

## **Covariance-controlled adaptive Langevin thermostat for large-scale Bayesian sampling**

Xiaocheng Shang, Brown University

Monte Carlo sampling for Bayesian posterior inference is a common approach used in machine learning. The Markov chain Monte Carlo procedures that are used are often discrete-time analogues of associated stochastic differential equations (SDEs). These SDEs are guaranteed to leave invariant the required posterior distribution. An area of current research addresses the computational benefits of stochastic gradient methods in this setting. Existing techniques rely on estimating the variance or covariance of the subsampling error, and typically assume constant variance. We propose a covariance-controlled adaptive Langevin thermostat that can effectively dissipate parameter-dependent noise while maintaining a desired target distribution. The proposed method achieves a substantial speedup over popular alternative schemes for large-scale machine learning applications.