Particle and ensemble filters for diffusions

Hans R. Künsch
Seminar für Statistik
ETH Zürich

We consider an unobservable nonlinear diffusion $X(t)$ and observations Y_j of the form

$$Y_j = HX(t_j) + \varepsilon_j.$$

The goal is to compute the filter distributions, i.e. conditional distributions of $X(t_j)$ given $Y_1, \dots, Y_j$. These are given by partial differential equations so that exact computation is not feasible. The particle filter and the ensemble filter both provide recursive Monte Carlo approximations. They consist of a propagation and an update step, and the two methods differ in the latter. Theoretically, the particle filter is preferable because it converges to the true filter distributions under general conditions, whereas the ensemble filter relies on a Gaussian assumption for the prediction density that usually is not valid. However, in many applications the ensemble filter performs much better.

The problems of the particle filter have two causes: First, a discrete approximation can be poor if the number of replicates (particles) is small, and second if the filter density is more concentrated than the prediction density, most of the particles are at the wrong place. The first problem can be solved by kernel smoothing, and we will explain how the ensemble filter can be viewed as a particle filter with maximal smoothing. Choosing the smoothing parameter adaptively has thus the potential to improve over both particle and ensemble filters. The second problem can in principle be adressed by modifying the update step and correcting by importance weights. However, finding a simple and efficient proposal distribution for the update seems to be difficult.

The problems and ideas will be illustrated by the double-well and the Lorenz model.