

A low dimensional tutorial on nonlinear dimension reduction

Elizabeth Qian, MIT and Greg Darnell, ICERM

Nonlinear dimension reduction is used for manifold learning, data compression and representation, and low-cost model approximation. In this tutorial, we explore two classes of methods for nonlinear dimension reduction, and show how general frameworks with particular choices (e.g., kernel or covariance prior) can be used for dimension reduction in diverse applications. First, we describe how principal component analysis can be extended to identify nonlinear low-dimensional structure using kernels that implicitly define a high-dimensional feature space. We relate these methods to autoencoders and variational autoencoders that can reduce to these more familiar methods with particular assumptions, but provide more representational expressiveness. We motivate our exploration with applications in statistical genetics and model reduction for dynamical systems. We present results and discuss challenges and opportunities particular to these applications.