


Fast Convergence of Exact Policy Gradient Methods

ABSTRACT - Policy gradient methods are among the most effective methods for reinforcement learning. Despite their long history and extensive recent activities, their convergence properties are not well understood even in the fundamental setting of finite Markov decision processes (the tabular setting). In this talk, we introduce several progresses on understanding exact policy gradient methods in the tabular setting. First, we derive an improved sublinear convergence rate for the projected policy gradient method based on the notion of weak gradient-mapping domination. Second, we show that a general class of policy mirror descent methods, including the natural policy gradient method, enjoy a linear convergence rate without relying on entropy regularization. Finally, we draw connection between preconditioned policy gradient methods and the classical policy iteration method, and comment on possible superlinear convergence.



Dr. Lin Xiao
Research Scientist
Facebook AI Research (FAIR)

SPEAKER BIO – Lin Xiao is a Research Scientist at Facebook AI Research (FAIR) in Seattle, Washington. His research interests include theory and algorithms for large-scale optimization and machine learning, reinforcement learning, and parallel and distributed computing. He received his PhD from Stanford University in 2004 and was a postdoctoral fellow at California Institute of Technology. Before joining Facebook in 2020, he spent 14 years as a Researcher at Microsoft Research. He was one of three winners for the Young Researcher competition at the first International Conference on Continuous Optimization in 2004, and won the Test of Time Award at NeurIPS 2019. He currently serves as an associate editor for the SIAM Journal on Optimization, and has served as area chairs for several machine learning conferences including NeurIPS, ICML and ICLR.