

Title: Noisy Monte Carlo

Abstract:

Monte Carlo algorithms often aim to draw from a distribution π by simulating a Markov chain with transition kernel \mathbf{P} such that π is invariant under \mathbf{P} . However, there are many situations for which it is impractical or impossible to draw from the transition kernel \mathbf{P} . For instance, this is the case with massive datasets, where it is prohibitively expensive to calculate the likelihood and is also the case for intractable likelihood models arising from, for example, Gibbs random fields, such as those found in spatial statistics and network analysis. A natural approach in these cases is to replace \mathbf{P} by an approximation $\hat{\mathbf{P}}$. Using theory from the stability of Markov chains we explore a variety of situations where it is possible to quantify how 'close' the chain given by the transition kernel $\hat{\mathbf{P}}$ is to the chain given by \mathbf{P} . We apply these results to several examples from spatial statistics and social network analysis.