

CS 491 Undergraduate Research

Sergio Servantez
Computer Science Department
Illinois Institute of Technology
Chicago, Illinois 60616
sservantez@hawk.iit.edu
Advisor: Dr. Zhiling Lan

Report 5: 4/4/19

We have successfully trained our first deep neural network on Theta using the MNIST dataset. Over the past couple weeks, the primary focus was on resolving two issues. First, modifying the code to access the dataset from a local directory instead of fetching it from an external Google server. And second, properly partitioning the dataset into the respective training and test data once read in. Both of these issues have been resolved. I have included a screenshot of some of the training output below which shows the loss value is around 2.3. This value by itself doesn't provide us with much information about the final accuracy of the DNN. For our next experiments, we will try to determine how to output the prediction accuracy of the network. The DNN is using a 2-layer convolution model with a densely connected layer with 1024 neurons. Now that we have successfully executing a training job, the focus will shift towards using the built-in Horovod profiler to analyze internode communication. Over the next couple weeks, the primary goal will be to run these experiments and collect the profiler data to better understand the communication patterns between the ranks.

```
INFO:tensorflow:loss = 2.296105, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.312255, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.3013809, step = 151 (5.100 sec)
INFO:tensorflow:loss = 2.3022878, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.298362, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.300405, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.3022506, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.2997, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.296441, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.3031778, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.3043177, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.3075616, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.2997036, step = 151 (5.100 sec)
INFO:tensorflow:loss = 2.3077056, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.3074558, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.2972646, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.3013003, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.3001208, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.3071516, step = 151 (5.099 sec)
INFO:tensorflow:loss = 2.304042, step = 151 (5.099 sec)
INFO:tensorflow:Saving checkpoints for 156 into ./checkpoints/model.ckpt.
```