GENERATIVE MODELING VIA MAXIMUM MEAN DISCREPANCY FLOWS

JOHANNES HERTRICH

We consider gradient flows with respect to the maximum mean discrepancy (MMD) with negative distance kernel, which is also known as energy distance. In order to achieve computational efficiency, we prove that for certain kernels the MMD coincides with its sliced version. Therefore, all computations can be performed in a one-dimensional setting, where the MMD with negative distance kernel can be evaluated by a simple sorting algorithm with improved computational complexity. This enables us to simulate MMD particle flows in high dimensions for a large number of particles. We approximate these particle flows by neural networks and apply them for generative modeling and posterior sampling in Bayesian inverse problems.
From a theoretical viewpoint, we study Wasserstein gradient flows with respect to our MMD functionals. Interestingly, particles might "explode" in this setting, i.e., the flow turns atomic measures into absolutely continuous ones and vice versa. We analytically derive the Wasserstein flows for some special cases and propose a numerical approximation of suitable forward and backward time discretizations by generative neural networks.