



TENSORFLOW

Module 1: Introduction to TensorFlow

This chapter provides a quick introduction to various features of TensorFlow, and some of the TensorFlow tools and projects that are included in the TensorFlow “family.” The intent of this chapter is to provide relevant TensorFlow information, how to visualize TensorFlow graphs via TensorBoard, how to invoke TensorFlow code in a browser via Jupyter notebooks, and where to train neural networks with free GPU support (Google Colaboratory). The material in this chapter will prepare you for Chapter 2, which provides a “foundation” of commonly used TensorFlow APIs (illustrated via short code samples) that are also in code samples in the remaining chapters of this book.

Module 2: Useful TensorFlow APIs

This chapter focuses on useful TensorFlow APIs that you will encounter in many TensorFlow code samples. In fact, you will use these APIs in code samples that involve more complex TensorFlow code that are beyond the scope of this book.

Module 3: TensorFlow Datasets

This chapter contains various code samples that illustrate how to use TensorFlow Datasets, which support a rich set of operators that can simplify your TensorFlow code and enable you to process very large datasets (i.e., datasets that are too large to fit in memory). You will learn about operators (such as `filter()` and `map()`) that you can specify via “method chaining” as part of the definition of a TF dataset. In addition, you’ll learn about TF estimators (in the `tf.estimator` namespace), TF layers (in the `tf.layers` namespace), and TFRecords

Module 4: Linear Regression

This chapter introduces linear regression, along with code samples of linear regression: first with NumPy APIs, and then with TensorFlow APIs. The code samples use an “incremental” approach, starting with simple examples that involve Python and NumPy code (often using the NumPy `linspace()` API). Then you will see comparable code samples involving TensorFlow code. In addition, the samples in the second half of the chapter usually involve concepts from the first half of the chapter.

Module 5: Logistic Regression

This chapter assumes that you have been exposed to activation functions (even if you do not fully understand them), and also that you have read the material in Chapter 4. In addition, you need a basic understanding of hidden layers in a neural network, which is not discussed in this book. Depending on your comfort level, you might need to read some preparatory material before diving into this chapter (there are many articles available online).

The rationale for the preceding assumptions is that this book is not for “absolute beginners” (as mentioned in the Preface), and providing a very detailed explanation for every concept in every code samples in every chapter would probably double the length of this book. However, this book belongs to the Pocket Primer series of books for advanced beginners, which assumes that readers have already been exposed (to a limited extent) to some of the concepts in this chapter.