DISSERTATION DEFENSE

Rui Liu

Sequential Decision Making for Large Scale Machine Learning

Monday, December 12, 2022
12:30 – 2:30pm
Virtual – Zoom

ABSTRACT: Large scale machine learning lies at the core of many artificial intelligence's recent successes exemplified by AlphaGo, BERT, DALL-E, GitHub Copilot, AlphaCode, and ChatGPT. These large scale models are undoubtedly powerful but very expensive to train. To make it less expensive, we incorporate sequential decision making into machine learning model training. Sequential decision making has long been the focus of stand-alone fields (e.g., reinforcement learning and multi-armed bandit). We observe that sequential decision making problems also appear in the context of training machine learning models under several different settings. We show that carefully resolving these problems leads to improved training efficiency, thus reducing the training time. Depending on the specific training settings, the decisions are about how to (a) select coordinates of the vectors, (b) select examples from the training set, or (c) route tokens across different machines. In this dissertation, we consider a variety of training settings including recommendation systems, distributed learning, curriculum learning and transformer models. We design decision making strategies tailored for each training setting, and demonstrate reduced training times and often better accuracies through extensive experiments.

CHAIR: Prof. Barzan Mozafari