Journées SMAI MODE 2018

Learning Generative Models with the Sinkhorn Loss

Aude Genevay
Université Paris Dauphine

Gabriel Peyré
CNRS - Ecole Normale Supérieure

Marco Cuturi
ENSAE

Résumé. Generative models are a class of statistical models that can easily be sampled from but whose likelihood is intractable because they are supported on low-dimensional manifolds in much higher-dimensional spaces. Inference for these models thus can't be performed with standard methods using maximum likelihood. That's why the ability to compare two degenerate probability distributions is a crucial factor in the estimation of these generative models. It is therefore no surprise that optimal transport (OT) metrics and their ability to handle measures with non-overlapping supports have emerged as a promising tool.
Yet, training generative models using OT raises formidable computational and statistical challenges, because of (i) the computational burden of evaluating OT losses, (ii) their instability and lack of smoothness, (iii) the difficulty to estimate them, as well as their gradients, in high dimension.
By introducing an OT-based loss called Sinkhorn loss, we can tackle these three issues by relying on two key ideas: (a) entropic smoothing, which turns the original OT loss into a differentiable and more robust quantity that can be computed using Sinkhorn fixed point iterations; (b) algorithmic (automatic) differentiation of these iterations with seamless GPU execution. Additionally, Entropic smoothing generates a family of losses interpolating between Wasserstein (OT) and Energy distance/Maximum Mean Discrepancy (MMD) losses, thus allowing to find a sweet spot leveraging the geometry of OT on the one hand, and the favorable high-dimensional sample complexity of MMD, which comes with unbiased gradient estimates.

Mots-clefs : Optimal transport, machine learning, generative models, neural networks

Références

[1] Chun-Liang LI, Wei-Cheng CHANG, Yu CHENG, et al.. *MMD GAN: Towards Deeper Understanding of Moment Matching Network*. arXiv preprint arXiv:1705.08584, 2017.

[2] Aude Genevay, Gabriel Peyré, Marco Cuturi. *Learning Generative Models with Sinkhorn Divergences* . arXiv preprint arXiv:1706.00292, 2017.