



HAL
open science

Interpolating between Optimal Transport and MMD using Sinkhorn Divergences

Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-Ichi Amari,
Alain Trouvé, Gabriel Peyré

► **To cite this version:**

Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-Ichi Amari, Alain Trouvé, et al.. Interpolating between Optimal Transport and MMD using Sinkhorn Divergences. 2018. hal-01898858

HAL Id: hal-01898858

<https://hal.science/hal-01898858v1>

Preprint submitted on 18 Oct 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Interpolating between Optimal Transport and MMD using Sinkhorn Divergences

Jean Feydy

DMA, École Normale Supérieure
CMLA, ENS Paris-Saclay

Thibault Séjourné

DMA, École Normale Supérieure

François-Xavier Vialard

LIGM, UPEM

Shun-ichi Amari

Brain Science Institute, RIKEN

Alain Trounev

CMLA, ENS Paris-Saclay

Gabriel Peyré

DMA, École Normale Supérieure

Abstract

Comparing probability distributions is a fundamental problem in data sciences. Simple norms and divergences such as the total variation and the relative entropy only compare densities in a point-wise manner and fail to capture the geometric nature of the problem. In sharp contrast, Maximum Mean Discrepancies (MMD) and Optimal Transport distances (OT) are two classes of distances between measures that take into account the geometry of the underlying space and metrize the convergence in law.

This paper studies the Sinkhorn divergences, a family of geometric divergences that interpolates between MMD and OT. Relying on a new notion of geometric entropy, we provide theoretical guarantees for these divergences: positivity, convexity and metrization of the convergence in law. On the practical side, we detail a numerical scheme that enables the large scale application of these divergences for machine learning: on the GPU, gradients of the Sinkhorn loss can be computed for batches of a million samples.

1 Introduction

Countless methods in machine learning and imaging sciences rely on comparisons between probability distributions. With applications ranging from shape matching (Vaillant and Glaunès, 2005; Kaltenmark et al., 2017) to classification (Frogner et al., 2015) and generative model training (Goodfellow et al., 2014), a common setting is that of *measure fitting*: given a unit-mass, positive empirical distribution $\beta \in \mathcal{M}_1^+(\mathcal{X})$ on a feature space \mathcal{X} , a loss function $L : \mathcal{M}_1^+(\mathcal{X}) \times \mathcal{M}_1^+(\mathcal{X}) \rightarrow \mathbb{R}$ and a model distribution $\alpha_\theta \in \mathcal{M}_1^+(\mathcal{X})$ parameterized by a vector θ , we strive to minimize

$\theta \mapsto L(\alpha_\theta, \beta)$ through gradient descent. Numerous papers focus on the construction of suitable models $\theta \mapsto \alpha_\theta$. But which loss function L should we use? If \mathcal{X} is endowed with a *ground distance* $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_+$, taking it into account can make sense and help descent algorithm to overcome spurious local minima.

Geometric divergences for Machine Learning.

Unfortunately, simple dissimilarities such as the Total Variation norm or the Kullback-Leibler relative entropy do not take into account the distance d on the feature space \mathcal{X} . As a result, they do not metrize the convergence in law (aka. the weak* topology of measures) and are unstable with respect to deformations of the distributions' supports. We recall that if \mathcal{X} is compact, α_n converges weak* towards α (denoted $\alpha_n \rightharpoonup \alpha$) if for all continuous test functions $f \in \mathcal{C}(\mathcal{X})$, $\langle \alpha_n, f \rangle \rightarrow \langle \alpha, f \rangle$ where $\langle \alpha, f \rangle \stackrel{\text{def.}}{=} \int_{\mathcal{X}} f d\alpha = \mathbb{E}[f(X)]$ for any random vector X with law α .

The two main classes of losses $L(\alpha, \beta)$ which avoid these shortcomings are Optimal Transport distances and Maximum Mean Discrepancies: they are continuous with respect to the convergence in law and metrize its topology. That is, $\alpha_n \rightharpoonup \alpha \Leftrightarrow L(\alpha_n, \alpha) \rightarrow 0$. The main purpose of this paper is to study the theoretical properties of a new class of *geometric* divergences which interpolates between these two families and thus offers an extra degree of freedom through a parameter ε that can be cross-validated in typical learning scenarios.

1.1 Previous works

OT distances and entropic regularization. A first class of geometric distances between measures is that of Optimal Transportation (OT) costs, which are computed as solutions of a linear program (Kantorovich, 1942) (see (1) below in the special case $\varepsilon = 0$). Enjoying many theoretical properties, these costs allow us to lift a “ground metric” on the fea-

ture space \mathcal{X} towards a metric on the space $\mathcal{M}_+^1(\mathcal{X})$ of probability distributions (Santambrogio, 2015). OT distances (sometimes referred to as Earth Mover’s Distances (Rubner et al., 2000)) are progressively being adopted as an effective tool in a wide range of situations, from computer graphics (Bonneel et al., 2016) to supervised learning (Frogner et al., 2015), unsupervised density fitting (Bassetti et al., 2006) and generative model learning (Montavon et al., 2016; Arjovsky et al., 2017; Salimans et al., 2018; Genevay et al., 2018; Sanjabi et al., 2018). However, in practice, solving the linear problem required to compute these OT distances is a challenging issue; many algorithms that leverage the properties of the underlying feature space (\mathcal{X}, d) have thus been designed to accelerate the computations, see (Peyré and Cuturi, 2017) for an overview.

Out of this collection of methods, entropic regularization has recently emerged as a computationally efficient way of approximating OT costs. For $\varepsilon > 0$, we define

$$\text{OT}_\varepsilon(\alpha, \beta) \stackrel{\text{def.}}{=} \min_{\pi_1=\alpha, \pi_2=\beta} \int_{\mathcal{X}^2} C \, d\pi + \varepsilon \text{KL}(\pi | \alpha \otimes \beta) \quad (1)$$

$$\text{where } \text{KL}(\pi | \alpha \otimes \beta) \stackrel{\text{def.}}{=} \int_{\mathcal{X}^2} \log \left(\frac{d\pi}{d\alpha d\beta} \right) d\pi,$$

where $C(x, y)$ is some symmetric positive cost function (we assume here that $C(x, x) = 0$) and where the minimization is performed over coupling measures $\pi \in \mathcal{M}_+^1(\mathcal{X}^2)$ as (π_1, π_2) denotes the two marginals of π . Typically, $C(x, y) = \|x - y\|^p$ on $\mathcal{X} \subset \mathbb{R}^D$ and setting $\varepsilon = 0$ in (1) allows us to retrieve the Earth Mover ($p = 1$) or the quadratic Wasserstein ($p = 2$) distances.

The idea of adding an entropic barrier $\text{KL}(\cdot | \alpha \otimes \beta)$ to the original linear OT program can be traced back to Schrödinger’s problem (Léonard, 2013) and has been used for instance in social sciences (Galichon and Salanié, 2010). Crucially, as highlighted in (Cuturi, 2013), the smooth problem (1) can be solved efficiently on the GPU as soon as $\varepsilon > 0$: the celebrated Sinkhorn algorithm (detailed in Section 3) allows us to compute efficiently a smooth, geometric loss OT_ε between sampled measures.

