

# Introduction To TensorFlow

Originalartikel

Backup

<html> <p>I had great fun writing neural network software in the 90s, and I have been anxious to try creating some using <a href=„<https://www.tensorflow.org/>“ target=„\_blank“>TensorFlow</a>.</p> <p>Google's machine intelligence framework is the new hotness right now. And when TensorFlow <a href=„<https://github.com/samjabrahams/tensorflow-on-raspberry-pi>“ target=„\_blank“>became installable on the Raspberry Pi</a>, working with it became very easy to do. In a short time I made a neural network that counts in binary. So I thought I'd pass on what I've learned so far. Hopefully this makes it easier for anyone else who wants to try it, or for anyone who just wants some insight into neural networks.</p> <p> <h2>What Is TensorFlow?</h2> <p>To quote the TensorFlow website, TensorFlow is an &#8220;open source software library for numerical computation using data flow graphs&#8221;.&#160;What do we mean by &#8220;data flow graphs&#8221;? Well, that's the really cool part. But before we can answer that, we'll need to talk a bit about the structure for a simple neural network.</p> <figure id=„attachment\_251874“ style=„width: 801px“ class=„wp-caption aligncenter“><img data-attachment-id=„251874“ data-permalink=„[https://hackaday.com/2017/04/11/introduction-to-tensorflow/bincounter\\_neural\\_network/](https://hackaday.com/2017/04/11/introduction-to-tensorflow/bincounter_neural_network/)“ data-orig-file=„[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_neural\\_network.jpg?w=801&h=441](https://hackadaycom.files.wordpress.com/2017/04/bincounter_neural_network.jpg?w=801&h=441)“ data-orig-size=„777,428“ data-comments-opened=„1“ data-image-meta=„{&quot;aperture&quot;:&quot;0&quot;,&quot;credit&quot;:&quot;?????&quot;,&quot;camera&quot;:&quot;,&quot;caption&quot;:&quot;,&quot;created\_timestamp&quot;:&quot;0&quot;,&quot;copyright&quot;:&quot;,&quot;focal\_length&quot;:&quot;0&quot;,&quot;iso&quot;:&quot;0&quot;,&quot;shutter\_speed&quot;:&quot;0&quot;,&quot;title&quot;:&quot;,&quot;orientation&quot;:&quot;0&quot;};“ data-image-title=„Binary counter neural network“ data-image-description=„“ data-medium-file=„[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_neural\\_network.jpg?w=801&h=441?w=400](https://hackadaycom.files.wordpress.com/2017/04/bincounter_neural_network.jpg?w=801&h=441?w=400)“ data-large-file=„[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_neural\\_network.jpg?w=801&h=441?w=777](https://hackadaycom.files.wordpress.com/2017/04/bincounter_neural_network.jpg?w=801&h=441?w=777)“ class=„wp-image-251874“ src=„[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_neural\\_network.jpg?w=801&h=441](https://hackadaycom.files.wordpress.com/2017/04/bincounter_neural_network.jpg?w=801&h=441)“ alt=„Binary counter neural network“ width=„801“ height=„441“ srcset=„[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_neural\\_network.jpg?w=777](https://hackadaycom.files.wordpress.com/2017/04/bincounter_neural_network.jpg?w=777), [https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_neural\\_network.jpg?w=250&h=138](https://hackadaycom.files.wordpress.com/2017/04/bincounter_neural_network.jpg?w=250&h=138) 250w, [https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_neural\\_network.jpg?w=400&h=220](https://hackadaycom.files.wordpress.com/2017/04/bincounter_neural_network.jpg?w=400&h=220) 400w, [https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_neural\\_network.jpg?w=768&h=423](https://hackadaycom.files.wordpress.com/2017/04/bincounter_neural_network.jpg?w=768&h=423) 768w“ sizes=„(max-width: 801px) 100vw, 801px“><figcaption class=„wp-caption-text“>Binary counter neural network</figcaption></figure> <h2>Basics of a Neural Network</h2> <p>A simple neural network has some input units where the input goes.&#160;It also has hidden units, so-called


