Lecture 2.2 (02:25-02:50)

Scaling of Piecewise Deterministic Monte Carlo for Anisotropic Targets

Kengo Kamatani (ISM)

Piecewise deterministic Markov processes (PDMPs) are a type of continuous-time Markov process that combine deterministic flows with jumps. Recently, PDMPs have garnered attention within the Monte Carlo community as a potential alternative to traditional Markov chain Monte Carlo (MCMC) methods. The Zig-Zag sampler and the Bouncy particle sampler are commonly used examples of the PDMP methodology which have also yielded impressive theoretical properties, but little is known about their robustness to extreme dependence or isotropy of the target density. It turns out that PDMPs may suffer from poor mixing due to anisotropy and this paper investigates this effect in detail in the stylised but important Gaussian case. To this end, we employ a multi-scale analysis framework in this paper. Our results show that when the Gaussian target distribution has two scales, of order 1 and ϵ, the computational cost of the Bouncy particle sampler is of order ϵ^{-1}, and the computational cost of the Zig-Zag sampler is either ϵ^{-1} or ϵ^{-2}, depending on the target distribution. In comparison, the cost of the traditional MCMC methods such as RWM or MALA is of order ϵ^{-2}, at least when the dimensionality of the small component is more than 1. Therefore, there is a robustness advantage to using PDMPs in this context.

This is joint work with Joris Bierkens (TU Delft) and Gareth O. Roberts (Warwick).