

Inference and Data Assimilation with Optimal Maps

Youssef Marzouk, Massachusetts Institute of Technology

We present a new approach to Bayesian inference that entirely avoids Markov chain simulation, by constructing a map that pushes forward the prior measure to the posterior measure. Existence and uniqueness of a suitable measure-preserving map is established by formulating the problem in the context of optimal transport theory. We discuss various means of explicitly parameterizing the map and computing it efficiently through solution of a stochastic optimization problem. The resulting algorithm overcomes many of the computational bottlenecks associated with Markov chain Monte Carlo. Advantages include analytical expressions for posterior moments, automatic evaluation of the marginal likelihood, clear convergence criteria, and the ability to generate independent uniformly-weighted posterior samples without additional model evaluations.

We also discuss extensions of the map approach to hierarchical Bayesian models and to problems of sequential data assimilation, i.e., filtering and smoothing. For the latter problems, one scheme involves pushing forward a fixed reference measure to each filtered state distribution, while an alternative scheme computes maps that push forward the filtering distribution from one stage to the other. We evaluate the performance of these schemes and show that the map approach inherently avoids the issues of sample impoverishment common among particle filters.

Numerical demonstrations thus include parameter inference in ordinary and partial differential equations and in spatial statistical models, as well as state estimation in nonlinear dynamical systems.