because from a user's perspective they're literally hidden. And there are output units, from which we get the results. Off to the side are also bias units, which are there to help control the values emitted from the hidden and output units. Connecting all of these units are a bunch of weights, which are just numbers, each of which is associated with two units. </p> <p>The way we instill intelligence into this neural network is to assign values to all those weights. That's what training a neural network does, find suitable values for those weights. Once trained, in our example, we'll set the input units to the binary digits 0, 0, and 0 respectively, TensorFlow will do stuff with everything in between, and the output units will magically contain the binary digits 0, 0, and 1 respectively. In case you missed that, it knew that the next number after binary 000 was 001. For 001, it should spit out 010, and so on up to 111, wherein it'll spit out 000. Once those weights are set appropriately, it'll know how to count. </p> <figure id="attachment\_251998" style="width: 410px" class="wp-caption alignright"><figcaption class="wp-caption-text">Binary counter neural network with matrices</figcaption></figure><p>One step in

&#8220;running&#8221; the neural network is to multiply the value of each weight by the value of its input unit, and then to store the result in the associated hidden unit. </p> <p>We can redraw the units and weights as arrays, or what are called lists in Python. From a math standpoint, they're matrices. We've redrawn only a portion of them in the diagram. Multiplying the input matrix with the weight matrix involves simple matrix multiplication resulting in the five element hidden matrix/list/array. </p> <h2

>From Matrices to Tensors</h2> <p>In TensorFlow, those lists are called


tensors. And the matrix multiplication step is called an operation, or op in programmer-speak, a term you'll have to get used to if you plan on reading the TensorFlow documentation. Taking it further, the whole neural network is a collection of tensors and the ops that operate on them. Altogether they make up a graph.

[https://hackaday.com/2017/04/11/introduction-to-tensorflow/tb\\_bincounter\\_full\\_graph/](https://hackaday.com/2017/04/11/introduction-to-tensorflow/tb_bincounter_full_graph/)




Binary counter's full graph

[https://hackaday.com/2017/04/11/introduction-to-tensorflow/tb\\_bincounter\\_expanded\\_graph/](https://hackaday.com/2017/04/11/introduction-to-tensorflow/tb_bincounter_expanded_graph/)



layer1 expanded

[https://hackaday.com/2017/04/11/introduction-to-tensorflow/tb\\_bincounter\\_expanded\\_graph/](https://hackaday.com/2017/04/11/introduction-to-tensorflow/tb_bincounter_expanded_graph/)



