Geophysical Research Abstracts
Vol. 14, EGU2012-13386, 2012
EGU General Assembly 2012

Optimal proposal densities for particle filters

P.J. van Leeuwen

University of Reading, DARC, Department of Meteorology, Reading, United Kingdom (p.j.vanleeuwen@reading.ac.uk, +441183788905)

Most data-assimilation problems in the geosciences are of very large dimension and nonlinear, either through nonlinear models, and/or through nonlinear observation operators. Most present-day data-assimilation methods for large-dimensional problems are based on linearisations, such as (Ensemble) Kalman filters and variational methods like 4DVar. There is a growing need for fully nonlinear data-assimilation methods, and particle filters could in principle serve this goal.

However, standard particle filters are notoriously inefficient in that they typically need millions or more model runs to represent the posterior pdf. It is easy to show that this large number is related to the number of independent observations that determine the weighting of the particles via the likelihood. These weights typically vary enormously, such that e.g. a weighted mean is effectively represented via one or a few particles. This problem has been reduced to some extend by using proposal densities that bring the particles closer to observations, and as such reduce the variance in the weights of particles. However, even if the so-called 'optimal proposal density' is used, in which the proposal density takes into account the future observations, the variance in the weights is so large that an astronomical number of particles is needed for any real-sized problem.

In this presentation we discuss new proposal densities that solve this weight degeneracy problem. The idea is simply that the proposal densities can be used not only to bring particles close to observations, but also to ensure that the weights of the particles are very similar. Two examples of such proposal densities are discussed, the equivalent-weights particle filter, and the new Gaussian-peak particle filter. The first scheme determines a target weight at the last time step before the observations come in, and moves the particles such that each obtains a weight very close to that target weight. The second new scheme slightly perturbs the particles such that while sampled from a very narrow pdf, their weights are determined from a very broad pdf, ensuring that the weights are nearly equal.

The methods are implemented and tested on a high-dimensional highly nonlinear one-dimensional problem and compared to the standard particle filter and the so-called 'optimal proposal density' scheme. It is shown that both the standard particle filter and the 'optimal proposal density' scheme are degenerate, while both new schemes properly represent the posterior pdf using a very small number of particles. These results show that particle filters can be made to work for real geophysical problems when the extra freedom in the proposal density related to the weights of the particles is explored.