
Sample Complexity of Sinkhorn Divergences

Aude Genevay
DMA,
ENS Paris

Lénaïc Chizat
Inria

Francis Bach
Inria and
DI, ENS Paris

Marco Cuturi
Google and
CREST ENSAE

Gabriel Peyré
CNRS and DMA,
ENS Paris

Abstract

Optimal transport (OT) distances and maximum mean discrepancies (MMD) are now routinely used in machine learning to compare probability measures. Our focus in this paper is on *Sinkhorn divergences* (SDs), a regularized variant of OT distances that can interpolate, depending on the regularization strength ε , between OT ($\varepsilon = 0$) and MMD ($\varepsilon = \infty$). Although the tradeoff induced by that regularization is by now well understood computationally (OT, SDs and MMD require respectively $O(n^3 \log n)$, $O(n^2)$ and n^2 operations to compare two samples of size n), much less is known in terms of the *sample complexity* of SDs, namely bounding the gap between the evaluation of SDs on two densities *vs.* samples from these densities. That complexity for OT and MMD stand at two extremes: $O(1/n^{1/d})$ for OT in dimension d and $O(1/\sqrt{n})$ for MMD. that for SDs has only been studied empirically. In this paper, we *(i)* derive a bound on the approximation error made with SDs when approximating OT as a function of the regularizer ε , *(ii)* prove that the optimizers of regularized OT are bounded in a Sobolev (RKHS) ball independent of the two measures and *(iii)* reformulate SDs as a maximization problem in a RKHS to obtain a sample complexity in $1/\sqrt{n}$ (as in MMD), with a constant that depends however on ε , making the bridge between OT and MMD complete.

1 Introduction

Optimal Transport (OT) has emerged in recent years as a powerful tool to compare probability distributions. Indeed, Wasserstein distances endow the space of probability measures with a rich Riemannian structure (Ambrosio et al., 2006), one that is able to capture meaningful geometric features between measures even when their supports do not overlap. OT has been, however, long neglected in data sciences for two main reasons, which could be loosely described as *computational*—computing OT is costly since it usually requires solving a network flow problem—and *statistical*—OT suffers from the curse-of-dimensionality, since has recalled later in this paper, the Wasserstein distance computed between two samples converges only very slowly to its population counterpart.

Recent years have witnessed significant advances on those computational aspects that hindered the application of OT. A recent wave of works have exploited entropic regularization, both to compare discrete measures with finite support (Cuturi, 2013) or measures that can be sampled from (Genevay et al., 2016). Among the many learning tasks performed with this regularization, one may cite domain adaptation (Courty et al., 2014), text retrieval (Huang et al., 2016) or multi-label classification (Frogner et al., 2015). The ability of OT to compare probability distributions with disjoint supports (as opposed to the Kullback-Leibler divergence) has also made it popular as a loss function to learn generative models (Genevay et al., 2018; Salimans et al., 2018; Beaumont et al., 2002) as an alternative to the approximation considered in (Arjovsky et al., 2017).

At the other end of the spectrum, the maximum mean discrepancy (MMD) (Gretton et al., 2006) is an integral probability metric (Sriperumbudur et al., 2012) on a reproducing kernel Hilbert space (RKHS) of test functions. The MMD is easy to compute, and has also been used in a very wide variety of applications, notably to estimate of generative models (Li et al., 2015; Dziugaite et al., 2015; Li et al., 2017).

OT and MMD differ, however, on a fundamental aspect: their sample complexity. The definition of sample complexity that we choose here is the convergence rate of a given metric between a measure and its empirical counterpart, as a function of the number of samples. This notion is crucial in machine learning, as bad sample complexity implies overfitting and high gradient variance when using these divergences for parameter estimation. In that context, it is well known that the sample complexity of MMD is independent of the dimension, scaling as $\frac{1}{\sqrt{n}}$ (Gretton et al., 2006) where n is the number of samples. In contrast, it is well known that standard OT suffers from the curse of dimensionality (Dudley, 1969): Its sample complexity is exponential in the dimension of the ambient space. Although it was recently shown that these results can be refined by considering the implicit dimension of data (Weed and Bach, 2017) or the densitie properties' Weed and Berthet (2019), the sample complexity of OT appears now to be the major bottleneck for the use of OT in high-dimensional machine learning problems.

A remedy to this problem may lie, again, in regularization. Divergences defined through regularized OT, known as Sinkhorn divergences, seem to be indeed less prone to over-fitting. Indeed, a certain amount of regularization seems to improve performance in simple learning tasks (Cuturi, 2013). Additionally, recent papers (Ramdas et al., 2017; Genevay et al., 2018) have pointed out the fact that Sinkhorn divergences are in fact interpolating between OT (when regularization goes to zero) and MMD (when regularization goes to infinity). However, aside from a recent central limit theorem in the case of measures supported on discrete spaces (Bigot et al., 2017), the convergence of empirical Sinkhorn divergences, and more generally their sample complexity, remains an open question.

Contributions. This paper provides three main contributions, which all exhibit theoretical properties of Sinkhorn divergences. Our first result is a bound on the speed of convergence of regularized OT to standard OT as a function of the regularization parameter, in the case of continuous measures. The second theorem proves that the optimizers of the regularized optimal transport problem lie in a Sobolev ball which is independent of the measures. This allows us to rewrite the Sinkhorn divergence as an expectation maximization problem in a RKHS ball and thus justify the use of kernel-SGD for regularized OT as advocated in (Genevay et al., 2016). As a consequence of this reformulation, we provide as our third contribution a sample complexity result. We focus on how the sample size and the regularization parameter affect the convergence of the empirical Sinkhorn divergence (i.e., computed from samples of two continuous measures) to the continuous

Sinkhorn divergence. We show that the Sinkhorn divergence benefits from the same sample complexity as MMD, scaling in $\frac{1}{\sqrt{n}}$ but with a constant that depends on the inverse of the regularization parameter. Thus sample complexity worsens when getting closer to standard OT, and there is therefore a tradeoff between a good approximation of OT (small regularization parameter) and fast convergence in terms of sample size (larger regularization parameter). We conclude this paper with a few numerical experiments to asses the dependence of the sample complexity on ε and d in very simple cases.

Notations. We consider \mathcal{X} and \mathcal{Y} two bounded subsets of \mathbb{R}^d and we denote by $|\mathcal{X}|$ and $|\mathcal{Y}|$ their respective diameter $\sup\{\|x - x'\|, x, x' \in \mathcal{X}(\text{resp. } \mathcal{Y})\}$. The space of positive Radon measures of mass 1 on \mathcal{X} is denoted $\mathcal{M}_+^1(\mathcal{X})$ and we use upper cases X, Y to denote random variables in these spaces. We use the notation $\varphi = O(1 + x^k)$ to say that $\varphi \in \mathbb{R}$ is bounded by a polynomial of order k in x with positive coefficients.

