TensorFlow for Deep Learning



Zurich University of Applied Sciences

Oliver Dürr

Datalab-Lunch Seminar Series Winterthur, 23 Nov, 2016

Code: github.com/oduerr/dl_tutorial/



4

Leftovers

- Notes and things I forgot
 - The Mandelbrot example now also includes loop (see Control_Flow/Madelbrot)
 - Extremely nice tool (DevDocs includes TF and many other useful libs)
- Today it's about using Tensorflow for deep learning. No theory of Deep Learning!

Outline

- Short Recap
- How to build networks
 - Scoping
- Using existing models
 - Accessing ops and tensors in existing networks
 - Fine-tuning adopt existing networks to a new task
- Debugging
 - Tensorboard
 - tf.Print()

Most important thing from last time: compute graph

- Edges are arrays with n indices (tensors of order n)
- Nodes are operations (ops)
- These tensors flow hence the name
- To steps process (allows e.g. for symbolic differentiation)
 - Build graph in a abstract fashion
 - Put values in and out with feeds and fetches



Photo credit TensorFlow documentation

The Tensors Flowing



Most important thing from last time: feeds and fetches





Libraries on top of TensorFlow

- There are lots of libraries on top of TensorFlow. Some of them are in the *tensorflow.contrib* package and are thus installed with TensorFlow
 - TF-Slim
 - nice to build networks
 - contains many pre-trained networks
 - skflow
 - scikit learn like interface (not used so far)
 - TF Learn (inside contrib)
 - I did not use it so far
- Notable exception is the TFLearn (http://tflearn.org/) library (outside TF)
 - Easy training
 - Can handle hdf5 files
 - Includes data augmentation

Since they are all build around TF they can be combined. Other libraries using TF as engine (e.g. Keras no experience so far)



Building Blocks for Networks

Photo credit CS231n







Name scopes

• Name like vgg16/conv1/conv1_1 allow to group complex networks



Creation of name scopes

You could simply name ops like 'conv1/Conv2d'. However, there is a nice mechanism to do so:

Out[3]: (u'conv1/Kernels:0', u'conv1/Conv2D:0')

With this mechanism it's also easy possible to create a network out of simple building blocks...

https://github.com/oduerr/dl_tutorial/blob/master/tensorflow/Building_Nice_Networks/Scoping.ipynb

Variable scopes as building blocks



conv3

Variable scopes to share variables

- Variable scoping is a mechanism to share the variables of (possible large) parts of a network, without the need to pass references.
- These shared variables are needed for example in Siamese Networks.
- Two function with go together:
 - tf.variable_scope() created the name-space or better context manager
 - tf.get_variable() gets or newly creates variables in the name scope
 - Here we do not use tf.Variable()

- See also
 - <u>https://www.tensorflow.org/versions/master/how_tos/variable_scope/index.html</u>
 - http://stackoverflow.com/questions/35919020/whats-the-difference-of-name-scope-and-a-variable-scope-in-tensorflow

Variable scoping (new variables)

```
tf.reset_default_graph()
with tf.variable_scope('var'):
    al = tf.get_variable('a', shape=(1))
    #This variable is used and thus this would result in an error
    #a1_1 = tf.get_variable('a', shape=(1))
    a2 = tf.get_variable('a2', shape=(1))
al.name,a2.name
```

```
(u'var/a:0', u'var/a2:0')
```

Variable scoping (shared variables)

```
tf.reset_default_graph()
with tf.variable_scope('var', reuse=False):
    a1 = tf.get_variable('a', shape=(1))
with tf.variable_scope('var', reuse=True):
    a1_1 = tf.get_variable('a', shape=(1)) #This variable is reused
    #This would give an error, since that variable has not been used before
    #a2 = tf.get_variable('a2', shape=(1))
al.name, al_1.name
```

```
(u'var/a:0', u'var/a:0')
```

https://github.com/oduerr/dl_tutorial/blob/master/tensorflow/Building_Nice_Networks/Scoping.ipynb

Notebook and sharing weights

Let's have a look at the notebook. In this notebook it is also explained how to **share variables** (e.g. for Siamese Networks) with variable scopes and $tf.get_variable('v', shape=(1,10))$.

https://github.com/oduerr/dl_tutorial/blob/master/ tensorflow/Building Nice Networks/Scoping.ipynb

