Feedback Particle Filter

Sean Meyn, University of Illinois at Urbana-Champaign

The objective of filtering is to compute or approximate a hidden state process based on a time-series of noisy measurements. The particle filter is a simulation-based technique that is intended to approximate the conditional distribution of the state process.

A control-oriented particle filter algorithm for nonlinear filtering will be described and illustrated with several examples. It is conjectured that the low variance of this simulation based algorithm is explained by its feedback structure. However, as always in Monte-Carlo, there is no free lunch: the feedback particle filter requires an up-front calculation to obtain a filter gain. This gain can be obtained from the solution to Poisson's equation for a Smoluchowski stochastic differential equation. This suggests approximation techniques, such as those used in TD-learning.

The talk will bring together concepts from nonlinear filtering, optimal control, reinforcement learning, and mean-field theory.

This is joint work with Tao Yang and Prashant Mehta at the University of Illinois.