90/templates',  
         '/Users/clueless_noob/templates', ]
```

However, there are a number of problems with this. First you have to add in the path for each setting, each time. Second, if you are running the app on different operating systems the black slashes have to be constructed differently (see the warning below).



Don't hard code Paths!

The road to Hell is paved with hard coded paths. [Hard-coding paths](#) is a [software engineering anti-pattern](#), and will make your project [less portable](#) - meaning that when you run it on another computer, it probably won't work!

Dynamic Paths

A better solution is to make use of built-in Python functions to work out the path of your templates directory automatically. This way, an absolute path can be obtained regardless of where you place your Django project's code. This in turn means that your project becomes more *portable*.

At the top of your `settings.py` file, there is a variable called `BASE_DIR`. This variable stores the path to the directory in which your project's `settings.py` module is contained. This is obtained by using the special Python `__file__` attribute, which is [set to the absolute path of your settings module](#). The call to `os.path.dirname()` then provides the reference to the absolute path of the directory containing the `settings.py` module. Calling `os.path.dirname()` again removes another layer, so that `BASE_DIR` contains `<workspace>/tango_with_django_project/`. You can see this process in action, if you are curious, by adding the following lines to your `settings.py` file.

```
print(__file__)  
print(os.path.dirname(__file__))  
print(os.path.dirname(os.path.dirname(__file__)))
```

Having access to the value of `BASE_DIR` makes it easy for you to reference other aspects of your Django project. As such, we can now create a new variable called `TEMPLATE_DIR` that will reference your new templates directory. We can make use of the `os.path.join()` function to join up multiple paths, leading to a variable definition like the example below.

```
TEMPLATE_DIR = os.path.join(BASE_DIR, 'templates')
```

Here we make use of `os.path.join()` to mash together the `BASE_DIR` variable and `'templates'`, which would yield `<workspace>/tango_with_django_project/templates/`. This means we can then use our new `TEMPLATE_DIR` variable to replace the hard coded path we defined earlier in `TEMPLATES`. Update the `DIRS` key/value pairing to look like the following.

```
'DIRS': [TEMPLATE_DIR, ]
```



Why TEMPLATE_DIR?

You've created a new variable called `TEMPLATE_DIR` at the top of your `settings.py` file because it's easier to access should you ever need to change it. For more complex Django projects, the `DIRS` list allows you to specify more than one template directory - but for this book, one location is sufficient to get everything working.



Concatenating Paths

When concatenating system paths together, always use `os.path.join()`. Using this built-in function ensures that the correct path separators are used. On a UNIX operating system (or derivative of), forward slashes (`/`) would be used to separate directories, whereas a Windows operating system would use backward slashes (`\`). If you manually append slashes to paths, you may end up with path errors when attempting to run your code on a different operating system, thus reducing your project's portability.

Adding a Template

With your template directory and path now set up, create a file called `index.html` and place it in the `templates/rango/` directory. Within this new file, add the following HTML code.

```
<!DOCTYPE html>
<html>

  <head>
    <title>Rango</title>
  </head>

  <body>
    <h1>Rango says...</h1>
    <div>
      hey there partner! <br />
      <strong>{{ boldmessage }}</strong> <br />
    </div>
    <div>
      <a href="/rango/about/">About</a> <br />
    </div>
  </body>

</html>
```

From this HTML code, it should be clear that a simple HTML page is going to be generated that greets a user with a *hello world* message. You might also notice some non-HTML in the form of `{{ boldmessage }}`. This is a *Django template variable*. We can set values to these variables so they are replaced with whatever we want when the template is rendered. We'll get to that in a moment.

To use this template, we need to reconfigure the `index()` view that we created earlier. Instead of dispatching a simple response, we will change the view to dispatch our template.

In `rango/views.py`, check to see if the following `import` statement exists at the top of the file. If it is not present, add it.

```
from django.shortcuts import render
```

You can then update the `index()` view function as follows. Check out the inline commentary to see what each line does.

```
def index(request):
    # Construct a dictionary to pass to the template engine as its context.
    # Note the key boldmessage is the same as {{ boldmessage }} in the template!
    context_dict = {'boldmessage': "Crunchy, creamy, cookie, candy, cupcake!"}

    # Return a rendered response to send to the client.
    # We make use of the shortcut function to make our lives easier.
    # Note that the first parameter is the template we wish to use.
    return render(request, 'rango/index.html', context=context_dict)
```

First, we construct a dictionary of key/value pairs that we want to use within the template. Then, we call the `render()` helper function. This function takes as input the user's request, the template filename, and the context dictionary. The `render()` function will take this data and mash it together with the template to produce a complete HTML page. This is then returned and dispatched to the user's web browser.




What is the Template Context?

When a template file is loaded with the Django templating system, a *template context* is created. In simple terms, a template context is essentially a Python dictionary that maps template variable names with Python variables. In the template we created earlier, we included a template variable name called `boldmessage`. In our updated `index(request)` view example, the string `Crunchy, creamy, cookie, candy, cupcake!` is mapped to template variable `boldmessage`. The string `Crunchy, creamy, cookie, candy, cupcake!` therefore replaces *any* instance of `{{ boldmessage }}` within the template.

Now that you have updated the view to employ the use of your template, start the Django development server and visit `http://127.0.0.1:8000/rango/`. You should see your simple HTML template rendered, just like the [example screenshot shown below](#).

If you don't, read the error message presented to see what the problem is, and then double check all the changes that you have made. One of the most common issues people have with templates is that the path is set incorrectly in `settings.py`. Sometimes it's worth adding a `print` statement to `settings.py` to report the `BASE_DIR` and `TEMPLATE_DIR` to make sure everything is correct.

This example demonstrates how to use templates within your views. However, we have only touched upon a fraction of the functionality provided by the Django templating engine. We will use templates in more sophisticated ways as you progress through this book. In the meantime, you can find out more about [templates from the official Django documentation](#).



What you should see when your first template is working correctly. Note the bold text - `Crunchy, creamy, cookie, candy, cupcake!` - which originates from the view, but is rendered in the template.

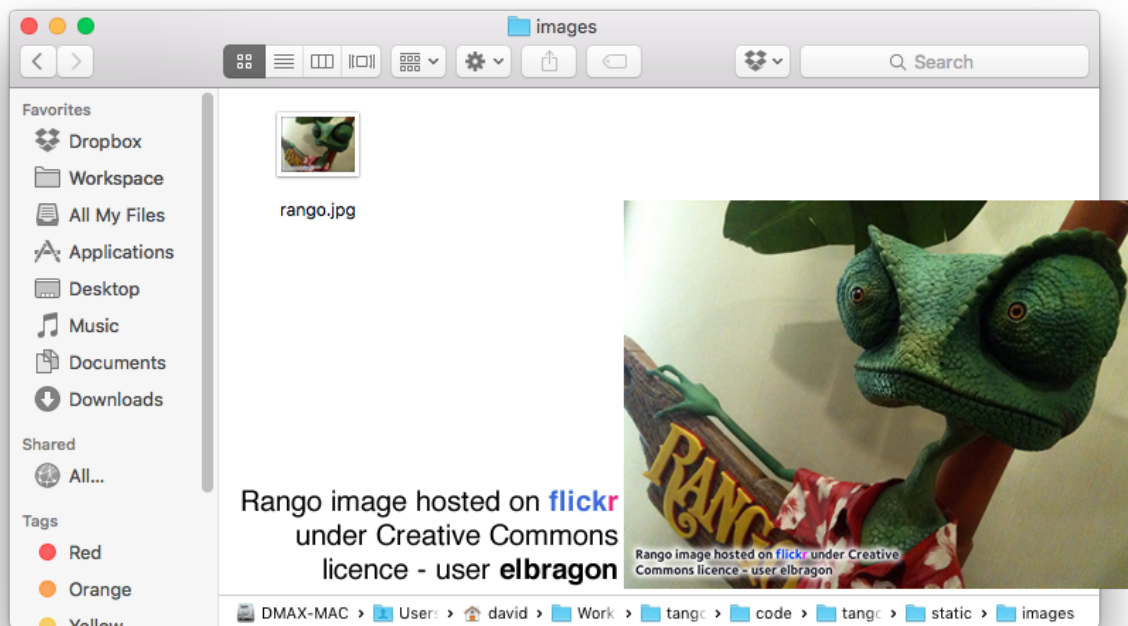
4.2 Serving Static Media

While you've got templates working, your Rango app is admittedly looking a bit plain right now - there's no styling or imagery. We can add references to other files in our HTML template such as *Cascading Style Sheets (CSS)*, *JavaScript* and images to improve the show. These are called *static files*, because they are not generated dynamically by a Web server; they are simply sent as is to a client's Web browser. This section shows you how to set Django up to serve static files, and shows you how to include an image within your simple template.

Configuring the Static Media Directory

To start, you will need to set up a directory in which static media files are stored. In your project directory (e.g. `<workspace>/tango_with_django_project/`), create a new directory called `static` and a new directory called `images` inside `static`. Check that the new `static` directory is at the same level as the `templates` directory you created earlier in this chapter.

Next, place an image inside the `images` directory. As shown in below, we chose a picture of [the chameleon Rango](#) - a fitting mascot, if ever there was one.



Rango the chameleon within our `static/images` media directory.

Just like the `templates` directory we created earlier, we need to tell Django about our new `static` directory. To do this, we once again need to edit our project's `settings.py` module. Within this file, we need to add a new variable pointing to our `static` directory, and a data structure that Django can parse to work out where our new directory is.

First of all, create a variable called `STATIC_DIR` at the top of `settings.py`, preferably underneath `BASE_DIR` and `TEMPLATES_DIR` to keep your paths all in the same place. `STATIC_DIR` should make use of the same `os.path.join` trick - but point to `static` this time around, just as shown below.

```
STATIC_DIR = os.path.join(BASE_DIR, 'static')
```

This will provide an absolute path to the location `<workspace>/tango_with_django_project/static/`. Once this variable has been created, we then need to create a new data structure called

STATICFILES_DIRS. This is essentially a list of paths with which Django can expect to find static files that can be served. By default, this list does not exist - **check** it doesn't before you create it. If you define it twice, you can start to confuse Django - and yourself.

For this book, we're only going to be using one location to store our project's static files - the path defined in STATIC_DIR. As such, we can simply set up STATICFILES_DIRS with the following.

```
STATICFILES_DIRS = [STATIC_DIR, ]
```



Keep settings.py Tidy!

It's in your best interests to keep your settings.py module tidy and in good order. Don't just put things in random places; keep it organised. Keep your DIRS variables at the top of the module so they are easy to find, and place STATICFILES_DIRS in the portion of the module responsible for static media (close to the bottom). When you come back to edit the file later, it'll be easier for you or other collaborators to find the necessary variables.

Finally, check that the STATIC_URL variable is defined within your settings.py module. If it is not, then define it as shown below. Note that this variable by default in Django 1.9 appears close to the end of the module, so you may have to scroll down to find it.

```
STATIC_URL = '/static/'
```

With everything required now entered, what does it all mean? Put simply, the first two variables STATIC_DIR and STATICFILES_DIRS refers to the locations on your computer where static files are stored. The final variable STATIC_URL then allows us to specify the URL with which static files can be accessed when we run our Django development server. For example, with STATIC_URL set to /static/, we would be able to access static content at `http://127.0.0.1:8000/static/`. *Think of the first two variables as server-side locations, and the third variable as the location with which clients can access static content.*



Test your Configuration

As a small exercise, test to see if everything is working correctly. Try and view the rango.jpg image in your browser when the Django development server is running. If your STATIC_URL is set to /static/ and rango.jpg can be found at images/rango.jpg, what is the URL you enter into your Web browser's window?

Don't proceed until you are sure your configuration is working!



Don't Forget the Slashes!

When setting `STATIC_URL`, check that you end the URL you specify with a forward slash (e.g. `/static/`, not `/static`). As per the [official Django documentation](#), not doing so can open you up to a world of pain. The extra slash at the end ensures that the root of the URL (e.g. `/static/`) is separated from the static content you want to serve (e.g. `images/rango.jpg`).



Serving Static Content

While using the Django development server to serve your static media files is fine for a development environment, it's highly unsuitable for a production environment. The [official Django documentation on deployment](#) provides further information about deploying static files in a production environment. We'll look at this issue in more detail however when we deploy Rango.

Static Media Files and Templates

Now that you have your Django project set up to handle static files, you can now make use of these files within your templates to improve their appearance and add additional functionality.

To demonstrate how to include static files, open up the `index.html` templates you created earlier, located in the `<workspace>/templates/rango/` directory. Modify the HTML source code as follows. The two lines that we add are shown with a HTML comment next to them for easy identification.

```
<!DOCTYPE html>

{% load staticfiles %} <!-- New line -->

<html>
  <head>
    <title>Rango</title>
  </head>

  <body>
    <h1>Rango says...</h1>

    <div>
      hey there partner! <br />
      <strong>{{ boldmessage }}</strong> <br />
    </div>
```



```
<div>
  <a href="/rango/about/">About</a><br />
   <!-- New line -->
</div>
</body>

</html>
```


The first new line added (i.e. `{% load staticfiles %}`) informs Django's template engine that we will be using static files with the template. This then enables us to access the media in the static directories via the use of the `static` [template tag](#). This indicates to Django that we wish to show the image located in the static media directory called `images/rango.jpg`. Template tags are denoted by curly brackets (e.g. `{% % }`), and calling `static` will combine the URL specified in `STATIC_URL` with `images/rango.jpg` to yield `/static/images/rango.jpg`. The HTML generated as a result would be:

```

```

If for some reason the image cannot be loaded, it is always a good idea to specify an alternative text tagline. This is what the `alt` attribute provides inside the `img` tag.

With these minor changes in place, start the Django development server once more and navigate to `http://127.0.0.1:8000/rango`. If everything has been done correctly, you will see a Webpage that looks similar to the [screenshot shown below](#).



Our first Rango template, complete with a picture of Rango the chameleon.



Templates and `<!DOCTYPE>`

When creating the HTML templates, always ensure that the `DOCTYPE` declaration appears on the **first line**. If you put the `{% load staticfiles %}` first then whitespace is added before the `DOCTYPE` declaration which can lead to your HTML code to **fail validation**.



Loading other Static Files

The `{% static %}` template tag call can be used whenever you wish to reference static files within a template. The code example below demonstrates how you could include JavaScript, CSS and images into your templates with correct HTML markup.

```
<!DOCTYPE html>
{% load staticfiles %}

<html>

  <head>
    <title>Rango</title>
    <!-- CSS -->
    <link rel="stylesheet" href="{% static "css/base.css" %}" />
    <!-- JavaScript -->
    <script src="{% static "js/jquery.js" %}"></script>
  </head>

  <body>
    <!-- Image -->
    
  </body>

</html>
```

Static files you reference will obviously need to be present within your static directory. If a requested file is not present or you have referenced it incorrectly, the console output provided by Django's development server will show a [HTTP 404 error](#). Try referencing a non-existent file and see what happens. Looking at the output snippet below, notice how the last entry's HTTP status code is 404.