MMD norms. Still, to define geometry-aware distances between measures, a simpler approach is to integrate a positive definite kernel $k(x, y)$ on the feature space \mathcal{X} . On a Euclidean feature space $\mathcal{X} \subset \mathbb{R}^D$, we typically use RBF kernels such as the Gaussian kernel $k(x, y) = \exp(-\|x - y\|^2/2\sigma^2)$ or the energy distance (conditionally positive) kernel $k(x, y) = -\|x - y\|$. The kernel loss is then defined, for $\xi = \alpha - \beta$, as

$$\text{L}_k(\alpha, \beta) \stackrel{\text{def.}}{=} \frac{1}{2} \|\xi\|_k^2 \stackrel{\text{def.}}{=} \frac{1}{2} \int_{\mathcal{X}^2} k(x, y) \, d\xi(x) d\xi(y). \quad (2)$$

If k is universal (Micchelli et al., 2006) (i.e. if the linear space spanned by functions $k(x, \cdot)$ is dense in $\mathcal{C}(\mathcal{X})$) we know that $\|\cdot\|_k$ metrizes the convergence in law. Such Euclidean norms, introduced for shape matching in (Glaunes et al., 2004), are often referred to as “Maximum Mean Discrepancies” (MMD) (Gretton et al., 2007). They have been extensively used for generative model (GANs) fitting in machine learning (Li et al., 2015; Dziugaite et al., 2015). MMD norms are cheaper to compute than OT and have a smaller *sample complexity* – i.e. approximation error when sampling a distribution.

1.2 Interpolating between OT and MMD using Sinkhorn divergences

Unfortunately though, the “flat” geometry that MMDs induce on the space of probability measures $\mathcal{M}_+^1(\mathcal{X})$ does not faithfully lift the ground distance on \mathcal{X} . For instance, on $\mathcal{X} = \mathbb{R}^D$, let us denote by α_τ the translation of α by $\tau \in \mathbb{R}^D$, defined through $\langle \alpha_\tau, f \rangle = \langle \alpha, f(\cdot + \tau) \rangle$ for continuous functions $f \in \mathcal{C}(\mathbb{R}^D)$. Wasserstein distance discrepancies defined for $C(x, y) = \|x - y\|^p$ are such that $\text{OT}_0(\alpha, \alpha_\tau)^{\frac{1}{p}} = \|\tau\|$.

In sharp contrast, MMD norms rely on *convolutions* that weigh the frequencies of $\xi = \alpha - \beta$ independently, according to the Fourier transform of the kernel function. For instance, up to a multiplicative constant of D , the energy distance $\text{L}_{-\|\cdot\|}(\alpha, \alpha_\tau)$ is given by $\int_{\mathbb{R}^D} (1 - \cos(\langle \omega, \tau \rangle)) |\hat{\alpha}(\omega)|^2 / \|\omega\|^{D+1} d\omega$ (Székely and Rizzo (2004), Lemma 1), where $\hat{\alpha}(\omega)$ is the Fourier transform of the probability measure α . Except for the trivial case of a Dirac mass $\alpha = \delta_{x_0}$ for some $x_0 \in \mathbb{R}^D$, we thus always have $\text{L}_{-\|\cdot\|}(\alpha, \alpha_\tau) < \|\tau\|$ with a value that strongly depends on the *smoothness* of the reference measure α . In practice, as evidenced in Figure 5, this theoretical shortcoming of MMD losses is reflected by *vanishing gradients* (or similar artifacts) next to the extreme points of the measures’ supports.

Sinkhorn divergences. On the one hand, OT losses have appealing *geometric* properties; on the other hand, cheap MMD norms scales up to large batches with a low sample complexity. Why not *interpolate* between them to get the best of both worlds?

Following (Genevay et al., 2018) (see also (Ramdas et al., 2017; Salimans et al., 2018; Sanjabi et al., 2018)) we consider a new cost built from OT_ε that we call a *Sinkhorn divergence*:

$$\text{S}_\varepsilon(\alpha, \beta) \stackrel{\text{def.}}{=} \text{OT}_\varepsilon(\alpha, \beta) - \frac{1}{2} \text{OT}_\varepsilon(\alpha, \alpha) - \frac{1}{2} \text{OT}_\varepsilon(\beta, \beta). \quad (3)$$

Such a formula satisfies $\text{S}_\varepsilon(\beta, \beta) = 0$ and interpolates between OT and MMD (Ramdas et al., 2017):

$$\text{OT}_0(\alpha, \beta) \stackrel{0 \leftarrow \varepsilon}{=} \text{S}_\varepsilon(\alpha, \beta) \stackrel{\varepsilon \rightarrow +\infty}{=} \frac{1}{2} \|\alpha - \beta\|_{-C}^2. \quad (4)$$

The entropic bias. Why bother with the auto-correlation terms $\text{OT}_\varepsilon(\alpha, \alpha)$ and $\text{OT}_\varepsilon(\beta, \beta)$? For positive values of ε , in general, $\text{OT}_\varepsilon(\beta, \beta) \neq 0$ so that minimizing $\text{OT}_\varepsilon(\alpha, \beta)$ with respect to α results in a biased solution: as evidenced by Figure 1, the gradient of OT_ε drives α towards a shrunk measure whose support is *smaller* than that of the target measure β . This is most evident as ε tends to infinity: $\text{OT}_\varepsilon(\alpha, \beta) \rightarrow \iint C(x, y) d\alpha(x) d\beta(y)$, a quantity that is minimized if α is a Dirac distribution located at the median (*resp.* the mean) value of β if $C(x, y) = \|x - y\|$ (*resp.* $\|x - y\|^2$).

In the literature, the formula (3) has been introduced more or less empirically to fix the *entropic bias* present in the OT_ε cost: with a structure that mimicks that of a squared kernel norm (2), it was assumed or conjectured that S_ε would define a positive definite loss function, suitable for applications in ML. This paper is all about *proving* that this is indeed what happens.

1.3 Contributions

The purpose of this paper is to show that the Sinkhorn divergences are convex, smooth, positive definite loss functions that metrize the convergence in law. Our main result is the theorem below, that ensures that one can indeed use S_ε as a *reliable* loss function for ML applications – whichever value of ε we pick.

Theorem 1. *Let \mathcal{X} be a compact metric space with a Lipschitz cost function $C(x, y)$ that induces, for $\varepsilon > 0$, a positive universal kernel $k_\varepsilon(x, y) \stackrel{\text{def.}}{=} \exp(-C(x, y)/\varepsilon)$. Then, S_ε defines a symmetric positive definite, smooth loss function that is convex in each of its input variables. It also metrizes the convergence in law: for all probability Radon measures α and $\beta \in \mathcal{M}_1^+(\mathcal{X})$,*

$$0 = S_\varepsilon(\beta, \beta) \leq S_\varepsilon(\alpha, \beta), \quad (5)$$

$$\alpha = \beta \iff S_\varepsilon(\alpha, \beta) = 0, \quad (6)$$

$$\alpha_n \rightarrow \alpha \iff S_\varepsilon(\alpha_n, \alpha) \rightarrow 0. \quad (7)$$

Notably, these results also hold for measures with bounded support on a Euclidean space $\mathcal{X} = \mathbb{R}^D$ endowed with ground cost functions $C(x, y) = \|x - y\|$ or $C(x, y) = \|x - y\|^2$ – which induce Laplacian and Gaussian kernels respectively.

