SAMPLING ADVANCES BY ADAPTIVE REGENERATIVE PROCESSES AND IMPORTANCE MONTE CARLO

CHRISTIAN ROBERT

This talk will cover two recent advances in sampling, achieved in collaboration with Arthur McKimm, Murray Pollock, Gareth Roberts, Andi Wang, and with Charly Andral, Randal Douc, Hugo Marival, respectively.

Enriching Brownian motion with regenerations from a fixed regeneration distribution μ at a particular regeneration rate κ results in a Markov process that has a target distribution π as its invariant distribution \cite{wang2021}. For the purpose of Monte Carlo inference, implementing such a scheme requires firstly selection of regeneration distribution μ, and secondly computation of a specific constant C. Both of these tasks can be very difficult in practice for good performance. In \cite{kimm2024}, We introduce a method for adapting the regeneration distribution, by adding point masses to it. This allows the process to be simulated with as few regenerations as possible and obviates the need to find said constant C. Moreover, the choice of fixed μ is replaced with the choice of the initial regeneration distribution, which is considerably less difficult. We establish convergence of this resulting self-reinforcing process and explore its effectiveness at sampling from a number of target distributions. The examples show that adapting the regeneration distribution guards against poor choices of fixed regeneration distribution and can reduce the error of Monte Carlo estimates of expectations of interest, especially when π is skewed.

The Importance Markov chain is a novel algorithm proposed by \cite{andral2024} bridging the gap between rejection sampling and importance sampling, moving from one to the other through a tuning parameter. Based on a modified sample of an instrumental Markov chain targeting an instrumental distribution (typically via a MCMC kernel), the Importance Markov chain produces an extended Markov chain where the marginal distribution of the first component converges to the target distribution. For example, when targeting a multimodal distribution, the instrumental distribution can be chosen as a tempered version of the target which allows the algorithm to explore its modes more efficiently. We obtain a Law of Large Numbers and a Central Limit Theorem as well as geometric ergodicity for this extended kernel under mild assumptions on the instrumental kernel. Computationally, the algorithm is easy to implement and pre-existing librairies can be used to sample from the instrumental distribution.