

*Lift and flow: manifold MCMC methods for efficient inference
in stiff posteriors*

Matt Graham @ (University College London), Au Khai Xiang, Alexandre
Thiery

R 3.07 Tue Z1 15:00-15:30

A challenging regime for employing Markov chain Monte Carlo (MCMC) methods to perform Bayesian inference is when the observed data tightly constrains only some directions in the latent space. Such ‘stiff’ posterior distributions have varying scales across the latent space and can exhibit complex geometries. The posterior mass may concentrate near lower-dimensional structures in the latent space and funnel-like pathologies can emerge where the distribution can vary between highly concentrated and highly diffuse regimes due to poor identification of scale parameters. As the step size parameter of MCMC methods typically needs to be set according to the smallest posterior scale to ensure chains can access all parts of the distribution, this can lead to slow exploration of less constrained directions or regions of the posterior. In this talk I will discuss an approach for constructing efficient Markov kernels targeting such posteriors when the underlying generative model is differentiable. The posterior distribution is lifted on to a manifold embedded in a higher dimensional space, and then a Hamiltonian flow, constrained to the manifold, simulated to generate proposed moves. In contrast to the original posterior the lifted distribution remains diffuse in the presence of highly informative observations, allowing the constrained Hamiltonian flow to be simulated with a large integrator step size and for chains to rapidly mix in the lifted distribution. As we demonstrate empirically, this can lead to substantial improvements in sampling efficiency over competing approaches.