2 Reminders on Sinkhorn Divergences

We consider two probability measures $\alpha \in \mathcal{M}_+^1(\mathcal{X})$ and β on $\mathcal{M}_+^1(\mathcal{Y})$. The Kantorovich formulation (1942) of optimal transport between α and β is defined by

$$W(\alpha, \beta) \stackrel{\text{def.}}{=} \min_{\pi \in \Pi(\alpha, \beta)} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\pi(x, y), \quad (\mathcal{P})$$

where the feasible set is composed of probability distributions over the product space $\mathcal{X} \times \mathcal{Y}$ with fixed marginals α, β :

$$\Pi(\alpha, \beta) \stackrel{\text{def.}}{=} \{ \pi \in \mathcal{M}_+^1(\mathcal{X} \times \mathcal{Y}) ; P_{1\#}\pi = \alpha, P_{2\#}\pi = \beta \},$$

where $P_{1\#}\pi$ (resp. $P_{2\#}\pi$) is the marginal distribution of π for the first (resp. second) variable, using the projection maps $P_1(x, y) = x; P_2(x, y) = y$ along with the push-forward operator $\#$.

The cost function c represents the cost to move a unit of mass from x to y . Through this paper, we will assume this function to be \mathcal{C}^∞ (more specifically, we need it to be $\mathcal{C}^{\frac{d}{2}+1}$). When $\mathcal{X} = \mathcal{Y}$ is endowed with a distance $d_{\mathcal{X}}$, choosing $c(x, y) = d_{\mathcal{X}}(x, y)^p$ where $p \geq 1$ yields the p -Wasserstein distance between probability measures.

We introduce regularized optimal transport, which consists in adding an entropic regularization to the optimal transport problem, as proposed in (Cuturi, 2013). Here we use the relative entropy of the transport plan with respect to the product measure $\alpha \otimes \beta$ following (Genevay

et al., 2016):

$$W_\varepsilon(\alpha, \beta) \stackrel{\text{def.}}{=} \min_{\pi \in \Pi(\alpha, \beta)} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\pi(x, y) + \varepsilon H(\pi \mid \alpha \otimes \beta), \quad (\mathcal{P}_\varepsilon)$$

where

$$H(\pi \mid \alpha \otimes \beta) \stackrel{\text{def.}}{=} \int_{\mathcal{X} \times \mathcal{Y}} \log \left(\frac{d\pi(x, y)}{d\alpha(x) d\beta(y)} \right) d\pi(x, y). \quad (1)$$

Choosing the relative entropy as a regularizer allows to express the dual formulation of regularized OT as the maximization of an expectation problem, as shown in (Genevay et al., 2016)

$$\begin{aligned} W_\varepsilon(\alpha, \beta) &= \max_{u \in \mathcal{C}(\mathcal{X}), v \in \mathcal{C}(\mathcal{Y})} \int_{\mathcal{X}} u(x) d\alpha(x) + \int_{\mathcal{Y}} v(y) d\beta(y) \\ &\quad - \varepsilon \int_{\mathcal{X} \times \mathcal{Y}} e^{\frac{u(x) + v(y) - c(x, y)}{\varepsilon}} d\alpha(x) d\beta(y) + \varepsilon \\ &= \max_{u \in \mathcal{C}(\mathcal{X}), v \in \mathcal{C}(\mathcal{Y})} \mathbb{E}_{\alpha \otimes \beta} [f_\varepsilon^{XY}(u, v)] + \varepsilon \end{aligned}$$

where $f_\varepsilon^{xy}(u, v) = u(x) + v(y) - \varepsilon e^{\frac{u(x) + v(y) - c(x, y)}{\varepsilon}}$. This reformulation as the maximum of an expectation will prove crucial to obtain sample complexity results. The existence of optimal dual potentials (u, v) is proved in the appendix. They are unique α - and β -a.e. up to an additive constant.

To correct for the fact that $W_\varepsilon(\alpha, \alpha) \neq 0$, (Genevay et al., 2018) propose Sinkhorn divergences, a natural normalization of that quantity defined as

$$\bar{W}_\varepsilon(\alpha, \beta) = W_\varepsilon(\alpha, \beta) - \frac{1}{2}(W_\varepsilon(\alpha, \alpha) + W_\varepsilon(\beta, \beta)). \quad (2)$$

This normalization ensures that $\bar{W}_\varepsilon(\alpha, \alpha) = 0$, but also has a noticeable asymptotic behavior as mentioned in (Genevay et al., 2018). Indeed, when $\varepsilon \rightarrow 0$ one recovers the original (unregularized) OT problem, while choosing $\varepsilon \rightarrow +\infty$ yields the maximum mean discrepancy associated to the kernel $k = -c/2$, where MMD is defined by:

$$\begin{aligned} MMD_k(\alpha, \beta) &= \mathbb{E}_{\alpha \otimes \alpha} [k(X, X')] + \mathbb{E}_{\beta \otimes \beta} [k(Y, Y')] \\ &\quad - 2\mathbb{E}_{\alpha \otimes \beta} [k(X, Y)]. \end{aligned}$$

In the context of this paper, we study in detail the sample complexity of $W_\varepsilon(\alpha, \beta)$, knowing that these results can be extended to $\bar{W}_\varepsilon(\alpha, \beta)$.

3 Approximating Optimal Transport with Sinkhorn Divergences

In the present section, we are interested in bounding the error made when approximating $W(\alpha, \beta)$ with $W_\varepsilon(\alpha, \beta)$.

Theorem 1. *Let α and β be probability measures on \mathcal{X} and \mathcal{Y} subsets of \mathbb{R}^d such that $|\mathcal{X}| = |\mathcal{Y}| \leq D$ and assume that c is L -Lipschitz w.r.t. x and y . It holds*

$$0 \leq W_\varepsilon(\alpha, \beta) - W(\alpha, \beta) \leq 2\varepsilon d \log \left(\frac{e^2 \cdot L \cdot D}{\sqrt{d} \cdot \varepsilon} \right) \quad (3)$$

$$\sim_{\varepsilon \rightarrow 0} 2\varepsilon d \log(1/\varepsilon). \quad (4)$$

Proof. For a probability measure π on $\mathcal{X} \times \mathcal{Y}$, we denote by $C(\pi) = \int c d\pi$ the associated transport cost and by $H(\pi)$ its relative entropy with respect to the product measure $\alpha \otimes \beta$ as defined in (1). Choosing π_0 a minimizer of $\min_{\pi \in \Pi(\alpha, \beta)} C(\pi)$, we will build our upper bounds using a family of transport plans with finite entropy that approximate π_0 . The simplest approach consists in considering block approximation. In contrast to the work of Carlier et al. (2017), who also considered this technique, our focus here is on quantitative bounds.