layer1 expanded

src=„[https://i2.wp.com/hackaday.com.files.wordpress.com/2017/04/tb\\_bincounter\\_expanded\\_graph.jpg?w=410&h=608&crop&ssl=1](https://i2.wp.com/hackaday.com.files.wordpress.com/2017/04/tb_bincounter_expanded_graph.jpg?w=410&h=608&crop&ssl=1)“ width=„410“ height=„608“ data-original-width=„410“ data-original-height=„608“ itemprop=„<http://schema.org/image>“ title=„layer1 expanded“ alt=„layer1 expanded“ style=„width: 410px; height: 608px;“></a> <div class=„tiled-gallery-caption“ itemprop=„caption description“> layer1 expanded </div> </div> </div> <!-- close group --> </div> <!-- close row --> </div> <p>Shown here are snapshots taken of <a href=„<https://hackaday.com/2017/03/24/ten-minute-tensorflow-speech-recognition/>“>TensorBoard, a tool for visualizing the graph as well as examining tensor values during and after training. The tensors are the lines, and written on the lines are the tensor&#8217;s&#8217;dimensions. Connecting the tensors are all the ops, though some of the things you see can be double-clicked on in order to expand for more detail, as we&#8217;ve done for layer1 in the second snapshot.</p> <p>At the very bottom is x, the name we&#8217;ve given for a placeholder op that allows us to provide values for the input tensor. The line going up and to the left from it is the input tensor. Continue following that line up and you&#8217;ll find the MatMul op, which does the matrix multiplication with that input tensor&#8217;and the tensor which is the other line leading into the MatMul op.&#8217;That tensor represents&#8217;the weights.</p> <p>All this was just to give you a feel for what a graph and its tensors and ops are, giving you a better idea of what we mean by&#8217;TensorFlow being a &#8220;software library for numerical computation using data flow graphs&#8221;. But why we would want to create these graphs?</p> <h2>Why Create Graphs?</h2> <p>The API that&#8217;s currently stable is one for Python, an interpreted language. Neural networks are compute intensive and a large one could&#8217;have thousands or even millions of weights. Computing by interpreting every step would take forever.</p> <p>So we instead create a graph made up of tensors and ops, describing the layout of the neural network, all mathematical operations, and&#8217;even initial values for variables. Only after we&#8217;ve created this graph do we then pass it to what TensorFlow calls a session. This is known as deferred execution. The session runs the graph using very efficient code. Not only that, but many of the operations, such as matrix multiplication, are ones that <a href=„[https://www.tensorflow.org/tutorials/using\\_gpu](https://www.tensorflow.org/tutorials/using_gpu)“ target=„\_blank“>can be done on a supported GPU</a> (Graphics Processing Unit) and the session will do that for you. Also, TensorFlow is built to be able to distribute the processing across multiple machines and/or GPUs. Giving it the complete graph allows it to do that.</p> <h2>Creating The Binary Counter Graph</h2> <p>And here&#8217;s the code for our binary counter neural network. You can find the full source code on <a href=„[https://github.com/Hack-a-Day/bincounter\\_TensorFlow\\_example/](https://github.com/Hack-a-Day/bincounter_TensorFlow_example/)“ target=„\_blank“>this GitHub page</a>. Note that&#8217;there&#8217;s additional code in it for saving information for use with&#8217;TensorBoard.</p> <p>We&#8217;ll start with the code for creating the graph of tensors and ops.</p> <pre style=„padding-left:30px;“>import tensorflow as tf<#13; sess = tf.InteractiveSession()<#13; &#13; NUM\_INPUTS = 3<#13; NUM\_HIDDEN = 5<#13; NUM\_OUTPUTS = 3<#13; </pre> <p>We first import the

## tensorflow

module, create a session for use&#8217;later, and, to make our code more understandable, we create a few variables containing the number of units in our network.</p> <pre style=„padding-left:30px;“>x = tf.placeholder(tf.float32, shape=[None, NUM\_INPUTS], name='x')<#13; y\_ = tf.placeholder(tf.float32, shape=[None, NUM\_OUTPUTS], name='y\_')<#13; </pre> <p>Then we create placeholders for our input and output units. A placeholder&#8217;is a TensorFlow op for things that we&#8217;ll provide values for later.&#8216;

## x

and

y\_

are now tensors in a new graph and each has a

placeholder

op associated with it. You might wonder why we define the shapes as

[None, NUM\_INPUTS]

and

[None, NUM\_OUTPUTS]

, two dimensional lists, and why

None

for the first dimension? In the overview of neural networks above it looks like we'll give it one input at a time and train it to produce a given output. It's more efficient though, if we give it multiple input/output pairs at a time, what's called a batch. The first dimension is for the number of input/output pairs in each batch. We won't know how many are in a batch until we actually give one later. And in fact, we're using the same graph for training, testing, and for actual usage so the batch size won't always be the same. So we use the Python placeholder object

None

for the size of the first dimension for now. 

```
W_fc1 = tf.truncated_normal([NUM_INPUTS, NUM_HIDDEN], mean=0.5, stddev=0.707); W_fc1 = tf.Variable(W_fc1, name='W_fc1'); b_fc1 = tf.truncated_normal([NUM_HIDDEN], mean=0.5, stddev=0.707); b_fc1 = tf.Variable(b_fc1, name='b_fc1'); h_fc1 = tf.nn.relu(tf.matmul(x, W_fc1) + b_fc1);
```

 That's followed by creating layer one of the neural network graph: the weights

W\_fc1

, the biases

b\_fc1

, and the hidden units

h\_fc1

. The &#8220;fc&#8221; is a convention meaning &#8220;fully connected&#8221;,&#160;since the weights connect every input unit to every hidden unit.

