

Particle filters for partially observed diffusions Hans R.

Künsch and Sybille Arnold
Seminar für Statistik
ETH Zürich

We consider a model that consists of an unobservable diffusion process (X_t) and observations Y_j 's that are conditionally independent given (X_t) and such that Y_j depends on X_{t_j} only. For the conditional distribution of (X_t) given the observations, one would have to solve partial differential equations which is difficult to do numerically. The particle filter provides recursive Monte Carlo approximations for the conditional distribution of X_t given observations up to time t . In the standard version, we only need to be able to simulate from the unconditional distribution of (X_t) , but it can perform poorly in cases where the information contained in the new observation is large. The auxiliary particle filter of Pitt and Shephard is in principle able to circumvent this problem. However, the construction of the ideal proposal is equivalent to computing the smoother with one observation and the evaluation of the weights involves the Girsanov formula. We discuss different approximations that come close to the ideal proposal. Finally, we show how the exact algorithms of Beskos et al. can be used to implement an accept/reject version of the particle filter.