Diagnosing and Resolving Lack of MCMC Convergence in a Dynamic System Model

David A. Campbell*
Department of Statistics and Actuarial Science
Simon Fraser University
Surrey, BC V3T 0A3
dac5@sfu.ca

1 Poster Abstract

When setting up a Metropolis-Hasting sampler, initially the users goal is adjusting the transition density to achieve a reasonable acceptance rate. After this has been achieved, the practitioner then runs the Markov Chain for a long time until convergence has been assessed and the sample is large enough for the desired precision of estimates.

We present a case study where targeting an ideal acceptance rate will not reliably achieve Markov chain convergence nor reasonable precision in a feasible number of iterations. The model in question is a common and simple two state compartment model, with only one observed state, such as is commonly used in Pharmacokinetic/Pharmacodynamic, epidemiological or chemical reaction models. The problem is that within a small neighbourhood of the mode, the marginal posterior is reasonably approximated by a narrow Gaussian distribution where, along one side, the marginal density stops decaying at a squared exponential rate and flattens out at a level in line with contours containing 99.9% of the highest density region for a Gaussian.

In assessing the adequacy of the transition density by the performance of the acceptance rate one could make the transition density very small and begin near the posterior mode. However eventually, the random walk will find the ridge, where a lack of slope in the posterior topology results in what appears to be a transition density that is much too narrow and a trace plot which shows transience rather than stationarity. On the other hand one could begin with a wide transition density variance and explore the flat ridge fully, but the relatively narrow peak around the posterior mode is likely to be missed, or if found, the transition variance will be too large to draw a well mixing sample near the mode that isn't disproportionately swamped with samples from the ridge.In this poster we show how to diagnose and resolve this problem while highlighting it's presence in a variety of commonly used models.

To diagnose the problem we advocate keeping track of the un-normalized posterior which is already computed within the Metropolis Hastings acceptance probability ratio. Plotted alongside or against the accepted parameter samples provides visual diagnostic insights and reveals posterior topology features. Comparing the un-normalized posterior with a density estimate from the Markov chain samples highlights cases of over-representation of sampled features. Even at the discovery of the unfavourable posterior features, what remains unclear is whether or not the posterior samples obtained from a long running MCMC sampler should include the ridge as a feature in the relevant highest density contours. Effectively, even after running the Metropolis-Hastings sampler for hundreds of thousands of iterations the resulting samples are not yet trustworthy. To fix this we again use the stored un-normalized posterior density values as sampling weights to complete a sampling importance re-sampling step to improve the quality of our posterior density estimate. In our case this revealed that the ridge was not an important feature for reasonable contours but instead was just a hazard for our random walk.

*Assistant Professor *http://www.stat.sfu.ca/∼dac5*