International Society for Bayesian Analysis, 9th World Meeting,
Hamilton Island, Australia, 2008.

Think locally, act globally: Combining the best features of particle methods and MCMC

Michael Johannes[1], Nicholas Polson[2], and Steven L. Scott[3⋆]

[1] Columbia University, New York, USA
[2] University of Chicago, Chicago, USA
[3] Capital One, USA
⋆ steve_the_bayesian@yahoo.com

This talk describes an algorithm for simulating from an arbitrary probability distribution π by combining features from Markov chain Monte Carlo (MCMC) and particle based methods. MCMC explores the sample space of π by a series of local moves. Particle methods simulate from π by first sampling many particles from a proposal distribution, then resampling the particle with appropriately chosen weights. Particle methods are limited by the need to choose an appropriate proposal distribution, but they have attractive global search properties that MCMC lacks (they have no difficulty locating multiple modes, for example).

Our algorithm, which we call Particle Posterior Sampling, runs n Markov chains in parallel. The state of each Markov chain at time t can be thought of as a particle, and the collection of particles define an empirical distribution which can be resampled. The result is an algorithm where particles move quickly to high probability regions, then use MCMC's local search capabilities to quickly explore those regions. The resulting trade off is that if n is large then t can be kept small (e.g. single digits). The worst-case behavior of the algorithm is that of n parallel Markov chains with a common starting point near the mode.

We illustrate the algorithm on several canonical problems where MCMC is known to struggle. These include carefully chosen examples from probit regression and Gaussian linear models with "spike and slab" prior distributions, as well as finite mixture models with an unknown number of components.