


SEBASTIAN RASCHKA



# Introduction to Artificial Neural Networks and Deep Learning

A Practical Guide  
with Applications in Python

# Introduction to Artificial Neural Networks and Deep Learning

A Practical Guide with Applications in Python

Sebastian Raschka

This book is for sale at <http://leanpub.com/ann-and-deeplearning>

This version was published on 2017-07-23



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 - 2017 Sebastian Raschka

# Contents

<b>Website</b> . . . . .	<b>i</b>
<b>Acknowledgments</b> . . . . .	<b>ii</b>
<b>Appendix G - TensorFlow Basics</b> . . . . .	<b>1</b>
TensorFlow in a Nutshell . . . . .	1
Installation . . . . .	3
Computation Graphs . . . . .	3
Variables . . . . .	5
Placeholder Variables . . . . .	8
Saving and Restoring Models . . . . .	9
Naming TensorFlow Objects . . . . .	13
CPU and GPU . . . . .	17
Control Flow . . . . .	18
TensorBoard . . . . .	21

# Website

Please visit the [GitHub repository](https://github.com/rasbt/deep-learning-book)<sup>1</sup> to download code examples used in this book.

If you like the content, please consider supporting the work by buying a copy of the book on [Leanpub](https://leanpub.com/ann-and-deeplearning)<sup>2</sup>.

I would appreciate hearing your opinion and feedback about the book! Also, if you have any questions about the contents, please don't hesitate to get in touch with me via [mail@sebastianraschka.com](mailto:mail@sebastianraschka.com) or join the [mailing list](https://groups.google.com/forum/#!forum/ann-and-dl-book)<sup>3</sup>.

Happy learning!

*Sebastian Raschka*

---

<sup>1</sup><https://github.com/rasbt/deep-learning-book>

<sup>2</sup><https://leanpub.com/ann-and-deeplearning>

<sup>3</sup><https://groups.google.com/forum/#!forum/ann-and-dl-book>

# Acknowledgments

I would like to give my special thanks to the readers, who provided feedback, caught various typos and errors, and offered suggestions for clarifying my writing.

- Appendix A: Artem Sobolev, Ryan Sun
- Appendix B: Brett Miller, Ryan Sun
- Appendix D: Marcel Blattner, Ignacio Campabadal, Ryan Sun
- Appendix F: Guillermo Moncecchi, Ged Ridgway, Ryan Sun, Patric Hindenberger
- Appendix H: Brett Miller, Ryan Sun

# Appendix G - TensorFlow Basics

This appendix offers a brief overview of TensorFlow, an open-source library for numerical computation and deep learning. This section is intended for readers who want to gain a basic overview of this library before progressing through the hands-on sections that are concluding the main chapters.

The majority of *hands-on* sections in this book focus on TensorFlow and its Python API, assuming that you have TensorFlow  $\geq 1.2$  installed if you are planning to execute the code sections shown in this book.

In addition to glancing over this appendix, I recommend the following resources from TensorFlow's official documentation for a more in-depth coverage on using TensorFlow:

- [Download and setup instructions](#)<sup>4</sup>
- [Python API documentation](#)<sup>5</sup>
- [Tutorials](#)<sup>6</sup>
- [TensorBoard, an optional tool for visualizing learning](#)<sup>7</sup>

## TensorFlow in a Nutshell

At its core, TensorFlow is a library for efficient multidimensional array operations with a focus on deep learning. Developed by the Google Brain Team, TensorFlow was open-sourced on November 9th, 2015. And augmented by its convenient Python API layer, TensorFlow has gained much popularity and wide-spread adoption in industry as well as academia.

TensorFlow shares some similarities with NumPy, such as providing data structures and computations based on multidimensional arrays. What makes TensorFlow particularly suitable for deep learning, though, are its primitives for defining functions on tensors, the ability of parallelizing tensor operations, and convenience tools such as automatic differentiation.

---

<sup>4</sup>[https://www.tensorflow.org/get\\_started/os\\_setup](https://www.tensorflow.org/get_started/os_setup)

<sup>5</sup>[https://www.tensorflow.org/api\\_docs/python/](https://www.tensorflow.org/api_docs/python/)

<sup>6</sup><https://www.tensorflow.org/tutorials/>

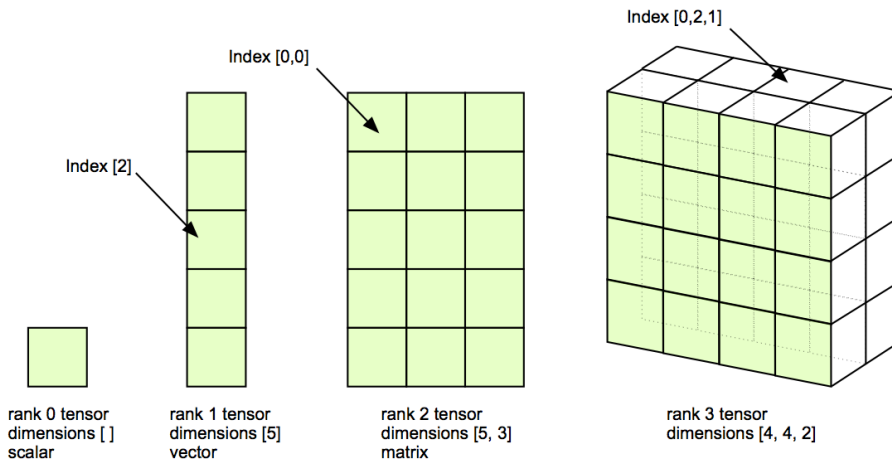
<sup>7</sup>[https://www.tensorflow.org/how\\_tos/summaries\\_and\\_tensorboard/](https://www.tensorflow.org/how_tos/summaries_and_tensorboard/)

While TensorFlow can be run entirely on a CPU or multiple CPUs, one of the core strength of this library is its support of GPUs (Graphical Processing Units) that are very efficient at performing highly parallelized numerical computations. In addition, TensorFlow also supports distributed systems as well as mobile computing platforms, including Android and Apple's iOS.

But what is a *tensor*? In simplifying terms, we can think of tensors as multidimensional arrays of numbers, as a generalization of scalars, vectors, and matrices.

1. Scalar:  $\mathbb{R}$
2. Vector:  $\mathbb{R}^n$
3. Matrix:  $\mathbb{R}^n \times \mathbb{R}^m$
4. 3-Tensor:  $\mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^p$
5. ...