Using pre-trained networks

Checkpointing (saving)

```
\dots #\leftarrow Definition of the network
epochs = 1000
saver = tf.train.Saver()
with tf.Session() as sess:
    sess.run(init op)
    for e in range(5):
        sess.run(train op, feed dict={x:x data, y:y data})
    res = sess.run([loss, a, b], feed dict={x:x data, y:y data})
    print(res)
    save path = saver.save(sess, "checkpoints/model.ckpt")
    print("Model saved in file: %s" % save path)
print('Finished all')
```

Checkpointing (restoring)

```
... # The network needs to be defined. It is not stored.
saver = tf.train.Saver()
with tf.Session() as sess:
    saver.restore(sess, "checkpoints/model.ckpt")
    res = sess.run([loss, a, b], feed_dict={x:x_data, y:y_data})
    print(res)
```

Weights and definition of the graph

Checkpointing just stores the weights. The definition of the network has to be defined or stored separately. If you what to do it all in one, you have to transform the weights in constants and save the network. This is referred to *freezing* the graph.

Using existing models [show]



 The notebook: <u>Using Trained Nets</u>* shows how to access trained networks



feed = tf.Graph.get_tensor_by_name(tf.get_default_graph(), 'Placeholder:0')
fetch = tf.Graph.get_tensor_by_name(tf.get_default_graph(), 'vgg_16/fc8/BiasAdd:0')
res = sess.run(fetch, feed_dict={feed:feed_vals})

This feeding and fetching is extremely useful for debugging, see later)

*https://github.com/oduerr/dl_tutorial/tree/master/tensorflow/stored_models

Transfer Learning [show]

Transfer Learning with CNNs





Debugging: run part of the graph feed and fetch

Tensorflow allows us to run parts of graph in isolation, i.e. only the relavant part of graph is executed (rather than executing *everything*)

```
x = tf.placeholder(tf.float32)
y = tf.placeholder(tf.float32)
bias = tf.Variable(1.0)
y pred = x ** 2 + bias # x -> x^2 + bias
loss = (y - y pred)**2  # 12 loss?
# Error: to compute loss, y is required as a dependency
print('Loss(x,y) = \%.3f' \% session.run(loss, {x: 3.0}))
# OK. print 1.000 = (3**2 + 1 - 9)**2
print('Loss(x,y) = %.3f' % session.run(loss, {x: 3.0, y: 9.0}))
# OK, print 10.000; for evaluating y pred only, input to y is not required
print('pred y(x) = \%.3f' % session.run(y pred, {x: 3.0}))
# OK, print 1.000 bias evaluates to 1.0
print('bias = %.3f' % session.run(bias))
```

Designing graphs for debugging

- You want to access the graph at different entry points
- You can get every Tensor in the graph (to feed or fetch)
 - E.g. tf.Graph.get_tensor_by_name(tf.get_default_graph(), 'Placeholder:0')
- However, it is much nicer to give the user handles to the tensors
- There are good and bad ways of doing so:

Designing graphs for debugging: bad

```
def alexnet(x):
    assert x.get shape().as list() == [224, 224, 3]
    conv1 = conv 2d(x, 96, 11, strides=4, activation='relu')
    pool1 = max pool 2d(conv1, 3, strides=2)
    conv2 = conv 2d(pool1, 256, 5, activation='relu')
    pool2 = max pool 2d(conv2, 3, strides=2)
    conv3 = conv 2d(pool2, 384, 3, activation='relu')
    conv4 = conv 2d(conv3, 384, 3, activation='relu')
    conv5 = conv 2d(conv4, 256, 3, activation='relu')
    pool5 = max pool 2d(conv5, 3, strides=2)
    fc6 = fully connected(pool5, 4096, activation='relu')
    fc7 = fully connected(fc6, 4096, activation='relu')
    output = fully connected(fc7, 1000, activation='softmax')
    return conv1, pool1, conv2, pool2, conv3, conv4, conv5, pool5, fc6, fc7
# At construction time
conv1, conv2, conv3, conv4, conv5, fc6, fc7, output = alexnet(images) # ?!
# During the training loop
, loss , conv1 , conv2 , conv3 , conv4 , conv5 , fc6 , fc7 = session.run(
        [train op, loss, conv1, conv2, conv3, conv4, conv5, fc6, fc7],
        feed dict = \{\ldots\}
```

Quite messy code!

