

LEARNING TO OPTIMIZE TRANSPORT PLANS

Giulia Luise

Microsoft Research

giulialuise@microsoft.com

Optimal transport distances and their regularized versions are a powerful tool to compare probability measures, that proved successful in many machine learning applications. In this talk, I will give a brief introduction on (entropy-regularized) optimal transport and dive into 'learning to optimize' transport plans leveraging amortized optimization.

Joint work with Brandon Amos (META), Samuel Cohen (UCL), Ievgen Redko (Aalto University).