(Bandit) Convex Optimization with Biased Noisy Gradient Oracles*

Csaba Szepesvári

Abstract

For bandit convex optimization we propose a model, where a gradient estimation oracle acts as an intermediary between a noisy function evaluation oracle and the algorithms. The algorithms can control the bias-variance tradeoff in the gradient estimates. We prove lower and upper bounds for the minimax error of algorithms that interact with the objective function by controlling this oracle. The upper bounds replicate many existing results (capturing the essence of existing proofs) while the lower bounds put a limit on the achievable performance in this setup. In particular, our results imply that no algorithm can achieve the optimal minimax error rate in stochastic bandit smooth convex optimization.

*Machine Learning External Seminar, Gatsby Unit, November 10, 2015.