```
[10/Apr/2016 15:12:48] "GET /rango/ HTTP/1.1" 200 374
[10/Apr/2016 15:12:48] "GET /static/images/rango.jpg HTTP/1.1" 304 0
[10/Apr/2016 15:12:52] "GET /static/images/not-here.jpg HTTP/1.1" 404 0
```

For further information about including static media you can read through the official [Django documentation on working with static files in templates](#).

4.3 Basic Workflow

With the chapter complete, you should now know how to setup and create templates, use templates within your views, setup and use the Django development server to serve static media files, *and*

include images within your templates. We've covered quite a lot!

Creating a template and integrating it within a Django view is a key concept for you to understand. It takes several steps, but will become second nature to you after a few attempts.

1. First, create the template you wish to use and save it within the `templates` directory you specified in your project's `settings.py` module. You may wish to use Django template variables (e.g. `{{ variable_name }}`) or [template tags](#) within your template. You'll be able to replace these with whatever you like within the corresponding view.
2. Find or create a new view within an application's `views.py` file.
3. Add your view specific logic (if you have any) to the view. For example, this may involve extracting data from a database and storing it within a list.
4. Within the view, construct a dictionary object which you can pass to the template engine as part of the [template's context](#).
5. Make use of the `render()` helper function to generate the rendered response. Ensure you reference the request, then the template file, followed by the context dictionary.
6. If you haven't already done so, map the view to a URL by modifying your project's `urls.py` file and the application specific `urls.py` file if you have one.

The steps involved for getting a static media file onto one of your pages is another important process you should be familiar with. Check out the steps below on how to do this.

1. Take the static media file you wish to use and place it within your project's `static` directory. This is the directory you specify in your project's `STATICFILES_DIRS` list within `settings.py`.
2. Add a reference to the static media file to a template. For example, an image would be inserted into an HTML page through the use of the `` tag.
3. Remember to use the `{% load staticfiles %}` and `{% static "<filename%" %}` commands within the template to access the static files. Replace `<filename>` with the path to the image or resource you wish to reference. **Whenever you wish to refer to a static file, use the static template tag!**



Exercises

Give the following exercises a go to reinforce what you've learnt from this chapter.

- Convert the about page to use a template as well, using a template called `about.html`.
- Within the new `about.html` template, add a picture stored within your project's static files.
- On the about page, include a line that says, `This tutorial has been put together by <your-name>`.

5. Models and Databases


When you think of databases, you will usually think of the *Structured Query Language (SQL)*, the common means with which we query the database for the data we require. With Django, querying an underlying database - which can store all sorts of data, such as your website's user details - is taken care of by the *object relational mapper (ORM)*. In essence, data stored within a database table can be encapsulated within a *model*. A model is a Python object that describes your database table's data. Instead of directly working on the database via SQL, you only need to manipulate the corresponding Python model object.

This chapter walks you through the basics of data management with Django and its ORM. You'll find it's incredibly easy to add, modify and delete data within your app's underlying database, and how straightforward it is to get data from the database to the Web browsers of your users.

5.1 Rango's Requirements

Before we get started, let's go over the data requirements for the Rango app that we are developing. Full requirements for the application are [provided in detail earlier on](#), but to refresh your memory, let's quickly summarise our client's requirements.

- Rango is essentially a *web page directory* - a site containing links to other websites.
- There are a number of different *webpage categories* with each category housing a number of links. [We assumed in the overview chapter](#) that this is a one-to-many relationship. Check out the [Entity Relationship diagram below](#).
- A category has a name, a number of visits, and a number of likes.
- A page refers to a category, has a title, URL and a number of views.



The Entity Relationship Diagram of Rango's two main entities.

5.2 Telling Django about Your Database

Before we can create any models, we need to set up our database with Django. In Django 1.9, a `DATABASES` variable is automatically created in your `settings.py` module when you set up a new project. It'll look similar to the following example.

```
DATABASES = {  
    'default': {  
        'ENGINE': 'django.db.backends.sqlite3',  
        'NAME': os.path.join(BASE_DIR, 'db.sqlite3'),  
    }  
}
```

We can pretty much leave this as is for our Rango app. You can see a `default` database that is powered by a lightweight database engine, [SQLite](#) (see the `ENGINE` option). The `NAME` entry for this database is the path to the database file, which is by default `db.sqlite3` in the root of your Django project.



Git Top Tip

If you are using Git, you might be tempted to add and commit the database file. This is not a good idea because if you are working on your app with other people, they are likely to change the database and this will cause endless conflicts.

Instead, add `db.sqlite3` to your `.gitignore` file so that it won't be added when you `git commit` and `git push`. You can also do this for other files like `*.pyc` and machine specific files.



Using other Database Engines

The Django database framework has been created to cater for a variety of different database backends, such as [PostgreSQL](#), [MySQL](#) and [Microsoft's SQL Server](#). For other database engines, other keys like `USER`, `PASSWORD`, `HOST` and `PORT` exist for you to configure the database with Django.

While we don't cover how to use other database engines in this book, there are guides online which show you how to do this. A good starting point is the [official Django documentation](#).

Note that SQLite is sufficient for demonstrating the functionality of the Django ORM. When you find your app has become viral and has accumulated thousands of users, you may want to consider [switching the database backend to something more robust](#).

5.3 Creating Models

With your database configured in `settings.py`, let's create the two initial data models for the Rango application. Models for a Django app are stored in the respective `models.py` module. This means that for Rango, models are stored within `rango/models.py`.

For the models themselves, we will create two classes - one class representing each model. Both must [inherit](#) from the `Model` base class, `django.db.models.Model`. The two Python classes will be the definitions for models representing *categories* and *pages*. Define the `Category` and `Page` model as follows.

```
class Category(models.Model):
    name = models.CharField(max_length=128, unique=True)

    def __str__(self): # For Python 2, use __unicode__ too
        return self.name

class Page(models.Model):
    category = models.ForeignKey(Category)
    title = models.CharField(max_length=128)
    url = models.URLField()
    views = models.IntegerField(default=0)

    def __str__(self): # For Python 2, use __unicode__ too
        return self.title
```



Check import Statements

At the top of the `models.py` module, you should see `from django.db import models`. If you don't see it, add it in.



`__str__()` or `__unicode__()`?

The `__str__()` and `__unicode__()` methods in Python generate a string representation of the class (similar to the `toString()` method in Java). In Python 2.x, strings are represented in ASCII format in the `__str__()` method. If you want [Unicode support](#), then you need to also implement the `__unicode__()` method.

In Python 3.x, strings are Unicode by default - so you only need to implement the `__str__()` method.

When you define a model, you need to specify the list of fields and their associated types, along with any required or optional parameters. By default, all models have an auto-increment integer x field called `id` which is automatically assigned and acts a primary key.

Django provides a [comprehensive series of built-in field types](#). Some of the most commonly used are detailed below.

- `CharField`, a field for storing character data (e.g. strings). Specify `max_length` to provide a maximum number of characters the field can store.
- `URLField`, much like a `CharField`, but designed for storing resource URLs. You may also specify a `max_length` parameter.
- `IntegerField`, which stores integers.
- `DateTimeField`, which stores a Python `datetime.date` object.



Other Field Types

Check out the [Django documentation on model fields](#) for a full listing of the Django field types you can use, along with details on the required and optional parameters that each has.

For each field, you can specify the `unique` attribute. If set to `True`, only one instance of a particular value in that field may exist throughout the entire database model. For example, take a look at our `Category` model defined above. The field name has been set to `unique` - thus every category name must be unique. This means that you can use the field like a primary key.

You can also specify additional attributes for each field, such as stating a default value with the syntax `default='value'`, and whether the value for a field can be blank (or `NULL`) (`null=True`) or not (`null=False`).

Django provides three mechanisms to related models in the database, which are:

- `ForeignKey`, a field type that allows us to create a [one-to-many relationship](#);
- `OneToOneField`, a field type that allows us to define a strict [one-to-one relationship](#); and
- `ManyToManyField`, a field type which allows us to define a [many-to-many relationship](#).

From our model examples above, the field `category` in model `Page` is of type `ForeignKey`. This allows us to create a one-to-many relationship with model/table `Category`, which is specified as an argument to the field's constructor.

Finally, it is good practice to implement the `__str__()` and/or `__unicode__()` methods. Without this method implemented when you go to print the object, it will show as `<Category: Category object>`. This isn't very useful when debugging or accessing the object - instead the code above will print, for example, `<Category: Python>` for the Python category. It is also helpful when we go to use the Admin Interface because Django will display the string representation of the object.

5.4 Creating and Migrating the Database

With our models defined in `models.py`, we can now let Django work its magic and create the tables in the underlying database. Django provides what is called a *migration tool* to help us set up and update the database to reflect any changes to your models. For example, if you were to add a new field then you can use the migration tools to update the database.

Setting up

First of all, the database must be *initialised*. This means creating it and all the associated tables within it so that data can then be stored within it. To do this, you must open a terminal or command prompt, and navigate to your project's root directory - where `manage.py` is stored. Run the following command.

```
$ python manage.py migrate
```

Operations to perform:

Apply all migrations: admin, contenttypes, auth, sessions

Running migrations:

```
Rendering model states... DONE
Applying contenttypes.0001_initial... OK
Applying auth.0001_initial... OK
Applying admin.0001_initial... OK
Applying admin.0002_logentry_remove_auto_add... OK
Applying contenttypes.0002_remove_content_type_name... OK
Applying auth.0002_alter_permission_name_max_length... OK
Applying auth.0003_alter_user_email_max_length... OK
Applying auth.0004_alter_user_username_opts... OK
Applying auth.0005_alter_user_last_login_null... OK
Applying auth.0006_require_contenttypes_0002... OK
Applying auth.0007_alter_validators_add_error_messages... OK
Applying sessions.0001_initial... OK
```

All apps that are installed in your Django project (check `INSTALLED_APPS` in `settings.py`) will update their database representations with this command. After this command is issued, you should then see a `db.sqlite3` file in your Django project's root.

Next, create a superuser to manage the database. Run the following command.

```
$ python manage.py createsuperuser
```

The superuser account will be used to access the Django admin interface, used later on in this chapter. Enter a username for the account, e-mail address and provide a password when prompted. Once completed, the script should finish successfully. Make sure you take a note of the username and password for your superuser account.

Creating and Updating Models/Tables

Whenever you make changes to your app's models, you need to *register* the changes via the `makemigrations` command in `manage.py`. Specifying the `rango` app as our target, we then issue the following command from our Django project's root directory.

```
$ python manage.py makemigrations rango
```

```
Migrations for 'rango':
```

```
0001_initial.py:
```

- Create model Category
- Create model Page

Upon the completion of this command, check the `rango/migrations` directory to see that a Python script has been created. It's called `0001_initial.py`, which contains all the necessary details to create your database schema at that particular migration.



Checking the Underlying SQL

If you want to check out the underlying SQL that the Django ORM issues to the database engine for a given migration, you can issue the following command.

```
$ python manage.py sqlmigrate rango 0001
```

In this example, `rango` is the name of your app, and `0001` is the migration you wish to view the SQL code for. Doing this allows you to get a better understanding of what exactly is going on at the database layer, such as what tables are created. You may find for complex database schemas including a many-to-many relationship that additional tables are created for you.

After you have created migrations for your app, you need to commit them to the database. Do so by once again issuing the `migrate` command.

```
$ python manage.py migrate
```

Operations to perform:

Apply all migrations: admin, rango, contenttypes, auth, sessions

Running migrations:

Rendering model states... DONE

Applying rango.0001_initial... OK

This output confirms that the database tables have been created in your database, and you are good to go.

However, you may have noticed that our `Category` model is currently lacking some fields that were specified in [Rango's requirements](#). Don't worry about this, as these will be added in later, allowing you to go through the migration process again.

5.5 Django Models and the Shell

Before we turn our attention to demonstrating the Django admin interface, it's worth noting that you can interact with Django models directly from the Django shell - a very useful tool for debugging purposes. We'll demonstrate how to create a `Category` instance using this method.

To access the shell, we need to call `manage.py` from within your Django project's root directory once more. Run the following command.

```
$ python manage.py shell
```

This will start an instance of the Python interpreter and load in your project's settings for you. You can then interact with the models, with the following terminal session demonstrating this functionality. Check out the inline commentary that we added to see what each command achieves.

```
# Import the Category model from the Rango application
>>> from rango.models import Category

# Show all the current categories
>>> print(Category.objects.all())
[] # Returns an empty list (no categories have been defined!)

# Create a new category object, and save it to the database.
>>> c = Category(name="Test")
>>> c.save()

# Now list all the category objects stored once more.
>>> print(Category.objects.all())
[<Category: test>] # We now have a category called 'Test' saved in the database!
```

```
# Quit the Django shell.  
>>> quit()
```

In the example, we first import the model that we want to manipulate. We then print out all the existing categories. As our underlying `Category` table is empty, an empty list is returned. Then we create and save a `Category`, before printing out all the categories again. This second print then shows the new `Category` just added. Note the name, `Test` appears in the second print - this is your `__str__()` or `__unicode__()` method at work!



Complete the Official Tutorial

The example above is only a very basic taster on database related activities you can perform in the Django shell. If you have not done so already, it's now a good time to complete [part two of the official Django Tutorial to learn more about interacting with models](#). Also check out the [official Django documentation on the list of available commands](#) for working with models.

5.6 Configuring the Admin Interface

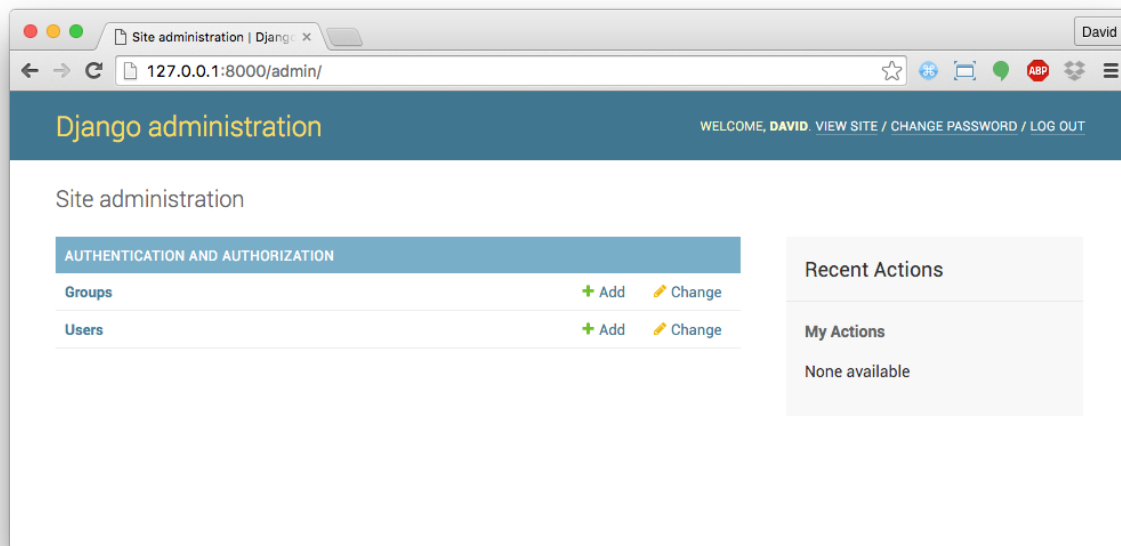
One of the eye-catching features of Django is the built-in, Web-based administrative interface that allows you to browse and edit data represented as model instances (from the corresponding database tables).

Setting everything up is relatively straightforward. In your project's `settings.py` module, you will notice that one of the preinstalled apps (within the `INSTALLED_APPS` list) is `django.contrib.admin`. Furthermore, there is a `urlpatterns` that matches `admin/` within your project's `urls.py` module.

By default, things are pretty much ready to go. Start the Django development server in the usual way with the following command.

```
$ python manage.py runserver
```

Navigate your Web browser to `http://127.0.0.1:8000/admin/`. You are then presented with a login prompt. Using the username and password you created previously with the `python manage.py createsuperuser` command, login. You are then presented with an interface looking [similar to the one shown below](#).



The Django admin interface, sans Rango models.

While this looks good, we are missing the `Category` and `Page` models that were defined for the Rango app. To include these models, we need to give Rango some help.


To do this, open the file `rango/admin.py`. With an `include` statement already present, modify the module so that you register each class you want to include. The example below registers both the `Category` and `Page` class to the admin interface.

```
from django.contrib import admin
from rango.models import Category, Page
```

```
admin.site.register(Category)
admin.site.register(Page)
```

Adding further classes which may be created in the future is as simple as adding another call to the `admin.site.register()` method.

With these changes saved, restart the Django development server and revisit the admin interface at `http://127.0.0.1:8000/admin/`. You will now see the `Category` and `Page` models, as shown below.



The Django admin interface, complete with Rango models.

Try clicking the Categorys link within the Rango section. From here, you should see the test category that we created earlier via the Django shell.



Experiment with the Admin Interface

You'll be using the admin interface quite a bit to verify data is stored correctly as you develop the Rango app. Experiment with it, and see how it all works. The interface is self-explanatory and straightforward to understand.

Delete the test category that was previously created. We'll be populating the database shortly with more example data.



User Management

The Django admin interface is your port of call for user management, through the Authentication and Authorisation section. Here, you can create, modify and delete user accounts, all with varying privilege levels.



Plural vs. Singular Spellings

Note the typo within the admin interface (Categorys, not Categories). This typo can be fixed by adding a nested Meta class into your model definitions with the `verbose_name_plural` attribute. Check out a modified version of the Category model below for an example, and [Django's official documentation on models](#) for more information about what can be stored within the Meta class.

