ID de Contribution: **25** Type: **Non spécifié**

Compressed and distributed least-squares regression: convergence rates with applications to Federated Learning

jeudi 20 juin 2024 09:00 (30 minutes)

We investigate the impact of compression on stochastic gradient algorithms for machine learning, a technique widely used in distributed and federated learning.
We underline differences in terms of convergence rates between several unbiased compression operators, that all satisfy the same condition on their variance, thus going beyond the classical worst-case analysis. To do so, we focus on the case of least-squares regression (LSR) and analyze a general stochastic approximation algorithm for minimizing quadratic functions relying on a random field. More particularly, we highlight the impact on the convergence of the covariance of the additive noise induced by the algorithm. We consider weak assumptions on the random field, tailored to the analysis (specifically, expected Hölder regularity), and on the noise covariance, enabling the analysis of various randomizing mechanisms, including compression. We then extend our results to the case of federated learning.

Orateur: PHILIPPENKO, Constantin (Ecole Polytechnique)

Classification de Session: Exposé court

Classification de thématique: Exposés courts