Ex 09 Tensorboard

Following tensorboard tutorial, we need to

- attach scalar summaries for scalar tensors like the learning rate, or loss
- attach a histogram summary to the output of a relu
- combine all summaries into a single op with tf.merge_all_summaries
- create a protocol buffer Summary object at each step
- use a train.SummaryWriter to write to disk
- Pass a Graph in the SummaryWriter constructor to see the computational graph

Code ex09_tensorboard.py

- train function calls new function: attach_tensorboard_summaries(model)
  - called after we make the train_op since that is when we add last ops to model
- attach_tensorboard_summaries does:
  - tf.scalar_summary on three things
  - tf.histogram_summary on one thing
  - returns the merge_all_summaries op
- then we make a summaryWriter
- We add the merge_all_summaries op to our training ops
- the result of sess.run on the merge_all_summaries, we pass it to our summaryWriter
- We use a directory on the network, not /tmp, which is on local machine
  - we want to read from the directory from other machines

Running

- launch the code on batch, even with a small batch size it can be slow on the interactive nodes
- maybe wait a few seconds for the code to read the data and print some output
  - then we know there is content for tensorboard to read
- go to your pslogin terminal
- cd to the mlearntut directory
- source the mlearntut-setup.sh if you haven't already
- you should see a tf_summaries_train directory. This is where the code writes the summaries
- execute:
  tensorboard --logdir=tf_summaries_train &
- That is run tensorboard in the background
- Now run a browser, the only one we have on pslogin is firefox
  firefox http://0.0.0.0:6006
- If we cannot all use the same port, then people should add --port PORT to the tensorboard run line, so we each run on different ports

Exercises

- Explore tensorboard
- Click on Graph, explore computational graph
- Add a new summary, maybe an image