Title: Optimal-transport distances between Markov chains

Abstract:

We propose a new framework for computing optimal-transport distances between Markov chains. We formulate the problem as a finite-dimensional linear program, which allows us to port several algorithmic ideas from other areas of optimal transport theory. In particular, our formulation makes it possible to introduce an appropriate notion of entropy regularization into the optimization problem, which in turn enables us to directly calculate optimal transport distances via a Sinkhorn-like method we call Sinkhorn Value Iteration (SVI). We show both theoretically and empirically that this method converges quickly to an optimal coupling, essentially at the same computational cost of running vanilla Sinkhorn in each pair of states. Along the way, we point out that our optimal transport distance exactly matches the common notion of bisimulation metrics between Markov chains, and thus our results also apply to computing such metrics, and in fact our algorithm turns out to be significantly more efficient than the best known methods developed so far for this purpose.

Based on joint work with Sergio Calo, Anders Jonsson, Ludovic Schwartz & Javier Segovia-Aguas