
Adaptive, delayed-acceptance MCMC for targets with expensive likelihoods

Chris Sherlock^{*1}, Andrew Golightly², and Daniel Henderson²

¹Lancaster University – Bailrigg, Lancaster. UK LA1 4YW, United Kingdom

²Newcastle University (UNITED KINGDOM) – United Kingdom

Abstract

When conducting Bayesian inference, delayed acceptance (DA) Metropolis-Hastings (MH) algorithms and DA pseudo-marginal MH algorithms can be applied when it is computationally expensive to calculate the true posterior or an unbiased estimate thereof, but a computationally cheap approximation is available. A first accept-reject stage is applied, with the cheap approximation substituted for the true posterior in the MH acceptance ratio. Only for those proposals which pass through the first stage is the computationally expensive true posterior (or unbiased estimate thereof) evaluated, with a second accept-reject stage ensuring that detailed balance is satisfied with respect to the intended true posterior. In some scenarios there is no obvious computationally cheap approximation. A weighted average of previous evaluations of the computationally expensive posterior provides a generic approximation to the posterior. If only the k -nearest neighbours have non-zero weights then evaluation of the approximate posterior can be made computationally cheap provided that the points at which the posterior has been evaluated are stored in a multi-dimensional binary tree, known as a KD-tree. The contents of the KD-tree are potentially updated after every computationally intensive evaluation. The resulting adaptive, delayed-acceptance [pseudo-marginal] Metropolis-Hastings algorithm is justified both theoretically and empirically. Guidance on tuning parameters is provided and the methodology is applied to a discretely observed Markov jump process characterising predator-prey interactions and an ODE system describing the dynamics of an autoregulatory gene network.

*Speaker