When we describe tensors, we refer to its “dimensions” as the *rank* (or *order*) of a tensor, which is not to be confused with the dimensions of a matrix. For instance, an  $m \times n$  matrix, where  $m$  is the number of rows and  $n$  is the number of columns, would be a special case of a rank-2 tensor. A visual explanation of tensors and their ranks is given in the figure below.



Tensors

## Installation

Code conventions in this book follow the Python 3.x syntax, and while the code examples should be backward compatible to Python 2.7, I highly recommend the use of Python  $\geq 3.5$ .

Once you have your Python Environment set up ([Appendix - Python Setup](#)), the most convenient ways for installing TensorFlow are via `pip` or `conda` – the latter only applies if you have the Anaconda/Miniconda Python distribution installed, which I prefer and recommend.

Since TensorFlow is under active development, I recommend you to consult the official “[Download and Setup](#)<sup>8</sup>” documentation for detailed installation instructions to install TensorFlow on your operating system, macOS, Linux, or Windows.

## Computation Graphs

In contrast to other tools such as NumPy, the numerical computations in TensorFlow can be categorized into two steps: a construction step and an execution step. Consequently, the typical workflow in TensorFlow can be summarized as follows:

- Build a computational graph
- Start a new *session* to evaluate the graph
  - Initialize variables
  - Execute the operations in the compiled graph

Note that the computation graph has no numerical values before we initialize and evaluate it. To see how this looks like in practice, let us set up a new graph for computing the column sums of a matrix, which we define as a constant tensor (`reduce_sum` is the TensorFlow equivalent of NumPy’s `sum` function).

In [1]:

---

<sup>8</sup>[https://www.tensorflow.org/get\\_started/os\\_setup](https://www.tensorflow.org/get_started/os_setup)



```
1 import tensorflow as tf
2
3 g = tf.Graph()
4 with g.as_default() as g:
5     tf_x = tf.constant([[1., 2.],
6                         [3., 4.],
7                         [5., 6.]], dtype=tf.float32)
8     col_sum = tf.reduce_sum(tf_x, axis=0)
9
10 print('tf_x:\n', tf_x)
11 print('\ncol_sum:\n', col_sum)
```

Out [1]:

```
1 tf_x:
2   Tensor("Const:0", shape=(3, 2), dtype=float32)
3
4 col_sum:
5   Tensor("Sum:0", shape=(2,), dtype=float32)
```

As we can see from the output above, the operations in the graph are represented as Tensor objects that require an explicit evaluation before the `tf_x` matrix is populated with numerical values and its column sum gets computed.

Now, we pass the graph that we created earlier to a new, active *session*, where the graph gets compiled and evaluated:

In [2]:

```
1 with tf.Session(graph=g) as sess:
2     mat, csum = sess.run([tf_x, col_sum])
3
4 print('mat:\n', mat)
5 print('\ncsum:\n', csum)
```

Out [2]:

```
1 mat:
2  [[ 1.  2.]
3   [ 3.  4.]
4   [ 5.  6.]]
5
6 csum:
7  [ 9. 12.]
```

Note that if we are only interested in the result of a particular operation, we don't need to run its dependencies – TensorFlow will automatically take care of that. For instance, we can directly fetch the numerical values of `col_sum_times_2` in the active session without explicitly passing `col_sum` to `sess.run(...)` as the following example illustrates:

**In [3]:**

```
1 g = tf.Graph()
2 with g.as_default() as g:
3     tf_x = tf.constant([[1., 2.],
4                        [3., 4.],
5                        [5., 6.]], dtype=tf.float32)
6     col_sum = tf.reduce_sum(tf_x, axis=0)
7     col_sum_times_2 = col_sum * 2
8
9
10 with tf.Session(graph=g) as sess:
11     csum_2 = sess.run(col_sum_times_2)
12
13 print('csum_2:\n', csum_2)
```

**Out [3]:**

```
1 csum_2:
2  [array([ 18., 24.], dtype=float32)]
```

## Variables

Variables are constructs in TensorFlow that allows us to store and update parameters of our models in the current session during training. To define a “variable” tensor, we use TensorFlow's `Variable()` constructor, which looks similar to the use of `constant`

that we used to create a matrix previously. However, to execute a computational graph that contains variables, we must initialize all variables in the active session first (using `tf.global_variables_initializer()`), as illustrated in the example below.

**In [1]:**

```
1 import tensorflow as tf
2
3 g = tf.Graph()
4 with g.as_default() as g:
5     tf_x = tf.Variable([[1., 2.],
6                        [3., 4.],
7                        [5., 6.]], dtype=tf.float32)
8     x = tf.constant(1., dtype=tf.float32)
9
10    # add a constant to the matrix:
11    tf_x = tf_x + x
12
13 with tf.Session(graph=g) as sess:
14     sess.run(tf.global_variables_initializer())
15     result = sess.run(tf_x)
16
17 print(result)
```

**Out [1]:**

```
1 [[ 2.  3.]
2  [ 4.  5.]
3  [ 6.  7.]
```

Now, let us do an experiment and evaluate the same graph twice:

**In [2]:**

```
1 with tf.Session(graph=g) as sess:
2     sess.run(tf.global_variables_initializer())
3     result = sess.run(tf_x)
4     result = sess.run(tf_x)
```

**Out [2]:**

```
1 [[ 2.  3.]
2  [ 4.  5.]
3  [ 6.  7.]]
```

As we can see, the result of running the computation twice did not affect the numerical values fetched from the graph. To update or to assign new values to a variable, we use TensorFlow's assign operation. The function syntax of assign is `assign(ref, val, ...)`, where 'ref' is updated by assigning 'value' to it:

**In [3]:**

```
1 g = tf.Graph()
2 with g.as_default() as g:
3     tf_x = tf.Variable([[1., 2.],
4                        [3., 4.],
5                        [5., 6.]], dtype=tf.float32)
6     x = tf.constant(1., dtype=tf.float32)
7
8     update_tf_x = tf.assign(tf_x, tf_x + x)
9
10
11 with tf.Session(graph=g) as sess:
12     sess.run(tf.global_variables_initializer())
13     result = sess.run(update_tf_x)
14     result = sess.run(update_tf_x)
15
16 print(result)
```

**Out [3]:**

```
1 [[ 3.  4.]
2  [ 5.  6.]
3  [ 7.  8.]]
```

As we can see, the contents of the variable `tf_x` were successfully updated twice now; in the active session we

- initialized the variable `tf_x`
- added a constant scalar 1. to `tf_x` matrix via `assign`

- added a constant scalar `1.` to the previously updated `tf_x` matrix via `assign`

Although the example above is kept simple for illustrative purposes, variables are an important concept in TensorFlow, and we will see throughout the chapters, they are not only useful for updating model parameters but also for saving and loading variables for reuse.

