Langevin Quasi-Monte Carlo

Sifan Liu
Stanford University, USA
sfliu@stanford.edu

Langevin Monte Carlo (LMC) and its stochastic gradient versions are powerful algorithms for sampling from complex high-dimensional distributions, which are common in data science and machine learning applications. To sample from the distribution with density $\pi(x) \propto \exp(-U(x))$, LMC generates the next sample by taking a step in the gradient direction ∇U with a Gaussian perturbation. Expectations w.r.t. the target distribution π are estimated by averaging over LMC samples. In ordinary Monte Carlo, it is well known that the estimation error can be substantially reduced by replacing independent random samples with quasi-random samples like low-discrepancy sequences. In this work, we show that the estimation error of LMC can also be reduced by using quasi-random samples. Specifically, we propose to use completely uniformly distributed sequences with certain low-discrepancy property to generate the Gaussian perturbations (and stochastic gradients). Under smoothness and convexity conditions, we prove that LMC with quasi-random samples achieves smaller errors than standard LMC. We provide rigorous theoretical analysis supported by compelling numerical experiments to demonstrate the effectiveness of our approach.

Joint work with Art Owen (Stanford University).