

11-695: Competitive Engineering  
Let's Code a Neural Net!

Spring 2018

## tf\_graph\_recap.py

```
1 import tensorflow as tf
2
3 def build_tf_graph():
4     # create ops, variables, etc.
5
6 def main(_args):
7     g = tf.Graph()
8     with g.as_default():
9         build_tf_graph()
10    with tf.Session() as sess:
11        output = sess.run([train_op, ...]) # execute the operation z
```

- Computational Graphs:
  - nodes or *ops*
  - edges are *tensors*
- Execution: `sess.run([a, a, b, c])`
  - *a, b, c and their parents* will be run
  - Two *a*s will be run *once*

## tf\_graph\_dict\_trick.py

```
1 import tensorflow as tf
2
3 def build_tf_graph():
4     # create ops, variables, etc.
5     ops = {
6         "train": train_op,
7         "grad_norm": gradient_norm,
8         "preds": get_predictions,
9     }
10    return ops
11
12 def main():
13     g = tf.Graph()
14     with g.as_default():
15         ops = build_tf_graph()
```


- Graph names are *different* from Python names
  - Can be lost when you build large graphs
  - You should use a Python dict to store these Python *handles*

## tf\_feed\_forward\_net.py

```
1 def feed_forward_net(x, dims=[256, 512, 128], num_classes=10):
2     # x is a tensor of size [N, H, W, C]
3     N, H, W, C = tf.unstack(tf.shape(x))
4
5 def build_tf_graph():
6     images, labels = get_data_ops()
7     logits = feed_forward_net(images)
8     # LATER: loss function, train_op, etc.
```

- What are images and labels?
  - They are *TF ops*
- What happens in `N, H, W, C = tf.unstack(tf.shape(x))`?
  - `tf.shape(x)` is a *TF ops*.
  - `tf.unstack(...)` is a *TF ops*.
  - So are `N, H, W, C`
- Get used to them. Everything in TF is an ops.

# Step 1: Build a Model in a TF Graph



## tf\_feed\_forward\_net.py

```
1 def feed_forward_net(x, dims=[256, 512, 128], num_classes=10):
2     # x is a tensor of size [N, H, W, C]
3     N, H, W, C = tf.unstack(tf.shape(x))
4     x = tf.reshape(x, [N, H * W * C]) # flatten
5     for layer_id, next_dim in enumerate(dims):
6         curr_dim = x.get_shape()[-1].value # get_shape() returns a <list>
7         with tf.variable_scope("layer_{}".format(layer_id)):
8             w = tf.get_variable("w", [curr_dim, next_dim]) # w's name: "layer_2/w"
9             x = tf.matmul(x, w)
10            x = tf.nn.relu(x)
11            curr_dim = x.get_shape()[-1].value # get_shape() returns a <list>
12            with tf.variable_scope("logits"):
13                w = tf.get_variable("w", [curr_dim, num_classes]) # w's name: "logits/w"
14                logits = tf.matmul(x, w)
15            return logits
```

- Flatten  $\rightarrow$  (Linear  $\rightarrow$  Nonlinear)  $\times$  100  $\rightarrow$  logits

### tf\_feed\_forward\_net.py

```
1 def feed_forward_net(x, dims=[256, 512, 128], num_classes=10):
2     # The mess we just discussed
3
4 def build_tf_graph():
5     images, labels = get_data_ops() # images, labels: [N, H, W, C], [N]
6     logits = feed_forward_net(images)
7     # cross entropy loss function
8     loss = tf.nn.sparse_softmax_cross_entropy_with_logits(
9         logits=logits, labels=labels)
10    loss = tf.reduce_mean(loss)
11    # LATER: train_op, etc.
```

- `tf.nn.sparse_softmax_cross_entropy_with_logits(logits, labels)` computes *all* losses.
  - Divide them by the batch size
- `tf.nn.sparse_softmax_cross_entropy_with_logits(...)` is *the correct way* to implement **cross entropy loss**

### tf\_feed\_forward\_net.py

```
1 def build_tf_graph():
2     images, labels = get_data_ops() # images, labels: [N, H, W, C], [N]
3     logits = feed_forward_net(images)
4     # cross entropy loss function
5     probs = tf.nn.softmax(logits)
6     label_probs = tf.gather(probs, labels, axis=1)
7     loss = tf.log(label_probs)
8     # LATER: train_op, etc.
```

- Here is one (out of many) wrong ways!
  - Correct values, but **much** slower!
- Why?
  - $\log \frac{\exp \{l_i\}}{\sum_j \exp \{l_j\}}$  is *fused* into a kernel.