## Placeholder Variables

Another important concept in TensorFlow is the use of placeholder variables, which allow us to feed the computational graph with numerical values in an active session at runtime.

In the following example, we will define a computational graph that performs a simple matrix multiplication operation. First, we define a placeholder variable that can hold 3x2-dimensional matrices. And after initializing the placeholder variable in the active session, we will use a dictionary, `feed_dict` we feed a NumPy array to the graph, which then evaluates the matrix multiplication operation.

In [1]:

```
1 import tensorflow as tf
2 import numpy as np
3
4 g = tf.Graph()
5 with g.as_default() as g:
6     tf_x = tf.placeholder(dtype=tf.float32,
7                           shape=(3, 2))
8
9     output = tf.matmul(tf_x, tf.transpose(tf_x))
10
11
12 with tf.Session(graph=g) as sess:
13     sess.run(tf.global_variables_initializer())
14     np_ary = np.array([[3., 4.],
15                       [5., 6.],
16                       [7., 8.]])
17     feed_dict = {tf_x: np_ary}
18     print(sess.run(output,
19                     feed_dict=feed_dict))
```

Out [1]:

```
1 [[ 25.  39.  53.]
2  [ 39.  61.  83.]
3  [ 53.  83. 113.]]
```

Throughout the main chapters, we will make heavy use of placeholder variables, which allow us to pass our datasets to various learning algorithms in the computational graphs.

## Saving and Restoring Models

Training deep neural networks requires a lot of computations and computational resources, and in practice, it would be infeasible to retrain our model each time we start a new TensorFlow session before we can use it to make predictions. In this section, we will go over the basics of saving and re-using the results of our TensorFlow models.

The most convenient way to store the main components of our model is to use TensorFlow's `Saver` class (`tf.train.Saver()`). To see how it works, let us reuse the simple example from the [Variables](#) section, where we added a constant `1.` to all elements in a `3x2` matrix:

**In [2]:**

```
1 g = tf.Graph()
2 with g.as_default() as g:
3
4     tf_x = tf.Variable([[1., 2.],
5                        [3., 4.],
6                        [5., 6.]], dtype=tf.float32)
7     x = tf.constant(1., dtype=tf.float32)
8
9     update_tf_x = tf.assign(tf_x, tf_x + x)
10
11     # initialize a Saver, which gets all variables
12     # within this computation graph context
13     saver = tf.train.Saver()
```

Now, after we initialized the graph above, let us execute its operations in a new session:

**In [3]:**

```
1 with tf.Session(graph=g) as sess:
2     sess.run(tf.global_variables_initializer())
3     result = sess.run(update_tf_x)
4
5     # save the model
6     saver.save(sess, save_path='./my-model.ckpt')
```

Notice the `saver.save` call above, which saves all variables in the graph to “checkpoint” files bearing the prefix `my-model.ckpt` in our local directory (`./`). Since we didn’t specify which variables we wanted to save when we instantiated a `tf.train.Saver()`, it saved all variables in the graph by default – here, we only have one variable, `tf_x`. Alternatively, if we are only interested in keeping particular variables, we can specify this by feeding `tf.train.Saver()` a dictionary or list of these variables upon instantiation. For example, if our graph contained more than one variable, but we were only interested in saving `tf_x`, we could instantiate a `saver` object as `tf.train.Saver([tf_x])`.

After we executed the previous code example, we should find the three `my-model.ckpt` files (in binary format) in our local directory:

- `my-model.ckpt.data-00000-of-00001`
- `my-model.ckpt.index`
- `my-model.ckpt.meta`

The file `my-model.ckpt.data-00000-of-00001` saves our main variable values, the `.index` file keeps track of the data structures, and the `.meta` file describes the structure of our computational graph that we executed.

Note that in our simple example above, we just saved our variable one single time. However, in real-world applications, we typically train models over multiple iterations or epochs, and it is useful to create intermediate checkpoint files during training so that we can pick up where we left off in case we need to interrupt our session or encounter unforeseen technical difficulties. For instance, by using the `global_step` parameter, we could save our results after each 10th iteration by making the following modification to our code:

**In [4]:**

```
1 g = tf.Graph()
2 with g.as_default() as g:
3
4     tf_x = tf.Variable([[1., 2.],
5                        [3., 4.],
6                        [5., 6.]], dtype=tf.float32)
7     x = tf.constant(1., dtype=tf.float32)
8
9     update_tf_x = tf.assign(tf_x, tf_x + x)
10
11     # initialize a Saver, which gets all variables
12     # within this computation graph context
13     saver = tf.train.Saver()
14
15 with tf.Session(graph=g) as sess:
16     sess.run(tf.global_variables_initializer())
17
18     for epoch in range(100):
19         result = sess.run(update_tf_x)
20         if not epoch % 10:
21             saver.save(sess,
22                       save_path='./my-model-multiple_ckpts.ckpt',
23                       global_step=epoch)
```

After we executed this code we find five `my-model.ckpt` files in our local directory:

- `my-model.ckpt-50` { `.data-00000-of-00001`, `.ckpt.index`, `.ckpt.meta`}
- `my-model.ckpt-60` { `.data-00000-of-00001`, `.ckpt.index`, `.ckpt.meta`}
- `my-model.ckpt-70` { `.data-00000-of-00001`, `.ckpt.index`, `.ckpt.meta`}
- `my-model.ckpt-80` { `.data-00000-of-00001`, `.ckpt.index`, `.ckpt.meta`}
- `my-model.ckpt-90` { `.data-00000-of-00001`, `.ckpt.index`, `.ckpt.meta`}

Although we saved our variables ten times, the saver only keeps the five most recent checkpoints by default to save storage space. However, if we want to keep more than five recent checkpoint files, we can provide an optional argument `max_to_keep=n` when we initialize the saver, where `n` is an integer specifying the number of the most recent checkpoint files we want to keep.

