

Finite Sample analysis of linear stochastic approximation and TD learning

Tuesday, June 18, 2024 9:30 AM (1 hour)

Abstract: In this talk, we consider the problem of obtaining sharp bounds for linear stochastic approximation. We then apply these results to temporal difference (TD) methods with linear functional approximation for policy evaluation in discounted Markov decision processes. We show that a simple algorithm with a universal and instance-independent step size together with Polyak-Ruppert tail averaging is sufficient to obtain near-optimal variance and bias terms. We also provide the respective sample complexity bounds. Our proof technique is based on refined error bounds for linear stochastic approximation together with the novel stability result for the product of random matrices that arise from the TD-type recurrence. We will also discuss how these results extend to the distributed / federated setting.

Session Classification: Keynote: Éric Moulines (École Polytechnique)