Geophysical Research Abstracts
Vol. 12, EGU2010-6198, 2010
EGU General Assembly 2010

Implicit particle filters and applications

Alexandre Chorin (1), Robert Miller (2), and Xuemin Tu (1)

(1) Department of Mathematics, University of California at Berkeley and Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, (2) College of Oceanic and Atmospheric Sciences, Oregon State University, Corvallis, OR 97331

Implicit particle filters and applications

Alexandre J. Chorin * Robert Miller † Xuemin Tu *

We have developed an efficient particle filter for nonlinear data assimilation, of striking simplicity and wide applicability [1,2].

Particle filters are often presented in the setting of an Ito stochastic differential equation (SDE): $dx = f(x,t)\,dt + g(x,t)\,dw$, where x is an m-dimensional vector, w is m-dimensional Brownian motion, f is an m-dimensional vector function, and $g(x,t)$ is an m by m diagonal matrix, supplemented by observations b^n at times $t^n, n = 1, 2, ...$ related to $x(t)$ by $b^n = h(x^n) + QW^n$, where h is a k-dimensional, generally nonlinear, vector function with $k \leq m$, Q is a k by k diagonal matrix, $x^n = x(t^n)$, and W^n is a Gaussian vector independent of the Brownian motion in the SDE. Bayesian particle filters follow "particles" X_i^n, where i enumerates the particles and n refers to the time, whose empirical density approximates the probability density functions (pdfs) P_n at times t^n. At each step, one first guesses a "prior" incorporating the information contained in the SDE; the prior is then corrected by sampling weights determined by the observations, yielding a "posterior" density. The catch is that, in common weighting schemes, most of the weights become very small very fast, leaving only a small number of significantly weighted particles.

Our "implicit" filter avoids the catch by reversing the procedure. Rather than find samples and then determine their probability, we first pick probabilities and then find samples that assume them. Specifically, we write the posterior density, proportional to $P(X^{n+1}|X^n)P(b^{n+1}|X^{n+1})$, as $\exp(-F(X^{n+1}))$ (this defines a function F); for each particle we sample a reference m-dimensional random variable ξ with a fixed pdf, say a Gaussian $\exp(-\xi^T\xi/2)/(2\pi)^{m/2}$ (T denotes a transpose). We then represent X^{n+1} as a function of ξ by solving the equation $F(X^{n+1}) - \phi = \xi^T\xi/2$, where ϕ is needed to make the equation solvable ($\phi = \min F$ will do the job but may not be the optimal choice). The resulting X^{n+1} is a sample of the posterior density, with sampling weight $\exp(-\phi)J$ (no explicit dependence on ξ!), where J is the Jacobian of the map $\xi \to X^{n+1}$. It is easy to see that these weights are well-distributed. In a test problem [4] devised to show the catastrophic collapse of the sampling weights in a Bayesian filter in a large number of dimensions, our algorithm produces equal weights for all the particles in any number of dimensions. Note that the function F depends also on X^n, so that we are not representing the posterior as a single function of a Gaussian, but as a sample of a large collection of functions of a Gaussian, a distinct function for every particle at every step.

If the observation function h is linear, one can choose ϕ so that the equation $F(X^{n+1}) - \phi = \xi^T\xi/2$ is directly solvable and the Jacobian J is a constant independent of particle; the method can then be viewed as an efficient implementation of the optimal importance sampler for sequential importance sampling [3]. If h is not linear, one can choose ϕ so that an iteration converges quickly and the Jacobian is easy to evaluate.

Each observation b has implications for the past as well as the future– it may reveal as unlikely what had previously appeared to be likely; a filter may need backward sampling for accuracy (a step often misleadingly motivated solely by the need to fight sample impoverishment in Bayesian filters). This we do along the same lines– define a suitable function F and represent each sample as a function of a reference variable.

We will present an application of our filter to the analysis of the circulation in the nearshore surf zone with a shallow-water model and synthetic data.

*Department of Mathematics, University of California at Berkeley and Lawrence Berkeley National Laboratory, Berkeley, CA, 94720.
†College of Oceanic and Atmospheric Sciences, Oregon State University, Corvallis, OR 97331.

References

[1] A. Chorin and X. Tu. Implicit sampling for particle filters. *Proc. Nat. Acad. Sc. USA*, 106:17249–17254, 2009.

[2] A. Chorin and X. Tu. Interpolation and iteration for nonlinear filters. 2009. Submitted to Math. Model Numer. Anal.

[3] A. Doucet, S. Godsill, and C. Andrieu. On sequential Monte Carlo sampling methods for Bayesian filtering. *Stat. Comp.*, 10:197–208, 2000.

[4] C. Snyder, T. Bengtsson, P. Bickel, and J. Anderson. Obstacles to high-dimensional particle filtering. *Mon. Wea. Rev.*, 136:4629–4640, 2008.