

Computing the bias of stochastic approximation with constant step-size via Stein's method.

Monday, June 17, 2024 1:30 PM (30 minutes)

Stochastic approximation algorithms are quite popular in reinforcement learning notably because they are powerful tools to study the convergence of algorithms based on stochastic gradient descent (like Q-learning or policy gradient). In this talk, I will focus on constant step-size stochastic approximation and present tools to compute its asymptotic bias, which is non-zero (both for Martingale noise or Markovian noise), contrary to the case of decreasing step-size. The analysis is based on a fine comparison of the generators of the stochastic system and its deterministic counterpart. It is similar to Stein's method.

Primary author: GAST, Nicolas (Inria, Univ. Grenoble Alpes)

Co-author: ALLMEIER, Sebastian (Inria, Univ. Grenoble Alpes)

Presenter: GAST, Nicolas (Inria, Univ. Grenoble Alpes)

Session Classification: Parallel session: Policy gradient methods: optimization and convergence