Mini-course: Adaptive Markov Chain Monte Carlo

Markov Chain Monte Carlo (MCMC) is a well-established technique for sampling from probability distributions. The method is particularly useful in high-dimensional Bayesian inference problems. Despite its success, MCMC algorithms are often difficult to calibrate. Adaptive MCMC methods have emerged as a development of MCMC where the design and optimization of the relevant Markov kernels becomes part of the sampling problem. More broadly, adaptive techniques in MCMC provide a framework for developing new algorithms where the Markov assumption can be lifted.

For example, a good implementation of the popular Random Walk Metropolis (RWM) algorithm requires the specification of the covariance matrix of the so-called proposal kernel, a tedious task in general. But an adaptive version of the RWM proposed by H. Haario and its collaborators easily handles this issue.

The Simulated Tempering (ST) algorithm of C. Geyer is a major contribution to the MCMC literature that allows to deal effectively with slow mixing and multi-modality. But ST does not mix well unless the "temperatures" and the pseudo-prior weights are properly set which is a difficult task. This problem can also be nicely handled adaptively using the Wang-Landau algorithm.

Another example of adaptation in Monte Carlo simulation is the Equi-Energy sampler which uses empirical measure approximations of the target distribution as proposal kernel.

This mini-course will introduce the students to these recent developments in Monte Carlo simulation. The course will have two parts. The first part will highlight the methodological developments through some stylized examples. The second part will introduce the students to a martingale approximation technique and show how this is used to study the asymptotic behavior of adaptive MCMC algorithms.

Prerequisites: statistical inference, Bayesian inference, measure-theoretic probability, martingales, Markov chains.