## Particle Methods for Hidden Markov Models

Olivier Cappé

Given a partially observed discrete-time dynamic model, the filtering task amounts to determining the successive distributions of the (unobservable) state of the system at each time index given the observations gathered up to that time. Although less obvious, solving the filtering problem is also required for maximum-likelihood parameter inference in such models. Particle filters, perhaps more properly called sequential Monte Carlo methods, are numerical simulation-based approaches to the filtering problem which build upon the importance sampling and sampling importance resampling methods.

In these lectures, I will first describe the structure of the sequence of successive filtering distributions in the context of general state space hidden Markov models. I will then review the fundamental elements of the sequential Monte Carlo approach, illustrating in particular the role of resampling. Finally, the use of particles methods for maximum-likelihood parameter estimation will be discussed using as an example the stochastic volatility model.

These lectures are based on the book Inference in Hidden Markov Models written with E. Moulines and T. Rydén (Springer-Verlag, to appear in 2005).