## tf.truncated\_normal

results in a number of ops and tensors which will later assign normalized, random numbers to all the weights.

### Variable

ops are given a value to do initialization with, random numbers; in this case, and keep their data across multiple runs. They're also handy for saving the neural network to a file, something you'll want to do once it's trained.

You can see where we'll be doing the matrix multiplication using the;

### matmul

op. We also insert an

### add

op which will add on the bias weights. The

### relu

op performs what we call an activation function. The matrix multiplication and the addition are linear operations. There's a very limited number of things a neural network can learn using just linear operations. The activation function provides some non-linearity. In the case of the relu activation function, it sets any values that are less than zero to zero, and all other values are left unchanged. Believe it or not, doing that opens up a

whole other world of things that can be learned.

```
W_fc2 = tf.truncated_normal([NUM_HIDDEN, NUM_OUTPUTS], mean=0.5, stddev=0.707); W_fc2 =
```

```
tf.Variable(W_fc2, name='W_fc2');
```

```
b_fc2 = tf.truncated_normal([NUM_OUTPUTS], mean=0.5, stddev=0.707); b_fc2 =
```

```
tf.Variable(b_fc2, name='b_fc2');
```

```
y = tf.matmul(h_fc1, W_fc2) + b_fc2;
```

The weights and biases for layer two are set up the same as for layer one but the output layer is different. We again will do a matrix multiplication, this time multiplying the weights and the hidden units, and then adding the bias weights.

We've left the activation function for the next bit of code.

```
results = tf.sigmoid(y, name='results');
```

```
cross_entropy = tf.reduce_mean(tf.nn.sigmoid_cross_entropy_with_logits(logits=y, labels=y_));
```

Sigmoid is another activation function, like the relu we encountered above, there to provide non-linearity. I used sigmoid here partly because

[the sigmoid equation](https://en.wikipedia.org/wiki/Sigmoid_function) results in values between 0 and 1, ideal for our binary counter example. I also used it because it's good for outputs where more than one output unit can have a large value. In our case, to represent the binary number 111, all the output units can have large values. When doing image classification we'd want something quite different, we'd want just one output unit to fire with a large value. For example, we'd want the output unit representing giraffes to have a large value if an image contains a giraffe. Something like

[softmax](https://en.wikipedia.org/wiki/Softmax_function) would be a good choice for image classification.

On close inspection, it looks like there's some duplication. We seem to be inserting sigmoid twice. We're actually creating two different, parallel outputs here. The

## cross\_entropy

&#160;tensor will be used during training of the neutral network.&#160;The

## results

tensor will be used when we run our trained neural network later for whatever purpose it&#8217;s created, for fun in our case.&#160;I don&#8217;t know if this is the best way of doing this, but it&#8217;s the&#160;way I came up with.</p> <pre style=„padding-left:30px;“>train\_step = tf.train.RMSPropOptimizer(0.25, momentum=0.5).minimize(cross\_entropy)&#13; </pre> <p>The last piece we add to our graph is the training. This is the&#160;op or ops that will adjust all the weights based on training data.&#160;Remember, we&#8217;re still just creating a graph here. The actual training will happen later when we run the graph.</p> <p>There are a few optimizers to chose from. I chose

## tf.train.RMSPropOptimizer

because,&#160;like the sigmoid, it works well for cases where all output values can be large. For classifying things as when doing image classification,&#160;

## tf.train.GradientDescentOptimizer

might be better.</p> <h2>Training And Using The Binary Counter</h2> <p>Having created the graph, it&#8217;s time to do the training. Once it&#8217;s&#160;trained, we can then use it.</p> <pre style=„padding-left:30px;“>inputvals = 13; [1, 1, 0], [1, 1, 1]&#13; targetvals = 13; [1, 1, 1], [0, 0,&#13; </pre> <p>First, we have some training data:

## inputvals

and

## targetvals

## inputvals

&#160;contains the inputs, and for each one there&#8217;s a corresponding

## targetvals

&#160;target value. For

## inputvals[0]

&#160;we have

## [0, 0, 0]

, and the expected output is

```
targetvals[0]
```

, which is

```
[0, 0, 1]
```

, and so on.</p> <pre style="padding-left:30px;">if do\_training == 1:&#13;

```
sess.run(tf.global_variables_initializer())&#13;
```

```
&#13;
```

```
for i in range(10001):&#13;
    if i%100 == 0:&#13;
        train_error = cross_entropy.eval(feed_dict={x: inputvals,
y_:targetvals})&#13;
        print("step %d, training error %g"%(i, train_error))&#13;
        if train_error < 0.0005:&#13;
            break&#13;
```

```
&#13;
```

```
sess.run(train_step, feed_dict={x: inputvals, y_: targetvals})&#13;
```

```
&#13;
```

```
if save_trained == 1:&#13;
    print("Saving neural network to %s.*"%(save_file))&#13;
    saver = tf.train.Saver()&#13;
    saver.save(sess, save_file)&#13;
```

</pre> <p>

```
do_training
```

and

```
save_trained
```

can be hardcoded, and changed for each use,&#160;or can be set using command line arguments.</p> <p>We first go through all those

```
Variable
```

ops and have them initialize their tensors.</p> <p>Then, for up to 10001 times we run the graph from the bottom up to&#160;the

**train\_step**

tensor, the last thing we added to our graph. We pass

**inputvals**

and

**targetvals**

to

**train\_step**

&#8216;s op or ops, which we’d added using

**RMSPropOptimizer**

. This is the step that adjusts all the weights such that the given inputs will result in something close to&#160;the corresponding target outputs. If the error between target outputs and actual outputs gets small enough sooner, then we break out of the loop. If you have thousands of input/output pairs then you could give it a subset of them at a time, the batch we spoke of earlier. But here we have only eight, and so we give all of them each time. If we want to, we can also save the network to a file. Once&#160;it&#8217;s trained well, we don&#8217;t need to train it again.

```
>else: # if we're not training then we must be loading from file
```

```
print("Loading neural network from %s"%(save_file))#13;
saver = tf.train.Saver()#13;
saver.restore(sess, save_file)#13;
# Note: the restore both loads and initializes the variables#13;
```

If we’re not training it then we instead load the trained network&#160;from a file. The file contains only the values for the tensors&#160;that have

**Variable**

ops. It doesn&#8217;t contain the structure of the&#160;graph. So even when running an already trained graph, we still need&#160;the code to create the graph. There is a way to save and load graphs&#160;from files using [https://www.tensorflow.org/programmers\\_guide/meta\\_graph](https://www.tensorflow.org/programmers_guide/meta_graph) but we’re not doing that here.

```
>print('\nCounting starting with: 0 0 0')#13;
res = sess.run(results, feed_dict={x: 0, 0, 0})#13;
print('%g %g %g'%(res[0][0], res[0][1], res[0][2]))#13;
for i in range(8):#13;
    res = sess.run(results, feed_dict={x: res})#13;
    print('%g %g %g'%(res[0][0], res[0][1], res[0][2]))#13;
```

In either case we try it out. Notice that we’re running it from&#160;the bottom of the graph up to the results tensor we’d talked about above, the duplicate output we’d created especially for when&#160;making use of the trained network.

We give it 000, and hope that it returns something close to 001. We pass what was returned, back in and run it again. Altogether we run it 9 times, enough times to count from 000 to 111 and then back to 000 again.

