Manifold Learning in Wasserstein Space

Matthew Thorpe

We give conditions sufficient for the linear approximation of a submanifold of the Wasserstein space to be equivalent (in the metric sense) to the restricted shortest-path Wasserstein distance on the submanifold. We then show how the latent manifold structure of the submanifold can be learned from samples and pairwise extrinsic Wasserstein distances. In particular, we show that the submanifold metric space can be asymptotically recovered in the sense of Gromov-Hausdorff from an appropriate graph. In addition, we demonstrate how tangent spaces can be asymptotically recovered via spectral analysis of a suitable "covariance operator" using optimal transport maps. This is joint work with Bernhard Schmitzer, Caroline Moosmueller and Keaton Hamm.