

MAY 11, 2016

Final Paper Rationale

BY KENNY FRIEDMAN

In this brief note, I explain the rationale behind my final paper. I have previously described much of this rationale in-person during the final presentations on the last day of class (available as a private Youtube video, here: <https://youtu.be/rbSq2DT-Opc>).

This semester, I took two AI-related classes: MAS.S63 (this class) and 6.803, Professor Patrick Winston's "Human Intelligence Enterprise" (homepage: <https://courses.csail.mit.edu/6.803/schedule.html>). While both classes covered the full spectrum of AI research, 6.803 put a much heavier focus on early/classic research, and MAS.S63 put a much heavier focus on deep learning and neuroscience.

During MAS.S63, it became more and more clear to me that connectionist/deep-learning/neural net architecture could provide a plausible cognitive model. Therefore, deep-learning could be the base-level substrate for all machine intelligence systems.

While reading the 6.803 papers, I saw that neural networks can not do many tasks that other systems have been able to do since the early 1960s. Moreover, there are certain tasks which neural nets do not seem even close to being able to accomplish.

MAS.S63 has enabled me to have a much deeper understanding of modern AI research, such as allowing me to explore many neural network implementations. Throughout the semester, I learned the basics of Torch and TensorFlow. I also enabled GPU training. Using this EC2 instance, I explore and experimented with many architectures that ranged from MNIST auto-encoders to style-net implementations.

At first, I planned to do a TensorFlow based final project. However, I had been grappling with many concepts between the two classes, and wanted to flesh out my understanding of their intersection.

This paper, "Escaping the Local Minimum", is the culmination of my exploration. I am particularly proud of the original research I did: I have accumulated over 40 citations with direct links, some of which have rarely been cited in their full form (such as a direct link to PDF scans of a 1970 Minsky quote or a NYTimes headline from 1988 on the state of the industry). I am also proud of the synthesis I've done between "Good Old Fashioned AI" and modern deep learning techniques.

I hope this paper will be both an interesting read and a useful resource for others.

Best,
Kenny Friedman