

Mirror and Preconditioned Gradient Descent in Wasserstein Space

As the problem of minimizing functionals on the Wasserstein space encompasses many applications in machine learning, different optimization algorithms on \mathbb{R}^d have received their counterpart analog on the Wasserstein space. We focus here on lifting two explicit algorithms: mirror descent and preconditioned gradient descent. These algorithms have been introduced to better capture the geometry of the function to minimize and are provably convergent under appropriate (namely relative) smoothness and convexity conditions. Adapting these notions to the Wasserstein space, we prove guarantees of convergence of some Wasserstein-gradient-based discrete-time schemes for new pairings of objective functionals and regularizers. The difficulty here is to carefully select along which curves the functionals should be smooth and convex. We illustrate the advantages of adapting the geometry induced by the regularizer on ill-conditioned optimization tasks, and showcase the improvement of choosing different discrepancies and geometries in a computational biology task of aligning single-cells.