Definition 1 (Block approximation). *For a resolution $\Delta > 0$, we consider the block partition of \mathbb{R}^d in hypercubes of side Δ defined as*

$$\{Q_k^\Delta = [k_1 \cdot \Delta, (k_1 + 1) \cdot \Delta] \times \dots \times [k_d \cdot \Delta, (k_d + 1) \cdot \Delta]; \quad k = (k_1, \dots, k_d) \in \mathbb{Z}^d\}.$$

To simplify notations, we introduce $Q_{ij}^\Delta \stackrel{\text{def.}}{=} Q_i^\Delta \times Q_j^\Delta$, $\alpha_i^\Delta \stackrel{\text{def.}}{=} \alpha(Q_i^\Delta)$, $\beta_j^\Delta \stackrel{\text{def.}}{=} \beta(Q_j^\Delta)$. The block approximation of π_0 of resolution Δ is the measure $\pi^\Delta \in \Pi(\alpha, \beta)$ characterized by

$$\pi^\Delta|_{Q_{ij}^\Delta} = \frac{\pi_0(Q_{ij}^\Delta)}{\alpha_i^\Delta \cdot \beta_j^\Delta} (\alpha|_{Q_i^\Delta} \otimes \beta|_{Q_j^\Delta})$$

for all $(i, j) \in (\mathbb{Z}^d)^2$, with the convention $0/0 = 0$.

π^Δ is nonnegative by construction. Observe also that for any Borel set $B \subset \mathbb{R}^d$, one has

$$\begin{aligned} \pi^\Delta(B \times \mathbb{R}^d) &= \sum_{(i, j) \in (\mathbb{Z}^d)^2} \frac{\pi_0(Q_{ij}^\Delta)}{\alpha_i^\Delta \cdot \beta_j^\Delta} \cdot \alpha(B \cap Q_i^\Delta) \cdot \beta_j^\Delta \\ &= \sum_{i \in \mathbb{Z}^d} \alpha(B \cap Q_i^\Delta) = \alpha(B), \end{aligned}$$

which proves, using the symmetric result in β , that π^Δ belongs to $\Pi(\alpha, \beta)$. As a consequence, for any $\varepsilon > 0$ one has $W_\varepsilon(\alpha, \beta) \leq C(\pi^\Delta) + \varepsilon H(\pi^\Delta)$. Recalling also that the relative entropy H is nonnegative over the set of probability measures, we have the bound

$$0 \leq W_\varepsilon(\alpha, \beta) - W(\alpha, \beta) \leq (C(\pi^\Delta) - C(\pi_0)) + \varepsilon H(\pi^\Delta).$$

We can now bound the terms in the right-hand side, and choose a value for Δ that minimizes these bounds.

The bound on $C(\pi^\Delta) - C(\pi_0)$ relies on the Lipschitz regularity of the cost function. Using the fact that $\pi^\Delta(Q_{ij}^\Delta) = \pi_0(Q_{ij}^\Delta)$ for all i, j , it holds

$$\begin{aligned} C(\pi^\Delta) - C(\pi_0) &= \sum_{(i,j) \in (\mathbb{Z}^d)^2} \pi_0(Q_{ij}^\Delta) \left(\sup_{x,y \in Q_{ij}^\Delta} c(x,y) \right. \\ &\quad \left. - \inf_{x,y \in Q_{ij}^\Delta} c(x,y) \right) \\ &\leq 2L\Delta\sqrt{d}, \end{aligned}$$

where L is the Lipschitz constant of the cost (separately in x and y) and $\Delta\sqrt{d}$ is the diameter of each set Q_i^Δ .

As for the bound on $H(\pi^\Delta)$, using the fact that $\pi_0(Q_{ij}^\Delta) \leq 1$ we get

$$\begin{aligned} H(\pi^\Delta) &= \sum_{(i,j) \in (\mathbb{Z}^d)^2} \log \left(\frac{\pi_0(Q_{ij}^\Delta)}{\alpha_i^\Delta \cdot \beta_j^\Delta} \right) \pi_0(Q_{ij}^\Delta) \\ &\leq \sum_{(i,j) \in (\mathbb{Z}^d)^2} \left(\log(1/\alpha_i^\Delta) + \log(1/\beta_j^\Delta) \right) \pi_0(Q_{ij}^\Delta) \\ &= -H^\Delta(\alpha) - H^\Delta(\beta), \end{aligned}$$

where we have defined $H^\Delta(\alpha) = \sum_{i \in \mathbb{Z}^d} \alpha_i^\Delta \log(\alpha_i^\Delta)$ and similarly for β . Note that in case α is a discrete measure with finite support, $H^\Delta(\alpha)$ is equal to (minus) the discrete entropy of α as long as Δ is smaller than the minimum separation between atoms of α . However, if α is not discrete then $H^\Delta(\alpha)$ blows up to $-\infty$ as Δ goes to 0 and we need to control how fast it does so. Considering α^Δ the block approximation of α with constant density α_i^Δ/Δ^d on each block Q_i^Δ and (minus) its differential entropy $H_{\mathcal{L}^d}(\alpha^\Delta) = \int_{\mathbb{R}^d} \alpha^\Delta(x) \log \alpha^\Delta(x) dx$, it holds $H^\Delta(\alpha) = H_{\mathcal{L}^d}(\alpha^\Delta) - d \cdot \log(1/\Delta)$. Moreover, using the convexity of $H_{\mathcal{L}^d}$, this can be compared with the differential entropy of the uniform probability on a hypercube containing \mathcal{X} of size $2D$. Thus it holds $H_{\mathcal{L}^d}(\alpha^\Delta) \geq -d \log(2D)$ and thus $H^\Delta(\alpha) \geq -d \cdot \log(2D/\Delta)$.