aligncenter" > <img data-attachment-id=„251879“  
data-  
permalink=„[https://hackaday.com/2017/04/11/introduction-to-tensorflow/bincounter\\_py\\_program\\_output\\_wide/](https://hackaday.com/2017/04/11/introduction-to-tensorflow/bincounter_py_program_output_wide/)“  
data-  
orig-  
file=„[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_py\\_program\\_output\\_wide.jpg?w=800&h=303](https://hackadaycom.files.wordpress.com/2017/04/bincounter_py_program_output_wide.jpg?w=800&h=303)“ data-orig-size=„832,315“ data-comments-opened=„1“ data-image-  
meta=„{“aperture“:0,“credit“:„?????“,“camera“:  
“:„,“caption“:„,“created\_timestamp“:„0“,“copyright“:„,“focal\_length“:„0“,“iso“:„0“,“orientation“:„0“,“shutter\_speed“:„0“,“title“:„,““orientation“:„0“}“ data-image-title=„Running the binary counter“ data-image-description=„  
data-  
medium-file=„[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_py\\_program\\_output\\_wide.jpg?w=800&h=303?w=400](https://hackadaycom.files.wordpress.com/2017/04/bincounter_py_program_output_wide.jpg?w=800&h=303?w=400)“  
data-  
large-  
file=„[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_py\\_program\\_output\\_wide.jpg?w=800&h=303?w=800](https://hackadaycom.files.wordpress.com/2017/04/bincounter_py_program_output_wide.jpg?w=800&h=303?w=800)“ class=„size-full wp-image-251879“  
src=„[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_py\\_program\\_output\\_wide.jpg?w=800&h=303](https://hackadaycom.files.wordpress.com/2017/04/bincounter_py_program_output_wide.jpg?w=800&h=303)“ alt=„Running the binary counter“ width=„800“ height=„303“  
srcset=„[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_py\\_program\\_output\\_wide.jpg?w=800&h=303](https://hackadaycom.files.wordpress.com/2017/04/bincounter_py_program_output_wide.jpg?w=800&h=303) 800w,  
[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_py\\_program\\_output\\_wide.jpg?w=250&h=95](https://hackadaycom.files.wordpress.com/2017/04/bincounter_py_program_output_wide.jpg?w=250&h=95) 250w,  
[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_py\\_program\\_output\\_wide.jpg?w=400&h=151](https://hackadaycom.files.wordpress.com/2017/04/bincounter_py_program_output_wide.jpg?w=400&h=151) 400w,  
[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_py\\_program\\_output\\_wide.jpg?w=768&h=291](https://hackadaycom.files.wordpress.com/2017/04/bincounter_py_program_output_wide.jpg?w=768&h=291) 768w,  
[https://hackadaycom.files.wordpress.com/2017/04/bincounter\\_py\\_program\\_output\\_wide.jpg](https://hackadaycom.files.wordpress.com/2017/04/bincounter_py_program_output_wide.jpg) 832w“  
sizes=„(max-width: 800px) 100vw, 800px“> <figcaption class=„wp-caption-text“>Running the  
binary counter</figcaption> </figure> <p>Here’s the output during successful training and  
subsequent counting. Notice that it trained within 200 steps through the loop. Very  
occasionally it does all 10001 steps without reducing the training error sufficiently, but once  
you’ve trained it successfully and saved it, that doesn’t matter.</p> <h2>The  
Next Step</h2> <p>As we said, the code for the binary counter neural network is on <a  
href=„[https://github.com/Hack-a-Day/bincounter\\_TensorFlow\\_example/](https://github.com/Hack-a-Day/bincounter_TensorFlow_example/)“ target=„\_blank“>our github  
page</a>. You can start with that, start from scratch, or use any of the many tutorials on the <a  
href=„<https://www.tensorflow.org/>“ target=„\_blank“>TensorFlow website</a>. Getting it to do  
something with hardware is definitely my next step, taking inspiration from <a  
href=„<https://hackaday.com/2016/10/09/tensorflow-robot-recognizes-objects/>“>this robot that [Lukas  
Biewald] made recognize objects around his workshop</a>.</p> <p>What are you using, or  
planning to use TensorFlow for? Let us know in the comments below and maybe we’ll give it a  
try in a future article!</p> </html>

From:

<https://schnipsl.qgelm.de/> - **Qgelm**



Permanent link:

<https://schnipsl.qgelm.de/doku.php?id=wallabag:introduction-to-tensorflow>

Last update: **2021/12/06 15:24**