Now that we learned how to save TensorFlow Variables, let us see how we can restore them. Assuming that we started a fresh computational session, we need to specify the graph first. Then, we can use the saver's `restore` method to restore our variables as shown below:



In [5]:

```
1 g = tf.Graph()
2 with g.as_default() as g:
3
4     tf_x = tf.Variable([[1., 2.],
5                         [3., 4.],
6                         [5., 6.]], dtype=tf.float32)
7     x = tf.constant(1., dtype=tf.float32)
8
9     update_tf_x = tf.assign(tf_x, tf_x + x)
10
11     # initialize a Saver, which gets all variables
12     # within this computation graph context
13     saver = tf.train.Saver()
14
15 with tf.Session(graph=g) as sess:
16     saver.restore(sess, save_path='./my-model.ckpt')
17     result = sess.run(update_tf_x)
18     print(result)
```

Out [5]:

```
1 [[ 3.  4.]
2  [ 5.  6.]
3  [ 7.  8.]]
```

Notice that the returned values of the `tf_x` Variable are now increased by a constant of two, compared to the values in the computational graph. The reason is that we ran the graph one time before we saved the variable,

```
1 with tf.Session(graph=g) as sess:
2     sess.run(tf.global_variables_initializer())
3     result = sess.run(update_tf_x)
4
5     # save the model
6     saver.save(sess, save_path='./my-model.ckpt')
```

and we ran it a second time when after we restored the session.

Similar to the example above, we can reload one of our checkpoint files by providing the desired checkpoint suffix (here: `-90`, which is the index of our last checkpoint):

**In [6]:**

```
1 with tf.Session(graph=g) as sess:
2     saver.restore(sess, save_path='./my-model-multiple_ckpts.ckpt-90')
3     result = sess.run(update_tf_x)
4     print(result)
```

**Out [6]:**

```
1 [[ 93.  94.]
2  [ 95.  96.]
3  [ 97.  98.]]
```

In this section, we merely covered the basics of saving and restoring TensorFlow models. If you want to learn more, please take a look at the official [API documentation](#)<sup>9</sup> of TensorFlow's `Saver` class.

## Naming TensorFlow Objects

When we create new TensorFlow objects like `Variables`, we can provide an optional argument for their `name` parameter – for example:

```
1 tf_x = tf.Variable([[1., 2.],
2                   [3., 4.],
3                   [5., 6.]],
4                   name='tf_x_0',
5                   dtype=tf.float32)
```

Assigning names to `Variables` explicitly is not a requirement, but I personally recommend making it a habit when building (more) complex models. Let us walk through a scenario to illustrate the importance of naming variables, taking the simple example from the previous section and add new variable `tf_y` to the graph:

**In [1]:**

---

<sup>9</sup>[https://www.tensorflow.org/api\\_docs/python/tf/train/Saver](https://www.tensorflow.org/api_docs/python/tf/train/Saver)

```
1 import tensorflow as tf
2
3 g = tf.Graph()
4 with g.as_default() as g:
5
6     tf_x = tf.Variable([[1., 2.],
7                         [3., 4.],
8                         [5., 6.]], dtype=tf.float32)
9
10    tf_y = tf.Variable([[7., 8.],
11                       [9., 10.],
12                       [11., 12.]], dtype=tf.float32)
13
14    x = tf.constant(1., dtype=tf.float32)
15    update_tf_x = tf.assign(tf_x, tf_x + x)
16    saver = tf.train.Saver()
17
18 with tf.Session(graph=g) as sess:
19     sess.run(tf.global_variables_initializer())
20     result = sess.run(update_tf_x)
21
22     saver.save(sess, save_path='./my-model.ckpt')
```

The variable `tf_y` does not do anything in the code example above; we added it for illustrative purposes, as we will see in a moment. Now, let us assume we started a new computational session and loaded our saved `my-model` into the following computational graph:

**In [2]:**

```
1 g = tf.Graph()
2 with g.as_default() as g:
3
4     tf_y = tf.Variable([[7., 8.],
5                         [9., 10.],
6                         [11., 12.]], dtype=tf.float32)
7
8     tf_x = tf.Variable([[1., 2.],
9                         [3., 4.],
10                        [5., 6.]], dtype=tf.float32)
11
12    x = tf.constant(1., dtype=tf.float32)
13    update_tf_x = tf.assign(tf_x, tf_x + x)
```

```
14     saver = tf.train.Saver()
15
16     with tf.Session(graph=g) as sess:
17         saver.restore(sess, save_path='./my-model.ckpt')
18         result = sess.run(update_tf_x)
19         print(result)
```

What results do you expect after running the code snippet above?

**Out [2]:**

```
1  [[ 8.  9.]
2   [10. 11.]
3   [12. 13.]]
```

Unless you paid close attention on how we initialized the graph above, this result above surely was not the one you expected. What happened? Intuitively, we expected our session to print

```
1  [[ 3.  4.]
2   [ 5.  6.]
3   [ 7.  8.]]
```

The explanation behind this unexpected result is that we reversed the order of `tf_y` and `tf_x` in the graph above. TensorFlow applies a default naming scheme to all operations in the computational graph, unless we use do it explicitly via the `name` parameter – or in other words, we confused TensorFlow by reversing the order of two similar objects, `tf_y` and `tf_x`.

To circumvent this problem, we could give our variables specific names – for example, `'tf_x_0'` and `'tf_y_0'`:

**In [3]:**

```
1 g = tf.Graph()
2 with g.as_default() as g:
3
4     tf_x = tf.Variable([[1., 2.],
5                        [3., 4.],
6                        [5., 6.]],
7                        name='tf_x_0',
8                        dtype=tf.float32)
9
10    tf_y = tf.Variable([[7., 8.],
11                      [9., 10.],
12                      ]],
13                      name='tf_y_0',
14                      dtype=tf.float32)
15
16    x = tf.constant(1., dtype=tf.float32)
17    update_tf_x = tf.assign(tf_x, tf_x + x)
18    saver = tf.train.Saver()
19
20 with tf.Session(graph=g) as sess:
21     sess.run(tf.global_variables_initializer())
22     result = sess.run(update_tf_x)
23
24     saver.save(sess, save_path='./my-model.ckpt')
```

Then, even if we flip the order of these variables in a new computational graph, TensorFlow knows which values to use for each variable when loading our model – assuming we provide the corresponding variable names:

**In [4]:**

```
1 g = tf.Graph()
2 with g.as_default() as g:
3
4     tf_y = tf.Variable([[7., 8.],
5                        [9., 10.],
6                        ]],
7                        name='tf_y_0',
8                        dtype=tf.float32)
9
10    tf_x = tf.Variable([[1., 2.],
11                      [3., 4.],
```

```
12         [5., 6.]),
13         name='tf_x_0',
14         dtype=tf.float32)
15
16     x = tf.constant(1., dtype=tf.float32)
17     update_tf_x = tf.assign(tf_x, tf_x + x)
18     saver = tf.train.Saver()
19
20     with tf.Session(graph=g) as sess:
21         saver.restore(sess, save_path='./my-model.ckpt')
22         result = sess.run(update_tf_x)
23         print(result)
```

Out [4]:

```
1  [[ 3.  4.]
2   [ 5.  6.]
3   [ 7.  8.]
```

## CPU and GPU

Please note that all code examples in this book, and all TensorFlow operations in general, can be executed on a CPU. If you have a GPU version of TensorFlow installed, TensorFlow will automatically execute those operations that have GPU support on GPUs and use your machine's CPU, otherwise. However, if you wish to define your computing device manually, for instance, if you have the GPU version installed but want to use the main CPU for prototyping, we can run an active section on a specific device using the `with` context as follows

```
1  with tf.Session() as sess:
2      with tf.device("/gpu:1"):
```

where

- “/cpu:0”: The CPU of your machine.
- “/gpu:0”: The GPU of your machine, if you have one.
- “/gpu:1”: The second GPU of your machine, etc.
- etc.

You can get a list of all available devices on your machine via

```
1 from tensorflow.python.client import device_lib
2
3 device_lib.list_local_devices()
```

For more information on using GPUs in TensorFlow, please refer to the GPU documentation at [https://www.tensorflow.org/how\\_tos/using\\_gpu/](https://www.tensorflow.org/how_tos/using_gpu/).



Another good way to check whether your current TensorFlow session runs on a GPU is to execute

```
>>> import tensorflow as tf
>>> tf.test.gpu_device_name()
```

In your current Python session. If a GPU is available to TensorFlow, it will return a non-empty string; for example,  `'/gpu:0'` . Otherwise, if no GPU can be found, the function will return an empty string.

## Control Flow

It is important to discuss TensorFlow's control flow mechanics, the way it handles control statements such as `if/else` and `while`-loops. Control flow in TensorFlow is not a complicated topic, but it can be quite unintuitive at first and a common pitfall for beginners – especially, in the context of how control flow in Python is handled.

To explain control flow in TensorFlow in a practical manner, let us consider a simple example first. The following graph is meant to add the value `1.0` to a placeholder variable `x` if `addition=True` and subtract `1.0` from `x` otherwise:

**In [1]:**

```
1 import tensorflow as tf
2
3 addition = True
4
5 g = tf.Graph()
6 with g.as_default() as g:
7     x = tf.placeholder(dtype=tf.float32, shape=None)
8     if addition:
9         y = x + 1.
10    else:
11        y = x - 1.
```

Now, let us create a new session and execute the graph by feeding a `1.0` to the placeholder. If everything works as expected, the session should return the value `2.0` since `addition=True` and  $1.0 + 1.0 = 2.0$ :

**In [2]:**

```
1 with tf.Session(graph=g) as sess:
2     result = sess.run(y, feed_dict={x: 1.})
3
4 print('Result:\n', result)
```

**Out [2]:**

```
1 Result:
2 2.0
```

It appears that the session did return the same value that it returned when `addition` was set to `True`. Why did this happen? The explanation for this is that the `if/else` statements in the previous code only apply to the graph construction step. Or in other words, we created a graph by visiting the code contained under the `if` statement, and since TensorFlow graphs are static, we have no way of running the code under the `else` statement – except for setting `addition=False` and creating a new graph.

However, we do not have to create a new graph each time we want to include control statements – TensorFlow implements a variety of helper functions that help with control flow inside a graph. For instance, to accomplish the little exercise of conditionally adding or subtracting a one from the placeholder variable `x`, we could use `tf.cond` as follows:

**In [3]:**



```
1 addition = True
2
3 g = tf.Graph()
4 with g.as_default() as g:
5     addition = tf.placeholder(dtype=tf.bool, shape=None)
6     x = tf.placeholder(dtype=tf.float32, shape=None)
7
8     y = tf.cond(addition,
9                 true_fn=lambda: tf.add(x, 1.),
10                false_fn=lambda: tf.subtract(x, 1.))
```

The basic use of `tf.cond` for conditional execution comes with three important arguments: a condition to check (here, if `addition` is `True` or `False`), a function that gets executed if the condition is `True` (`true_fn`) and a function that gets executed if the condition is `False` (`false_fn`), respectively.

Next, let us repeat the little exercise from earlier and see if toggling the `addition` value between `True` and `False` affects the conditional execution that is now part of the graph:

**In [4]:**

```
1 with tf.Session(graph=g) as sess:
2     result = sess.run(y, feed_dict={addition:True,
3                                     x: 1.})
4
5 print('Result:\n', result)
```

**Out [4]:**

```
1 Result:
2 2.0
```

**In [5]:**

```
1 with tf.Session(graph=g) as sess:
2     result = sess.run(y, feed_dict={addition:False,
3                                     x: 1.})
4
5 print('Result:\n', result)
```

**Out [5]:**

```
1 Result:
2 0.0
```

Finally, we get the expected results

- “1.0 + 1.0 = 2.0” if addition=True
- “1.0 - 1.0 = 0.0” if addition=False

While this section provides you with the most important concept behind control flow in Python versus TensorFlow, there are many control statements (and logical operators) that we have not covered. Since the use of other control statements is analogous to `tf.cond`, I recommend you to visit [TensorFlow’s API documentation](#)<sup>10</sup>, which provides an overview of all the different operators for control flow and links to useful examples.

## TensorBoard

TensorBoard is one of the coolest features of TensorFlow, which provides us with a suite of tools to visualize our computational graphs and operations before and during runtime. Especially, when we are implementing large neural networks, our graphs can be quite complicated, and TensorBoard is only useful to visually track the training cost and performance of our network, but it can also be used as an additional tool for debugging our implementation. In this section, we will go over the basic concepts of TensorBoard, but make sure you also check out the [official documentation](#)<sup>11</sup> for more details.

