

University of Chicago and Toyota Technological Institute at Chicago Machine Learning Seminar Series

PRESENTS:



Lin Yang
Princeton University

Title: Learn Policy Optimally via Efficiently Utilizing Data

Abstract:

Recent years have witnessed increasing empirical successes in reinforcement learning. Nevertheless, it is an irony that many theoretical problems in this field are not well understood even in the most basic setting. For instance, the optimal sample and time complexities of policy learning in finite-state Markov decision process still remain unclear. Given a state-transition sampler, we develop a novel algorithm that learns an approximate-optimal policy in near-optimal time and using a minimal number of samples. The algorithm makes updates by processing samples in a “streaming” fashion, which requires small memory and naturally adapts to large-scale data. Our result resolves the long-standing open problem on the sample complexity of Markov decision process and provides new insights on how to use data efficiently in learning and optimization. The algorithm and analysis can be extended to solve two-person stochastic games and feature-based Markov decision problems while achieving near-optimal sample complexity. We further illustrate several other examples of learning and optimization over streaming data, with applications in accelerating Astrophysical discoveries and improving network securities.

Bio:

Lin Yang is currently a postdoctoral researcher at Princeton University working with Prof. Mengdi Wang. He obtained two Ph.D. degrees simultaneously in Computer Science and in Physics & Astronomy from Johns Hopkins University in 2017. Prior to that, he obtained a bachelor’s degree from Tsinghua University. His research focuses on developing fast algorithms for large-scale optimization and machine learning. This includes reinforcement learning and streaming methods for optimization and function approximations. His algorithms have been applied to real-world applications including accelerating astrophysical discoveries and improving network security. He has published numerous papers in top Computer Science conferences including NeurIPS, ICML, STOC, and PODS. At Johns Hopkins, he was a recipient of the Dean Robert H. Roy Fellowship.

Host: Rebecca Willett

Wednesday, April 17, 1 – 2 pm
Harper Center (Booth) Room C04
Pizza provided by UChicago CS Department