Correlated pseudo marginal schemes for partially observed diffusion processes

Andrew Golightly[1], Emma Bradley[1]

[1]*School of Mathematics & Statistics, Newcastle University, UK*

Abstract

Stochastic differential equations (SDEs) provide a natural framework for modelling intrinsic stochasticity inherent in many continuous-time physical processes. Performing fully Bayesian inference for such models, using discrete time data that may be incomplete and subject to measurement error is a challenging problem. One widely used approach is to replace unavailable transition densities with an Euler-Maruyama approximation, made sufficiently accurate through the introduction of intermediate times between observations. Integrating over the uncertainty in the process at these intermediate times typically necessitates the use of Markov chain Monte Carlo (MCMC). Pseudo marginal MCMC schemes are increasingly used, since for a given discretisation level, the observed data likelihood can be unbiasedly estimated using a particle filter. When observations are particularly informative, a diffusion bridge construct can be used to drive the particle filter. Recent work in state-space settings has shown how the pseudo marginal approach can be made much more efficient by correlating the underlying pseudo random numbers used to form the estimate of likelihood at the current and proposed values of the unknown parameters. We extend this approach to the discretised diffusion process framework by correlating the Gaussian innovations that drive the diffusion bridge construct used inside particle filter. We find that the resulting approach offers substantial gains in efficiency over a standard implementation.