

Conservative Hamiltonian Monte Carlo

Dr. Geoffrey McGregor¹

Markov Chain Monte Carlo (MCMC) methods enable us to extract meaningful statistics from complex distributions which frequently appear in parameter estimation, Bayesian statistics, statistical mechanics and machine learning. Similar to how flipping a coin, or rolling a dice, allows us to sample from the corresponding distributions underlying these processes, MCMC methods enable us to sample from more complex distributions. The sample statistics of the sequence generated by MCMC will converge to those of the target distribution, or “stationary distribution”, provided certain acceptance and rejection criteria are satisfied. However, as the dimensionality of the stationary distribution increases, the acceptance rate of traditional MCMC methods inevitably diminishes and their convergence slows down substantially. This has led to recent developments in computational techniques, such as Hamiltonian Monte Carlo (HMC), to improve the performance in convergence and acceptance rate. Specifically, HMC proposes samples for acceptance or rejection by solving a Hamiltonian system of differential equations using volume-preserving numerical methods. In this talk, we introduce the Conservative Hamiltonian Monte Carlo (CHMC) method, which instead 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 on acceptance rates.

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

¹Mathematics and Statistics, University of Northern British Columbia, Prince George, B.C., V2N4Z9, Canada (Geoffrey.McGregor@unbc.ca).