```
class Category(models.Model):
    name = models.CharField(max_length=128, unique=True)

    class Meta:
        verbose_name_plural = 'categories'

    def __str__(self):
        return self.name
```



Expanding admin.py

It should be noted that the example `admin.py` module for your Rango app is the most simple, functional example available. However you can customise the Admin interface in a number of ways. Check out the [official Django documentation on the admin interface](#) for more information if you're interested.

5.7 Creating a Population Script

Entering test data into your database tends to be a hassle. Many developers will add in some bogus test data by randomly hitting keys, just like `wTFzm8j3z7`. Rather than do this, it is better to write a script so that you and your collaborators works from the same tests data. Furthermore, this approach would guarantee that you have useful and pseudo realistic data rather than random junk. It's therefore good practice to create what we call a *population script* for your app. This script is designed to automatically populate your database with test data for you

To create a population script for Rango, start by creating a new Python module within your Django project's root directory (e.g. `<workspace>/tango_with_django_project/`). Create the `populate_rango.py` file and add the following code.

```
1 import os
2 os.environ.setdefault('DJANGO_SETTINGS_MODULE',
3                       'tango_with_django_project.settings')
4
5 import django
6 django.setup()
7 from rango.models import Category, Page
8
9 def populate():
10     # First, we will create lists of dictionaries containing the pages
11     # we want to add into each category.
12     # Then we will create a dictionary of dictionaries for our categories.
13     # This might seem a little bit confusing, but it allows us to iterate
14     # through each data structure, and add the data to our models.
15
16     python_pages = [
17         {"title": "Official Python Tutorial",
18          "url": "http://docs.python.org/2/tutorial/"},
19         {"title": "How to Think like a Computer Scientist",
20          "url": "http://www.greenteapress.com/thinkpython/"},
21         {"title": "Learn Python in 10 Minutes",
22          "url": "http://www.korokithakis.net/tutorials/python/" } ]
23
24     django_pages = [
25         {"title": "Official Django Tutorial",
26          "url": "https://docs.djangoproject.com/en/1.9/intro/tutorial01/"},
27         {"title": "Django Rocks",
28          "url": "http://www.djangorocks.com/"},
29         {"title": "How to Tango with Django",
30          "url": "http://www.tangowithdjango.com/" } ]
31
32     other_pages = [
33         {"title": "Bottle",
34          "url": "http://bottlepy.org/docs/dev/"},
35         {"title": "Flask",
36          "url": "http://flask.pocoo.org" } ]
37
38     cats = {"Python": {"pages": python_pages},
39            "Django": {"pages": django_pages},
40            "Other Frameworks": {"pages": other_pages} }
41
42     # If you want to add more categories or pages,
```



```

43     # add them to the dictionaries above.
44
45     # The code below goes through the cats dictionary, then adds each category,
46     # and then adds all the associated pages for that category.
47     # if you are using Python 2.x then use cats.iteritems() see
48     # http://docs.quantifiedcode.com/python-anti-patterns/readability/
49     # for more information about how to iterate over a dictionary properly.
50
51     for cat, cat_data in cats.items():
52         c = add_cat(cat)
53         for p in cat_data["pages"]:
54             add_page(c, p["title"], p["url"])
55
56     # Print out the categories we have added.
57     for c in Category.objects.all():
58         for p in Page.objects.filter(category=c):
59             print("- {0} - {1}".format(str(c), str(p)))
60
61 def add_page(cat, title, url, views=0):
62     p = Page.objects.get_or_create(category=cat, title=title)[0]
63     p.url=url
64     p.views=views
65     p.save()
66     return p
67
68 def add_cat(name):
69     c = Category.objects.get_or_create(name=name)[0]
70     c.save()
71     return c
72
73 # Start execution here!
74 if __name__ == '__main__':
75     print("Starting Rango population script...")
76     populate()

```



Understand this Code!

To reiterate what we wrote earlier, don't simply copy, paste and leave. Add the code to your new module, and then step through line by line to work out what is going on. It'll help with your understanding.

Below we have provided explanations - hopefully you'll learn something new!

While this looks like a lot of code, what is going on is essentially a series of function calls to two small functions, `add_page()` and `add_cat()` defined towards the end of the module. Reading through the code, we find that execution starts at the *bottom* of the module - look at lines 75 and 76. This is because above this point, we define functions, these are not executed unless we call them. When the interpreter hits `if __name__ == '__main__':`, we call the `populate()` function.



What does `__name__ == '__main__'` Represent?

The `__name__ == '__main__'` trick is a useful one that allows a Python module to act as either a reusable module or a standalone Python script. Consider a reusable module as one that can be imported into other modules (e.g. through an `import` statement), while a standalone Python script would be executed from a terminal/Command Prompt by entering `python module.py`.

Code within a conditional `if __name__ == '__main__':` statement will therefore only be executed when the module is run as a standalone Python script. Importing the module will not run this code; any classes or functions will however be fully accessible to you.



Importing Models

When importing Django models, make sure you have imported your project's settings by importing `django` and setting the environment variable `DJANGO_SETTINGS_MODULE` to be your project's setting file, as demonstrated in lines 1 to 6 above. You then call `django.setup()` to import your Django project's settings.

If you don't do this crucial step, **an exception will be raised when you attempt to import your models as the necessary infrastructure has not yet been initialised.** This is why we import `Category` and `Page` *after* the settings have been loaded on line 8.

The `for` loop occupying lines 51-54 is responsible for the calling the `add_cat()` and `add_page()` functions repeatedly. These functions are in turn responsible for the creation of new categories and pages. `populate()` keeps tabs on categories that are created. As an example, a reference to a new category is stored in local variable `c` - check line 52 above. This is done as a `Page` requires a `Category` reference. After `add_cat()` and `add_page()` are called in `populate()`, the function concludes by looping through all new `Category` and associated `Page` objects, displaying their names on the terminal.



Creating Model Instances

We make use of the convenience `get_or_create()` method for creating model instances in the population script above. As we don't want to create duplicates of the same entry, we can use `get_or_create()` to check if the entry exists in the database for us. If it doesn't exist, the method creates it. It does, then a reference to the specific model instance is returned.

This helper method can remove a lot of repetitive code for us. Rather than doing this laborious check ourselves, we can make use of code that does exactly this for us.

The `get_or_create()` method returns a tuple of `(object, created)`. The first element `object` is a reference to the model instance that the `get_or_create()` method creates if the database entry was not found. The entry is created using the parameters you pass to the method - just like `category`, `title`, `url` and `views` in the example above. If the entry already exists in the database, the method simply returns the model instance corresponding to the entry. `created` is a boolean value; `True` is returned if `get_or_create()` had to create a model instance.

This explanation therefore means that the `[0]` at the end of our call to the `get_or_create()` returns the object reference only. Like most other programming language data structures, Python tuples use [zero-based numbering](#).

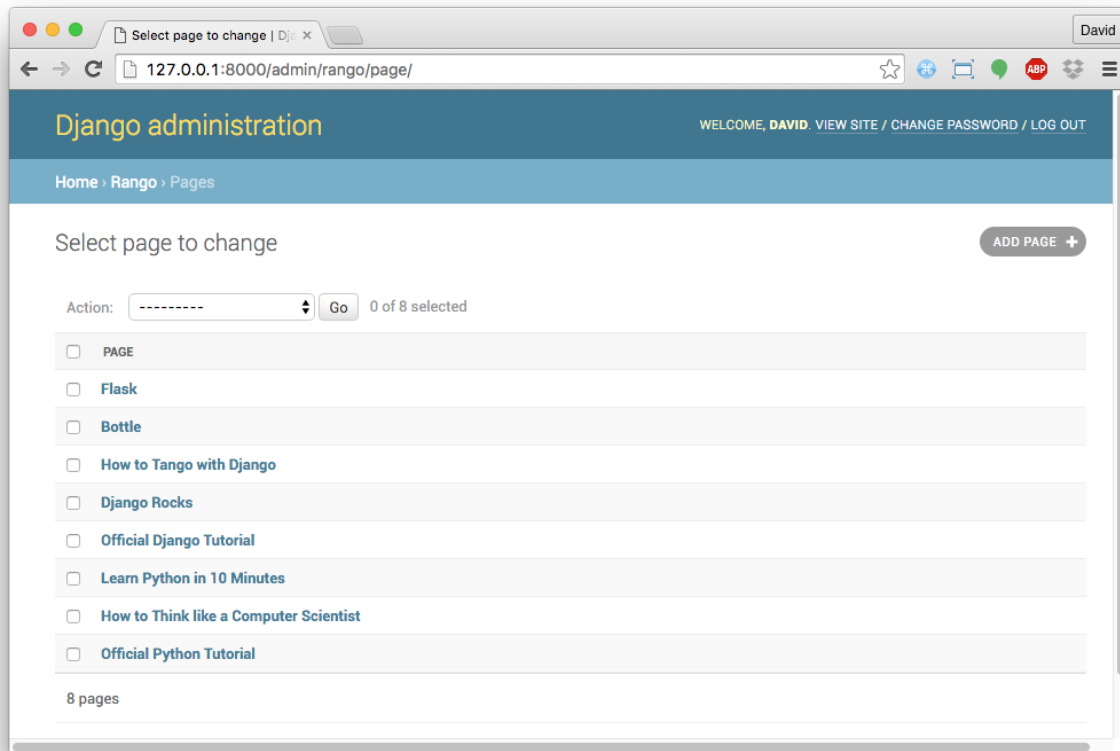
You can check out the [official Django documentation](#) for more information on the handy `get_or_create()` method.

When saved, you can then run your new populations script by changing the present working directory in a terminal to the Django project's root. It's then a simple case of executing the command `$ python populate_rango.py`. You should then see output similar to that shown below.

```
$ python populate_rango.py
```

```
Starting Rango population script...
- Python - Official Python Tutorial
- Python - How to Think like a Computer Scientist
- Python - Learn Python in 10 Minutes
- Django - Official Django Tutorial
- Django - Django Rocks
- Django - How to Tango with Django
- Other Frameworks - Bottle
- Other Frameworks - Flask
```

Next, verify that the population script actually populated the database. Restart the Django development server, navigate to the admin interface (at `http://127.0.0.1:8000/admin/`) and check that you have some new categories and pages. Do you see all the pages if you click Pages, like in the figure shown below?



The Django admin interface, showing the `Page` model populated with the new population script. Success!

While creating a population script may take time, you will save yourself time in the long run. When deploying your app elsewhere, running the population script after setting everything up means you can start demonstrating your app straight away. You'll also find it very handy when it comes to [unit testing your code](#).

5.8 Workflow: Model Setup

Now that we've covered the core principles of dealing with Django's ORM, now is a good time to summarise the processes involved in setting everything up. We've split the core tasks into separate sections for you. Check this section out when you need to quickly refresh your mind of the different steps.

Setting up your Database

With a new Django project, you should first [tell Django about the database you intend to use](#) (i.e. configure `DATABASES` in `settings.py`). You can also register any models in the `admin.py` module of your app to make them accessible via the admin interface.

Adding a Model

The workflow for adding models can be broken down into five steps.

1. First, create your new model(s) in your Django application's `models.py` file.
2. Update `admin.py` to include and register your new model(s).
3. Perform the migration `$ python manage.py makemigrations`.
4. Apply the changes `$ python manage.py migrate`. This will create the necessary infrastructure within the database for your new model(s).
5. Create/edit your population script for your new model(s).

Invariably, there will be times when you will have to delete your database. When this happens, run the following commands from the `manage.py` module.


1. `migrate` your database - this will set everything up in the new database. Ensure that your app is listed in the migrations that are committed. If it is not, run the `makemigrations <appname>` command, where `<appname>` is the name of your app.
2. Create a new administrative account with the `createsuperuser` command.



Exercises

Now that you've completed this chapter, try out these exercises to reinforce and practice what you have learnt. **Once again, note that the following chapters will have expected you to have completed these exercises!**

- Update the `Category` model to include the additional attributes `views` and `likes` where the `default` values for each are both zero (0).
- Make the migrations for your app, then migrate your database to commit the changes.
- Update your population script so that the Python category has 128 views and 64 likes, the Django category has 64 views and 32 likes, and the Other Frameworks category has 32 views and 16 likes.
- Delete and recreate your database, populating it with your updated population script.
- Complete parts [two](#) and [seven](#) of the official Django tutorial. These sections will reinforce what you've learnt on handling databases in Django, and show you additional techniques to customising the Django admin interface.
- Customise the admin interface. Change it in such a way so that when you view the `Page` model, the table displays the category, the name of the page and the `url` - just [like in the screenshot shown below](#).



The updated admin interface Page view, complete with columns for category and URL.



Exercise Hints

If you require some help or inspiration to get these exercises done, these hints will hopefully help you out.

- Modify the Category model by adding in the fields, view and likes as IntegerFields.
- Modify the add_cat function in the populate.py script, to take the views and likes. Once you get the Category c, then you can update the number of views with c.views, and similarly with likes. Don't forget to save() the instance!
- To customise the admin interface, you will need to edit rango/admin.py and create a PageAdmin class that inherits from admin.ModelAdmin.
- Within your new PageAdmin class, add list_display = ('title', 'category', 'url').
- Finally, register the PageAdmin class with Django's admin interface. You should modify the line admin.site.register(Page). Change it to admin.site.register(Page, PageAdmin) in Rango's admin.py file.



Tests

We have written a few tests to check if you have completed the exercises. To check your work so far, [download the tests.py script](#) from our [GitHub repository](#), and save it within your rango app directory.

To run the tests, issue the following command in the terminal or Command Prompt.

