Sequential inference for parameters and hidden states using particle methods

Omiros Papaspiliopoulos

Department of Economics, Universitat Pompeu Fabra, Barcelona, Spain

A central theme in the talk will be that of sequential parameter estimation and model choice. A golden standard would be to apply the iterated batch importance sampling (IBIS) algorithm of Chopin (2002). This is a sequential Monte Carlo algorithm in the θ-dimension, that samples values of θ, reweights iteratively these values using the likelihood increments $p(y_t|y_{1:t-1}, \theta)$, and rejuvenates the θ-particles through a resampling step and a MCMC update step.

In state-space models these likelihood increments are intractable in most cases, but they may be unbiasedly estimated by a particle filter in the x-dimension, for any fixed θ. This motivates the algorithm proposed in this talk: a sequential Monte Carlo algorithm, defined in the θ-dimension, which propagates and resamples many particle filters in the x-dimension. The filters in the x-dimension are an example of the random weight particle filter as in Fearnhead et al. (2010). On the other hand, the particle Markov chain Monte Carlo (PMCMC) framework developed in Andrieu et al. (2010) allows us to design appropriate MCMC rejuvenation steps. Thus, the θ-particles target the correct posterior distribution at each iteration t, despite the intractability of the likelihood increments.

We explore the applicability of our algorithm in both sequential and non-sequential applications and consider various degrees of freedom, as for example increasing dynamically the number of x-particles. We contrast our approach to various competing methods, both conceptually and empirically.