This theorem legitimizes the use of the *unbiased* Sinkhorn divergences S_ε instead of OT_ε in model-fitting applications. Indeed, computing S_ε is roughly as expensive as OT_ε (the computation of the corrective factors being cheap, as detailed in Section 3) and the “debiasing” formula (3) allows us to guarantee that the unique minimizer of $\alpha \mapsto S_\varepsilon(\alpha, \beta)$ is the target distribution β (see Figure 1). Section 3 details how

to implement these divergences efficiently: our algorithms scale up to millions of samples thanks to freely available GPU routines. To conclude, we showcase in Section 4 the typical behavior of S_ε compared with OT_ε and standard MMD losses.

2 Proof of Theorem 1

We now give the proof of Theorem 1. Our argument relies on a new Bregman divergence derived from a weak* continuous entropy that we call the *Sinkhorn entropy* (see Section 2.2). We believe this (convex) entropy function to be of independent interest. Note that all this section is written under the assumptions of Theorem 1; the proof of some intermediate results can be found in the appendix.

2.1 Properties of the OT_ε loss

First, let us recall some standard results of regularized OT theory (Peyr e and Cuturi, 2017). Thanks to the Fenchel-Rockafellar theorem, we can rewrite Cuturi’s loss (1) as

$$\text{OT}_\varepsilon(\alpha, \beta) \stackrel{\text{def.}}{=} \max_{(f, g) \in \mathcal{C}(\mathcal{X})^2} \langle \alpha, f \rangle + \langle \beta, g \rangle - \varepsilon \langle \alpha \otimes \beta, \exp(\frac{1}{\varepsilon}(f \oplus g - C)) - 1 \rangle, \quad (8)$$

where $f \oplus g$ is the tensor sum $(x, y) \in \mathcal{X}^2 \mapsto f(x) + g(y)$. The primal-dual relationship linking an optimal transport plan π solving (1) to an optimal dual pair (f, g) that solves (8) is

$$\pi = \exp(\frac{1}{\varepsilon}(f \oplus g - C)) \cdot (\alpha \otimes \beta). \quad (9)$$

Crucially, the first order optimality conditions for the dual variables are equivalent to the primal’s marginal constraints ($\pi_1 = \alpha, \pi_2 = \beta$) on (9). They read

$$f = T(\beta, g) \quad \alpha\text{-a.e.} \quad \text{and} \quad g = T(\alpha, f) \quad \beta\text{-a.e.}, \quad (10)$$




Figure 1 – **Removing the entropic bias.** Solution α (in red) of the fitting problem $\min_\alpha L(\alpha, \beta)$ for some β shown in blue. Here, $C(x, y) = \|x - y\|$ on the unit square \mathcal{X} in \mathbb{R}^2 and $\varepsilon = .1$. The positions of the red dots were optimized by gradient descent, starting from a normal Gaussian sample.

where the ‘‘Sinkhorn mapping’’ $T : \mathcal{M}_1^+(\mathcal{X}) \times \mathcal{C}(\mathcal{X}) \rightarrow \mathcal{C}(\mathcal{X})$ is defined through

$$T : (\alpha, f) \mapsto \left(y \in \mathcal{X} \mapsto \min_{x \sim \alpha} [C(x, y) - f(x)] \right), \quad (11)$$

with a SoftMin operator of strength ε defined through

$$\min_{x \sim \alpha} \varphi(x) \stackrel{\text{def.}}{=} -\varepsilon \log \int_{\mathcal{X}} \exp \left(-\frac{1}{\varepsilon} \varphi(x) \right) d\alpha(x). \quad (12)$$

Dual potentials. The following proposition recalls some important properties of OT_ε and the associated dual potentials. Its proof can be found in Section B.1.

Proposition 1 (Properties of OT_ε). *The optimal potentials (f, g) exist and are unique (α, β) -a.e. up to an additive constant, i.e. $\forall K \in \mathbb{R}, (f + K, g - K)$ is also optimal. At optimality, we get*

$$\text{OT}_\varepsilon(\alpha, \beta) = \langle \alpha, f \rangle + \langle \beta, g \rangle. \quad (13)$$

We recall that a function $F : \mathcal{M}_1^+(\mathcal{X}) \rightarrow \mathbb{R}$ is said to be *differentiable* if there exists $\nabla F(\alpha) \in \mathcal{C}(\mathcal{X})$ such that for any displacement $\xi = \beta - \beta'$ with $(\beta, \beta') \in \mathcal{M}_1^+(\mathcal{X})^2$, we have

$$F(\alpha + t\xi) = F(\alpha) + t\langle \xi, \nabla F(\alpha) \rangle + o(t). \quad (14)$$

The following proposition, whose proof is detailed in Section B.2, shows that the dual potentials are the gradients of OT_ε .

Proposition 2. *OT_ε is weak* continuous and differentiable. Its gradient reads*

$$\nabla \text{OT}_\varepsilon(\alpha, \beta) = (f, g) \quad (15)$$

where (f, g) satisfies $f = T(\beta, g)$ and $g = T(\alpha, f)$ on the whole domain \mathcal{X} and T is the Sinkhorn mapping (11).

Let us stress that even though the solutions of the dual problem (8) are defined (α, β) -a.e., the gradient (15) is defined on the whole domain \mathcal{X} . Fortunately, an optimal dual pair (f_0, g_0) defined (α, β) -a.e. satisfies the optimality condition (10) and can be *extended* in a canonical way: to compute the ‘‘gradient’’ pair $(f, g) \in \mathcal{C}(\mathcal{X})^2$ associated to a pair of measures (α, β) , using $f = T(\beta, g_0)$ and $g = T(\alpha, f_0)$ is enough.

2.2 Sinkhorn and Hausdorff divergences

Having recalled some standard properties of OT_ε , let us now state a few *original* facts about the corrective, symmetric term $-\frac{1}{2}\text{OT}_\varepsilon(\alpha, \alpha)$ used in (3). We still suppose that (\mathcal{X}, d) is a compact set endowed with a symmetric, Lipschitz cost function $C(x, y)$. For $\varepsilon > 0$, the associated *Gibbs kernel* is defined through

$$k_\varepsilon : (x, y) \in \mathcal{X} \times \mathcal{X} \mapsto \exp \left(-C(x, y)/\varepsilon \right). \quad (16)$$

Crucially, we now assume that k_ε is a *positive universal kernel* on the space of signed Radon measures.