```
$ python manage.py test rango
```

If you are interested in learning about automated testing, now is a good time to check out the [chapter on testing](#). The chapter runs through some of the basics on testing that can perform in Django.

6. Automated Testing

It is good practice to get into the habit of writing and developing tests. A lot of software engineering is about writing and developing tests and test suites in order to ensure the software is robust. Of course, most of the time, we are too busy trying to build things to bother about making sure that they work. Or too arrogant to believe it would fail.

According to the [Django Tutorial](#), there are numerous reasons why you should include tests:

- Test will save you time: a change in a complex system can cause failures in unpredictable places.
- Tests don't just identify problems, they prevent them: tests show where the code is not meeting expectations.
- Tests make your code more attractive: “Code without tests is broken by design”, Jacob Kaplan-Moss, One of Django's original developers.
- Tests help teams work together: they make sure your team doesn't inadvertently break your code.

According to the [Python Guide] (<http://docs.python-guide.org/en/latest/writing/tests/>), there are a number of general rules you should try to follow when writing tests. Below are some main rules:

- Tests should focus on one small bit of functionality
- Tests should have a clear purpose
- Tests should be independent.
- Run your tests, before you code, and before your commit and push your code.
- Even better create a hook that tests code on push.
- Use long and descriptive names for tests.



Testing in Django

Currently this chapter provides the very basics of testing and follows a similar format to the [Django Tutorial](#), with some additional notes. We hope to expand this further in the future.

6.1 Running Tests

In built in Django is machinery to test the applications built. You can do this by issuing the following command:


```
$ python manage.py test rango
```

```
Creating test database for alias 'default'...
```

```
-----  
Ran 0 tests in 0.000s
```

```
OK
```

```
Destroying test database for alias 'default'...
```

This will run through the tests associated with the rango application. At the moment, nothing much happens. That is because you may have noticed the file `rango/tests.py` only contains an import statement. Every time you create an application, Django automatically creates such a file to encourage you to write tests.

From this output, you might also notice that a database called `default` is referred to. When you run tests, a temporary database is constructed, which your tests can populate, and perform operations on. This way your testing is performed independently of your live database.

Testing the models in Rango

Ok, lets create a test. In the `Category` model, we want to ensure that views are either zero or positive, because the number of views, let's say, can never be less than zero. To create a test for this we can put the following code into `rango/tests.py`:

```
from django.test import TestCase
from rango.models import Category

class CategoryMethodTests(TestCase):
    def test_ensure_views_are_positive(self):
        """
        ensure_views_are_positive should results True for categories
        where views are zero or positive
        """
        cat = Category(name='test', views=-1, likes=0)
        cat.save()
        self.assertEqual((cat.views >= 0), True)
```

The first thing you should notice, if you have not written tests before, is that we have to inherit from `TestCase`. The naming over the method in the class also follows a convention, all tests start with `test_` and they also contain some type of assertion, which is the test. Here we are checking if the values are equal, with the `assertEqual` method, but other types of assertions are also possible.

See the Python Documentation on unit tests, <https://docs.python.org/2/library/unittest.html> for other commands (i.e. `assertItemsEqual`, `assertListEqual`, `assertDictEqual`, etc). Django's testing machinery is derived from Python's but also provides a number of other asserts and specific test cases.

Now lets run test:

```
1 $ python manage.py test rango
2
3
4 Creating test database for alias 'default'...
5 F
6 =====
7 FAIL: test_ensure_views_are_positive (rango.tests.CategoryMethodTests)
8 -----
9 Traceback (most recent call last):
10     File "/Users/leif/Code/tango_with_django_project_19/rango/tests.py",
11     line 12, in test_ensure_views_are_positive
12         self.assertEqual((cat.views>=0), True)
13     AssertionError: False != True
14
15 -----
16 Ran 1 test in 0.001s
17
18 FAILED (failures=1)
```

As we can see this test fails. This is because the model does not check whether the value is less than zero or not. Since we really want to ensure that the values are non-zero, we will need to update the model, to ensure that this requirement is fulfilled. Do this now by adding some code to the `Category` models, `save()` method, that checks the value of `views`, and updates it accordingly.

Once you have updated your model, you can now re-run the test, and see if your code now passes it. If not, try again.

Let's try adding another test, that ensures an appropriate slug line is created i.e. one with dashes, and in lowercase. Add the following code to `rango/tests.py`:

```

1 def test_slug_line_creation(self):
2     """
3     slug_line_creation checks to make sure that when we add a category an appropriate
4     slug line is created
5     i.e. "Random Category String" -> "random-category-string"
6     """
7     cat = cat('Random Category String')
8     cat.save()
9     self.assertEqual(cat.slug, 'random-category-string')

```

Does your code still work?

Testing Views

So far we have written tests that focus on ensuring the integrity of the data housed in the models. Django also provides testing mechanisms to test views. It does this with a mock client, that internally makes a call to a Django view via the url. In the test you have access to the response (including the html) and the context dictionary.

Let's create a test that checks that when the index page loads, it displays the message that There are no categories present, when the Category model is empty.

```

from django.core.urlresolvers import reverse

class IndexViewTests(TestCase):

    def test_index_view_with_no_categories(self):
        """
        If no questions exist, an appropriate message should be displayed.
        """
        response = self.client.get(reverse('index'))
        self.assertEqual(response.status_code, 200)
        self.assertContains(response, "There are no categories present.")
        self.assertQuerysetEqual(response.context['categories'], [])

```

First of all, the django TestCase has access to a client object, which can make requests. Here, it uses the helper function reverse to look up the url of the index page. Then it tries to get that page, where the response is stored. The test then checks a number of things: did the page load ok? Does the response, i.e. the html contain the phrase "There are no categories present.", and does the context dictionary contain an empty categories list. Recall that when you run tests, a new database is created, which by default is not populated.

Let's now check the resulting view when categories are present. First add a helper method.

```
from rango.models import Category

def add_cat(name, views, likes):
    c = Category.objects.get_or_create(name=name)[0]
    c.views = views
    c.likes = likes
    c.save()
    return c
```

Then add another method to the class `IndexViewTests(TestCase)`:

```
def test_index_view_with_categories(self):
    """
    If no questions exist, an appropriate message should be displayed.
    """

    add_cat('test',1,1)
    add_cat('temp',1,1)
    add_cat('tmp',1,1)
    add_cat('tmp test temp',1,1)

    response = self.client.get(reverse('index'))
    self.assertEqual(response.status_code, 200)
    self.assertContains(response, "tmp test temp")

    num_cats = len(response.context['categories'])
    self.assertEqual(num_cats , 4)
```

In this test, we populate the database with four categories, and then check if the page loads, if it contains the text `tmp test temp` and if the number of categories is equal to 4.

Testing the Rendered Page

It is also possible to perform tests that load up the application and programmatically interact with the DOM elements on the HTML pages by using either Django's test client and/or Selenium, which is an "in-browser" framework to test the way the HTML is rendered in a browser.

6.2 Coverage Testing

Code coverage measures how much of your code base has been tested, and how much of your code has been put through its paces via tests. You can install a package called `coverage` via `pip install coverage` which automatically analyses how much code coverage you have. Once you have `coverage` installed, run the following command:

```
$ coverage run --source='.' manage.py test rango
```

This will run through all the tests and collect the coverage data for the rango application. To see the coverage report you need to then type:

```
$ coverage report
```

Name	Stmts	Miss	Cover

manage	6	0	100%
populate	33	33	0%
rango/__init__	0	0	100%
rango/admin	7	0	100%
rango/forms	35	35	0%
rango/migrations/0001_initial	5	0	100%
rango/migrations/0002_auto_20141015_1024	5	0	100%
rango/migrations/0003_category_slug	5	0	100%
rango/migrations/0004_auto_20141015_1046	5	0	100%
rango/migrations/0005_userprofile	6	0	100%
rango/migrations/__init__	0	0	100%
rango/models	28	3	89%
rango/tests	12	0	100%
rango/urls	12	12	0%
rango/views	110	110	0%
tango_with_django_project/__init__	0	0	100%
tango_with_django_project/settings	28	0	100%
tango_with_django_project/urls	9	9	0%
tango_with_django_project/wsgi	4	4	0%

TOTAL	310	206	34%

We can see from the above report that critical parts of the code have not been tested, ie. rango/views. For more details about using the package coverage visit: <http://nedbatchelder.com/code/coverage/>



Exercises

Lets say that we want to extend the Page to include two additional fields, `last_visit` and `first_visit` which will be of type `timedate`.

- Update the model to include these two fields
- Update the add page functionality, and the goto functionality.
- Add in a test to ensure the last visit or first visit is not in the future
- Add in a test to ensure that the last visit equal to or after the first visit.
- Run through [Part Five of the official Django Tutorial](#) to learn more about testing.
- Check out the [tutorial on test driven development by Harry Percival](#).

7. Deploying Your Project

This chapter provides a step-by-step guide on how to deploy your Django applications. We'll be looking at deploying applications on [PythonAnywhere](#), an online IDE and web hosting service. The service provides in-browser access to the server-based Python and Bash command line interfaces, meaning you can interact with PythonAnywhere's servers just like you would with a regular terminal instance on your own computer. Currently, PythonAnywhere are offering a free account which sets you up with an adequate amount of storage space and CPU time to get a Django application up and running.



Go Git It!

You can do this chapter independently, assuming you have some working knowledge of Git. If you don't, refer to the [chapter on using Git](#).

7.1 Creating a PythonAnywhere Account

First, [sign up for a Beginner PythonAnywhere account](#). If your application takes off and becomes popular, you can always upgrade your account at a later stage to gain more storage space and CPU time along with a number of other benefits - such as hosting specific domains and SSH abilities, for example.

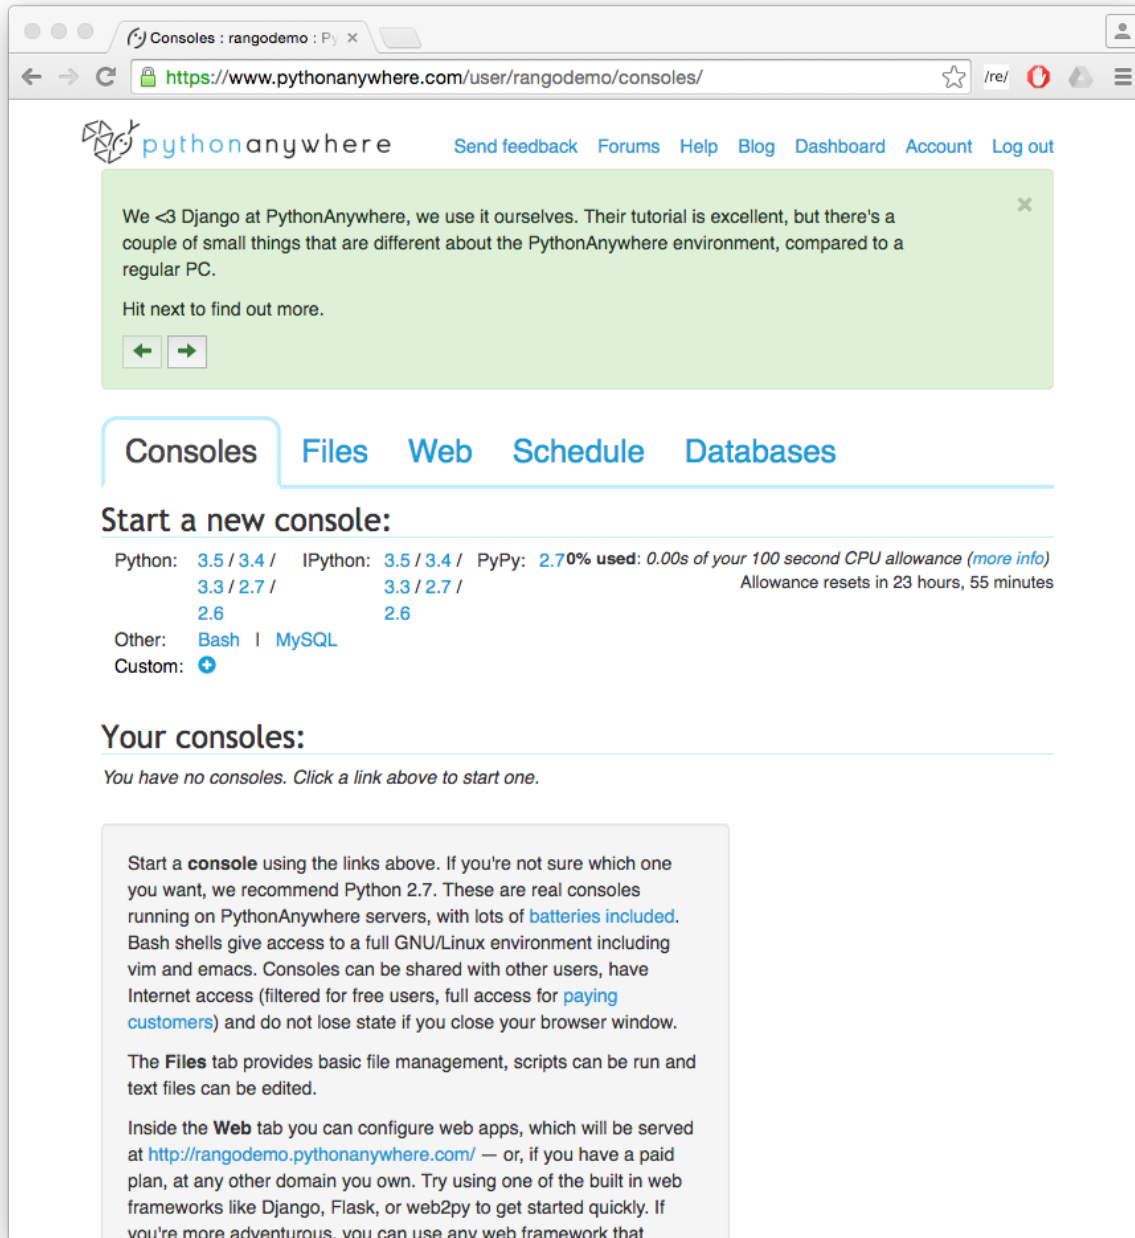
Once your account has been created, you will have your own little slice of the World Wide Web at `http://<username>.pythonanywhere.com`, where `<username>` is your PythonAnywhere username. It is from this URL that your hosted application will be available.

7.2 The PythonAnywhere Web Interface

The PythonAnywhere web interface contains a *dashboard* which in turn provides a series of tabs allowing you to manage your application. The tabs as [illustrated in the figure below](#) include:

- a *consoles* tab, allowing you to create and interact with Python and Bash console instances;
- a *files* tab, which allows you to upload to and organise files within your disk quota;
- a *web* tab, allowing you to configure settings for your hosted web application;
- a *schedule* tab, allowing you to setup tasks to be executed at particular times; and
- a *databases* tab, which allows you to configure a MySQL instance for your applications should you require it.

Of the five tabs provided, we'll be working primarily with the *consoles* and *web* tabs. The [PythonAnywhere Wiki](#) provides a series of detailed explanations on how to use the other tabs.



The PythonAnywhere dashboard, showing the *Consoles* tab.

7.3 Creating a Virtual Environment

As part of its standard default Bash environment, PythonAnywhere comes with Python 2.7.6 and a number of pre-installed Python Packages (including *Django 1.3.7* and *Django-Registration 0.8*). Since we are using a different setup, we need to select a particular Python version and setup a virtual environment for our application.

First, open a Bash console from the PythonAnywhere *Consoles* tab by clicking the *Bash* link. When the terminal is ready for you to use, enter the following commands.

```
$ mkvirtualenv --python=<python-version> rango
```

If you've coded up the tutorial using Python 3.x, then change `<python-version>` to either `python3.4` or `python3.5`. If you are using Python 2.7.x, then change `<python-version>` to `python2.7`. The command you enter creates a new virtual environment called `rango` using the version of Python that you specified. For example, below is the output for when we created a Python 2.7 virtual environment.

```
13:38 ~ $ mkvirtualenv --python=python2.7 rango
Running virtualenv with interpreter /usr/bin/python2.7
New python executable in /home/rangodemo/.virtualenvs/rango/bin/python2.7
Also creating executable in /home/rangodemo/.virtualenvs/rango/bin/python
Installing setuptools, pip, wheel...done.
virtualenvwrapper creating /home/rangodemo/.virtualenvs/.../predeactivate
virtualenvwrapper creating /home/rangodemo/.virtualenvs/.../postdeactivate
virtualenvwrapper creating /home/rangodemo/.virtualenvs/.../preactivate
virtualenvwrapper creating /home/rangodemo/.virtualenvs/.../postactivate
virtualenvwrapper creating /home/rangodemo/.virtualenvs/.../get_env_details
```

Note in the example above, the PythonAnywhere username used is `rangodemo` - this will be replaced with your own username. The process of creating the virtual environment will take a little while to complete, after which you will be presented with a slightly different prompt.

```
(rango)13:40 ~ $
```

Note the inclusion of `(rango)` compared to the previous command prompt. This signifies that the `rango` virtual environment has been activated, so any package installations will be done within that virtual environment, leaving the wider system setup alone. If you issue the command `ls -la`, you will see that a directory called `.virtualenvs` has been created. This is the directory in which all of your virtual environments and associated packages will be stored. To confirm the setup, issue the command `which pip`. This will print the location in which the active `pip` binary is located - hopefully within `.virtualenvs` and `rango`, as shown in the example below.

```
/home/<username>/.virtualenvs/test/bin/pip
```

To see what packages are already installed, enter `pip list`. Now we can customise the virtual environment by installing the required packages for our Rango application. Install all the required packages, by issuing the following commands.

```
$ pip install -U django==1.9.5
$ pip install pillow
$ pip install django-registration-redux
$ pip install django-bootstrap-toolkit
```

Alternatively, you could use `pip freeze > requirements.txt` to save your current development environment, and then on PythonAnywhere, run `pip install -r requirements.txt` to install all the packages in one go.



Waiting for the Bits to Download

Installing all these packages may take some time, so you can relax, call a friend, or tweet about our tutorial [@tangowithdjango!](#)

Once installed, check if Django has been installed with the command `which django-admin.py`. You should receive output similar to the following example.

