

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

PRESENTS:



Zhaoran Wang
Northwestern

Title: On the Computational and Statistical Efficiency of Policy Optimization in (Deep) Reinforcement Learning

Abstract: Coupled with powerful function approximators such as neural networks, policy optimization plays a key role in the tremendous empirical successes of deep reinforcement learning. In sharp contrast, the theoretical understandings of policy optimization remain rather limited from both the computational and statistical perspectives. From the perspective of computational efficiency, it remains unclear whether policy optimization converges to the globally optimal policy in a finite number of iterations, even given infinite data. From the perspective of statistical efficiency, it remains unclear how to attain the globally optimal policy with a finite regret or sample complexity.

To address the computational question, I will show that, under suitable conditions, natural policy gradient/proximal policy optimization/trust-region policy optimization (NPG/PPO/TRPO) converges to the globally optimal policy at a sublinear rate, even when it is coupled with neural networks. To address the statistical question, I will present an optimistic variant of NPG/PPO/TRPO, namely OPPO, which incorporates exploration in a principled manner and attains a \sqrt{T} -regret.

(Joint work with Qi Cai, Chi Jin, Jason Lee, Boyi Liu, Zhuoran Yang)

Bio: Zhaoran Wang is an assistant professor at Northwestern University, working at the interface of machine learning, statistics, and optimization. He is the recipient of the AISTATS (Artificial Intelligence and Statistics Conference) notable paper award, ASA (American Statistical Association) best student paper in statistical learning and data mining, INFORMS (Institute for Operations Research and the Management Sciences) best student paper finalist in data mining, and the Microsoft fellowship.

Host: Mladen Kolar

Friday, December 6, 10:30 – 11:30 am
TTIC, 6045 S. Kenwood Avenue,
5th Floor, Room 526
(Remote broadcast at JCL, Room 390)