Definition 1 (Sinkhorn negentropy). *Under the assumptions above, we define the Sinkhorn negentropy of a probability Radon measure $\alpha \in \mathcal{M}_1^+(\mathcal{X})$ through*

$$F_\varepsilon(\alpha) \stackrel{\text{def.}}{=} -\frac{1}{2}\text{OT}_\varepsilon(\alpha, \alpha). \quad (17)$$

The following proposition is the cornerstone of our approach to prove the positivity of S_ε , providing an alternative expression of F_ε . Its proof relies on a change of variables $\mu = \exp(f/\varepsilon)\alpha$ in (8) that is detailed in the Section B.3 of the appendix.

Proposition 3. *Let (\mathcal{X}, d) be a compact set endowed with a symmetric, Lipschitz cost function $C(x, y)$ that induces a positive kernel k_ε . Then, for $\varepsilon > 0$ and $\alpha \in \mathcal{M}_1^+(\mathcal{X})$, one has*

$$\frac{1}{\varepsilon}F_\varepsilon(\alpha) + \frac{1}{2} = \min_{\mu \in \mathcal{M}^+(\mathcal{X})} \langle \alpha, \log \frac{d\alpha}{d\mu} \rangle + \frac{1}{2}\|\mu\|_{k_\varepsilon}^2. \quad (18)$$

The following proposition, whose proof can be found in the Section B.4 of the appendix, leverages the alternative expression (18) to ensure the convexity of F_ε .

Proposition 4. *Under the same hypotheses as Proposition 3, F_ε is a strictly convex functional on $\mathcal{M}_1^+(\mathcal{X})$.*

We now define an auxiliary ‘‘Hausdorff’’ divergence that can be interpreted as an OT_ε loss with *decoupled* dual potentials.

Definition 2 (Hausdorff divergence). *Thanks to Proposition 2, the Sinkhorn negentropy F_ε is differentiable in the sense of (14). For any probability measures $\alpha, \beta \in \mathcal{M}_1^+(\mathcal{X})$ and regularization strength $\varepsilon > 0$, we can thus define*

$$H_\varepsilon(\alpha, \beta) \stackrel{\text{def.}}{=} \frac{1}{2}\langle \alpha - \beta, \nabla F_\varepsilon(\alpha) - \nabla F_\varepsilon(\beta) \rangle \geq 0.$$

It is the symmetric Bregman divergence induced by the strictly convex functional F_ε (Bregman, 1967) and is therefore a positive definite quantity.

2.3 Proof of the Theorem

We are now ready to conclude. First, remark that the dual expression (8) of $\text{OT}_\varepsilon(\alpha, \beta)$ as a maximization of linear forms ensures that $\text{OT}_\varepsilon(\alpha, \beta)$ is *convex* with respect to α and with respect to β (but *not* jointly convex if $\varepsilon > 0$). S_ε is thus convex with respect to both inputs α and β as a sum of the functions OT_ε and F_ε – see Proposition 4.

Convexity also implies that,

$$\begin{aligned} \text{OT}_\varepsilon(\alpha, \alpha) + \langle \beta - \alpha, \nabla_2 \text{OT}_\varepsilon(\alpha, \alpha) \rangle &\leq \text{OT}_\varepsilon(\alpha, \beta), \\ \text{OT}_\varepsilon(\beta, \beta) + \langle \alpha - \beta, \nabla_1 \text{OT}_\varepsilon(\beta, \beta) \rangle &\leq \text{OT}_\varepsilon(\alpha, \beta). \end{aligned}$$

Using (15) to get $\nabla_2 \text{OT}_\varepsilon(\alpha, \alpha) = -\nabla F_\varepsilon(\alpha)$, $\nabla_1 \text{OT}_\varepsilon(\beta, \beta) = -\nabla F_\varepsilon(\beta)$ and summing the above inequalities, we show that $H_\varepsilon \leq S_\varepsilon$, which implies (5).

To prove (6), note that $S_\varepsilon(\alpha, \beta) = 0 \Rightarrow H_\varepsilon(\alpha, \beta) = 0$, which implies that $\alpha = \beta$ since F_ε is a *strictly* convex functional.

Finally, we show that S_ε metrizes the convergence in law (7) in the Section B.5 of the appendix.

3 Computational scheme

We have shown that Sinkhorn divergences (3) are positive definite, convex loss functions on the space of probability measures. Let us now detail their *implementation* on modern hardware.

Encoding measures. For the sake of simplicity, we focus on discrete, *sampled* measures on a Euclidean feature space $\mathcal{X} \subset \mathbb{R}^D$. Our input measures α and $\beta \in \mathcal{M}_1^+(\mathcal{X})$ are represented as sums of weighted Dirac atoms

$$\alpha = \sum_{i=1}^N \alpha_i \delta_{\mathbf{x}_i}, \quad \beta = \sum_{j=1}^M \beta_j \delta_{\mathbf{y}_j} \quad (19)$$

and encoded as two pairs (α, \mathbf{x}) and (β, \mathbf{y}) of float arrays. Here, $\alpha \in \mathbb{R}_+^N$ and $\beta \in \mathbb{R}_+^M$ are *non-negative* vectors of shapes $[N]$ and $[M]$ that sum up to 1, whereas $\mathbf{x} \in (\mathbb{R}^D)^N$ and $\mathbf{y} \in (\mathbb{R}^D)^M$ are real-valued tensors of shapes $[N, D]$ and $[M, D]$ – if we follow python’s convention.

3.1 The Sinkhorn algorithm(s)

Working with dual vectors. Proposition 1 is key to the modern theory of regularized Optimal Transport: it allows us to compute the OT_ε cost – and thus the Sinkhorn divergence S_ε , thanks to (3) – using dual variables that have the same *memory footprint* as the input measures: solving (8) in our discrete setting, we only need to store the *sampled values* of the dual potentials f and g on the measures’ supports.

We can thus work with *dual vectors* $\mathbf{f} \in \mathbb{R}^N$ and $\mathbf{g} \in \mathbb{R}^M$, defined through $\mathbf{f}_i = f(\mathbf{x}_i)$ and $\mathbf{g}_j = g(\mathbf{y}_j)$, which encode an *implicit* transport plan π from α to β (9). Crucially, the optimality condition (10) now reads:

$$\forall i \in [1, N], \forall j \in [1, M],$$

$$\mathbf{f}_i = -\varepsilon \text{LSE}_{k=1}^M \left(\log(\beta_k) + \frac{1}{\varepsilon} \mathbf{g}_k - \frac{1}{\varepsilon} C(\mathbf{x}_i, \mathbf{y}_k) \right) \quad (20)$$

$$\mathbf{g}_j = -\varepsilon \text{LSE}_{k=1}^N \left(\log(\alpha_k) + \frac{1}{\varepsilon} \mathbf{f}_k - \frac{1}{\varepsilon} C(\mathbf{x}_k, \mathbf{y}_j) \right) \quad (21)$$

$$\text{where } \text{LSE}_{k=1}^N(V_k) = \log \sum_{k=1}^N \exp(V_k) \quad (22)$$

denotes a (stabilized) log-sum-exp reduction.