```
/home/<username>/.virtualenvs/rango/bin/django-admin.py
```



Virtual Environments on PythonAnywhere

PythonAnywhere also provides instructions on how to setup virtual environments. [Check out their Wiki documentation for more information.](#)

Virtual Environment Switching

Moving between virtual environments can be done pretty easily. For this to work, you need to make sure that `virtualenvwrapper.sh` has been loaded by running `source virtualenvwrapper.sh`.

Rather than doing this each time you open up a console, you can add it to your `.bashrc` profile which is located in your home directory. Doing so will ensure the command is executed automatically for you every time you start a new Bash console instance. Any Bash consoles active will need to be closed for the changes to take effect.

With this done, you can then launch into a pre-existing virtual environment with the `workon` command. To load up the `rango` environment, enter:

```
16:48 ~ $ workon rango
```

where `rango` can be replaced with the name of the virtual environment you wish to use. Your prompt should then change to indicate you are working within a virtual environment.

```
(rango) 16:49 ~ $
```

You can then leave the virtual environment using the `deactivate` command. Your prompt should then be missing the `(rango)` prefix, with an example shown below.

```
(rango) 16:49 ~ $ deactivate  
16:51 ~ $
```

Cloning your Git Repository

Now that your virtual environment for Rango is all setup, you can now clone your Git repository to obtain a copy of your project's files. Clone your repository by issuing the following command from your home directory:

```
$ git clone https://<USERNAME>:<PASSWORD>@github.com/<OWNER>/<REPO_NAME>.git
```

where you replace - `<USERNAME>` with your GitHub username; - `<PASSWORD>` with your GitHub password; - `<OWNER>` with the username of the person who owns the repository; and - `<REPO_NAME>` with the name of your project's repository.

If you haven't put your code in a Git repository, you can clone the version we have made, by issuing the following command:

```
16:54 ~ $ git clone https://github.com/leifos/tango_with_django_19.git
```

Setting Up the Database

With your files cloned, you must then prepare your database. We'll be using the `populate_rango.py` module that we created earlier in the book. As we'll be running the module, you must ensure that you are using the `rango` virtual environment (i.e. you see `(rango)` as part of your prompt - if not, invoke `workon rango`). From your home directory, move into the `tango_with_django_19` directory, then to the `code` directory. Finally, `cd` into the `tango_with_django_project` directory - the directory with `manage.py` in it. Now issue the following commands.

```
(rango) 16:55 ~/tango_with_django $ python manage.py makemigrations rango
(rango) 16:55 ~/tango_with_django $ python manage.py migrate
(rango) 16:56 ~/tango_with_django $ python populate_rango.py
(rango) 16:57 ~/tango_with_django $ python manage.py createsuperuser
```

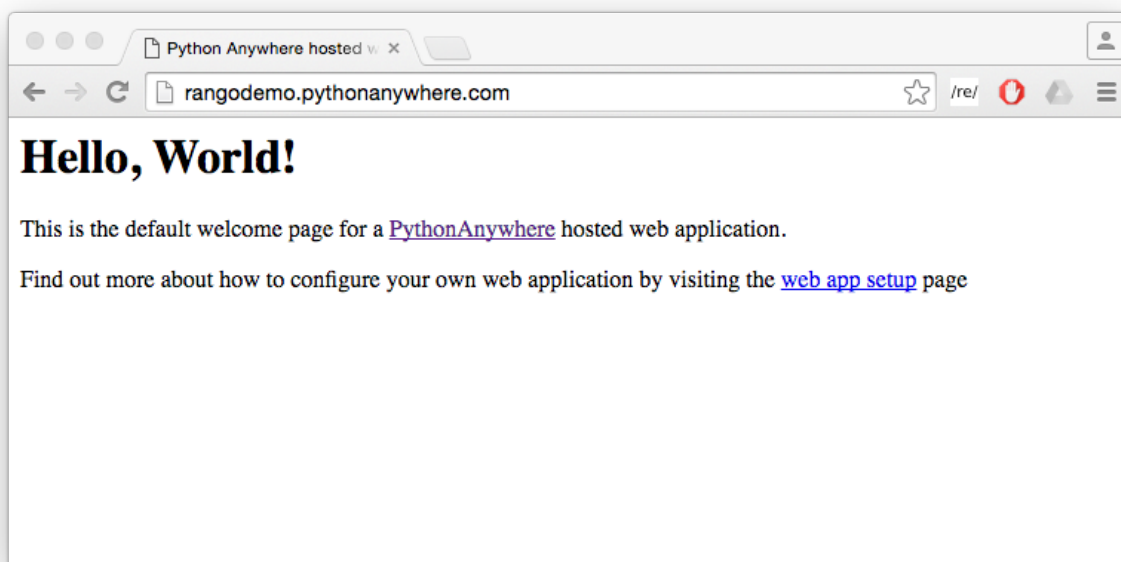
As discussed earlier in the book, the first command creates the migrations for the rango app, then the `migrate` command creates the *SQLite3* database. Once the database is created, the database can be populated and a superuser created.

7.4 Setting up Your Web Application

Now that the database is setup, we need to configure the PythonAnywhere *NGINX* Web server to serve up your application. Within PythonAnywhere's Web interface, navigate to your *dashboard* and click on the *Web* tab. On the left of the page that appears, click *Add a new web app*.

A popup box will then appear. Follow the instructions on-screen, and when the time comes, select the *manual configuration* option and complete the wizard. Make sure you select the same Python version as the one you selected earlier.

In a new tab or window in your Web browser, go visit your PythonAnywhere subdomain at the address `http://<username>.pythonanywhere.com`. You should be presented with the *default Hello, World! webpage, as shown below*. This is because the WSGI script is currently serving up this page, and not your Django application. This is what we need to change next.



The default PythonAnywhere *hello world* webpage.

Configure the Virtual Environment

To set the virtual environment for your app, navigate to the *Web* tab in PythonAnywhere's dashboard. From there, scroll all the way down under you see the heading *Virtualenv*.

Enter in the path to your virtual environment. Assuming you created a virtual environment called *rango* the path would be:

```
/home/<username>/.virtualenvs/rango
```

You can start a console to check if it is successful.

Now in the *Code* section, you can set the path to your web applications source code.

```
/home/<username>/<path-to>/tango_with_django_project/
```

If you have checked out code from our GitHub account, then the path will be something like:

```
/home/<username>/tango_with_django_19/code/tango_with_django_project/
```

Configuring the WSGI Script

The [Web Server Gateway Interface](#), a.k.a. *WSGI* provides a simple and universal interface between Web servers and Web applications. PythonAnywhere uses WSGI to bridge the server-application link and map incoming requests to your subdomain to your web application.

To configure the WSGI script, navigate to the *Web* tab in PythonAnywhere's dashboard. From there, click the *Web* tab. Under the *Code* heading you can see a link to the WSGI configuration file in the *Code* section: e.g. `/var/www/<username>_pythonanywhere_com_wsgi.py`

The good people at PythonAnywhere have set up a sample WSGI file for us with several possible configurations. For your Web application, you'll need to configure the Django section of the file. The example below demonstrates a possible configuration for you application.

```
import os
import sys

# ADD YOUR PROJECT TO THE PYTHONPATH FOR THE PYTHON INSTANCE
path = '/home/<username>/<path-to>/tango_with_django_project/'
if path not in sys.path:
    sys.path.append(path)

# IMPORTANTLY GO TO THE PROJECT DIR
```

```
os.chdir(path)

# TELL DJANGO WHERE YOUR SETTINGS MODULE IS LOCATED
os.environ.setdefault('DJANGO_SETTINGS_MODULE',
                     'tango_with_django_project.settings')

# IMPORT THE DJANGO SETUP
import django
django.setup()

# IMPORT THE DJANGO WSGI HANDLER TO TAKE CARE OF REQUESTS
import django.core.handlers.wsgi
application = django.core.handlers.wsgi.WSGIHandler()
```

Ensure that you replace <username> with your PythonAnywhere username, and update any other path settings to suit your application. You should also remove all other code from the WSGI configuration script to ensure no conflicts take place.

The script adds your project's directory to the PYTHONPATH for the Python instance that runs your web application. This allows Python to access your project's modules. If you have additional paths to add, you can easily insert them here. You can then specify the location of your project's `settings.py` module. The final step is to include the Django WSGI handler and invoke it for your application.

When you have completed the WSGI configuration, click the *Save* button at the top-right of the webpage. Navigate back to the *Web* tab within the PythonAnywhere dashboard, and click the *Reload* button at the top of the page. When the application is reloaded, you can then revisit your PythonAnywhere subdomain at `http://<username>.pythonanywhere.com`. Hopefully, if all went well, you should see your application up and running. If not, check through your scripts and paths carefully. Double check your paths by actually visiting the directories, and use `pwd` to confirm the path.



Bad Gateway Errors??

During testing, we noted that you can sometimes receive HTTP 502 - Bad Gateway errors instead of your application. Try reloading your application again, and then waiting a longer. If the problem persists, try reloading again. If the problem still persists, [check out your log files](#) to see if any accesses/errors are occurring, before contacting the PythonAnywhere support.

Assigning Static Paths

We're almost there. One issue which we still have to address is to sort out paths for our application. Doing so will allow PythonAnywhere's servers to serve your static content, for example From the

PythonAnywhere dashboard, click the *Web* tab and choose the subdomain hosting your application from the list on the left.

Underneath the *Static files* header, perform the following.

Click the `Enter` path text. Set this to (all on one line):

```
/home/<username>/virtualenvs/rango/lib/<python-version>/site-packages/django/  
contrib/admin/static/admin
```

where `<username>` should be replaced with your PythonAnywhere username. `<python-version>` should also be replaced with `2.7`, `3.4`, etc., depending on which Python version you selected. You may also need to change `rango` if this is not the name of your application's virtual environment. Remember to hit return to confirm the path. Then click `Enter URL` and enter `/static/admin`, followed by hitting return.

Repeat the two steps above for the URL `/static/` and path `/home/<username>/<path-to>/tango_ -with_django_project/static`, with the path setting pointing to the `static` directory of your Web application.

With these changes saved, reload your web application by clicking the *Reload* button at the top of the page. Don't forget about potential HTTP 502 - Bad Gateway errors, as discussed previously. Setting the static folders means that when you visit the admin interface, it has the predefined Django style sheets, and that you can access images and scripts. Reload your Web application, and you should now notice that your images are present.

Bing API Key

[Add your Bing API key](#) to `bing.key` to enable the search functionality in Rango.

Turning off DEBUG Mode

When your application is ready to go, it's a good idea to instruct Django that your application is now hosted on a production server. To do this, open your project's `settings.py` file and change `DEBUG = True` to `DEBUG = False`. This disables [Django's debug mode](#), and removes explicit error messages.

Changing the value of `DEBUG` also means you should set the `ALLOWED_HOSTS` property. Failing to perform this step will make Django return HTTP 400 Bad Request errors. Alter `ALLOWED_HOSTS` so that it includes your PythonAnywhere subdomain like in the example below.

```
ALLOWED_HOSTS = ['<username>.pythonanywhere.com']
```

Again, ensure `<username>` is changed to your PythonAnywhere username. Once complete, save the file and reload the application via the PythonAnywhere Web interface.

7.5 Log Files

Deploying your Web application to an online environment introduces another layer of complexity. It is likely that you will encounter new and bizarre errors due to unsuspecting problems. When facing such errors, vital clues may be found in one of the three log files that the Web server on PythonAnywhere creates.

Log files can be viewed via the PythonAnywhere web interface by clicking on the *Web* tab, or by viewing the files in `/var/log/` within a Bash console instance. The files provided are:

- `access.log`, which provides a log of requests made to your subdomain;
- `error.log`, which logs any error messages produced by your web application; and
- `server.log`, providing log details for the UNIX processes running your application.

Note that the names for each log file are prepended with your subdomain. For example, `access.log` will have the name `<username>.pythonanywhere.com.access.log`.

When debugging, you may find it useful to delete or move the log files so that you don't have to scroll through a huge list of previous attempts. If the files are moved or deleted, they will be recreated automatically when a new request or error arises.



Exercises

Congratulations, you've successfully deployed Rango!

- Tweet a link of your application to [@tangowithdjango](#).
- E-mail us to let us know, and let us know your thoughts on the book.

8. Final Thoughts

In this book, we have gone through the process of web development from specification to deployment. Along the way we have shown how to use the Django framework to construct the models, views and templates associated with a web application. We have also demonstrated how toolkits and services like Bootstrap, JQuery, Bing Search, PythonAnywhere, etc can be integrated within an application. However, the road doesn't stop here. While, as we have only painted the broad brush strokes of a web application - as you have probably noticed there are lots of improvements that could be made to Rango - and these finer details often take a lot more time to complete as you polish the application. By developing your application on a firm base and good setup you will be able to construct up to 80% of your site very rapidly and get a working demo online.

In future versions of this book we intend to provide some more details on various aspects of the framework, along with covering the basics of some of the other fundamental technologies associated with web development. If you have any suggestions or comments about how to improve the book please get in touch.

Please report any bugs, problems, etc, or submit change requests via GitHub: https://github.com/leifos/tango_with_django_19/

9. Setting up your System

9.1 Installing the Software

Now that you have a decent understanding of how to interact with the terminal, you can begin to install the software required for this tutorial.

Installing Python

So, how do you go about installing Python 2.7/3.0 on your computer? You may already have Python installed on your computer - and if you are using a Linux distribution or OS X, you will definitely have it installed. Some of your operating system's functionality is [implemented in Python](#), hence the need for an interpreter!

Unfortunately, nearly all modern operating systems utilise a version of Python that is older than what we require for this tutorial. There's many different ways in which you can install Python, and many of them are sadly rather tricky to accomplish. We demonstrate the most commonly used approaches, and provide links to additional reading for more information.

warning

This section will detail how to run Python 2.7.5 *alongside* your current Python installation. It is regarded as poor practice to remove your operating system's default Python installation and replace it with a newer version. Doing so could render aspects of your operating system's functionality broken!

Apple OS X

The most simple way to get Python 2.7.5 installed on your Mac is to download and run the simple installer provided on the official Python website. You can download the installer by visiting the webpage at <http://www.python.org/getit/releases/2.7.5/>.

warning

Ensure that you download the .dmg file that is relevant to your particular OS X installation!

1. Once you have downloaded the .dmg file, double-click it in the Finder.
2. The file mounts as a separate disk and a new Finder window is presented to you.

3. Double-click the file `Python.mpkg`. This will start the Python installer.
4. Continue through the various screens to the point where you are ready to install the software. You may have to provide your password to confirm that you wish to install the software.
5. Upon completion, close the installer and eject the Python disk. You can now delete the downloaded `.dmg` file.

You should now have an updated version of Python installed, ready for Django! Easy, huh?

Linux Distributions

Unfortunately, there are many different ways in which you can download, install and run an updated version of Python on your Linux distribution. To make matters worse, methodologies vary from distribution to distribution. For example, the instructions for installing Python on [Fedora](#) may differ from those to install it on an [Ubuntu](#) installation.

However, not all hope is lost. An awesome tool (or a *Python environment manager*) called [pythonbrew](#) can help us address this difficulty. It provides an easy way to install and manage different versions of Python, meaning you can leave your operating system's default Python installation alone. Hurrah!

Taken from the instructions provided from [the pythonbrew GitHub page](#) and [this Stack Overflow question and answer page](#), the following steps will install Python 2.7.5 on your Linux distribution.

1. Open a new terminal instance.
2. Run the command `curl -kL http://xrl.us/pythonbrewinstall | bash`. This will download the installer and run it within your terminal for you. This installs pythonbrew into the directory `~/.pythonbrew`. Remember, the tilde (`~`) represents your home directory!
3. You then need to edit the file `~/.bashrc`. In a text editor (such as `gedit`, `nano`, `vi` or `emacs`), add the following to a new line at the end of `~/.bashrc`: `[[-s $HOME/.pythonbrew/etc/bashrc]] && source $HOME/.pythonbrew/etc/bashrc`
4. Once you have saved the updated `~/.bashrc` file, close your terminal and open a new one. This allows the changes you make to take effect.
5. Run the command `pythonbrew install 2.7.5` to install Python 2.7.5.
6. You then have to *switch* Python 2.7.5 to the *active* Python installation. Do this by running the command `pythonbrew switch 2.7.5`.
7. Python 2.7.5 should now be installed and ready to go.

note

Directories and files beginning with a period or dot can be considered the equivalent of *hidden files* in Windows. [Dot files](#) are not normally visible to directory-browsing tools, and are commonly used for configuration files. You can use the `ls` command to view hidden files by adding the `-a` switch to the end of the command, giving the command `ls -a`.

Windows

By default, Microsoft Windows comes with no installations of Python. This means that you do not have to worry about leaving existing versions be; installing from scratch should work just fine. You can download a 64-bit or 32-bit version of Python from [the official Python website](#). If you aren't sure which one to download, you can determine if your computer is 32-bit or 64-bit by looking at the instructions provided [on the Microsoft website](#).

1. When the installer is downloaded, open the file from the location to which you downloaded it.
2. Follow the on-screen prompts to install Python.
3. Close the installer once completed, and delete the downloaded file.

Once the installer is complete, you should have a working version of Python ready to go. By default, Python 2.7.5 is installed to the folder `C:\Python27`. We recommend that you leave the path as it is.

Upon the completion of the installation, open a Command Prompt and enter the command `python`. If you see the Python prompt, installation was successful. However, in certain circumstances, the installer may not set your Windows installation's `PATH` environment variable correctly. This will result in the `python` command not being found. Under Windows 7, you can rectify this by performing the following:

1. Click the *Start* button, right click *My Computer* and select *Properties*.
2. Click the *Advanced* tab.
3. Click the *Environment Variables* button.
4. In the *System variables* list, find the variable called *Path*, click it, then click the *Edit* button.
5. At the end of the line, enter `;C:\python27;C:\python27\scripts`. Don't forget the semicolon - and certainly *do not* add a space.
6. Click OK to save your changes in each window.
7. Close any Command Prompt instances, open a new instance, and try run the `python` command again.

This should get your Python installation fully working. Windows XP, [has slightly different instructions](#), and [so do Windows 8 installationsthis](#).

Setting Up the PYTHONPATH

With Python now installed, we now need to check that the installation was successful. To do this, we need to check that the `PYTHONPATH` [environment variable](#) is setup correctly. `PYTHONPATH` provides the Python interpreter with the location of additional Python [packages and modules](#) which add extra functionality to the base Python installation. Without a correctly set `PYTHONPATH`, we'll be unable to install and use Django!

First, let's verify that our `PYTHONPATH` variable exists. Depending on the installation technique that you chose, this may or may not have been done for you. To do this on your UNIX-based operating system, issue the following command in a terminal.