Summing up, we have for all $\Delta > 0$

$$W_\varepsilon(\alpha, \beta) - W(\alpha, \beta) \leq 2L\Delta\sqrt{d} + 2\varepsilon d \cdot \log(2D/\Delta).$$

The above bound is convex in Δ , minimized with $\Delta = 2\sqrt{d} \cdot \varepsilon/L$. This yields

$$W_\varepsilon(\alpha, \beta) - W(\alpha, \beta) \leq 4\varepsilon d + 2\varepsilon d \log \left(\frac{L \cdot D}{\sqrt{d} \cdot \varepsilon} \right). \quad \square$$

4 Properties of Sinkhorn Potentials

We prove in this section that Sinkhorn potentials are bounded in the Sobolev space $\mathbf{H}^s(\mathbb{R}^d)$ regardless of the marginals α and β . For $s > \frac{d}{2}$, $\mathbf{H}^s(\mathbb{R}^d)$ is a reproducing kernel Hilbert space (RKHS): This property will be crucial to establish sample complexity results later on, using standard tools from RKHS theory.

Definition 2. The Sobolev space $\mathbf{H}^s(\mathcal{X})$, for $s \in \mathbb{N}^*$, is the space of functions $\varphi : \mathcal{X} \subseteq \mathbb{R}^d \rightarrow \mathbb{R}$ such that for every multi-index k with $|k| \leq s$ the mixed partial derivative $\varphi^{(k)}$ exists and belongs to $L^2(\mathcal{X})$. It is endowed with the following inner-product

$$\langle \varphi, \psi \rangle_{\mathbf{H}^s(\mathcal{X})} = \sum_{|k| \leq s} \int_{\mathcal{X}} \varphi^{(k)}(x) \psi^{(k)}(x) dx. \quad (5)$$

Theorem 2. When \mathcal{X} and \mathcal{Y} are two compact sets of \mathbb{R}^d and the cost c is \mathcal{C}^∞ , then the Sinkhorn potentials (u, v) are uniformly bounded in the Sobolev space $\mathbf{H}^s(\mathbb{R}^d)$ and their norms satisfy

$$\|u\|_{\mathbf{H}^s} = O\left(1 + \frac{1}{\varepsilon^{s-1}}\right) \text{ and } \|v\|_{\mathbf{H}^s} = O\left(1 + \frac{1}{\varepsilon^{s-1}}\right),$$

with constants that only depend on $|\mathcal{X}|$ (or $|\mathcal{Y}|$ for v), d , and $\|c^{(k)}\|_\infty$ for $k = 0, \dots, s$. In particular, we get the following asymptotic behavior in ε : $\|u\|_{\mathbf{H}^s} = O(1)$ as $\varepsilon \rightarrow +\infty$ and $\|u\|_{\mathbf{H}^s} = O(\frac{1}{\varepsilon^{s-1}})$ as $\varepsilon \rightarrow 0$.

To prove this theorem, we first need to state some regularity properties of the Sinkhorn potentials.

Proposition 1. If \mathcal{X} and \mathcal{Y} are two compact sets of \mathbb{R}^d and the cost c is \mathcal{C}^∞ , then

- $u(x) \in [\min_y v(y) - c(x, y), \max_y v(y) - c(x, y)]$ for all $x \in \mathcal{X}$
- u is L -Lipschitz, where L is the Lipschitz constant of c
- $u \in \mathcal{C}^\infty(\mathcal{X})$ and $\|u^{(k)}\|_\infty = O(1 + \frac{1}{\varepsilon^{k-1}})$

and the same results also stand for v (inverting u and v in the first item, and replacing \mathcal{X} by \mathcal{Y}).

Proof. The proofs of all three claims exploit the optimality condition of the dual problem:

$$\exp\left(\frac{-u(x)}{\varepsilon}\right) = \int \exp\left(\frac{v(y) - c(x, y)}{\varepsilon}\right) \beta(y) dy. \quad (6)$$

Since β is a probability measure, $e^{\frac{-u(x)}{\varepsilon}}$ is a convex combination of $\varphi : x \mapsto e^{\frac{v(x) - c(x, y)}{\varepsilon}}$ and thus $e^{\frac{-u(x)}{\varepsilon}} \in [\min_y \varphi(y), \max_y \varphi(y)]$. We get the desired bounds by taking the logarithm. The two other points use the following lemmas:

Lemma 1. The derivatives of the potentials are given by the following recurrence

$$u^{(n)}(x) = \int g_n(x, y) \gamma_\varepsilon(x, y) \beta(y) dy, \quad (7)$$

where

$$\begin{aligned} g_{n+1}(x, y) &= g'_n(x, y) + \frac{u'(x) - c'(x, y)}{\varepsilon} g_n(x, y), \\ g_1(x, y) &= c'(x, y) \text{ and } \gamma_\varepsilon(x, y) = \exp\left(\frac{u(x) + v(y) - c(x, y)}{\varepsilon}\right). \end{aligned}$$

Lemma 2. *The sequence of auxiliary functions $(g_k)_{k=0\dots}$ verifies $\|u^{(k)}\|_\infty \leq \|g_k\|_\infty$. Besides, for all $j = 0, \dots, k$, for all $k = 0, \dots, n - 2$, $\|g_{n-k}^{(j)}\|_\infty$ is bounded by a polynomial in $\frac{1}{\varepsilon}$ of order $n - k + j - 1$.*

The detailed proofs of the lemmas can be found in the appendix. We give here a sketch in the case where $d = 1$. Lemma 1 is obtained by a simple recurrence, consisting in differentiating both sides of the dual optimality condition. Differentiating under the integral is justified with the usual domination theorem, bounding the integrand thanks to the Lipschitz assumption on c , and this bound is integrable thanks to the marginal constraint. Differentiating once and rearranging terms gives:

$$u'(x) = \int c'(x, y) \gamma_\varepsilon(x, y) \beta(y) dy. \quad (8)$$

where γ_ε is defined in Lemma 1. One can easily see that $\gamma'_\varepsilon(x, y) = \frac{u'(x) - c'(x, y)}{\varepsilon} \gamma_\varepsilon(x, y)$ and this allows to conclude the recurrence, by differentiating both sides of the equality. From the primal constraint, we have that $\int_Y \gamma_\varepsilon(x, y) \beta(y) dy = 1$. Thus thanks to Lemma 1 we immediately get that $\|u^{(n)}\|_\infty \leq \|g_n\|_\infty$. For $n = 1$, since $g_1 = c'$ we get that $\|u'\|_\infty = \|c'\|_\infty = L$ and this proves the second point of Proposition 1. For higher values of n , we need the result from Lemma 2. This property is also proved by recurrence, but requires a bit more work. To prove the induction step, we need to go from bounds on $g_{n-k}^{(i)}$, for $k = 0, \dots, n - 2$ and $i = 0, \dots, k$ to bounds on $g_{n+1-k}^{(i)}$, for $k = 0, \dots, n - 1$ and $i = 0, \dots, k$. Hence only new quantities that we need to bound are $g_{n+1-k}^{(k)}$, $k = 0, \dots, n - 1$. This is done by another (backwards) recurrence on k which involves some tedious computations, based on Leibniz formula, that are detailed in the appendix. \square

Combining the bounds of the derivatives of the potentials with the definition of the norm in \mathbf{H}^s , is enough to complete the proof of Theorem 2.