If (\mathbf{f}, \mathbf{g}) is an optimal pair of dual vectors that satisfies

Equations (20-21), we deduce from (13) that

$$OT_\varepsilon(\alpha_i, \mathbf{x}_i, \beta_j, \mathbf{y}_j) = \sum_{i=1}^N \alpha_i \mathbf{f}_i + \sum_{j=1}^M \beta_j \mathbf{g}_j. \quad (23)$$

But how can we solve this *coupled* system of equations given $\alpha, \mathbf{x}, \beta$ and \mathbf{y} as input data?

The Sinkhorn algorithm. One simple answer: by enforcing (20) and (21) alternatively, updating the vectors \mathbf{f} and \mathbf{g} until convergence (Cuturi, 2013). Starting from null potentials $\mathbf{f}_i = 0 = \mathbf{g}_j$, this numerical scheme is nothing but a block-coordinate ascent on the dual problem (8). One step after another, we are enforcing null derivatives on the dual cost with respect to the \mathbf{f}_i ’s and the \mathbf{g}_j ’s.

Convergence. The ‘‘Sinkhorn loop’’ converges quickly towards its unique optimal value: it enjoys a linear convergence rate (Peyr e and Cuturi, 2017) that can be improved with some heuristics (Thibault et al., 2017). When computed through the *dual* expression (23), OT_ε and its gradients (26-27) are *robust* to small perturbations of the values of \mathbf{f} and \mathbf{g} : monitoring convergence through the L^1 norm of the updates on \mathbf{f} and breaking the loop as we reach a set tolerance level is thus a sensible stopping criterion. In practice, if ε is large enough – say, $\varepsilon \geq .05$ on the unit square with an Earth Mover’s cost $C(x, y) = \|x - y\|$ – waiting for 10 or 20 iterations is more than enough.

Symmetric OT_ε problems. All in all, the baseline Sinkhorn loop provides an efficient way of solving the discrete problem $OT_\varepsilon(\alpha, \beta)$ for generic input measures. But in the specific case of the (symmetric) corrective terms $OT_\varepsilon(\alpha, \alpha)$ and $OT_\varepsilon(\beta, \beta)$ introduced in (3), we can do better.

The key here is to remark that if $\alpha = \beta$, the dual problem (8) becomes a concave maximization problem that is *symmetric* with respect to its two variables f and g . Hence, there exists a (unique) optimal dual pair $(f, g = f)$ on the diagonal which is characterized in the discrete setting by the symmetric optimality condition:

$$\forall i \in [1, N],$$

$$\mathbf{f}_i = -\varepsilon \text{LSE}_{k=1}^N \left[\log(\alpha_k) + \frac{1}{\varepsilon} \mathbf{f}_k - \frac{1}{\varepsilon} C(\mathbf{x}_i, \mathbf{x}_k) \right]. \quad (24)$$

Fortunately, given α and \mathbf{x} , the optimal vector \mathbf{f} that solves this equation can be computed by iterating a *well-conditioned* fixed-point update:

$$\mathbf{f}_i \leftarrow \frac{1}{2} \left(\mathbf{f}_i - \varepsilon \text{LSE}_{k=1}^N \left[\log(\alpha_k) + \frac{1}{\varepsilon} \mathbf{f}_k - \frac{1}{\varepsilon} C(\mathbf{x}_i, \mathbf{x}_k) \right] \right). \quad (25)$$

This symmetric variant of the Sinkhorn algorithm can be shown to converge much faster than the standard

loop applied to a pair $(\alpha, \beta = \alpha)$ on the diagonal, and three iterations are usually enough to compute accurately the optimal dual vector.

3.2 Computing the Sinkhorn divergence and its gradients

Given two pairs (α, \mathbf{x}) and (β, \mathbf{y}) of float arrays that encode the probability measures α and β (19), we can now implement the Sinkhorn divergence $S_\varepsilon(\alpha, \beta)$:

The *cross-correlation* dual vectors $\mathbf{f} \in \mathbb{R}^N$ and $\mathbf{g} \in \mathbb{R}^M$ associated to the discrete problem $\text{OT}_\varepsilon(\alpha, \beta)$ can be computed using the Sinkhorn iterations (20-21).

The *autocorrelation* dual vectors $\mathbf{p} \in \mathbb{R}^N$ and $\mathbf{q} \in \mathbb{R}^M$, respectively associated to the symmetric problems $\text{OT}_\varepsilon(\alpha, \alpha)$ and $\text{OT}_\varepsilon(\beta, \beta)$, can be computed using the *symmetric* Sinkhorn update (25).

The Sinkhorn *loss* can be computed using (3) and (23):

$$S_\varepsilon(\alpha_i, \mathbf{x}_i, \beta_j, \mathbf{y}_j) = \sum_{i=1}^N \alpha_i(\mathbf{f}_i - \mathbf{p}_i) + \sum_{j=1}^M \beta_j(\mathbf{g}_j - \mathbf{q}_j).$$

What about the gradients? In this day and age, we could be tempted to rely on the *automatic differentiation* engines provided by modern libraries, which let us differentiate the result of twenty or so Sinkhorn iterations as a mere composition of elementary operations (Genevay et al., 2018). But beware: this loop has a lot more *structure* than a generic feed forward network. Taking advantage of it is key to a x2-x3 gain in performances, as we now describe.

Crucially, we must remember that the Sinkhorn loop is a *fixed point* iterative solver: at convergence, its solution satisfies an equation given by the implicit function theorem. Thanks to (15), using the very definition of gradients in the space of probability measures (14) and the intermediate variables in the computation of $S_\varepsilon(\alpha, \beta)$, we get that

$$\partial_{\alpha_i} S_\varepsilon(\alpha_i, \mathbf{x}_i, \beta_j, \mathbf{y}_j) = \mathbf{f}_i - \mathbf{p}_i \quad (26)$$

$$\text{and} \quad \partial_{\mathbf{x}_i} S_\varepsilon(\alpha_i, \mathbf{x}_i, \beta_j, \mathbf{y}_j) = \nabla \varphi(\mathbf{x}_i), \quad (27)$$

where $\varphi : \mathcal{X} \rightarrow \mathbb{R}$ is equal to $\mathbf{f}_i - \mathbf{p}_i$ on the \mathbf{x}_i 's and is defined through

$$\begin{aligned} \varphi(x) = & -\varepsilon \log \sum_{j=1}^M \exp \left[\log(\beta_j) + \frac{1}{\varepsilon} \mathbf{g}_j - \frac{1}{\varepsilon} C(x, \mathbf{y}_j) \right] \\ & + \varepsilon \log \sum_{i=1}^N \exp \left[\log(\alpha_i) + \frac{1}{\varepsilon} \mathbf{p}_i - \frac{1}{\varepsilon} C(x, \mathbf{x}_i) \right]. \end{aligned}$$

Graph surgery with PyTorch. Assuming convergence in the Sinkhorn loops, it is thus possible to compute the gradients of S_ε *without having to backprop*

through the twenty or so iterations of the Sinkhorn algorithm: we only have to differentiate the expression above with respect to x . But does it mean that we should differentiate C or the log-sum-exp operation by hand? Fortunately, no!

