

Sinkhorn Distributionally Robust Optimization

We study distributionally robust optimization (DRO) with Sinkhorn distance—a variant of Wasserstein distance based on entropic regularization. We derive convex programming dual reformulation for general nominal distributions, transport costs, and loss functions. Compared with Wasserstein DRO, our proposed approach offers enhanced computational tractability for a broader class of loss functions, and the worst-case distribution exhibits greater plausibility in practical scenarios. To solve the dual reformulation, we develop a stochastic mirror descent algorithm with biased gradient oracles. Remarkably, this algorithm achieves near-optimal sample complexity for both smooth and nonsmooth loss functions, nearly matching the sample complexity of the Empirical Risk Minimization counterpart. Finally, we provide numerical examples using synthetic and real data to demonstrate its superior performance.