Particle Filters for High Dimension Spatial Systems

Jonathan Briggs

Background - Particle Filters:

Particle filters are a well established Monte Carlo methodology for estimating through time the distribution of the state vector that completely defines a system using the state-space model framework [4]. At each discrete time-point there are two steps: First a prediction ensemble is generated by applying the state space model evolution equation to each member of the previous time point filter ensemble. Second a filter ensemble is generated by applying the likelihood defined by the state-space model observation equation to any available observations in a Bayesian update using the prediction ensemble as a prior. Usually this update is carried out via importance sampling. This process is initialised by sampling from an assumed filter distribution immediately before the first time point. Particle filters are a generalisation of the Kalman filter which requires normal prediction and filter distributions. The sequential nature of particle filter estimation means that state estimates only use observations taken at or before each estimated time point.

The problem:

When the observations are high dimension, the filter ensemble collapses to a small number of distinct points, providing very poor estimates [1][2].

For example in a model of the low trophic level marine eco-system at the Bermuda Atlantic Timeseries (BATS) [3] site, that has 5 state variables resolved at 350 locations and 52 observations the filter distribution has 1 distinct point:

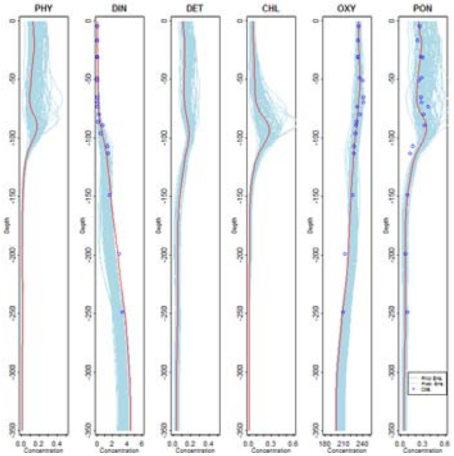

A Solution:

Particle smoothers are similar to particle filters except that they use observations available before and after the current time point in making their state estimates [5].

It is possible to use a particle smoother defined on a sequence of locations (rather than the traditional sequence of time points) to carry out the Bayesian update. Considering only one location at a time in the smoother reduces the dimensionality of the problem, avoiding filter ensemble collapse.

The location evolution equation and initial distributions can be taken from the prediction ensemble and particle filter evolution equation

The location observation equation can be taken directly from the particle filter observation equation.

Results:

A particle smoother update was applied to the same BATS model and observations as has been shown to lead to filter ensemble collapse.

The filter ensemble has 99 out of 100 distinct points at the surface location, attenuating to 17 out of 100 distinct points at the deepest location – a vast improvement over the 1 out of 100 distinct points at all locations using standard methods.

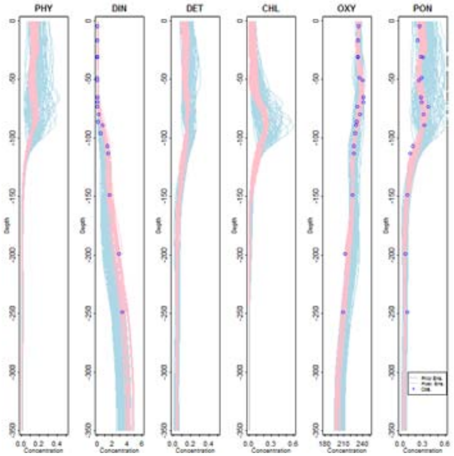

References:

[1] Van Leeuwen, P.J. Particle filtering in geophysical systems. Monthly Weather Review, 137 :4089-4114, 2009

[2] T Bengtsson, P. Bickel, and B. Li. Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems. Probability and Statistics: Essays in Honor of David A. Freedman,2:316-334, 2008.

[3] Steinberg, D. K. et al. Overview of the {US} {JGOFS} Bermuda Atlantic Time-series Study (BATS): a decade-scale look at ocean biology and biogeochemistry. Deep Sea Research part II, 48: 1405-1447, 2001.

[4] Arulampalam, M., S. et al. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. IEEE Transactions on Signal Processing, 50: 174-188, 2002

[5] Fearnhead et al. A sequential smoothing algorithm with linear computational cost. Biometrika, 97: 447-464, 2010