KILIAN FATRAS, Mila - Québec AI Institute, McGill University
Minibatch Optimal Transport distances meets Deep Learning

Optimal transport distances have found many applications in machine learning for their capacity to compare non-parametric probability distributions. Yet their algorithmic complexity generally prevents their direct use on large scale datasets. Among the possible strategies to alleviate this issue, practitioners can rely on computing estimates of these distances over minibatches of data. In this talk, we present an analysis of this practice. We notably argue that it is equivalent to an implicit regularization of the original problem, with appealing properties such as unbiased estimators, gradients and a concentration bound around the expectation. We also highlight in this talk some limits of this strategy, arguing it is not a distance and it can lead to undesirable smoothing effects. As an alternative, we suggest that the same minibatch strategy coupled with unbalanced optimal transport can yield more robust behaviours while preserving the same theoretical properties. Our experimental study shows that in challenging problems associated to domain adaptation, the use of unbalanced optimal transport leads to significantly better results, competing with or surpassing recent baselines.