Scaling particle MCMC methods with local proposals

Alexander Shestopaloff
Queen Mary University of London

Particle MCMC methods such as iterated conditional SMC are commonly used for state inference in non-linear, non-Gaussian state space models. Standard implementations of iterated cSMC provide an efficient way to sample state sequences in low-dimensional state space models. However, efficiently scaling iterated cSMC methods to perform well in models with a high-dimensional state remains a challenge. One reason for this is the use of a global proposal, without reference to the current state sequence. In high dimensions, such a proposal will typically not be well-matched to the posterior and impede efficient sampling. I will describe a novel approach based on a *local* embedded HMM (Hidden Markov Model) to construct efficient proposals in high dimensions that are local relative to the current state sequence. By considering several examples, I will demonstrate that both strategies improve the performance of iterated cSMC for state sequence sampling in high-dimensional state space models.