Modern libraries such as PyTorch (Paszke et al., 2017) are flexible enough to let us “hack” the naive `autograd` algorithm, and act as though the optimal dual vectors $\mathbf{f}_i, \mathbf{p}_i, \mathbf{g}_j$ and \mathbf{q}_j did not depend on the input variables of S_ε . As documented in our reference code,

`github.com/jeanfeydy/global-divergences`,

an appropriate use of the `.detach()` method in PyTorch is enough to get the best of both worlds: an *automatic* differentiation engine that computes our gradients using the formula *at convergence* instead of the baseline backpropagation algorithm. All in all, as evidenced by the benchmarks provided Figure 3, this trick allows us to divide by a factor 2-3 the time needed to compute a Sinkhorn divergence and its gradient with respect to the \mathbf{x}_i 's.

3.3 Scaling up to large datasets

The Sinkhorn iterations rely on a single non-trivial operation: the log-sum-exp reduction (22). In the ML literature, this *SoftMax* operator is often understood as a row- or column-wise reduction that acts on $[N, M]$ matrices. But as we strive to implement the update rules (20-21) and (25) on the GPU, we can go further.

Batch computation. First, if the number of samples N and M in both measures is small enough, we can optimize the GPU usage by computing Sinkhorn divergences *by batches* of size B . In practice, this can be achieved by encoding the cost function C as a 3D tensor of size $[B, N, M]$ made up of stacked matrices $(C(\mathbf{x}_i, \mathbf{y}_j))_{i,j}$, while \mathbf{f} and \mathbf{g} become $[B, N]$ and $[B, M]$ tensors, respectively. Thanks to the broadcasting syntax supported by modern libraries, we can then seamlessly compute, in parallel, loss values $S_\varepsilon(\alpha_k, \beta_k)$ for k in $[1, B]$.

The KeOps library. Unfortunately though, tensor-centric methods such as the one presented above cannot scale to measures sampled with large numbers N and M of Dirac atoms: as these numbers exceed 10,000, huge $[N, M]$ matrices stop fitting into GPU memories. To alleviate this problem, we leveraged the KeOps library (Charlier et al., 2018) that provides *online* map-reduce routines on the GPU with full PyTorch integration. Performing online log-sum-exp reductions with a *running maximum*, the KeOps primitives allow us to compute Sinkhorn divergences with a *linear* memory footprint. As evidenced by the




Figure 2 – **The KeOps library allows us to break the memory bottleneck.** Using CUDA routines that sum kernel values without storing them in memory, we can outperform baseline, tensorized, implementations of the energy distance. Experiments performed on $\mathcal{X} = \mathbb{R}^3$ with a cheap laptop’s GPU (GTX 960M).




Figure 3 – **Sinkhorn divergences scale up to finely sampled distributions.** As a rule of thumb, Sinkhorn divergences take 20-50 times as long to compute as a baseline MMD – even though the explicit gradient formula (26-27) lets us win a factor 2-3 compared with a *naive* autograd implementation. For the sake of this benchmark, we ran the Sinkhorn and symmetric Sinkhorn loops with fixed numbers of iterations: 20 and 3 respectively, which is more than enough for measures on the unit (hyper)cube if $\varepsilon \geq .05$ – here, we work in \mathbb{R}^3 .

benchmarks of Figures 2-3, computing the gradient of a Sinkhorn loss with 100,000 samples per measure is then a matter of seconds.

4 Numerical illustration

In the previous sections, we have provided theoretical guarantees on top of a comprehensive implementation guide for the family of *Sinkhorn divergences* S_ε . Let us now describe the *geometry* induced by these new loss functions on the space of probability measures.

Gradient flows. To compare MMD losses L_k with Cuturi’s original cost OT_ε and the de-biased Sinkhorn divergence S_ε , a simple yet relevant experiment is to let a *model* distribution $\alpha(t)$ flow with time t along the “Wasserstein-2” gradient flow of a loss functional $\alpha \mapsto L(\alpha, \beta)$ that drives it towards a target distribution β (Santambrogio, 2015). This corresponds to the “non-parametric” version of the data fitting problem evoked in Section 1, where the parameter θ is nothing but the vector of positions \mathbf{x} that encodes the support of a measure $\alpha = \frac{1}{N} \sum_{i=1}^N \delta_{\mathbf{x}_i}$. Understood as a “model free” idealization of fitting problems in machine learning, this experiment allows us to grasp the typical behavior of the loss function as we discover the deformations of the support that it favors.




Figure 4 – Gradient flows for 1-D measures sampled with $N = M = 5000$ points – we display $\alpha(t)$ (in red) and β (in blue) through kernel density estimations on the segment $[0, 1]$. The legend on the left indicates the function that is minimized with respect to α . Here $k(x, y) = -\|x - y\|$, $C(x, y) = \|x - y\|$ and $\varepsilon = .10$ on the second and fourth lines, $\varepsilon = .01$ on the third. In 1D, the optimal transport problem can be solved using a sort algorithm: for the sake of comparison, we can thus display the “true” dynamics of the Earth Mover’s Distance in the fifth line.




Figure 5 – **Gradient flows for 2-D measures.** The setting is the same as in Figure 4, but the measures $\alpha(t)$ (in red) and β (in blue) are now directly displayed as point clouds of $N = M = 500$ points. The evolution of the support of $\alpha(t)$ is thus made apparent, and we display as a green vector field the descent direction $-\nabla_{\mathbf{x}_i} L(\alpha, \beta)$.

In Figures 4 and 5, $\beta = \frac{1}{M} \sum_{j=1}^M \delta_{\mathbf{y}_j}$ is a fixed target measure while $\alpha = \frac{1}{N} \sum_{i=1}^N \delta_{\mathbf{x}_i(t)}$ is parameterized by a time-varying point cloud $\mathbf{x}(t) = (\mathbf{x}_i(t))_{i=1}^N \in (\mathbb{R}^D)^N$ in dimension $D = 1$ or 2 . Starting from a set initial condition at time $t = 0$, we simply integrate the ODE

$$\dot{\mathbf{x}}(t) = -N \nabla_{\mathbf{x}} \left[L \left(\frac{1}{N} \sum_{i=1}^N \delta_{\mathbf{x}_i}, \beta \right) \right] (\mathbf{x}_i(t))$$

with a Euler scheme and display the evolution of $\alpha(t)$ up to time $t = 5$.

Interpretation. In both figures, the fourth line highlights the entropic bias that is present in the OT_ϵ loss: $\alpha(t)$ is driven towards a minimizer that is a “shrunk” version of β . As showed in Theorem 1, the de-biased loss S_ϵ does not suffer from this issue: just like MMD norms, it can be used as a *reliable*, positive-definite divergence.

Going further, the dynamics induced by the Sinkhorn divergence interpolates between that of an MMD ($\epsilon = +\infty$) and Optimal Transport ($\epsilon = 0$), as shown in (4). Here, $C(x, y) = \|x - y\|$ and we can indeed remark that the second and third lines bridge the gap between the flow of the energy distance $L_{-\|\cdot\|}$ (in the first line) and that of the Earth Mover’s cost OT_0 which moves

particles according to an optimal transport plan.