Proof. (Theorem 2) The norm of u in $\mathbf{H}^s(\mathcal{X})$ is

$$\|u\|_{\mathbf{H}^s} = \left(\sum_{|k| \leq s} \int_{\mathcal{X}} (u^{(k)})^2 \right)^{\frac{1}{2}} \leq |\mathcal{X}| \left(\sum_{|k| \leq s} \|u^{(k)}\|_\infty^2 \right)^{\frac{1}{2}}.$$

From Proposition 1 we have that $\forall k, \|u^{(k)}\|_\infty = O(1 + \frac{1}{\varepsilon^{k-1}})$ and thus we get that $\|u\|_{\mathbf{H}^s} = O(1 + \frac{1}{\varepsilon^{s-1}})$. We just proved the bound in $\mathbf{H}^s(\mathcal{X})$ but we actually want to have a bound on $\mathbf{H}^s(\mathbb{R}^d)$. This is immediate thanks to the Sobolev extension theorem (Calderón, 1961) which guarantees that $\|u\|_{\mathbf{H}^s(\mathbb{R}^d)} \leq C \|u\|_{\mathbf{H}^s(\mathcal{X})}$ under the assumption that \mathcal{X} is a bounded Lipschitz domain. \square

This result, aside from proving useful in the next section to obtain sample complexity results on the Sinkhorn divergence, also proves that kernel-SGD can be used to solve continuous regularized OT. This idea introduced in Genevay et al. (2016) consists in assuming the potentials are in the ball of a certain RKHS, to write them as a linear combination of kernel functions and then perform stochastic gradient descent on these coefficients. Knowing the radius of the ball and the kernel associated with the RKHS (here the Sobolev or Matérn kernel) is crucial to obtain good numerical performance and ensure the convergence of the algorithm.

5 Approximation from Samples

In practice, measures α and β are only known through a finite number of samples. Thus, what can be actually computed in practice is the Sinkhorn divergence between the empirical measures $\hat{\alpha}_n \stackrel{\text{def.}}{=} \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ and $\hat{\beta}_n \stackrel{\text{def.}}{=} \frac{1}{n} \sum_{i=1}^n \delta_{Y_i}$, where (X_1, \dots, X_n) and (Y_1, \dots, Y_n) are n -samples from α and β , that is

$$\begin{aligned} W_\varepsilon(\hat{\alpha}_n, \hat{\beta}_n) &= \max_{u, v} \sum_{i=1}^n u(X_i) + \sum_{i=1}^n v(Y_i) \\ &\quad - \varepsilon \sum_{i=1}^n \exp\left(\frac{u(X_i) + v(Y_i) - c(X_i, Y_i)}{\varepsilon}\right) + \varepsilon \\ &= \max_{u, v} \frac{1}{n} \sum_{i=1}^n f_\varepsilon^{X_i Y_i}(u, v) + \varepsilon, \end{aligned}$$

where $(X_i, Y_i)_{i=1}^n$ are i.i.d random variables distributed according to $\alpha \otimes \beta$. On actual samples, these quantities can be computed using Sinkhorn's algorithm (Cuturi, 2013).

Our goal is to quantify the error that is made by approximating α, β by their empirical counterparts $\hat{\alpha}_n, \hat{\beta}_n$, that is bounding the following quantity:

$$\begin{aligned} |W_\varepsilon(\alpha, \beta) - W_\varepsilon(\hat{\alpha}_n, \hat{\beta}_n)| &= \\ |\mathbb{E} f_\varepsilon^{XY}(u^*, v^*) - \frac{1}{n} \sum_{i=1}^n f_\varepsilon^{X_i Y_i}(\hat{u}, \hat{v})|, \quad (9) \end{aligned}$$

where (u^*, v^*) are the optimal Sinkhorn potentials associated with (α, β) and (\hat{u}, \hat{v}) are their empirical counterparts.

Theorem 3. *Consider the Sinkhorn divergence between two measures α and β on \mathcal{X} and \mathcal{Y} two bounded subsets of \mathbb{R}^d , with a C^∞ , L -Lipschitz cost c . One has*

$$\mathbb{E} |W_\varepsilon(\alpha, \beta) - W_\varepsilon(\hat{\alpha}_n, \hat{\beta}_n)| = O\left(\frac{e^{\frac{\kappa}{\varepsilon}}}{\sqrt{n}} \left(1 + \frac{1}{\varepsilon^{\lfloor d/2 \rfloor}}\right)\right)$$

where $\kappa = 2L|\mathcal{X}| + \|c\|_\infty$ and constants only depend on $|\mathcal{X}|, |\mathcal{Y}|, d$, and $\|c^{(k)}\|_\infty$ for $k = 0 \dots \lfloor d/2 \rfloor$. In particular, we get the following asymptotic behavior in

ε :

$$\begin{aligned} \mathbb{E}|W_\varepsilon(\alpha, \beta) - W_\varepsilon(\hat{\alpha}_n, \hat{\beta}_n)| &= O\left(\frac{e^{\frac{\kappa}{\varepsilon}}}{\varepsilon^{\lfloor d/2 \rfloor} \sqrt{n}}\right) \text{ as } \varepsilon \rightarrow 0 \\ \mathbb{E}|W_\varepsilon(\alpha, \beta) - W_\varepsilon(\hat{\alpha}_n, \hat{\beta}_n)| &= O\left(\frac{1}{\sqrt{n}}\right) \text{ as } \varepsilon \rightarrow +\infty. \end{aligned}$$

An interesting feature from this theorem is the fact when ε is large enough, the convergence rate does not depend on ε anymore. This means that at some point, increasing ε will not substantially improve convergence. However, for small values of ε the dependence is critical.

