Teleportation MCMC

Supervisors: Murray Pollock (Warwick)*, Gareth Roberts (Warwick) †
Collaborators: Cyril Chimisov (Google), Andi Wang (Oxford OxWaSP)

Interested students should schedule a meeting to discuss this project prior to selection
Find updated project listing / availability at `https://warwick.ac.uk/mpollock/projects`

Overview

Simple MCMC algorithms can often be characterised as 'local', whereby the algorithm moves around using a succession of relatively small, local moves; or 'global', in which more ambitious moves to entirely different parts of the state space are attempted. For example, local methods would include the Random Walk Metropolis algorithm and the Metropolis-adjusted Langevin algorithm, whereas the Independence Metropolis algorithm would be a global method. Although it seems more intuitive to prefer global moves, such algorithms are known to scale badly in dimension, often having exponentially small acceptance probabilities (as a function of dimension) so that local MCMC methods are often required.

One strategy which has been attempted is to combine local and global algorithms (resulting in so called 'hybrid' schemes). Choosing the local and global dynamics jointly in order to target the correct invariant distribution is difficult, and so the natural approach to constructing a hybrid scheme is to choose both the local and global dynamics to separately target the correct invariant distribution, which ensures that the resulting hybrid algorithm also targets the correct invariant distribution. Although this is a pragmatic approach for selection of the dynamics, it is also one of the primary reasons such approaches have not been widely adopted – hybrid schemes suffer from the shortcomings of both the constituent local and global approaches, for instance in high dimensional situations the global moves will almost certainly be rejected leading to considerable inefficiency.

Recent work by the supervisors (in conjunction with Andi Wang (Oxford OxWaSP)) has been to develop approaches in which the local and global moves separately have the incorrect invariant distribution (indeed, may not have an invariant distribution at all!), but when hybridised jointly target the correct invariant distribution. This new and substantial flexibility opens considerable scope for the development of new part-local part-global methods, with the possibility of very desirable theoretical and methodological properties. For instance, it transpires that the resulting MCMC schemes allow for a practical 'Coupling from the Past' [Propp and Wilson, 1996] argument to be implemented, enabling perfect (iid) simulation from the invariant distribution. The methodology has a striking resemblance to Google's celebrated 'Page Rank' algorithm [Brin and Page, 1998], in which a network of websites is explored by a combination of local moves which move to 'adjacent' (referencing) websites, and frequent 'teleportation' (ie global) moves which begin exploring other parts of the network (so chosen to induce stability).

Current work by the supervisors has focussed on developing methodology where the local moves are driven by diffusions, however there is no particular restriction to diffusion local moves, and instead one could consider local PDMP moves (Piecewise Deterministic Stochastic Processes) moves (such as the Bouncy Particle Sampler [Bouchard-Côté et al., 2018] or Zig-Zag [Bierkens et al., 2016]) or embedding other local MCMC

*Email: *m.pollock@warwick.ac.uk*
†Email: *gareth.o.roberts@warwick.ac.uk*

schemes. Directions for doctoral projects in this area could include: designing new hybrid schemes; investigating the scaling properties of these approaches; optimisation of these methods; efficient implementation; or, taking a fresh look at applications such as Page Rank and applying our newly developed theory and methodology.

Selected References

[Bierkens et al., 2016] Bierkens, J., Fearnhead, P., and Roberts, G. (2016). The zig-zag process and super-efficient sampling for bayesian analysis of big data. *arXiv preprint arXiv:1607.03188*.

[Bouchard-Côté et al., 2018] Bouchard-Côté, A., Vollmer, S. J., and Doucet, A. (2018). The bouncy particle sampler: A nonreversible rejection-free markov chain monte carlo method. *Journal of the American Statistical Association*, pages 1–13.

[Brin and Page, 1998] Brin, S. and Page, L. (1998). The anatomy of a large-scale hypertextual web search engine. *Computer networks and ISDN systems*, 30(1-7):107–117.

[Propp and Wilson, 1996] Propp, J. G. and Wilson, D. B. (1996). Exact sampling with coupled markov chains and applications to statistical mechanics. *Random Structures & Algorithms*, 9(1-2):223–252.