

Sociotechnical Synthesis

STS 4600
Spring 2021

Brian Yu
B.S. Computer Science

Signature  _____ Date 03 May 2021

Brian Yu

Approved  _____ Date 03 May 2021
Professor Richard Jacques, Department of Engineering and Society

Deep neural networks are complex machine learning algorithms trained on very large datasets. The efficacy of deep learning models increases with the amount of data that the models are trained on and the size of the model. These models are transforming a wide array of fields. In science and medicine, deep learning allows us to understand how protein folding works and rapidly discover new drugs. In transportation and robotics, deep learning is the main driver behind advances in computer vision and planning, allowing for advances in self-driving cars and other autonomous robots. My technical project focused on the distributed training of deep learning models, which allows us to use many computers to train the same model, enabling faster training of larger, more powerful models. My STS research centered on the ethics of deep learning and uses the STS frameworks of user studies, the politics of artifacts, and sub politics to investigate the sociotechnical impacts of deep learning.

The technical portion of my project was focused on leveraging networks of computers, known as distributed systems, to train deep learning models in a process known as distributed training. There are many different types of deep learning models, each with its own unique set of computational demands. Thus, distributed training is a complex problem. Furthermore, given a cluster of computers and an arbitrary number of models, it is challenging to train all of these models on the cluster so that all cluster resources are fully utilized and model training finishes quickly. My project involved implementing state-of-the-art deep learning scheduling algorithms on both simulated and real computer clusters. These implementations provide a baseline that researchers in my lab can use to compare their scheduling algorithms to.

In my STS research, I used several STS frameworks to investigate the sociotechnical impacts of deep learning. Using user studies, I learned about how deep learning powered recommendation systems learn from users' behavior and then in turn influence that behavior,

forming a feedback loop. By studying the politics of deep learning, I learned how deep learning models are being used to make decisions with weighty ethical and political implications, such as determining the likelihood of criminal recidivism, which directly impacts an inmate's chance of getting parole. Using the sociotechnical framework of sub politics, I learned how companies are trying to self-govern their use of deep learning by creating private AI ethics boards. These ethics boards are often opaque to the public and have incentives that align more closely with the companies that sponsor them than the public for which they claim to benefit. Using these STS frameworks, I concluded that we need more regulation and oversight to ensure that deep learning models are being used ethically and responsibly.

Deep learning models are powerful tools that have already helped us solve many important problems and will continue to do so in the future. In the technical portion of my project, I dove deeply into the technical details of distributed training and learned about distributed systems, task scheduling, and deep learning model training. In my STS research, I investigated the ethics of deep learning and demonstrated the need for increased regulation, oversight, and interpretability of deep learning models. Deep learning is responsible for many incredible advances in science and technology and is being used in increasingly important contexts. We need to ensure that such powerful technology is used ethically and responsibly.

Acknowledgments

Thank you to Professor Haiying Shen, Professor Tolu Odumosu, and Professor Richard Jacques for their guidance and assistance during my technical project and STS research.