

Project Milestone 4

07-400

Justin Zhang

<https://www.andrew.cmu.edu/user/justinz/>

1 Progress Report

Things are going good. I have been collecting data for the different approximations on CPU/GPU/TPU, and the results have been pretty expected. Using CPU has been slowest by far, with GPU and TPU having comparable performance. TPUs are optimized for large matrix operations, so their performance scales better with larger amounts of data. Surprisingly, TPUs have slower times when low rank approximation is used; I believe this may be due to the former observation and the added overhead of computing an approximation.

Currently, I have most of the results for datasets of size up to 1 million, but I will compile results for 25 million sized datasets to see if my previous observations hold. Moreover, I will try to vary the batch sizes, which I believe GPUs will deal with smaller batch sizes better and vice versa for TPUs.

I have an adequate amount of results to start doing an analysis of how different parameters affect TPU performance, so I will start making cool graphs and magic statistical numbers.

Additionally, I will compile my literature review and prepare for the poster presentation. I will continue to communicate with Rashmi and Jason on how to put together my semester's research.

2 Major Changes

No major changes, I am making good progress towards the end of this project.

3 Meeting your Milestone

I have met my milestone, and I am currently compiling my final poster presentation for next week.

4 Surprises

Using TPU architecture is a headache due to proprietary reasons. I have spent a significant amount of time this week figuring out why Google Colab was so slow when using the TPU. For one, you have to manually connect to the TPU cluster in code (thanks stackoverflow). Also, TPUs can't access local storage that CPUs and GPUs usually can, so I had to open a Google Cloud Bucket™ to store my data. Whether this is by design or not, Google is making it very unclear whether storing my datasets in their buckets is free or not. I guess I will find out at the end of the month.

5 Revisions to your 07-400 Milestones

I have no revisions to make at this time.

6 Resources Needed

The current resources I need are Google Cloud TPU credits which the google people will apparently supply to me. I am currently running on a free account which is good enough so far.

References