

Eric Moulines

When Moreau meets Langevin: Sampling from high-dimensional log-concave non-smooth densities

Recently, the problem of designing MCMC sampler adapted to high-dimensional distributions and with sensible theoretical guarantees has received a lot of interest. The applications are numerous, including large-scale inference in machine learning, Bayesian nonparametrics, Bayesian inverse problem, aggregation of experts among others. When the density is L -smooth (the log-density is continuously differentiable and its derivative is Lipschitz), we will advocate the use of a “rejection-free” algorithm, based on the discretization of the Euler diffusion with either constant or decreasing stepsizes. We will present several new results allowing convergence to stationarity under different conditions for the log-density (from the weakest, bounded oscillations on a compact set and super-exponential in the tails to the strong concavity). When the log-density is not smooth (a problem which typically appears when using sparsity inducing priors for example), we still suggest to use a Euler discretization but of the Moreau envelope of the non-smooth part of the log-density. An importance sampling correction may be later applied to correct the target. Several numerical illustrations will be presented to show that this algorithm (named MYULA) can be practically used in a high dimensional setting. Finally, non-asymptotic bounds convergence bounds (in total variation and Wasserstein distances) are derived.