## tf\_feed\_forward\_net.py

```
1 def build_tf_graph():
2     images, labels = get_data_ops() # images, labels: [N, H, W, C], [N]
3     logits = feed_forward_net(images)
4     loss = tf.nn.sparse_softmax_cross_entropy_with_logits(logits=logits, labels=labels)
5     # train_op
6     optimizer = tf.train.GradientDescentOptimizer(learning_rate=1.0)
7     train_op = optimizer.minimize(loss)
8     # LATER: predictions and accuracies
```

- optimizer is *not* the train op
- optimizer.minimize(loss) is the train op
- If you call sess.run([train\_op]), it will perform
  - Forward pass images and labels
  - Back-propagation
  - Update all variables using gradient descents.



## tf\_feed\_forward\_net.py

```
1 def build_tf_graph():
2     images, labels = get_data_ops() # images, labels: [N, H, W, C], [N]
3     logits = feed_forward_net(images)
4     loss = tf.nn.sparse_softmax_cross_entropy_with_logits(logits=logits, labels=labels)
5     optimizer = tf.train.GradientDescentOptimizer(learning_rate=1.0)
6     train_op = optimizer.minimize(loss)
7     # predictions and accuracies
8     preds = tf.argmax(logits, axis=1)
9     accus = tf.equal(preds, labels) # TF boolean tensor
10    accus = tf.cast(accus, dtype=tf.int32)
11    accus = tf.reduce_sum(accus)
12    ops = {
13        "loss": loss,
14        "train_op": train_op,
15        "preds": preds,
16        "accus": accus,
17    }
18    return ops
```

- Everything is a TF ops
  - Get used to this concept / idea / paradigm!
  - I will repeat this until you are used to it!

## tf\_feed\_forward\_net.py

```
1 def build_tf_graph():
2     # The mess we discussed
3     ops = {
4         "loss": loss, "train_op": train_op, "preds": preds, "accus": accus
5     }
6     return ops
7
8 def main(_args):
9     g = tf.Graph()
10    with g.as_default():
11        ops = build_tf_graph()
12        with tf.Session() as sess:
13            sess.run(tf.global_variables_initializer()) # this is juts an ops!
14            for train_step in range(10000):
15                output = sess.run([ops["train_op"]])
```

- We built a TF graph with
  - A feed forward network
  - Essential ops *to train* and *to use* the network

## tf\_conv\_net.py

```
1 def feed_forward_net(x, dims=[256, 512, 128], num_classes=10):
2     # some mess you have seen
3
4 def conv_net(x, kernel_sizes=[3, 5, 7], num_channels=[128, 256, 512], num_classes=10):
5     # some mess you will fill in
6
7 def build_tf_graph():
8     images, labels = get_data_ops()
9     logits_feed_forward = feed_forward_net(images)
10    logits_conv = conv_net(images)
11    # some more mess
12    return ops
13
14 def main(_args):
15     with tf.Graph().as_default:
16         ops = build_tf_graph()
17         # even more mess
```

- Everything is the same, but
  - Replace `feed_forward_net` with `conv_net`

## tf\_conv\_net.py

```
1 def conv_net(x, kernel_sizes=[3, 5, 7], num_channels=[128, 256, 512], num_classes=10):
2     # x: tensor of size [N, H, W, C]
3     N, H, W, C = tf.unstack(tf.shape(x))
4     for layer_id, (k_size, next_c) in enumerate(zip(kernel_sizes, num_channels)):
5         curr_c = x.get_shape()[-1].value # get_shape() returns a <list>
6         with tf.variable_scope("layer_{}".format(layer_id)):
7             # w's name: "layer_2/w"
8             w = tf.get_variable("w", [k_size, k_size, curr_c, next_c])
9             x = tf.nn.conv2d(x, w, padding="SAME")
10            x = tf.nn.relu(x)
11            x = tf.reshape(x, [N, -1])
12            curr_c = x.get_shape()[-1].value # get_shape() returns a <list>
13            with tf.variable_scope("logits"):
14                w = tf.get_variable("w", [curr_c, num_classes]) # w's name: "logits/w"
15                logits = tf.matmul(x, w)
16            return logits
```

- No more flatten!
- Just change your `tf.matmul` into `tf.nn.conv2d()`
  - `padding="SAME"`