Measure transport, variational inference, and low-dimensional maps

Youssef Marzouk, Massachusetts Institute of Technology

We will discuss how transport maps, i.e., deterministic couplings between probability measures, can enable useful new approaches to Bayesian computation. A first use involves a combination of optimal transport and Metropolis correction; here, we use continuous transportation to transform typical MCMC proposals into adapted non-Gaussian proposals, both local and global. Second, we discuss a variational approach to Bayesian inference that constructs a deterministic transport map from a reference distribution to the posterior, without resorting to MCMC. Independent and unweighted posterior samples can then be obtained by pushing forward reference samples through the map.

Making either approach efficient in high dimensions, however, requires identifying and exploiting low-dimensional structure. We present new results relating sparsity of transport maps to the conditional independence structure of the target distribution, and discuss how this structure can be revealed through the analysis of certain average derivative functionals. This connection between transport and graphical models yields many useful algorithms for efficient ordering and decomposition---here, generalized to the continuous and non-Gaussian setting. The resulting inference algorithms involve either the direct identification of sparse maps or the composition of low-dimensional maps and rotations. We demonstrate our approaches on Bayesian inference problems arising in spatial statistics and in partial differential equations.

This is joint work with Alessio Spantini and Daniele Bigoni.