GEOFFREY MCGREGOR, University of Northern British Columbia
Conservative Hamiltonian Monte Carlo

Markov Chain Monte Carlo (MCMC) methods enable meaningful extraction of statistics from complex distributions, frequently appearing in applications such as parameter estimation, Bayesian statistics, statistical mechanics and machine learning. The main goal of MCMC is to generate a sequence of samples which converges to a specified stationary distribution. However, as the dimensionality of the target distribution increases, the convergence rate of typical MCMC sequences toward the stationary distribution slows down dramatically. This has led to recent developments in computational techniques, such as Hamiltonian Monte Carlo (HMC), to improve on the performance in convergence and acceptance rate of proposed samples by solving specific Hamiltonian systems using symplectic methods. Nonetheless, maintaining high acceptance rates using HMC in high dimensions remains a significant challenge. In this talk, we introduce the Conservative Hamiltonian Monte Carlo (CHMC) method, which alternatively utilizes an energy-preserving numerical method, known as the Discrete Multiplier Method. We show that CHMC converges to the correct stationary distribution under appropriate conditions, and provide numerical examples showcasing improvements in acceptance rates and in scaling for high dimensional problems.

This is joint work with Andy Wan from the University of Northern British Columbia.