Please note that in both experiments, the gradient of the energy distance with respect to the \mathbf{x}_i ’s vanishes at the extreme points of α ’s support. Crucially, for small enough values of ϵ , S_ϵ recovers the translation-aware geometry of OT and we observe a *clean* convergence of $\alpha(t)$ to β as no sample lags behind.

5 Conclusion

Recently introduced in the ML literature, the Sinkhorn divergences were designed to interpolate between MMD and OT. We have now shown that they also come with a bunch of desirable properties: positivity, convexity, metrization of the convergence in law and scalability to large datasets.

To the best of our knowledge, it is the first time that a loss derived from the theory of entropic Optimal Transport is shown to stand on such a firm ground. As the foundations of this theory are progressively being settled, we now hope that researchers will be free to focus on one of the major open problems in the field: the interaction of *geometric* loss functions with concrete machine learning models.

Bibliography

- Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein GAN. *arXiv preprint arXiv:1701.07875*.
- Bassetti, F., Bodini, A., and Regazzini, E. (2006). On minimum Kantorovich distance estimators. *Statistics & probability letters*, 76(12):1298–1302.
- Bonneel, N., Peyr e, G., and Cuturi, M. (2016). Wasserstein barycentric coordinates: Histogram regression using optimal transport. *ACM Transactions on Graphics*, 35(4).
- Bregman, L. M. (1967). The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics*, 7(3):200–217.
- Charlier, B., Feydy, J., and Glaun es, J. (2018). Kernel operations on the gpu, with autodiff, without memory overflows. <http://www.kernel-operations.io>. Accessed: 2018-10-04.
- Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. In *Adv. in Neural Information Processing Systems*, pages 2292–2300.
- Dziugaite, G. K., Roy, D. M., and Ghahramani, Z. (2015). Training generative neural networks via maximum mean discrepancy optimization. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pages 258–267.
- Franklin, J. and Lorenz, J. (1989). On the scaling of multidimensional matrices. *Linear Algebra and its applications*, 114:717–735.
- Frogner, C., Zhang, C., Mobahi, H., Araya, M., and Poggio, T. A. (2015). Learning with a Wasserstein loss. In *Advances in Neural Information Processing Systems*, pages 2053–2061.
- Galichon, A. and Salani e, B. (2010). Matching with trade-offs: Revealed preferences over competing characteristics. *Preprint hal-00473173*.
- Genevay, A., Peyr e, G., and Cuturi, M. (2018). Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617.
- Glaun es, J., Trouv e, A., and Younes, L. (2004). Diffeomorphic matching of distributions: A new approach for unlabelled point-sets and sub-manifolds matching. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, volume 2, pages II–II. Ieee.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680.
- Gretton, A., Borgwardt, K. M., Rasch, M., Sch olkopf, B., and Smola, A. J. (2007). A kernel method for the two-sample-problem. In *Advances in neural information processing systems*, pages 513–520.
- Kaltenmark, I., Charlier, B., and Charon, N. (2017). A general framework for curve and surface comparison and registration with oriented varifolds. In *Computer Vision and Pattern Recognition (CVPR)*.
- Kantorovich, L. (1942). On the transfer of masses (in Russian). *Doklady Akademii Nauk*, 37(2):227–229.
- L eonard, C. (2013). A survey of the Schr odinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*.
- Li, Y., Swersky, K., and Zemel, R. (2015). Generative moment matching networks. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 1718–1727.
- Micchelli, C. A., Xu, Y., and Zhang, H. (2006). Universal kernels. *Journal of Machine Learning Research*, 7(Dec):2651–2667.
- Montavon, G., M uller, K.-R., and Cuturi, M. (2016). Wasserstein training of restricted boltzmann machines. In *Advances in Neural Information Processing Systems*, pages 3718–3726.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. (2017). Automatic differentiation in pytorch.
- Peyr e, G. and Cuturi, M. (2017). Computational optimal transport. *arXiv:1610.06519*.
- Ramdas, A., Trillos, N. G., and Cuturi, M. (2017). On wasserstein two-sample testing and related families of nonparametric tests. *Entropy*, 19(2).
- Rubner, Y., Tomasi, C., and Guibas, L. J. (2000). The earth mover’s distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99–121.
- Salimans, T., Zhang, H., Radford, A., and Metaxas, D. (2018). Improving GANs using optimal transport. *arXiv preprint arXiv:1803.05573*.
- Sanjabi, M., Ba, J., Razaviyayn, M., and Lee, J. D. (2018). On the convergence and robustness of training GANs with regularized optimal transport. *arXiv preprint arXiv:1802.08249*.
- Santambrogio, F. (2015). *Optimal Transport for applied mathematicians*, volume 87 of *Progress in Nonlinear Differential Equations and their applications*. Springer.

- Székely, G. J. and Rizzo, M. L. (2004). Hierarchical clustering via joint between-within distances: Extending ward's minimum variance method. *J. Classification*, 22:151–183.
- Thibault, A., Chizat, L., Dossal, C., and Papadakis, N. (2017). Overrelaxed sinkhorn-knopp algorithm for regularized optimal transport. *arXiv preprint arXiv:1711.01851*.
- Vaillant, M. and Glaunès, J. (2005). Surface matching via currents. In *Biennial International Conference on Information Processing in Medical Imaging*, pages 381–392. Springer.

A Standard results

Before detailing our proofs, we first recall some well-known results regarding the Kullback-Leibler divergence and the SoftMin operator defined in (12).

A.1 The Kullback-Leibler divergence

First properties. For any pair of Radon measures $\alpha, \beta \in \mathcal{M}^+(\mathcal{X})$ on the compact metric set (\mathcal{X}, d) , the Kullback-Leibler divergence is defined through

$$\text{KL}(\alpha, \beta) \stackrel{\text{def.}}{=} \begin{cases} \langle \alpha, \log \frac{d\alpha}{d\beta} - 1 \rangle + \langle \beta, 1 \rangle & \text{if } \alpha \ll \beta \\ +\infty & \text{otherwise.} \end{cases}$$

It can be rewritten as an f -divergence associated to

$$\psi : x \in \mathbb{R}_{\geq 0} \mapsto x \log(x) - x + 1 \in \mathbb{R}_{\geq 0},$$

with $0 \cdot \log(0) = 0$, as

$$\text{KL}(\alpha, \beta) = \begin{cases} \langle \beta, \psi(\frac{d\alpha}{d\beta}) \rangle & \text{if } \alpha \ll \beta \\ +\infty & \text{otherwise.} \end{cases} \quad (28)$$

Since ψ is a strictly convex function with a unique global minimum at $\psi(1) = 0$, we thus get that $\text{KL}(\alpha, \beta) \geq 0$ with equality iff. $\alpha = \beta$.

