

Scalable Bayesian Inference with Hamiltonian Monte Carlo

Michael Betancourt, University of Warwick

Despite the promise of big data, inferences are often limited not by sample size but rather by systematic effects. Only by carefully modeling these effects can we take full advantage of the data -- big data must be complemented with big models and the algorithms that can fit them. One such algorithm is Hamiltonian Monte Carlo, which exploits the inherent geometry of the posterior distribution to admit full Bayesian inference that scales to the complex models of practical interest. In this talk I will discuss the theoretical foundations of Hamiltonian Monte Carlo, elucidating the geometric nature of its scalable performance and stressing the properties critical to a robust implementation.