Speaker: Axel Finke (Loughborough University)

Date: 19/10/2021 at 13:15 on Microsoft Teams

Title: Conditional sequential Monte Carlo in high dimensions

Abstract:
The iterated conditional sequential Monte Carlo (i-CSMC) algorithm from Andrieu, Doucet and Holenstein (2010) is an MCMC approach for efficiently sampling from the joint posterior distribution of the T latent states in challenging time-series models, e.g. in non-linear or non-Gaussian state-space models. It is also the main ingredient in particle Gibbs samplers which infer unknown model parameters alongside the latent states. In this work, we first prove that the i-CSMC algorithm suffers from a curse of dimension in the dimension of the states, D: it breaks down unless the number of samples ("particles"), N, proposed by the algorithm grows exponentially with D. Then, we present a novel "local" version of the algorithm which proposes particles using Gaussian random-walk moves that are suitably scaled with D. We prove that this iterated random-walk conditional sequential Monte Carlo (i-RW-CSMC) algorithm avoids the curse of dimension: for arbitrary N, its acceptance rates and expected squared jumping distance converge to non-trivial limits as $D \to \infty$. If $T = N = 1$, our proposed algorithm reduces to a Metropolis–Hastings or Barker's algorithm with Gaussian random-walk moves and we recover the well known scaling limits for such algorithms.