```
1 $ echo $PYTHONPATH
```

On a Windows-based machine, open a Command Prompt and issue the following.

```
1 $ echo %PYTHONPATH%
```

If all works, you should then see output that looks something similar to the example below. On a Windows-based machine, you will obviously see a Windows path, most likely originating from the C drive.

```
1 /opt/local/Library/Frameworks/Python.framework/Versions/2.7/lib/python2.7/site-p\
2 ackages:
```

This is the path to your Python installation's `site-packages` directory, where additional Python packages and modules are stored. If you see a path, you can continue to the next part of this tutorial. If you however do not see anything, you'll need to do a little bit of detective work to find out the path. On a Windows installation, this should be a trivial exercise: `site-packages` is located within the `lib` folder of your Python installation directory. For example, if you installed Python to `C:\Python27`, `site-packages` will be at `C:\Python27\Lib\site-packages\`.

UNIX-based operating systems however require a little bit of detective work to discover the path of your `site-packages` installation. To do this, launch the Python interpreter. The following terminal session demonstrates the commands you should issue.

```
1 $ python
2
3 Python 2.7.5 (v2.7.5:ab05e7dd2788, May 13 2013, 13:18:45)
4 [GCC 4.2.1 (Apple Inc. build 5666) (dot 3)] on darwin
5 Type "help", "copyright", "credits" or "license" for more information.
6
7 >>> import site
8 >>> print site.getsitepackages()[0]
9
10 '/Library/Frameworks/Python.framework/Versions/2.7/lib/python2.7/site-packages'
11
12 >>> quit()
```

Calling `site.getsitepackages()` returns a list of paths that point to additional Python package and module stores. The first typically returns the path to your `site-packages` directory - changing the list index position may be required depending on your installation. If you receive an error stating that `getsitepackages()` is not present within the `site` module, verify you're running the correct version of Python. Version 2.7.5 should include this function. Previous versions of the language do not include this function.

The string which is shown as a result of executing `print site.getsitepackages()[0]` is the path to your installation's `site-packages` directory. Taking the path, we now need to add it to your configuration. On a UNIX-based or UNIX-derived operating system, edit your `.bashrc` file once more, adding the following to the bottom of the file.

```
1 export PYTHONPATH=$PYTHONPATH:<PATH_TO_SITE-PACKAGES>
```

Replace `<PATH_TO_SITE-PACKAGES>` with the path to your `site-packages` directory. Save the file, and quit and reopen any instances of your terminal.

On a Windows-based computer, you must follow the instructions shown in Section `requirements-install-python-windows` to bring up the environment variables settings dialog. Add a `PYTHONPATH` variable with the value being set to your `site-packages` folder, which is typically `C:\Python27\Lib\site-packages\`.

Using Setuptools and Pip

Installing and setting up your development environment is a really important part of any project. While it is possible to install Python Packages such as Django separately, this can lead to numerous problems and hassles later on. For example, how would you share your setup with another developer? How would you set up the same environment on your new machine? How would you upgrade to the latest version of the package? Using a package manager removes much of the hassle involved in setting up and configuring your environment. It will also ensure that the package you install is the correct for the version of Python you are using, along with installing any other packages that are dependent upon the one you want to install.

In this book, we will be using *Pip*. Pip is a user-friendly wrapper over the *Setuptools* Python package manager. Because Pip depends on Setuptools, we are required to ensure that both are installed on your computer.

To start, we should download Setuptools from the [official Python package website](#). You can download the package in a compressed `.tar.gz` file. Using your favourite file extracting program, extract the files. They should all appear in a directory called `setuptools-1.1.6` - where `1.1.6` represents the Setuptools version number. From a terminal instance, you can then change into the directory and execute the script `ez_setup.py` as shown below.

```
1 $ cd setuptools-1.1.6
2 $ sudo python ez_setup.py
```

In the example above, we also use `sudo` to allow the changes to become system-wide. The second command should install Setuptools for you. To verify that the installation was successful, you should be able to see output similar to that shown below.

```
1 Finished processing dependencies for setuptools==1.1.6
```

Of course, `1.1.6` is substituted with the version of Setuptools you are installing. If this line can be seen, you can move onto installing Pip. This is a trivial process, and can be completed with one simple command. From your terminal instance, enter the following.

```
1 $ sudo easy_install pip
```

This command should download and install Pip, again with system-wide access. You should see the following output, verifying Pip has been successfully installed.

```
1 Finished processing dependencies for pip
```

Upon seeing this output, you should be able to launch Pip from your terminal. To do so, just type `pip`. Instead of an unrecognised command error, you should be presented with a list of commands and switches that Pip accepts. If you see this, you're ready to move on!

note

With Windows-based computers, follow the same basic process. You won't need to enter the `sudo` command, however.

9.2 Virtual Environments

We're almost all set to go! However, before we continue, it's worth pointing out that while this setup is fine to begin with, there are some drawbacks. What if you had another Python application that requires a different version to run? Or you wanted to switch to the new version of Django, but still wanted to maintain your Django 1.7 project?

The solution to this is to use [virtual environments](#). Virtual environments allow multiple installations of Python and their relevant packages to exist in harmony. This is the generally accepted approach to configuring a Python setup nowadays.

They are pretty easy to setup, once you have pip installed, and you know the right commands. You need to install a couple of additional packages.

```
1 $ pip install virtualenv
2 $ pip install virtualenvwrapper
```

The first package provides you with the infrastructure to create a virtual environment. See a [non-magical introduction to Pip and Virtualenv for Python Beginners](#) by Jamie Matthews for details about using virtualenv. However, using just *virtualenv* alone is rather complex. The second package provides a wrapper to the functionality in the virtualenv package and makes life a lot easier.

If you are using a linux/unix based OS, then to use the wrapper you need to call the following shell script from your command line: :

```
1 $ source virtualenvwrapper.sh
```

It is a good idea to add this to your bash/profile script. So you dont have to run it each and every time you want to use virtualenvironments.

However, if you are using windows, then install the [virtualenvwrapper-win](#) package:

```
1 $ pip install virtualenvwrapper-win
```

Now you should be all set to create a virtual environment:

```
1 $ mkvirtualenv rango
```

You can list the virtual environments created with `lsvirtualenv`, and you can activate a virtual environment as follows:

```
1 $ workon rango
2 (rango)$
```

Your prompt will change and the current virtual environment will be displayed, i.e. rango. Now within this environment you will be able to install all the packages you like, without interfering with your standard or other environments. Try `pip list` to see you dont have Django or Pillow installed in your virtual environment. You can now install them with pip so that they exist in your virtual environment.

Later on when we go to deploy the application, we will go through a similar process see Chapter [Deploying your Application<virtual-environment>](#) and set up a virtual environment on PythonAnywhere.

Code Repository

We should also point out that when you develop code, you should always house your code within a version-controlled repository such as [SVN](#) or [GIT](#). We won't be going through this right now so that we can get stuck into developing an application in Django. We have however provided a crash course on GIT [<git-crash-course>](#). We highly recommend that you set up a GIT repository for your own projects. Doing so could save you from disaster.

9.3 Exercises

To get comfortable with your environment, try out the following exercises.

- Install Python 2.7.5+ and Pip.
- Play around with your CLI and create a directory called code, which we use to create our projects in.
- Install the Django and Pillow packages.
- Setup your Virtual Environment
- Setup your account on GitHub
- Download and setup a Integrated Development Environemnt (like PyCharm) -

We have made the code for the book and application that you build available on GitHub, see [Tango W](#)

If you spot any errors or problem with the book, you can make a change request!

- If you have any problems with the exercises, you can check out the repository and see how we completed them.

IDE

Installing Git

Setup Git - the Git crash course expects a small section in this chapter detailing what you have to do to install git. Perhaps move the Windows information bit in the setting up Git part to here.

10. A Crash Course in UNIX-based Commands

Depending on your computing background, you may or may not have encountered a UNIX based system, or a derivative of. This small crash course focuses on getting you up to speed with the *terminal*, an application in which you issue commands for the computer to execute. This differs from a point-and-click *Graphical User Interface (GUI)*, the kind of interface that has made computing so much more accessible. A terminal based interface may be more complex to use, but the benefits of using such an interface include getting things done quicker, and more accurately, too.



Not for Windows!

Note that we're focusing on the Bash shell, a shell for UNIX-based operating systems and their derivatives, including OS X and Linux distributions. If you're a Windows user, you can use the [Windows Command Prompt](#) or [Windows PowerShell](#). Users of Windows 10 will [also be able to issue Bash commands directly to the Command Prompt](#). You could also experiment by [installing Cygwin](#) to bring Bash commands to Windows.

10.1 Using the Terminal

UNIX based operating systems and derivatives - such as OS X and Linux distributions - all use a similar looking terminal application, typically using the [Bash shell](#). All possess a core set of commands which allow you to navigate through your computer's filesystem and launch programs - all without the need for any graphical interface.

Upon launching a new terminal instance, you'll be typically presented with something resembling the following.

```
sibu:~ david$
```

What you see is the *prompt*, and indicates when the system is waiting to execute your every command. The prompt you see varies depending on the operating system you are using, but all look generally very similar. In the example above, there are three key pieces of information to observe:

- your username and computer name (username of david and computer name of sibu);
- your *present working directory* (the tilde, or ~); and

- the privilege of your user account (the dollar sign, or \$).



What is a Directory?

Above, we refer to your present working directory. But what exactly is a *directory*? If you have used a Windows computer up until now, you'll probably know a directory as a *folder*. The concept of a folder is analogous to a directory - it is a cataloguing structure which can contain references to other files and directories.

The dollar sign (\$) typically indicates that the user is a standard user account. Conversely, a hash symbol (#) may be used to signify the user logged in has [root privileges](#). Whatever symbol is present is used to signify that the computer is awaiting your input.



Prompts can Differ

The information presented by the prompt on your computer may differ from the example shown above. For example, some prompts may display the current date and time, or any other information. It all depends how your computer is set up.

When you are using the terminal, it is important to know where you are in the file system. To find out where you are, you can issue the command `pwd`. This will display your *Present Working Directory* (hence `pwd`). For example, check the example terminal interactions below.

```
Last login: Wed Mar 23 15:01:39 2016
sibu:~ david$ pwd
/users/grad/david
sibu:~ david$
```

You can see that the present working directory in this example is `/users/grad/david`.

You'll also note that the prompt indicates that the present working directory is a tilde `~`. The tilde is used a special symbol which represents your *home directory*. The base directory in any UNIX based file system is the *root directory*. The path of the root directory is denoted by a single forward slash (`/`). As folders (or directories) are separated in UNIX paths with a `/`, a single `/` denotes the root!

If you are not in your home directory, you can *Change Directory* (`cd`) by issuing the following command:

```
sibu:/ david$ cd ~
sibu:~ david$
```

Note how the present working directory switches from / to ~ upon issuing the `cd ~` command.



Path Shortcuts

UNIX shells have a number of different shorthand ways for you to move around your computer's filesystem. You've already seen that a forward slash (/) represents the **root directory**, and the tilde (~) represents your home directory in which you store all your personal files. However, there are a few more special characters you can use to move around your filesystem in conjunction with the `cd` command.

- Issuing `cd ~` will always return you to your home directory. On some UNIX or UNIX derivatives, simply issuing `cd` will return you to your home directory, too.
- Issuing `cd ..` will move your present working directory **up one level** of the filesystem hierarchy. For example, if you are currently in `/users/grad/david/code/`, issuing `cd ..` will move you to `/users/grad/david/`.
- Issuing `cd -` will move you to the **previous directory you were working in**. Your shell remembers where you were, so if you were in `/var/tmp/` and moved to `/users/grad/david/`, issuing `cd -` will move you straight back to `/var/tmp/`. This command obviously only works if you've move around at least once in a given terminal session.

Now, let's create a directory within the home directory called `code`. To do this, you can use the *Make Directory* command, called `mkdir`.

```
sibu:~ david$ mkdir code
sibu:~ david$
```

There's no confirmation that the command succeeded. We can check that the `code` directory was created by changing the present working directory with the `cd` command.

```
sibu:~ david$ cd code
sibu:code david$
```

Issuing a subsequent `pwd` command to confirm our present working directory yields `/users/grad/david/code` - our home directory, with `code` appended to the end. You can also see from the prompt in the example above that the present working directory changes from `~` to `code`.



Change Back

Now issue the command to change back to your home directory. What command do you enter?

From your home directory, let's now try out another command to see what files and directories exist. This new command is called `ls`, shorthand for *list*. Issuing `ls` in your home directory will yield something similar to the following.

```
sibu:~ david$ ls
code
```

This shows us that there's something present our home directory called `code`, as we would expect. We can obtain more detailed information by adding a `l` switch to the end of the `ls` command - with `l` standing for *list*.

```
sibu:~ david$ ls -l
drwxr-xr-x  2 david  grad   68  2 Apr 11:07 code
```

This provides us with additional information, such as the modification date (2 Apr 11:07), who the file belongs to (user `david` of group `grad`), the size of the entry (68 bytes), and the file permissions (`drwxr-xr-x`). While we don't go into file permissions here, the key thing to note is the `d` at the start of the string which denotes the entry is a directory. If we then add some files to our home directory and reissue the `ls -l` command, we then can observe differences in the way files are displayed as opposed to directories.