We prove this result in the rest of this section. The main idea is to exploit standard results from PAC-learning in RKHS. Our theorem is an application of the following result from Bartlett and Mendelson (2002) (combining Theorem 12,4) and Lemma 22 in their paper):

Proposition 2. (Bartlett-Mendelson '02) Consider α a probability distribution, ℓ a B-lipschitz loss and \mathcal{G} a given class of functions. Then

$$\mathbb{E}_\alpha \left[\sup_{g \in \mathcal{G}} \mathbb{E}_\alpha \ell(g, X) - \frac{1}{n} \sum_{i=1}^n \ell(g, X_i) \right] \leq 2B \mathbb{E}_\alpha \mathcal{R}(\mathcal{G}(X_1^n))$$

where $\mathcal{R}(\mathcal{G}(X_1^n))$ is the Rademacher complexity of class \mathcal{G} defined by $\mathcal{R}(\mathcal{G}(X_1^n)) = \sup_{g \in \mathcal{G}} \mathbb{E}_\sigma \frac{1}{n} \sum_{i=1}^n \sigma_i g(X_i)$ where $(\sigma_i)_i$ are iid Rademacher random variables. Besides, when \mathcal{G} is a ball of radius λ in a RKHS with kernel k the Rademacher complexity is bounded by

$$\mathcal{R}(\mathcal{G}_\lambda(X_1^n)) \leq \frac{\lambda}{n} \sqrt{\sum_{i=1}^n k(X_i, X_i)}.$$

Our problem falls in this framework thanks to the following lemma:

Lemma 3. Let $\mathcal{H}_\lambda^s \stackrel{\text{def}}{=} \{u \in \mathbf{H}^s(\mathbb{R}^d) \mid \|u\|_{\mathbf{H}^s(\mathbb{R}^d)} \leq \lambda\}$, then there exists λ such that:

$$|W_\varepsilon(\alpha, \beta) - W_\varepsilon(\hat{\alpha}_n, \hat{\beta}_n)| \leq$$

$$3 \sup_{(u,v) \in (\mathcal{H}_\lambda^s)^2} |\mathbb{E} f_\varepsilon^{XY}(u, v) - \frac{1}{n} \sum_{i=1}^n f_\varepsilon^{X_i Y_i}(u, v)|.$$

Proof. Inserting $\mathbb{E} f_\varepsilon^{XY}(\hat{u}, \hat{v})$ and using the triangle inequality in (9) gives

$$\begin{aligned} |W_\varepsilon(\alpha, \beta) - W_\varepsilon(\hat{\alpha}_n, \hat{\beta}_n)| &\leq |\mathbb{E} f_\varepsilon^{XY}(u^*, v^*) - \mathbb{E} f_\varepsilon^{XY}(\hat{u}, \hat{v})| \\ &+ |\mathbb{E} f_\varepsilon^{XY}(\hat{u}, \hat{v}) - \frac{1}{n} \sum_{i=1}^n f_\varepsilon^{X_i Y_i}(\hat{u}, \hat{v})|. \end{aligned}$$

From Theorem 2, we know that the all the dual potentials are bounded in $\mathbf{H}^s(\mathbb{R}^d)$ by a constant λ

which doesn't depend on the measures. Thus the second term is bounded by $\sup_{(u,v) \in (\mathcal{H}_\lambda^s)^2} |\mathbb{E} f_\varepsilon(u, v) -$

The first quantity needs to be broken down further. Notice that it is non-negative since (u^*, v^*) is the maximizer of $\mathbb{E} f_\varepsilon(\cdot, \cdot)$ so we can leave out the absolute value. We have:

$$\begin{aligned} \mathbb{E} f_\varepsilon^{XY}(u^*, v^*) - \mathbb{E} f_\varepsilon^{XY}(\hat{u}, \hat{v}) &\leq \\ \mathbb{E} f_\varepsilon^{XY}(u^*, v^*) - \frac{1}{n} \sum_{i=1}^n f_\varepsilon^{X_i Y_i}(u^*, v^*) &\quad (10) \end{aligned}$$

$$+ \frac{1}{n} \sum_{i=1}^n f_\varepsilon^{X_i Y_i}(u^*, v^*) - \frac{1}{n} \sum_{i=1}^n f_\varepsilon^{X_i Y_i}(\hat{u}, \hat{v}) \quad (11)$$

$$+ \frac{1}{n} \sum_{i=1}^n f_\varepsilon^{X_i Y_i}(\hat{u}, \hat{v}) - \mathbb{E} f_\varepsilon^{XY}(\hat{u}, \hat{v}) \quad (12)$$

Both (10) and (12) can be bounded by $\sup_{(u,v) \in (\mathcal{H}_\lambda^s)^2} |\mathbb{E} f_\varepsilon^{XY}(u, v) - \frac{1}{n} \sum_{i=1}^n f_\varepsilon^{X_i Y_i}(u, v)|$ while (11) is non-positive since (\hat{u}, \hat{v}) is the maximizer of $\frac{1}{n} \sum_{i=1}^n f_\varepsilon^{X_i Y_i}(\cdot, \cdot)$. \square

To apply Proposition 2 to Sinkhorn divergences we need to prove that (a) the optimal potentials are in a RKHS and (b) our loss function f_ε is Lipschitz in the potentials.

The first point has already been proved in the previous section. The RKHS we are considering is $\mathbf{H}^s(\mathbb{R}^d)$ with $s = \lfloor \frac{d}{2} \rfloor + 1$. It remains to prove that f_ε is Lipschitz in (u, v) on a certain subspace that contains the optimal potentials.

Lemma 4. Let $\mathcal{A} = \{(u, v) \mid u \oplus v \leq 2L|\mathcal{X}| + \|c\|_\infty\}$. We have:

(i) the pairs of optimal potentials (u^*, v^*) such that $u^*(0) = 0$ belong to \mathcal{A} ,

(ii) f_ε is B-Lipschitz in (u, v) on \mathcal{A} with $B \leq 1 + \exp(2 \frac{L|\mathcal{X}| + \|c\|_\infty}{\varepsilon})$.

Proof. Let us prove that we can restrict ourselves to a subspace on which f_ε is Lipschitz in (u, v) .

$$f_\varepsilon^{xy}(u, v) = u(x) + v(y) - \varepsilon \exp\left(\frac{u(x) + v(y) - c(x, y)}{\varepsilon}\right)$$

$$\nabla f_\varepsilon(u, v) = 1 - \exp\left(\frac{u + v - c}{\varepsilon}\right).$$

To ensure that f_ε is Lipschitz, we simply need to ensure that the quantity inside the exponential is upper-bounded at optimality and then restrict the function to all (u, v) that satisfy that bound.




Figure 1: $\bar{W}_\varepsilon(\hat{\alpha}_n, \hat{\alpha}'_n)$ as a function of n in log-log space : Influence of ε for fixed d on two uniform distributions on the hypercube with quadratic cost.




Figure 2: $\bar{W}_\varepsilon(\hat{\alpha}_n, \hat{\alpha}'_n)$ as a function of n in log-log space : Influence of d for fixed ε on two uniform distributions on the hypercube with quadratic cost.

