Learning with the Sinkhorn Loss

Aude Genevay
Université de Paris Dauphine

Optimal Transport is a powerful tool to compare probability distributions, but it suffers from a curse of dimensionality and a computational burden that make it impractical for high dimensional problems usually encountered in machine learning settings. Entropic regularization, as introduced by Marco Cuturi, allows to define a regularized version of the Wasserstein distance, that we call Sinkhorn loss. Aside from alleviating the computational cost of OT, entropy also reduces its sample complexity, by interpolating between OT and MMD distances (as introduced by Gretton et al.). We will discuss practical and theoretical properties of the Sinkhorn loss, and illustrate them on machine learning problems.