

Variational Hamiltonian Monte Carlo

Babak Shababa, The Center for Statistical Consulting

Traditionally, the field of computational Bayesian statistics has been divided into two main subfields: variational methods and Markov chain Monte Carlo (MCMC). In recent years, however, several methods have been proposed based on combining variational Bayesian inference and MCMC simulation in order to improve their overall accuracy and computational efficiency. This marriage of fast evaluation and flexible approximation provides a promising means of designing scalable Bayesian inference methods. In this talk, we explore the possibility of incorporating variational approximation into a state-of-the-art MCMC method, Hamiltonian Monte Carlo (HMC), to reduce the required expensive computation involved in the sampling procedure, which is the bottleneck for many applications of HMC in big data problems. To this end, we exploit the regularity in parameter space to construct a free-form approximation of the target distribution by a fast and flexible surrogate function using an optimized additive model of proper random basis. The surrogate provides sufficiently accurate approximation while allowing for fast computation in the sampling procedure, resulting in an efficient approximate inference algorithm. We demonstrate the advantages of our method on both synthetic and real data problems.