```
sibu:~ david$ ls -l
drwxr-xr-x  2 david  grad     68  2 Apr 11:07 code
-rw-r--r--@ 1 david  grad 303844 1 Apr 16:16 document.pdf
-rw-r--r--  1 david  grad    14  2 Apr 11:14 readme.md
```

One final useful switch to the `ls` command is the `a` switch, which displays *all* files and directories. This is useful because some directories and files can be *hidden* by the operating system to keep things looking tidy. Issuing the command yields more files and directories!

```
sibu:~ david$ ls -la
-rw-r--r--  1 david  grad    463 20 Feb 19:58 .profile
drwxr-xr-x 16 david  grad    544 25 Mar 11:39 .virtualenvs
drwxr-xr-x  2 david  grad     68  2 Apr 11:07 code
-rw-r--r--@ 1 david  grad 303844 1 Apr 16:16 document.pdf
-rw-r--r--  1 david  grad    14  2 Apr 11:14 readme.md
```

This command shows a hidden directory `.virtualenvs` and a hidden file `.profile`. Note that hidden files on a UNIX based computer (or derivative) start with a period (`.`). There's no special hidden file attribute you can apply, unlike on Windows computers.



Combining `ls` Switches

You may have noticed that we combined the `l` and `a` switches in the above `ls` example to force the command to output a list displaying all hidden files. This is a valid command - and there are [even more switches you can use](#) to customise the output of `ls`.

Creating files is also easy to do, straight from the terminal. The `touch` command creates a new, blank file. If we wish to create a file called `new.txt`, issue `touch new.txt`. If we then list our directory, we then see the file added.

```
sibu:~ david$ ls -l
drwxr-xr-x  2 david  grad      68  2 Apr 11:07 code
-rw-r--r--@ 1 david  grad  303844  1 Apr 16:16 document.pdf
-rw-r--r--  1 david  grad       0  2 Apr 11:35 new.txt
-rw-r--r--  1 david  grad      14  2 Apr 11:14 readme.md
```

Note the filesize of `new.txt` - it is zero bytes, indicating an empty file. We can start editing the file using one of the many available text editors that are available for use directly from a terminal, such as `nano` or `vi`. While we don't cover how to use these editors here, you can [have a look online for a simple how-to tutorial](#). We suggest starting with `nano` - while there are not as many features available compared to other editors, using `nano` is much simpler.

10.2 Core Commands

In the short tutorial above, you've covered a few of the core commands such as `pwd`, `ls` and `cd`. There are however a few more standard UNIX commands that you should familiarise yourself with before you start working for real. These are listed below for your reference, with most of them focusing upon file management. The list comes with an explanation of each, and an example of how to use them.

- `pwd`: As explained previously, this command displays your *present working directory* to the terminal. The full path of where you are presently is displayed.
- `ls`: Displays a list of files in the current working directory to the terminal.
- `cd`: In conjunction with a path, allows you to change your present working directory. For example, the command `cd /users/grad/david/` changes the current working directory to `/users/grad/david/`. You can also move up a directory level without having to provide the [absolute path](#) by using two dots, e.g. `cd ..`
- `cp`: Copies files and/or directories. You must provide the *source* and the *target*. For example, to make a copy of the file `input.py` in the same directory, you could issue the command `cp input.py input_backup.py`.

- `mv`: Moves files/directories. Like `cp`, you must provide the *source* and *target*. This command is also used to rename files. For example, to rename `numbers.txt` to `letters.txt`, issue the command `mv numbers.txt letters.txt`. To move a file to a different directory, you would supply either an absolute or relative path as part of the target - like `mv numbers.txt /home/david/numbers.txt`.
- `mkdir`: Creates a directory in your current working directory. You need to supply a name for the new directory after the `mkdir` command. For example, if your current working directory was `/home/david/` and you ran `mkdir music`, you would then have a directory `/home/david/music/`. You will need to then `cd` into the newly created directory to access it.
- `rm`: Shorthand for *remove*, this command removes or deletes files from your filesystem. You must supply the filename(s) you wish to remove. Upon issuing a `rm` command, you will be prompted if you wish to delete the file(s) selected. You can also remove directories [using the recursive switch](#). Be careful with this command - recovering deleted files is very difficult, if not impossible!
- `rmdir`: An alternative command to remove directories from your filesystem. Provide a directory that you wish to remove. Again, be careful: you will not be prompted to confirm your intentions.
- `sudo`: A program which allows you to run commands with the security privileges of another user. Typically, the program is used to run other programs as `root` - the [superuser](#) of any UNIX-based or UNIX-derived operating system.



There's More!

This is only a brief list of commands. Check out Ubuntu's documentation on [Using the Terminal](#) for a more detailed overview, or the [Cheat Sheet](#) by FOSSwire for a quick, handy reference guide. Like anything else, the more you practice, the more comfortable you will feel working with the terminal.

11. Virtual Environments

Virtual environments allow multiple installations of Python and their relevant packages to exist in harmony. This is the generally accepted approach to configuring a Python setup nowadays.

They are pretty easy to setup. Assuming you have `pip` installed, you can install the following packages:

```
$ pip install virtualenv
$ pip install virtualenvwrapper
```

The first package provides you with the infrastructure to create a virtual environment. See [a non-magical introduction to Pip and Virtualenv for Python Beginners](#) by Jamie Matthews for details about using `virtualenv`. However, using just `virtualenv` alone is rather complex. The second package provides a wrapper to the functionality in the `virtualenv` package and makes life a lot easier. The wrapper provides a series of extensions by [Doug Hellman](#) to the original `virtualenv` tool, making it easier for us to create, delete and use virtual environments.

If you are using a linux/unix based OS, then to use the wrapper you need to call the following shell script from your command line:

```
$ source virtualenvwrapper.sh
```

It is a good idea to add this to your `bash`/profile script. So you don't have to run it each and every time you want to use virtual environments. However, if you are using windows, then install the [virtualenvwrapper-win](#) package:

```
$ pip install virtualenvwrapper-win
```

Now you should be all set to create a virtual environment:

```
$ mkvirtualenv rango
```

You can list the virtual environments created with `lsvirtualenv`, and you can activate a virtual environment as follows:


```
$ workon rango  
(rango)$
```

Your prompt will change and the current virtual environment will be displayed, i.e. `rango`. Now within this environment you will be able to install all the packages you like, without interfering with your standard or other environments. Try `pip list` to see you don't have Django or Pillow installed in your virtual environment. You can now install them with `pip` so that they exist in your virtual environment.

In our [chapter on deployment](#), we will go through a similar process when deploying your application to [PythonAnywhere](#).

12. A Git Crash Course

We strongly recommend that you spend some time familiarising yourself with a [version control](#) system for your application's codebase. This chapter provides you with a crash course in how to use [Git](#), one of the many version control systems available. Originally developed by [Linus Torvalds](#), Git is today [one of the most popular version control systems in use](#), and is used by open-source and closed-source projects alike.

This tutorial demonstrates at a high level how Git works, explains the basic commands that you can use, and provides an explanation of Git's workflow. By the end of this chapter, you'll be able to make contributions to a Git repository, enabling you to work solo, or in a team.

12.1 Why Use Version Control?

As your software engineering skills develop, you will find that you are able to plan and implement solutions to ever more complex problems. As a rule of thumb, the larger the problem specification, the more code you have to write. The more code you write, the greater the emphasis you should put on software engineering practices. Such practices include the use of design patterns and the *DRY (Don't Repeat Yourself)* principle.

Think about your experiences with programming thus far. Have you ever found yourself in any of these scenarios?

- Made a mistake to code, realised it was a mistake and wanted to go back?
- Lost code (through a faulty drive), or had a backup that was too old?
- Had to maintain multiple versions of a product (perhaps for different organisations)?
- Wanted to see the difference between two (or more) versions of your codebase?
- Wanted to show that a particular change broke or fixed a piece of code?
- Wanted to submit a change (patch) to someone else's code?
- Wanted to see how much work is being done (where it was done, when it was done, or who did it)?

Using a version control system makes your life easier in *all* of the above cases. While using version control systems at the beginning may seem like a hassle it will pay off later - so it's good to get into the habit now!

We missed one final (and important) argument for using version control. With ever more complex problems to solve, your software projects will undoubtedly contain a large number of files containing source code. It'll also be likely that you *aren't working alone on the project; your project will probably have more than one contributor*. In this scenario, it can become difficult to avoid conflicts when working on files.

12.2 How Git Works

Essentially, Git comprises of four separate storage locations: your **workspace**, the **local index**, the **local repository** and the **remote repository**. As the name may suggest, the remote repository is stored on some remote server, and is the only location stored on a computer other than your own. This means that there are two copies of the repository - your local copy, and the remote copy. Having two copies is one of the main selling points of Git over other version control systems. You can make changes to your local repository when you may not have Internet access, and then apply any changes to the remote repository at a later stage. Only once changes are made to the remote repository can other contributors see your changes.



What is a Repository?

We keep repeating the word *repository*, but what do we actually mean by that? When considering version control, a repository is a data structure which contains metadata (a set of data that describes other data, hence *meta*) concerning the files which you are storing within the version control system. The kind of metadata that is stored can include aspects such as the historical changes that have taken place within a given file, so that you have a record of all changes that take place.

If you want to learn more about the metadata stored by Git, there is a [technical tutorial available](#) for you to read through.

For now though, let's provide an overview of each of the different aspects of the Git system. We'll recap some of the things we've already mentioned just to make sure it makes sense to you.

- As already explained, the **remote repository** is the copy of your project's repository stored on some remote server. This is particularly important for Git projects that have more than one contributor - you require a central place to store all the work that your team members produce. You could set up a Git server on a computer with Internet access and a properly configured firewall (check out [this Git server tutorial](#), for example), or simply use one of many services providing free Git repositories. One of the most widely used services available today is [GitHub](#). In fact, this book has a Git [repository](#) on GitHub!
- The **local repository** is a copy of the remote repository stored on your computer (locally). This is the repository to which you make all your additions, changes and deletions. When you reach a particular milestone, you can then *push* all your local changes to the remote repository. From there, you can instruct your team members to retrieve your changes. This concept is known as *pulling* from the remote repository. We'll subsequently explain pushing and pulling in a bit more detail.
- The **local index** is technically part of the local repository. The local index stores a list of files that you want to be managed with version control. This is explained in more detail [later in this chapter](#). You can have a look [here](#) to see a discussion on what exactly a Git index contains.

- The final aspect of Git is your **workspace**. Think of this folder or directory as the place on your computer where you make changes to your version controlled files. From within your workspace, you can add new files or modify or remove previously existing ones. From there, you then instruct Git to update the repositories to reflect the changes you make in your workspace. This is important - *don't modify code inside the local repository - you only ever edit files in your workspace*.

Next, we'll be looking at how to [get your Git workspace set up and ready to go](#). We'll also [discuss the basic workflow](#) you should use when using Git.

12.3 Setting up Git

We assume that you've got Git installed with the software to go. One easy way to test the software out is to simply issue `git` to your terminal or Command Prompt. If you don't see a command not found error, you're good to go. Otherwise, have a look at how to install Git to your system.



Using Git on Windows

Like Python, Git doesn't come as part of a standard Windows installation. However, Windows implementations of the version control system can be downloaded and installed. You can download the official Windows Git client from the [Git website](#). The installer provides the `git` command line program, which we use in this crash course. You can also download a program called *TortoiseGit*, a graphical extension to the Windows Explorer shell. The program provides a really nice right-click Git context menu for files. This makes version control really easy to use. You can [download TortoiseGit](#) for free. Although we do not cover how to use TortoiseGit in this crash course, many tutorials exist online for it. Check [this tutorial](#) if you are interested in using it.

We recommend however that you stick to the command line program. We'll be using the commands in this crash course. Furthermore, if you switch to a UNIX/Linux development environment at a later stage, you'll be glad you know the commands!

Setting up your Git workspace is a straightforward process. Once everything is set up, you will begin to make sense of the directory structure that Git uses. Assume that you have signed up for a new account on [GitHub](#) and [created a new repository on the service](#) for your project. With your remote repository setup, follow these steps to get your local repository and workspace setup on your computer. We'll assume you will be working from your `<workspace>` directory.

1. Open a terminal and navigate to your home directory (e.g. `$ cd ~`).
2. *Clone* the remote repository - or in other words, make a copy of it. Check out how to do this below.
3. Navigate into the newly created directory. That's your workspace in which you can add files to be version controlled!

How to Clone a Remote Repository

Cloning your repository is a straightforward process with the `git clone` command. Supplement this command with the URL of your remote repository - and if required, authentication details, too. The URL of your repository varies depending on the provider you use. If you are unsure of the URL to enter, it may be worth querying it with your search engine or asking someone in the know.

For GitHub, try the following command, replacing the parts below as appropriate:

```
$ git clone https://<USER>:<PASS>@github.com/<OWNER>/<REPO_NAME>.git <workspace>
```

where you replace

- `<USER>` with your GitHub username;
- `<PASS>` with your GitHub password;
- `<OWNER>` with the username of the person who owns the repository;
- `<REPO_NAME>` with the name of your project's repository; and
- `<workspace>` with the name for your workspace directory. This is optional; leaving this option out will simply create a directory with the same name as the repository.

If all is successful, you'll see some text like the example shown below.

```
$ git clone https://github.com/leifos/tango_with_django_19
Cloning into 'tango_with_django_19'...
remote: Counting objects: 18964, done.
remote: Total 18964 (delta 0), reused 0 (delta 0), pack-reused 18964
Receiving objects: 100% (18964/18964), 99.69 MiB | 3.51 MiB/s, done.
Resolving deltas: 100% (13400/13400), done.
Checking connectivity... done.
```

If the output lines end with `done`, everything should have worked. Check your filesystem to see if the directory has been created.



Not using GitHub?

There are many websites that provide Git repositories - some free, some paid. While this chapter uses GitHub, you are free to use whatever service you wish. Other providers include [Atlassian Bitbucket](#) and [Unfuddle](#). You will of course have to change the URL from which you clone your repository if you use a service other than GitHub.

The Directory Structure

Once you have cloned your remote repository onto your local computer, navigate into the directory with your terminal, Command Prompt or GUI file browser. If you have cloned an empty repository the workspace directory should appear empty. This directory is therefore your blank workspace with which you can begin to add your project's files.

However, the directory isn't blank at all! On closer inspection, you will notice a hidden directory called `.git`. Stored within this directory are both the local repository and local index. **Do not alter the contents of the `.git` directory.** Doing so could damage your Git setup and break version control functionality. *Your newly created workspace therefore actually contains within it the local repository and index.*

Final Tweaks

With your workspace setup, now would be a good time to make some final tweaks. Here, we discuss two cool features you can try which could make your life (and your team members') a little bit easier.

When using your Git repository as part of a team, any changes you make will be associated with the username you use to access your remote Git repository. However, you can also specify your full name and e-mail address to be included with changes that are made by you on the remote repository. Simply open a Command Prompt or terminal and navigate to your workspace. From there, issue two commands: one to tell Git your full name, and the other to tell Git your e-mail address.

