The Monge Gap, A Regularizer to Learn Optimal Transport Maps

Marco Cuturi
Apple / CREST-ENSAE, France
cuturi@apple.com

Optimal transport (OT) theory has been been used in machine learning to study and characterize maps that can push-forward efficiently a probability measure onto another. Recent works have drawn inspiration from Brenier's theorem, which states that when the ground cost is the squared-Euclidean distance, the "best" map to morph a continuous measure in $\mathcal{P}()$ into another must be the gradient of a convex function. To exploit that result, Makkuva et al. 2020, Korotin et. al. 2020 consider maps $T = \nabla f_\theta$, where f_θ is an input convex neural network (ICNN), as defined by Amos et al. 2017, and fit θ with SGD using samples. Despite their mathematical elegance, fitting OT maps with ICNNs raises many challenges, due notably to the many constraints imposed on θ; the need to approximate the conjugate of f_θ; or the limitation that they only work for the squared-Euclidean cost. More generally, we question the relevance of using Brenier's result, which only applies to densities, to constrain the architecture of candidate maps fitted on samples. Motivated by these limitations, we propose a radically different approach to estimating OT maps: Given a cost c and a reference measure ρ, we introduce a regularizer, the Monge gap $\mathcal{M}_\rho^c(T)$ of a map T. That gap quantifies how far a map T deviates from the ideal properties we expect from a c-OT map. In practice, we drop all architecture requirements for T and simply minimize a distance (e.g., the Sinkhorn divergence) between $T\sharp\mu$ and ν, regularized by $\mathcal{M}_\rho^c(T)$. We study $\mathcal{M}_\rho^c$, and show how our simple pipeline outperforms significantly other baselines in practice.

https://arxiv.org/abs/2302.04953

Joint work with Théo Uscidda (CREST-ENSAE).