To visualize a computational graph via TensorBoard, let us create a simple graph with two Variables, the tensors `tf_x` and `tf_y` with shape `[2, 3]`. The first operation is to add these two tensors together. Second, we transpose `tf_x` and multiply it with `tf_y`:

**In [1]:**

---

<sup>10</sup>[https://www.tensorflow.org/api\\_guides/python/control\\_flow\\_ops](https://www.tensorflow.org/api_guides/python/control_flow_ops)

<sup>11</sup>[https://www.tensorflow.org/how\\_tos/summaries\\_and\\_tensorboard/](https://www.tensorflow.org/how_tos/summaries_and_tensorboard/)

```
1 import tensorflow as tf
2
3 g = tf.Graph()
4 with g.as_default() as g:
5
6     tf_x = tf.Variable([[1., 2.],
7                        [3., 4.],
9                        [5., 6.]],
10                       name='tf_x_0',
11                       dtype=tf.float32)
12
13     tf_y = tf.Variable([[7., 8.],
14                        [9., 10.],
15                        [11., 12.]],
16                       name='tf_y_0',
17                       dtype=tf.float32)
18
19     output = tf_x + tf_y
20     output = tf.matmul(tf.transpose(tf_x), output)
```

If we want to visualize the graph via TensorBoard, we need to instantiate a new `FileWriter` object in our session, which we provide with a `logdir` and the graph itself. The `FileWriter` object will then write a `protobuf`<sup>12</sup> file to the `logdir` path that we can load into TensorBoard:

In [2]:

```
1 with tf.Session(graph=g) as sess:
2     sess.run(tf.global_variables_initializer())
3
4     # create FileWriter object that writes the logs
5     file_writer = tf.summary.FileWriter(logdir='logs/1', graph=g)
6     result = sess.run(output)
7     print(result)
```

Out [2]:

---

<sup>12</sup><https://developers.google.com/protocol-buffers/docs/overview>

```

1 [[ 124.  142.]
2  [ 160.  184.]]

```

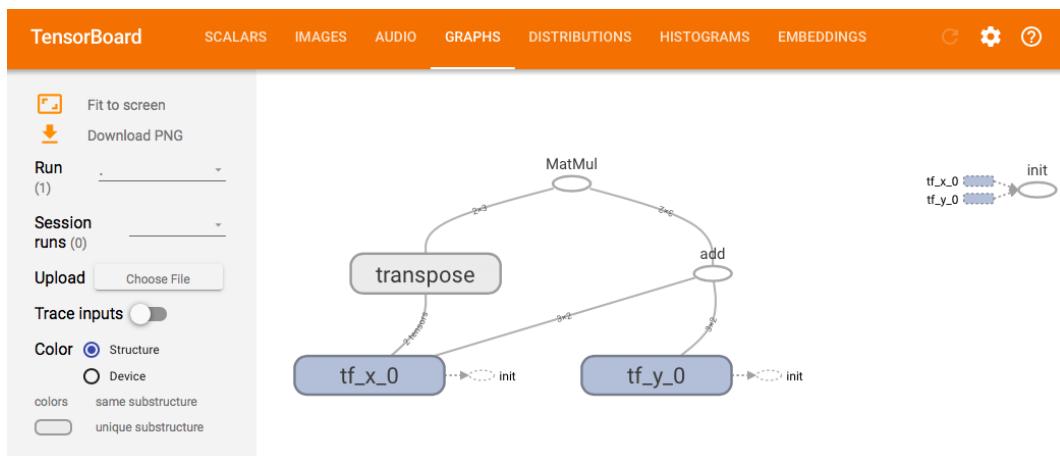
If you installed TensorFlow via pip, the `tensorboard` command should be available from your command line terminal. So, after running the preceding code examples for defining the graph and running the session, you just need to execute the command `tensorboard --logdir logs/1`. You should then see an output similar to the following:

```

1 Desktop Sebastian{$$} tensorboard --logdir logs/1
2 Starting TensorBoard b'41' on port 6006
3 (You can navigate to http://xxx.xxx.x.xx:6006)

```

Copy and paste the `http` address from the terminal and open it in your favorite web browser to open the TensorBoard window. Then, click on the Graph tab at the top, to visualize the computational graph as shown in the figure below:



TensorBoard

In our TensorBoard window, we can now see a visual summary of our computational graph (as shown in the screenshot above). The dark-shaded nodes labeled as `tf_x_0` and `tf_y_0` are the two variables we initialized, and following the connective lines, we can track the flow of operations. We can see the graph edges that are connecting `tf_x_0` and `tf_y_0` to an `add` node, with is the addition we defined in the graph, followed by the multiplication with the transpose of and the result of `add`.

Next, we are introducing the concept of `name_scopes`, which lets us organize different parts in our graph. In the following code example, we are going to take the initial code snippets and add with `tf.name_scope(...)` contexts as follows:

**In [3]:**

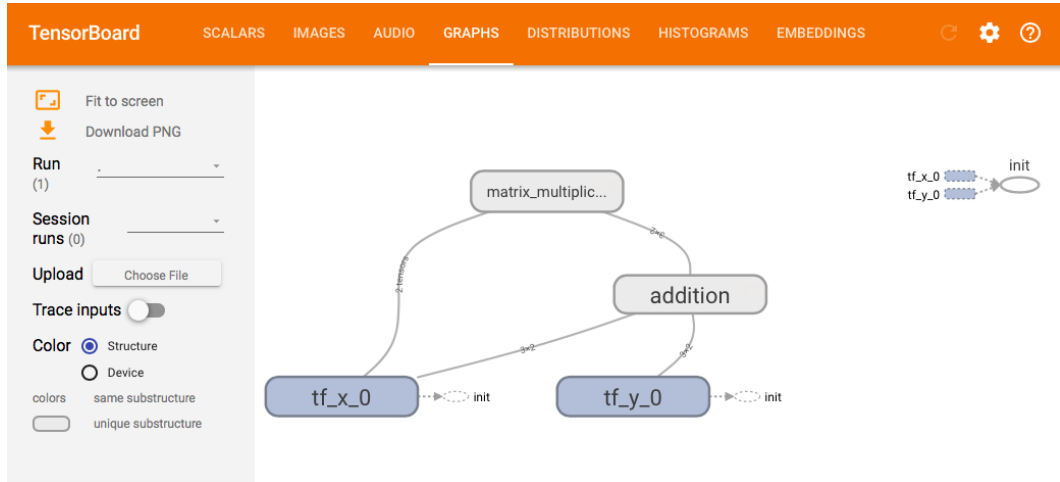
```
1 g = tf.Graph()
2 with g.as_default() as g:
3
4     tf_x = tf.Variable([[1., 2.],
5                         [3., 4.],
6                         [5., 6.]],
7                         name='tf_x_0',
8                         dtype=tf.float32)
9
10    tf_y = tf.Variable([[7., 8.],
11                       [9., 10.],
12                       [11., 12.]],
13                       name='tf_y_0',
14                       dtype=tf.float32)
15
16    # add custom name scope
17    with tf.name_scope('addition'):
18        output = tf_x + tf_y
19
20    # add custom name scope
21    with tf.name_scope('matrix_multiplication'):
22        output = tf.matmul(tf.transpose(tf_x), output)
23
24    with tf.Session(graph=g) as sess:
25        sess.run(tf.global_variables_initializer())
26        file_writer = tf.summary.FileWriter(logdir='logs/2', graph=g)
27        result = sess.run(output)
28        print(result)
```