```
$ git config user.name "John Doe"
$ git config user.email "johndoe123@me.com"
```

Obviously, replace the example name and e-mail address with your own - unless your name actually is John Doe.

Git also provides you with the capability to stop - or ignore - particular files from being added to version control. For example, you may not wish a file containing unique keys to access web services from being added to version control. If the file were to be added to the remote repository, anyone could theoretically access the file by cloning the repository. With Git, files can be ignored by including them in the `.gitignore` file, which resides in the root of `<workspace>`. When adding files to version control, Git parses this file. If a file that is being added to version control is listed within `.gitignore`, the file is ignored. Each line of `.gitignore` should be a separate file entry.

Check out the following example of a `.gitignore` file:

```
1 `config/api_keys.py`
2 `*.pyc`
```

In this example file, there are two entries - one on each line. The first entry prompts Git to ignore the file `api_keys.py` residing within the `config` directory of your repository. The second entry then prompts Git to ignore *all* instance of files with a `.pyc` extension, or compiled Python files. This is a really nice feature: you can use *wildcards* to make generic entries if you need to!



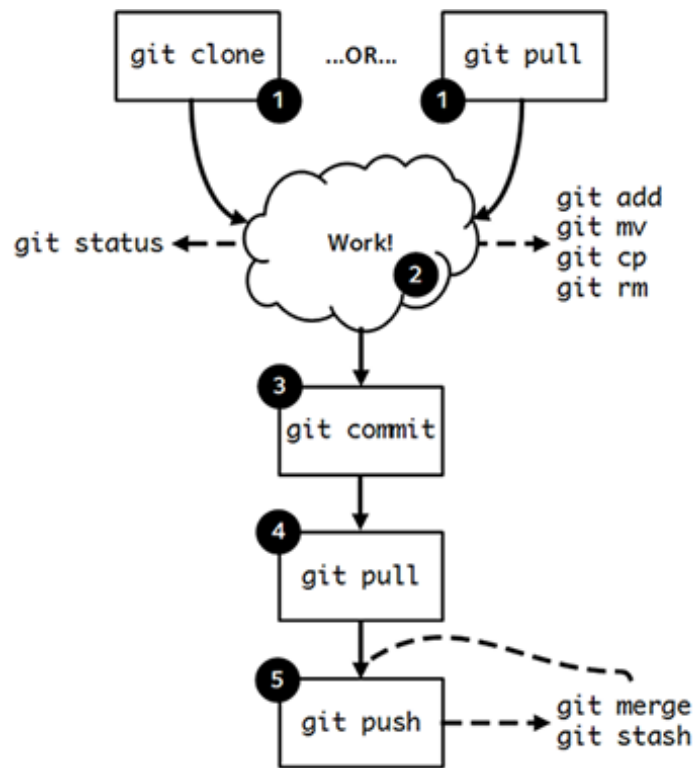
.gitignore - What else should I ignore?

There are many kinds of files you could safely ignore from being committed and pushed to your Git repositories. Examples include temporary files, databases (that can easily be recreated) and operating system-specific files. Operating system-specific files include configurations for the appearance of the directory when viewed in a given file browser. Windows computers create `thumbs.db` files, while OS X creates `.DS_Store` files.

When you create a new repository on GitHub, the service can offer to create a `.gitignore` file based upon the languages you will use in your project, which can save you some time setting everything up.

12.4 Basic Commands and Workflow

With your repository cloned and ready to go on your local computer, you're ready to get to grips with the Git workflow. This section shows you the basic Git workflow - and the associated Git commands you can issue.



We have provided a pictorial representation of the basic Git workflow as shown above. Match each of the numbers in the black circles to the numbered descriptions below to read more about each stage. Refer to this diagram whenever you're unsure about the next step you should take - it's very useful!

1. Starting Off

Before you can start work on your project, you must prepare Git. If you haven't yet sorted out your project's Git workspace, you'll need to [clone your repository to set it up](#).

If you've already cloned your repository, it's good practice to get into the habit of updating your local copy by using the `git pull` command. This *pulls* the latest changes from the remote repository onto your computer. By doing this, you'll be working from the same page as your team members. This will reduce the possibility of conflicting versions of files, which really does make your life a bit of a nightmare.

To perform a `git pull`, first navigate to your `<workspace>` directory within your Command Prompt or terminal, then issue `git pull`. Check out the snippet below from a Bash terminal to see exactly what you need to do, and what output you should expect to see.


```
$ cd <workspace>
$ git pull
remote: Counting objects: 3, done.
remote: Compressing objects: 100% (2/2), done.
remote: Total 3 (delta 0), reused 0 (delta 0), pack-reused 0
Unpacking objects: 100% (3/3), done.
From https://github.com/someuser/somerepository
   86a0b3b..a7cec3d  master    -> origin/master
Updating 86a0b3b..a7cec3d
Fast-forward
 README.md | 1 +
 1 file changed, 1 insertion(+)
 create mode 100644 README.md
```

This example shows that a `README.md` file has been updated or created from the latest pull.



Getting an Error?

If you receive `fatal: Not a git repository (or any of the parent directories): .git`, you're not in the correct directory. You need `cd` to your workspace directory - the one in which you cloned your repository to. A majority of Git commands only work when you're in a Git repository.



Pull before you Push!

Always `git pull` on your local copy of your repository before you begin work. **Always!**

2. Doing Some Work!

Once your workspace has been cloned or updated with the latest changes, it's time for you to get some work done! Within your workspace directory, you can take existing files and modify them. You can delete them too, or add new files to be version controlled.

When you modify your repository in any way, you need to keep Git up-to-date of any changes. Doing so allows Git to update your local index. The list of files stored within the local index are then used to perform your next *commit*, which we'll be discussing in the next step. To keep Git informed, there are several Git commands which let you update the local index. Three of the commands are near-identical to those that were discussed in the [Unix Crash Course](#) (e.g. `cp`, `mv`), with the addition of a `git` prefix.

- The first command `git add` allows you to request Git to add a particular file to the next commit for you. A common newcomer mistake is to assume that `git add` is used for adding new files to your repository only - *this is not the case. You must tell Git what modified files you wish to commit, too.* The command is invoked by typing `git add <filename>`, where `<filename>` is the name of the file you wish to add to your next commit. Multiple files and directories can be added with the command `git add .` - **but be careful with this.**
- `git mv` performs the same function as the Unix `mv` command - it moves files. The only difference between the two is that `git mv` updates the local index for you before moving the file. Specify the filename with the syntax `git mv <current_filename> <new_filename>`. For example, with this command you can move files to a different directory within your repository. This will be reflected in your next commit. The command is also used to rename files - from the old filename to the new.
- `git cp` allows you to make a copy of a file or directory while adding references to the new files into the local index for you. The syntax is the same as `git mv` above where the filename or directory name is specified thus: `git cp <current_filename> <copied_filename>`.
- The command `git rm` adds a file or directory delete request into the local index. While the `git rm` command does not delete the file straight away, the requested file or directory is removed from your filesystem and the Git repository upon the next commit. The syntax is similar to the `git add` command, where a filename can be specified thus: `git rm <filename>`. Note that you can add a large number of requests to your local index in one go, rather than removing each file manually. For example, `git rm -rf media/` creates delete requests in your local index for the `media/` directory. The `r` switch enables Git to *recursively* remove each file within the `media/` directory, while `f` allows Git to *forcibly* remove the files. Check out the [Wikipedia page](#) on the `rm` command for more information.

Lots of changes between commits can make things pretty confusing. You may easily forget what files you've already instructed Git to remove, for example. Fortunately, you can run the `git status` command to see a list of files which have been modified from your current working directory, but haven't been added to the local index for processing. Check out typical output from the command below to get a taste of what you can see.



Working with `.gitignore`

If you have [set up your `.gitignore` file correctly](#), you'll notice that files matching those specified within the `.gitignore` file are...ignored when you `git add` them. This is the intended behaviour - these files are not supposed to be committed to version control! If you however do need a file to be included that is in `.gitignore`, you can *force* Git to include it if necessary with the `git add -f <filename>` command.

```
$ git status
On branch master
Your branch is up-to-date with 'origin/master'.
Changes to be committed:
  (use "git reset HEAD <file>..." to unstage)
    modified:   chapter-unix.md
Changes not staged for commit:
  (use "git add <file>..." to update what will be committed)
  (use "git checkout -- <file>..." to discard changes in working directory)
    modified:   chapter-git.md
```

From this example above, we can see that the file `chapter-unix.md` has been added to the latest commit, and will therefore be updated in the next `git push`. The file `chapter-git.md` has been updated, but `git add` hasn't been run on the file, so the changes won't be applied to the repository.



Checking Status

For further information on the `git status` command, check out the [official Git documentation](#).

3. Committing your Changes

We've mentioned *committing* several times in the previous step - but what does it mean? Committing is when you save changes - which are listed in the local index - that you have made within your workspace. The more often you commit, the greater the number of opportunities you'll have to revert back to an older version of your code if things go wrong. Make sure you commit often, but don't commit an incomplete or broken version of a particular module or function. There's a lot of discussion as to when the ideal time to commit is. [Have a look at this Stack Overflow page](#) for the opinions of several developers. It does however make sense to commit only when everything is working. If you find you need to roll back to a previous commit only to find nothing works, you won't be too happy.

To commit, you issue the `git commit` command. Any changes to existing files that you have indexed will be saved to version control at this point. Additionally, any files that you've requested to be copied, removed, moved or added to version control via the local index will be undertaken at this point. When you commit, you are updating the *HEAD* of your local repository.



Commit Requirements

In order to successfully commit, you need to modify at least one file in your repository and instruct Git to commit it, through the `git add` command. See the previous step for more information on how to do this.

As part of a commit, it's incredibly useful to your future self and others to explain why you committed when you did. You can supply an optional message with your commit if you wish to do so. Instead of simply issuing `git commit`, run the following amended command.

```
$ git commit -m "Updated helpers.py to include a Unicode conversion function, str\
r_to_unicode()."
```

From the example above, you can see that using the `-m` switch followed by a string provides you with the opportunity to append a message to your commit. Be as explicit as you can, but don't write too much. People want to see at a glance what you did, and do not want to be bored or confused with a long essay. At the same time, don't be too vague. Simply specifying `Updated helpers.py` may tell a developer what file you modified, but they will require further investigation to see exactly what you changed.



Sensible Commits

Although frequent commits may be a good thing, you will want to ensure that what you have written actually *works* before you commit. This may sound silly, but it's an incredibly easy thing to not think about. To reiterate, committing code which doesn't actually work can be infuriating to your team members if they then rollback to a version of your project's codebase which is broken!

4. Synchronising your Repository



Important when Collaborating

Synchronising your local repository before making changes is crucial to ensure you minimise the chance for conflicts occurring. Make sure you get into the habit of doing a `pull` before you push.

After you've committed your local repository and committed your changes, you're just about ready to send your commit(s) to the remote repository by *pushing* your changes. However, what if someone within your group pushes their changes before you do? This means your local repository will be out of sync with the remote repository, meaning that any `git push` command that you issue will fail.

It's therefore always a good idea to check whether changes have been made on the remote repository before updating it. Running a `git pull` command will pull down any changes from the remote repository, and attempt to place them within your local repository. If no changes have been made, you're clear to push your changes. If changes have been made and cannot be easily rectified, you'll need to do a little bit more work.

In scenarios such as this, you have the option to *merge* changes from the remote repository. After running the `git pull` command, a text editor will appear in which you can add a comment explaining why the merge is necessary. Upon saving the text document, Git will merge the changes from the remote repository to your local repository.



Editing Merge Logs

If you do see a text editor on your Mac or Linux installation, it's probably the `vi` text editor. If you've never used `vi` before, check out [this helpful page containing a list of basic commands](#) on the Colorado State University Computer Science Department website. If you don't like `vi`, [you can change the default text editor](#) that Git calls upon. Windows installations most likely will bring up Notepad.

5. Pushing your Commit(s)

Pushing is the phrase used by Git to describe the sending of any changes in your local repository to the remote repository. This is the way in which your changes become available to your other team members, who can then retrieve them by running the `git pull` command in their respective local workspaces. The `git push` command isn't invoked as often as committing - *you require one or more commits to perform a push*. You could aim for one push per day, when a particular feature is completed, or at the request of a team member who is after your updated code.

To push your changes, the simplest command to run is:

```
$ git push origin master
```

As explained on [this Stack Overflow question and answer page](#) this command instructs the `git push` command to push your local master branch (where your changes are saved) to the *origin* (the remote server from which you originally cloned). If you are using a more complex setup involving [branching and merging](#), alter `master` to the name of the branch you wish to push.



Important Push?

If your `git push` is particularly important, you can also alert other team members to the fact they should really update their local repositories by pulling your changes. You can do this through a *pull request*. Issue one after pushing your latest changes by invoking the command `git request-pull master`, where `master` is your branch name (this is the default value). If you are using a service such as GitHub, the web interface allows you to generate requests without the need to enter the command. Check out [the official GitHub website's tutorial](#) for more information.

12.5 Recovering from Mistakes

This section presents a solution to a coder's worst nightmare: what if you find that your code no longer works? Perhaps a refactoring went terribly wrong, or something was changed by another team member without discussion. Whatever the reason, using a form of version control always gives you a last resort: rolling back to a previous commit. This section details how to do just that. We follow the information given from [this Stack Overflow](#) question and answer page.



Changes may be Lost!

You should be aware that this guide will rollback your workspace to a previous iteration. Any uncommitted changes that you have made will be lost, with a very slim chance of recovery! Be wary. If you are having a problem with only one file, you could always view the different versions of the files for comparison. Have a look [at this Stack Overflow page](#) to see how to do that.

Rolling back your workspace to a previous commit involves two steps: determining which commit to roll back to, and performing the rollback. To determine what commit to rollback to, you can make use of the `git log` command. Issuing this command within your workspace directory will provide a list of recent commits that you made, your name and the date at which you made the commit. Additionally, the message that is stored with each commit is displayed. This is where it is highly beneficial to supply commit messages that provide enough information to explain what is going on. Check out the following output from a `git log` invocation below to see for yourself.

```
commit 88f41317640a2b62c2c63ca8d755feb9f17cf16e          <- Commit hash
Author: John Doe <someaddress@domain.com>              <- Author
Date:   Mon Jul 8 19:56:21 2013 +0100                  <- Date/time
    Nearly finished initial version of the requirements chapter <- Message
commit f910b7d557bf09783b43647f02dd6519fa593b9f
Author: John Doe <someaddress@domain.com>
Date:   Wed Jul 3 11:35:01 2013 +0100
    Added in the Git figures to the requirements chapter.
commit c97bb329259ee392767b87cfe7750ce3712a8bdf
Author: John Doe <someaddress@domain.com>
Date:   Tue Jul 2 10:45:29 2013 +0100
    Added initial copy of Sphinx documentation and tutorial code.
commit 2952efa9a24dbf16a7f32679315473b66e3ae6ad
Author: John Doe <someaddress@domain.com>
Date:   Mon Jul 1 03:56:53 2013 -0700
    Initial commit
```

From this list, you can choose a commit to rollback to. For the selected commit, you must take the commit hash - the long string of letters and numbers. To demonstrate, the top (or HEAD) commit hash

in the example output above is `88f41317640a2b62c2c63ca8d755feb9f17cf16e`. You can select this in your terminal and copy it to your computer's clipboard.

With your commit hash selected, you can now rollback your workspace to the previous revision. You can do this with the `git checkout` command. The following example command would rollback to the commit with hash `88f41317640a2b62c2c63ca8d755feb9f17cf16e`.

```
$ git checkout 88f41317640a2b62c2c63ca8d755feb9f17cf16e .
```

Make sure that you run this command from the root of your workspace, and do not forget to include the dot at the end of the command! The dot indicates that you want to apply the changes to the entire workspace directory tree. After this has completed, you should then immediately commit with a message indicating that you performed a rollback. Push your changes and alert your collaborators - perhaps with a pull request. From there, you can start to recover from the mistake by putting your head down and getting on with your project.



Exercises

If you haven't undertaken what we've been discussing in this chapter already, you should go through everything now to ensure your Git repository is ready to go. To try everything out, you can create a new file `README.md` in the root of your `<workspace>` directory. The file [will be used by GitHub](#) to provide information on your project's GitHub homepage.

- Create the file, and write some introductory text to your project.
- Add the file to the local index upon completion of writing, and commit your changes.
- Push the new file to the remote repository and observe the changes on the GitHub website.

Once you have completed these basic steps, you can then go back and edit the readme file some more. Add, commit and push - and then try to revert to the initial version to see if it all works as expected.



There's More!

There are other more advanced features of Git that we have not covered in this chapter. Examples include **branching** and **merging**, which are useful for projects with different release versions, for example. There are many fantastic tutorials available online if you are interested in taking your super-awesome version control skills a step further. For more details about such features take a look at this [tutorial on getting started with Git](#), the [Git Guide](#) or [Learning about Git Branching](#).

However, if you're only using this chapter as a simple guide to getting to grips with Git, everything that we've covered should be enough. Good luck!