Interpretable Optimal Transport in High-Dimensions with Feature-Sparse Maps

Marco Cuturi
Apple, France
cuturi@apple.com

Optimal transport (OT) theory focuses, among all maps that can morph a probability measure onto another, on those that are the "thriftiest", i.e. such that the average cost between and its image is as small as possible. Many computational approaches have been proposed to estimate such Monge maps when that cost is the squared-Euclidean distance, using for instance neural networks. Such methods have been successfully applied to single-cell genomics, but have limitations in that they do not work well in high-dimensions, and require typically a massive drop in data dimensionality before being applied (e.g. using PCA to lower dimension of 35k dimensional gene counts to a few 50 directions). After recalling how these methods operate, I will present a recent method that allows the recovery of feature-sparse maps, e.g. such that the displacement vector $T(x) - x$ be sparse. I will detail its construction, and explain how this method can recover interpretable OT maps in high dimensions.

Joint work with Michal Klein (Apple) and Pierre Ablin (Apple).