**Out [3]:**

```
1 [[ 124.  142.]
2  [ 160.  184.]]
```

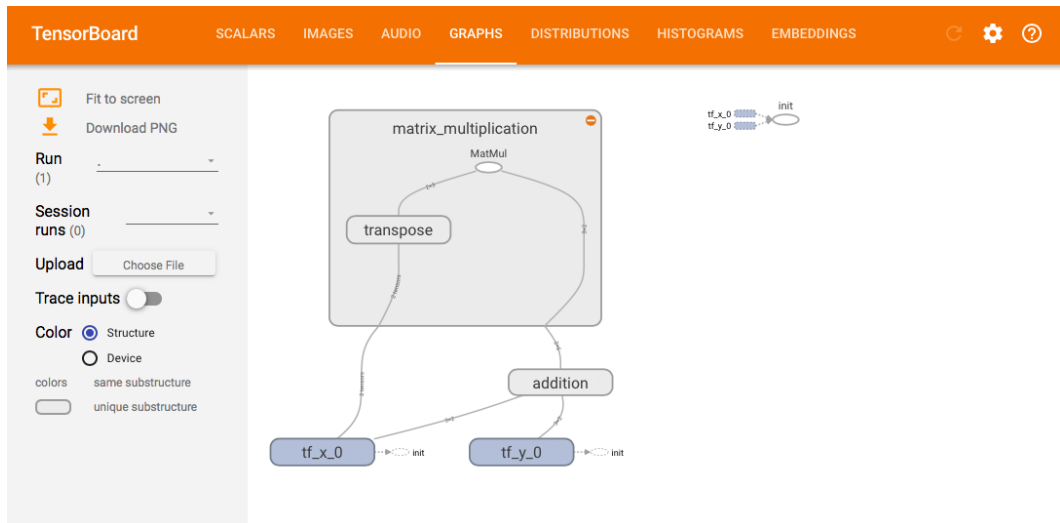
After executing the code example above, quit your previous TensorBoard session by pressing CTRL+C in the command line terminal and launch a new TensorBoard session via `tensorboard`

--logdir logs/2. After you refreshed your browser window, you should see the following graph:



TensorBoard

Comparing this visualization to our initial one, we can see that our operations have been grouped into our custom name scopes. If we double-click on one of these name scope summary nodes, we can expand it and inspect the individual operations in more details as shown for the `matrix_multiplication` name scope in the screenshot below:



TensorBoard

So far, we have only been looking at the computational graph itself. However, TensorBoard implements many more useful features. In the following example, we will make use of the “Scalar” and “Histogram” tabs. The “Scalar” tab in TensorBoard allows us to track scalar values over time, and the “Histogram” tab is useful for displaying the distribution of value in our tensor Variables (for instance, the model parameters during training). For simplicity, let us take our previous code snippet and modify it to demonstrate the capabilities of TensorBoard:

In [4]:

```

1 g = tf.Graph()
2 with g.as_default() as g:
3
4     some_value = tf.placeholder(dtype=tf.int32,
5                               shape=None,
6                               name='some_value')
7
8     tf_x = tf.Variable([[1., 2.],
9                       [3., 4.],
10                      [5., 6.]],
11                      name='tf_x_0',
12                      dtype=tf.float32)
13
14     tf_y = tf.Variable([[7., 8.],

```

```

15         [9., 10.],
16         [11., 12.]],
17         name='tf_y_0',
18         dtype=tf.float32)
19
20     with tf.name_scope('addition'):
21         output = tf_x + tf_y
22
23     with tf.name_scope('matrix_multiplication'):
24         output = tf.matmul(tf.transpose(tf_x), output)
25
26     with tf.name_scope('update_tensor_x'):
27         tf_const = tf.constant(2., shape=None, name='some_const')
28         update_tf_x = tf.assign(tf_x, tf_x * tf_const)
29
30     # create summaries
31     tf.summary.scalar(name='some_value', tensor=some_value)
32     tf.summary.histogram(name='tf_x_values', values=tf_x)
33
34     # merge all summaries into a single operation
35     merged_summary = tf.summary.merge_all()

```

Notice that we added an additional placeholder to the graph which later receives a scalar value from the session. We also added a new operation that updates our `tf_x` tensor by multiplying it with a constant 2.:

```

1     with tf.name_scope('update_tensor_x'):
2         tf_const = tf.constant(2., shape=None, name='some_const')
3         update_tf_x = tf.assign(tf_x, tf_x * tf_const)

```

Finally, we added the lines

```

1     # create summaries
2     tf.summary.scalar(name='some_value', tensor=some_value)
3     tf.summary.histogram(name='tf_x_values', values=tf_x)

```

at the end of our graph. These will create the “summaries” of the values we want to display in TensorBoard later. The last line of our graph is



```
1 merged_summary = tf.summary.merge_all()
```

which summarizes all the `tf.summary` calls to one single operation, so that we only have to fetch one variable from the graph when we execute the session. When we executed the session, we simply fetched this merged summary from `merged_summary` as follows:

```
1 result, summary = sess.run([update_tf_x, merged_summary],
2                             feed_dict={some_value: i})
```

Next, let us add a `for`-loop to our session that runs the graph five times, and feeds the counter of the range iterator to the `some_value` placeholder variable:

**In [5]:**

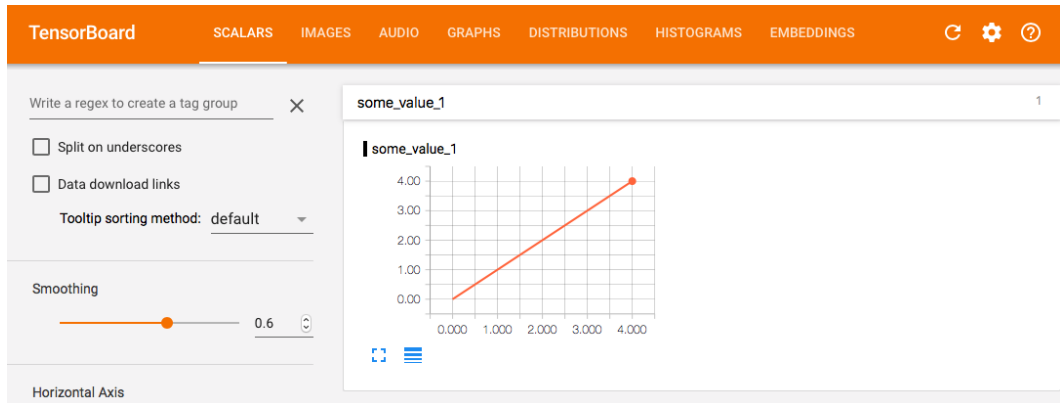
```
1 with tf.Session(graph=g) as sess:
2
3     sess.run(tf.global_variables_initializer())
4
5     # create FileWriter object that writes the logs
6     file_writer = tf.summary.FileWriter(logdir='logs/3', graph=g)
7
8     for i in range(5):
9         # fetch the summary from the graph
10        result, summary = sess.run([update_tf_x, merged_summary],
11                                   feed_dict={some_value: i})
12        # write the summary to the log
13        file_writer.add_summary(summary=summary, global_step=i)
14        file_writer.flush()
```

The two lines at the end of the preceding code snippet,

```
1     file_writer.add_summary(summary=summary, global_step=i)
2     file_writer.flush()
```

will write the summary data to our log file and the `flush` method updates TensorBoard. Executing `flush` explicitly is usually not necessary in real-world applications, but since the computations in our graph are so simple and “cheap” to execute, TensorBoard may not fetch the updates in real time. To visualize the results, quit your previous TensorBoard session (via `CTRL+C`) and execute `tensorboard --logdir logs/3` from the command line.

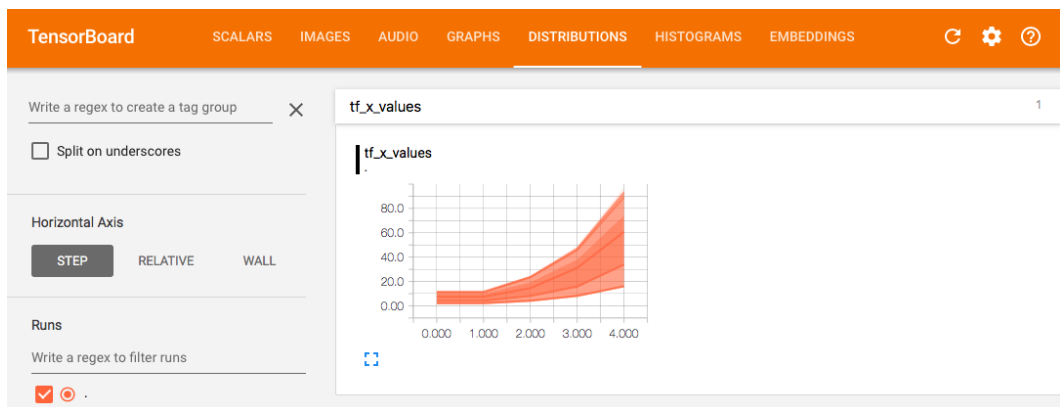
In the TensorBoard window under the tab “Scalar,” you should now see an entry called “some\_value\_1,” which refers to our `placeholder` in the graph that we called `some_value`. Since we just fed it the iteration index of our `for`-loop, we expect to see a linear graph with the iteration index on the x- and y-axis:



TensorBoard

Keep in mind that this is just a simple demonstration of how `tf.summary.scalar` works. For instance, more useful applications include the tracking of the training loss and the predictive performance of a model on training and validation sets throughout the different training rounds or epochs.

Next, let us go to the “Distributions” tab:

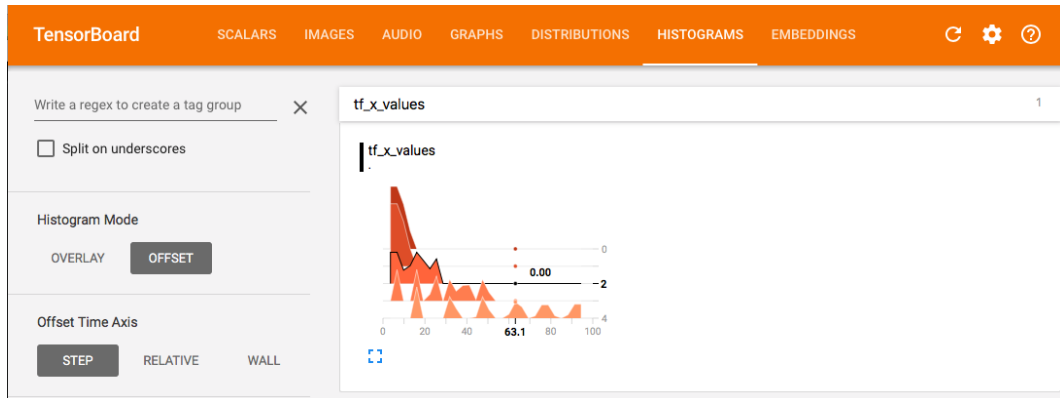


TensorBoard

The “Distributions” graph above shows us the distribution of values in `tf_x` for each step in

the for-loop. Since we doubled the value in the tensor after each for-loop iteration, we the distribution graph grows wider over time.

Finally, let us head over to the “Histograms” tab, which provides us with an individual histogram for each for-loop step that we can scroll through. Below, I selected the 3rd for-loop step that highlights the histogram of values in `tf_x` during this step:



TensorBoard

Since TensorBoard is such a highly visual tool with graph and data exploration in mind, I highly recommend you to take it for a test drive and explore it interactively. Also, there are several features that we haven't covered in this simple introduction to TensorBoard, so be sure to check out the [official documentation](https://www.tensorflow.org/how_tos/summaries_and_tensorboard/)<sup>13</sup> for more information.

---

<sup>13</sup>[https://www.tensorflow.org/how\\_tos/summaries\\_and\\_tensorboard/](https://www.tensorflow.org/how_tos/summaries_and_tensorboard/)