

Scaling Optimal Transport for High dimensional Learning

Gabriel Peyré

CNRS and ENS, PSL University

`gabriel.peyre@ens.fr`

Optimal transport (OT) has recently gained lot of interest in machine learning. It is a natural tool to compare in a geometrically faithful way probability distributions. It finds applications in both supervised learning (using geometric loss functions) and unsupervised learning (to perform generative model fitting). OT is however plagued by the curse of dimensionality, since it might require a number of samples which grows exponentially with the dimension. In this talk, I will explain how to leverage entropic regularization methods to define computationally efficient loss functions, approximating OT with a better sample complexity. More information and references can be found on the website of our book “Computational Optimal Transport”.