Designing graphs for debugging: good with dictionary

```
def alexnet(x, net={}):
    assert x.get_shape().as_list() == [224, 224, 3]
    net['conv1'] = conv_2d(x, 96, 11, strides=4, activation='relu')
    net['pool1'] = max_pool_2d(net['conv1'], 3, strides=2)
    net['conv2'] = conv_2d(net['pool1'], 256, 5, activation='relu')
    net['pool2'] = max_pool_2d(net['conv2'], 3, strides=2)
    net['conv3'] = conv_2d(net['pool2'], 384, 3, activation='relu')
    net['conv4'] = conv_2d(net['conv3'], 384, 3, activation='relu')
    net['conv5'] = conv_2d(net['conv4'], 256, 3, activation='relu')
    net['pool5'] = max_pool_2d(net['conv5'], 3, strides=2)
    net['fc6'] = fully_connected(net['pool5'], 4096, activation='relu')
    net['fc7'] = fully_connected(net['fc7'], 1000, activation='relu')
    net['output']
```

```
net = {}
output = alexnet(images, net)
# access intermediate layers like net['conv5'], net['fc7'], etc.
```

Better

Designing graphs for debugging: good with class

```
class AlexNetModel():
    # ...
    def build model(self, x):
        assert x.get shape().as list() == [224, 224, 3]
        self.conv1 = conv 2d(x, 96, 11, strides=4, activation='relu')
        self.pool1 = max pool 2d(self.conv1, 3, strides=2)
        self.conv2 = conv 2d(self.pool1, 256, 5, activation='relu')
        self.pool2 = max pool 2d(self.conv2, 3, strides=2)
        self.conv3 = conv 2d(self.pool2, 384, 3, activation='relu')
        self.conv4 = conv 2d(self.conv3, 384, 3, activation='relu')
        self.conv5 = conv 2d(self.conv4, 256, 3, activation='relu')
        self.pool5 = max pool 2d(self.conv5, 3, strides=2)
        self.fc6 = fully connected(self.pool5, 4096, activation='relu')
        self.fc7 = fully connected(self.fc6, 4096, activation='relu')
        self.output = fully connected(self.fc7, 1000, activation='softmax')
        return self.output
model = AlexNetModel()
```

```
output = model.build_model(images)
# access intermediate layers like self.conv5, self.fc7, etc.
```

Better

Summaries

Taken from: LinearRegression/02 Inspecting the graph.ipynb

```
resi = a*x + b - y
loss = tf.reduce_sum(tf.square(resi), name='loss')
...
#Definition of ops to be stored
loss_summary = tf.scalar_summary("loss_summary", loss) #<--- creates op!
resi_summart = tf.histogram_summary("resi_summart", resi)
merged_summary op = tf.merge_all_summaries()#<----- Combine all ops to be stored</pre>
```

```
sess.run(init_op)
#Where to store
writer = tf.train.SummaryWriter("/tmp/dumm/run1", tf.get_default_graph(),
'graph.pbtxt')
for e in range(epochs): #Fitting the data for 10 epochs
...
#Running the graph to produce output
sum_str = sess.run(merged_summary_op, feed_dict={x:x_vals, y:y_vals})
```

```
writer.add_summary(sum_str, e) #<--- writing out the output</pre>
```

Print- a bit non-trivial at the first sight



More Debugging

- tf.py func()

- tf.Assert()
 - Creates runtime assertions
- Possibility to use python code as tensorflow op.

```
def my_func(x):
    # x will be a numpy array with the contents of the placeholder below
    return np.sinh(x)
inp = tf.placeholder(tf.float32, [...])
y = py_func(my_func, [inp], [tf.float32])
```

The above snippet constructs a tf graph which invokes a numpy sinh(x) as an op in the graph.

For a detailed explanation of the above concepts see: <u>https://wookayin.github.io/TensorflowKR-2016-talk-debugging/</u>

Debugging with embedded python code

```
def _debug_plot(a_val, b_val, x_val, y_val):
    plt.scatter(x_val,y_val)
    ablineValues = [a_val * x_ + b_val for i, x_ in enumerate(x_val)]
    plt.plot(x_val, ablineValues)
    return False
```

Decorating the loss function

```
debug_op = tf.py_func(_debug_plot, [a, b, x, y], [tf.bool])
with tf.control_dependencies(debug_op):
    loss = tf.identity(loss, name='out')
```