Recall the bounds on the optimal potentials from Proposition 1. We have that $\forall x \in \mathcal{X}, y \in \mathcal{Y}$,

$$u(x) \leq L|x| \quad \text{and} \quad v(y) \leq \max_x u(x) - c(x, y).$$

Since we assumed \mathcal{X} to be a bounded set, denoting by $|\mathcal{X}|$ the diameter of the space we get that at optimality $\forall x \in \mathcal{X}, y \in \mathcal{Y}$

$$u(x) + v(y) \leq 2L|\mathcal{X}| + \|c\|_\infty.$$

Let us denote $\mathcal{A} = \{(u, v) \in (\mathbf{H}^s(\mathbb{R}^d))^2 \mid u \oplus v \leq 2L|\mathcal{X}| + \|c\|_\infty\}$, we have that $\forall (u, v) \in \mathcal{A}$,

$$|\nabla f_\varepsilon(u, v)| \leq 1 + \exp\left(2 \frac{L|\mathcal{X}| + \|c\|_\infty}{\varepsilon}\right). \quad \square$$

We now have all the required elements to prove our sample complexity result on the Sinkhorn loss, by applying Proposition 2.

Proof. (Theorem 3) Since f_ε is Lipschitz and we are optimizing over $\mathbf{H}^s(\mathbb{R}^d)$ which is a RKHS, we can apply Proposition 2 to bound the sup in Lemma 3. We get:

$$\mathbb{E}|W_\varepsilon(\alpha, \beta) - W_\varepsilon(\hat{\alpha}_n, \hat{\beta}_n)| \leq 3 \frac{2B\lambda}{n} \mathbb{E} \sqrt{\sum_{i=1}^n k(X_i, X_i)}$$

where $B \leq 1 + \exp(2 \frac{L|\mathcal{X}| + \|c\|_\infty}{\varepsilon})$ (Lemma 4), $\lambda = O(\max(1, \frac{1}{\varepsilon^{d/2}}))$ (Theorem 2). We can further bound

$\sqrt{\sum_{i=1}^n k(X_i, X_i)}$ by $\sqrt{n \max_{x \in \mathcal{X}} k(x, x)}$ where k is the kernel associated to $H^s(\mathbb{R}^d)$ (usually called Matern or Sobolev kernel) and thus $\max_{x \in \mathcal{X}} k(x, x) = k(0, 0) := K$ which doesn't depend on n or ε . Combining all these bounds, we get the convergence rate in $\frac{1}{\sqrt{n}}$ with different asymptotic behaviors in ε when it is large or small. \square

Using similar arguments, we can also derive a concentration result:

Corollary 1. *With probability at least $1 - \delta$,*

$$|W_\varepsilon(\alpha, \beta) - W_\varepsilon(\hat{\alpha}_n, \hat{\beta}_n)| \leq 6B \frac{\lambda K}{\sqrt{n}} + C \sqrt{\frac{2 \log \frac{1}{\delta}}{n}}$$

where B, λ, K are defined in the proof above, and $C = \kappa + \varepsilon \exp(\frac{\kappa}{\varepsilon})$ with $\kappa = 2L|\mathcal{X}| + \|c\|_\infty$.

Proof. We apply the bounded differences (Mc Diarmid) inequality to $g : (x_1, \dots, x_n) \mapsto \sup_{u, v \in \mathcal{H}_\lambda^s} (\mathbb{E} f_\varepsilon^{XY} - \frac{1}{n} f_\varepsilon^{X_i, Y_i})$. From Lemma 4 we get that $\forall x, y, f_\varepsilon^{xy}(u, v) \leq \kappa + \varepsilon e^{\kappa/\varepsilon} \stackrel{\text{def}}{=} C$, and thus, changing one of the variables in g changes the value of the function by at most $2C/n$. Thus the bounded differences inequality gives

$$\mathbb{P}(|g(X_1, \dots, X_n) - \mathbb{E}g(X_1, \dots, X_n)| > t) \leq 2 \exp\left(-\frac{t^2 n}{2C^2}\right)$$

Choosing $t = C\sqrt{\frac{2\log\frac{1}{\delta}}{n}}$ yields that with probability at least $1 - \delta$

$$g(X_1, \dots, X_n) \leq \mathbb{E}g(X_1, \dots, X_n) + C\sqrt{\frac{2\log\frac{1}{\delta}}{n}}$$

and from Theorem 3 we already have

$$\mathbb{E}g(X_1, \dots, X_n) = \mathbb{E} \sup_{u, v \in \mathcal{H}_\lambda^s} (\mathbb{E}f_\varepsilon^{XY} - \frac{1}{n}f_\varepsilon^{X_i, Y_i}) \leq \frac{2B\lambda K}{\sqrt{n}}.$$

□

6 Experiments

We conclude with some numerical experiments on the sample complexity of Sinkhorn Divergences. Since there are no explicit formulas for W_ε in general, we consider $\bar{W}_\varepsilon(\hat{\alpha}_n, \hat{\alpha}'_n)$ where $\hat{\alpha}_n \stackrel{\text{def.}}{=} \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$, $\hat{\alpha}'_n \stackrel{\text{def.}}{=} \frac{1}{n} \sum_{i=1}^n \delta_{X'_i}$ and (X_1, \dots, X_n) and (X'_1, \dots, X'_n) are two independent n -samples from α . Note that we use in this section the normalized Sinkhorn Divergence as defined in (2), since we know that $\bar{W}_\varepsilon(\alpha, \alpha) = 0$ and thus $\bar{W}_\varepsilon(\hat{\alpha}_n, \hat{\alpha}'_n) \rightarrow 0$ as $n \rightarrow +\infty$.

Each of the experiments is run 300 times, and we plot the average of $\bar{W}_\varepsilon(\hat{\alpha}_n, \hat{\alpha}'_n)$ as a function of n in log-log space, with shaded standard deviation bars.

First, we consider the uniform distribution over a hypercube with the standard quadratic cost $c(x, y) = \|x - y\|_2^2$, which falls within our framework, as we are dealing with a C^∞ cost on a bounded domain. Figure 1 shows the influence of the dimension d on the convergence, while Figure 2 shows the influence of the regularization ε on the convergence for a given dimension. The influence of ε on the convergence rate increases with the dimension: the curves are almost parallel for all values of ε in dimension 2 but they get further apart as dimension increases. As expected from our bound, there is a cutoff which happens here at $\varepsilon = 1$. All values of $\varepsilon \geq 1$ have similar convergence rates, and the dependence on $\frac{1}{\varepsilon}$ becomes clear for smaller values. The same cutoff appears when looking at the influence of the dimension on the convergence rate for a fixed ε . The curves are parallel for all dimensions for $\varepsilon \geq 1$ but they have very different slopes for smaller ε .

