DIMENSION REDUCTION FOR STATISTICAL INFERENCE VIA (DIMENSIONAL)

MATTHEW LI

Statistical inference in high dimensions --- whether by measure transport, by sampling, or by other methods --- is often made possible only by identifying and exploiting low-dimensional structure. For Bayesian inference in particular, one useful class of structure is to model the high-dimensional target measure as a low-dimensional update of a dominating reference measure (e.g., the prior). Determining optimal ``dimension reduction'' in this sense is computationally intractable, but we discuss how the logarithmic Sobolev and Poincar\'e inequalities, and generalizations thereof, allow us to derive approximations with certifiable error properties that are typically good enough in practice.

As a byproduct, for certain reference measures, we identify linear low-dimensional subspaces which exert universal control over the worst-case approximation error with respect to the KL divergence, the squared Hellinger metric, and everything in between. In the latter part of the talk, we then discuss: why should such tools from Markov semigroup theory have any bearing on dimension reduction at all? We answer this by showing that the optimal choice of functional inequality for the KL divergence is the \emph{dimensional} logarithmic Sobolev inequality which, in some sense, implicitly captures the low-dimensional update structure that we seek. We demonstrate some applications of these ideas to modern Bayesian inverse problems with GAN-based generative priors. Based on joint work with Olivier Zahm, Youssef Marzouk, Tiangang Cui, and Fengyi Li.