Posterior Sampling with the Proximal Stochastic Gradient Langevin Algorithm

Nicolas Papadakis

We investigate the challenge of sampling from distributions defined by non-convex potentials using the Unadjusted Langevin Algorithm (ULA). Our analysis establishes the stability of the discrete-time ULA under drift approximations, provided the potential is non-convex and strongly convex at infinity. To this end, we focus on the Proximal Stochastic Gradient Langevin Algorithm (PSGLA), which integrates the forward-backward optimization framework with a ULA step. By leveraging our stability results and properties of the Moreau envelope, we present the first convergence proof for PSGLA in the context of non-convex potentials. We validate our theoretical findings empirically using synthetic data and real-world imaging inverse problems. Our experiments illustrate that PSGLA not only preserves the restoration quality of posterior sampling but also achieves faster convergence rates compared to the standard Stochastic Gradient Langevin Algorithm.