We relax next some of the assumptions needed in our theorem to see how the Sinkhorn divergence behaves empirically. First we relax the regularity assumption on the cost, using $c(x, y) = \|x - y\|_1$. As seen on the two left images in figure 3 the behavior is very similar to the quadratic cost but with a more pronounced influence of ε , even for small dimensions. The fact that the convergence rate gets slower as ε gets smaller is already

very clear in dimension 2, which wasn't the case for the quadratic cost. The influence of the dimension for a given value of ε is not any different however.

We also relax the bounded domain assumption, considering a standard normal distribution over \mathbb{R}^d with a quadratic cost. While the influence of ε on the convergence rate is still obvious, the influence of the dimension is less clear. There is also a higher variance, which can be expected as the concentration bound from Corollary 1 depends on the diameter of the domain.

For all curves, we observe that d and ε impact variance, with much smaller variance for small values of ε and high dimensions. From the concentration bound, the dependency on ε coming from the uniform bound on f_ε is of the form $\varepsilon \exp(\kappa/\varepsilon)$, suggesting higher variance for small values of ε . This could indicate that our uniform bound on f_ε is not tight, and we should consider other methods to get tighter bounds in further work.

7 Conclusion

We have presented two convergence theorems for SDs: a bound on the approximation error of OT and a sample complexity bound for empirical Sinkhorn divergences. The $1/\sqrt{n}$ convergence rate is similar to MMD, but with a constant that depends on the inverse of the regularization parameter, which nicely complements the interpolation property of SDs pointed out in recent papers. Furthermore, the reformulation of SDs as the maximization of an expectation in a RKHS ball also opens the door to a better use of kernel-SGD for the computation of SDs.

Our numerical experiments suggest some open problems. It seems that the convergence rate still holds for unbounded domains and non-smooth cost functions. Besides, getting tighter bounds in our theorem might allow us to derive a sharp estimate on the optimal ε to approximate OT for a given n , by combining our two convergence theorems together.

References

- L. Ambrosio, N. Gigli, and G. Savaré. *Gradient flows in metric spaces and in the space of probability measures*. Springer, 2006.
- M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- P. L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3 (Nov):463–482, 2002.
- M. A. Beaumont, W. Zhang, and D. J. Balding. Approx-




Figure 3: $\bar{W}_\varepsilon(\hat{\alpha}_n, \hat{\alpha}'_n)$ as a function of n in log-log space - cost $c(x, y) = \|x - y\|_1$ with uniform distributions (two leftmost figures) and quadratic cost $c(x, y) = \|x - y\|_2^2$ with standard normal distributions (right figure).

- imate bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002.
- J. Bigot, E. Cazelles, and N. Papadakis. Central limit theorems for sinkhorn divergence between probability distributions on finite spaces and statistical applications. *arXiv preprint arXiv:1711.08947*, 2017.
- P. J. Bushell. Hilbert’s metric and positive contraction mappings in a Banach space. *Archive for Rational Mechanics and Analysis*, 52(4):330–338, 1973.
- A. Calderón. Lebesgue spaces of differentiable functions. In *Proc. Sympos. Pure Math*, volume 4, pages 33–49, 1961.
- G. Carlier, V. Duval, G. Peyré, and B. Schmitzer. Convergence of entropic schemes for optimal transport and gradient flows. *SIAM Journal on Mathematical Analysis*, 49(2):1385–1418, 2017.
- Y. Chen, T. Georgiou, and M. Pavon. Entropic and displacement interpolation: a computational approach using the hilbert metric. *SIAM Journal on Applied Mathematics*, 76(6):2375–2396, 2016.
- N. Courty, R. Flamary, and D. Tuia. Domain adaptation with regularized optimal transport. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 274–289. Springer, 2014.
- M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Adv. in Neural Information Processing Systems*, pages 2292–2300, 2013.
- R. Dudley. The speed of mean glivenko-cantelli convergence. *The Annals of Mathematical Statistics*, 40(1): 40–50, 1969.
- G. Dziugaite, D. Roy, and Z. Ghahramani. Training generative neural networks via maximum mean discrepancy optimization. In *Uncertainty in Artificial Intelligence-Proceedings of the 31st Conference, UAI 2015*, pages 258–267, 2015.
- J. Franklin and J. Lorenz. On the scaling of multidimensional matrices. *Linear Algebra and its applications*, 114:717–735, 1989.
- C. Frogner, C. Zhang, H. Mobahi, M. Araya, and T. Poggio. Learning with a Wasserstein loss. In *Adv. in Neural Information Processing Systems*, pages 2044–2052, 2015.
- A. Genevay, M. Cuturi, G. Peyré, and F. Bach. Stochastic optimization for large-scale optimal transport. In D. D. Lee, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Proc. NIPS’16*, pages 3432–3440. Curran Associates, Inc., 2016.
- A. Genevay, G. Peyre, and M. Cuturi. Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617, 2018.
- A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. Smola. A kernel method for the two-sample problem. In *Adv. in Neural Information Processing Systems*, pages 513–520, 2006.
- G. Huang, C. Guo, M. J. Kusner, Y. Sun, F. Sha, and K. Q. Weinberger. Supervised word mover’s distance. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 4862–4870. Curran Associates, Inc., 2016.
- L. Kantorovich. On the transfer of masses (in Russian). *Doklady Akademii Nauk*, 37(2):227–229, 1942.
- C.-L. Li, W.-C. Chang, Y. Cheng, Y. Yang, and B. Póczos. MMD GAN: Towards deeper understanding of moment matching network. *arXiv preprint arXiv:1705.08584*, 2017.
- Y. Li, K. Swersky, and R. Zemel. Generative moment matching networks. In *International Conference on Machine Learning*, pages 1718–1727, 2015.
- A. Ramdas, N. G. Trillos, and M. Cuturi. On wasserstein two-sample testing and related families of non-parametric tests. *Entropy*, 19(2):47, 2017.
- T. Salimans, H. Zhang, A. Radford, and D. Metaxas. Improving GANs using optimal transport. In *International Conference on Learning Representations*, 2018.

- B. Sriperumbudur, K. Fukumizu, A. Gretton, B. Schoelkopf, and G. Lanckriet. On the empirical estimation of integral probability metrics. *Electronic Journal of Statistics*, 6:1550–1599, 2012.
- J. Weed and F. Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in wasserstein distance. *arXiv preprint arXiv:1707.00087*, 2017.
- J. Weed and Q. Berthet. Estimation of smooth densities in wasserstein distance. *arXiv preprint arXiv:1902.01778*, 2019.