Dual formulation. The convex conjugate of ψ is defined for $u \in \mathbb{R}$ by

$$\begin{aligned} \psi^*(u) &\stackrel{\text{def.}}{=} \sup_{x>0} (xu - \psi(x)) \\ &= e^u - 1, \end{aligned}$$

$$\text{and we have } \psi(x) + \psi^*(u) \geq xu \quad (29)$$

for all $(x, u) \in \mathbb{R}_{\geq 0} \times \mathbb{R}$, with equality if $x > 0$ and $u = \log(x)$. This allows us to rewrite the Kullback-Leibler divergence as the solution of a dual concave problem:

Proposition 5 (Dual formulation of KL). *Under the assumptions above,*

$$\text{KL}(\alpha, \beta) = \sup_{h \in \mathcal{F}_b(\mathcal{X}, \mathbb{R})} \langle \alpha, h \rangle - \langle \beta, e^h - 1 \rangle \quad (30)$$

where $\mathcal{F}_b(\mathcal{X}, \mathbb{R})$ is the space of bounded measurable functions from \mathcal{X} to \mathbb{R} .

Proof. Lower bound on the sup. If α is not absolutely continuous with respect to β , there exists a Borel set A such that $\alpha(A) > 0$ and $\beta(A) = 0$. Consequently, for $h = \lambda \mathbf{1}_A$,

$$\langle \alpha, h \rangle - \langle \beta, e^h - 1 \rangle = \lambda \alpha(A) \xrightarrow{\lambda \rightarrow +\infty} +\infty.$$

Otherwise, if $\alpha \ll \beta$, we define $h_* = \log \frac{d\alpha}{d\beta}$ and see that

$$\langle \alpha, h_* \rangle - \langle \beta, e^{h_*} - 1 \rangle = \text{KL}(\alpha, \beta).$$

If $h_n = \log(\frac{d\alpha}{d\beta}) \mathbf{1}_{1/n \leq d\alpha/d\beta \leq n} \in \mathcal{F}_b(\mathcal{X}, \mathbb{R})$, the monotone and dominated convergence theorems then allow us to show that

$$\langle \alpha, h_n \rangle - \langle \beta, e^{h_n} - 1 \rangle \xrightarrow{n \rightarrow +\infty} \text{KL}(\alpha, \beta).$$

Upper bound on the sup. If $h \in \mathcal{F}_b(\mathcal{X}, \mathbb{R})$ and $\alpha \ll \beta$, combining (28) and (29) allow us to show that

$$\begin{aligned} &\text{KL}(\alpha, \beta) - \langle \alpha, h \rangle + \langle \beta, e^h - 1 \rangle \\ &= \langle \beta, \psi(\frac{d\alpha}{d\beta}) + \psi^*(h) - h \frac{d\alpha}{d\beta} \rangle \geq 0. \end{aligned}$$

The optimal value of $\langle \alpha, h \rangle - \langle \beta, e^h - 1 \rangle$ is bounded above and below by $\text{KL}(\alpha, \beta)$: we get (30). \square

Since $\langle \alpha, h \rangle - \langle \beta, e^h - 1 \rangle$ is a convex function of (α, β) , taking the supremum over test functions $h \in \mathcal{F}_b(\mathcal{X}, \mathbb{R})$ defines a *convex* divergence:

Proposition 6. *The KL divergence is a (jointly) convex function on $\mathcal{M}^+(\mathcal{X}) \times \mathcal{M}^+(\mathcal{X})$.*

Going further, the density of continuous functions in the space of bounded measurable functions allows us to restrict the optimization domain:

Proposition 7. *Under the same assumptions,*

$$\text{KL}(\alpha, \beta) = \sup_{h \in \mathcal{C}(\mathcal{X}, \mathbb{R})} \langle \alpha, h \rangle - \langle \beta, e^h - 1 \rangle \quad (31)$$

where $\mathcal{C}(\mathcal{X}, \mathbb{R})$ is the space of (bounded) continuous functions on the compact set \mathcal{X} .

Proof. Let $h = \sum_{i \in I} h_i \mathbf{1}_{A_i}$ be a simple Borel function on \mathcal{X} , and let us choose some error margin $\delta > 0$. Since α and β are Radon measures, for any i in the finite set of indices I , there exists a compact set K_i and an open set V_i such that $K_i \subset A_i \subset V_i$ and

$$\sum_{i \in I} \max[\alpha(V_i \setminus K_i), \beta(V_i \setminus K_i)] \leq \delta.$$

Moreover, for any $i \in I$, there exists a continuous function φ_i such that $\mathbf{1}_{K_i} \leq \varphi_i \leq \mathbf{1}_{V_i}$. The continuous function $g = \sum_{i \in I} h_i \varphi_i$ is then such that

$$|\langle \alpha, g - h \rangle| \leq \|h\|_\infty \delta \quad \text{and} \quad |\langle \beta, e^g - e^h \rangle| \leq \|e^h\|_\infty \delta$$

so that

$$\begin{aligned} &|(\langle \alpha, h \rangle - \langle \beta, e^h - 1 \rangle) - (\langle \alpha, g \rangle - \langle \beta, e^g - 1 \rangle)| \\ &\leq (\|h\|_\infty + \|e^h\|_\infty) \delta. \end{aligned}$$

As we let our simple function approach any measurable function in $\mathcal{F}_b(\mathcal{X}, \mathbb{R})$, choosing δ arbitrarily small, we then get (31) through (30). \square

We can then show that the Kullback-Leibler divergence is weakly lower semi-continuous:

Proposition 8. *If $\alpha_n \rightharpoonup \alpha$ and $\beta_n \rightharpoonup \beta$ are weakly converging sequences in $\mathcal{M}^+(\mathcal{X})$, we get*

$$\liminf_{n \rightarrow +\infty} \text{KL}(\alpha_n, \beta_n) \geq \text{KL}(\alpha, \beta).$$

Proof. According to (31), the KL divergence is defined as a pointwise supremum of weakly continuous applications

$$\varphi_h : (\alpha, \beta) \mapsto \langle \alpha, h \rangle - \langle \beta, e^h - 1 \rangle,$$

for $h \in \mathcal{C}(\mathcal{X}, \mathbb{R})$. It is thus lower semi-continuous for the convergence in law. \square

A.2 SoftMin Operator

Proposition 9 (The SoftMin interpolates between a minimum and a sum). *Under the assumptions of the definition (12), we get that*

$$\begin{aligned} \min_{x \sim \alpha} \varphi(x) &\xrightarrow{\varepsilon \rightarrow 0} \min_{x \in \text{Supp}(\alpha)} \varphi(x) \\ &\xrightarrow{\varepsilon \rightarrow +\infty} \langle \alpha, \varphi \rangle. \end{aligned}$$

If φ and ψ are two continuous functions in $\mathcal{C}(\mathcal{X})$ such that $\varphi \leq \psi$,

$$\min_{x \sim \alpha} \varphi(x) \leq \min_{x \sim \alpha} \psi(x). \quad (32)$$

Finally, if $K \in \mathbb{R}$ is constant with respect to x , we have that

$$\min_{x \sim \alpha} [K + \varphi(x)] = K + \min_{x \sim \alpha} [\varphi(x)]. \quad (33)$$

Proposition 10 (The SoftMin operator is continuous). *Let (α_n) be a sequence of probability measures converging weakly towards α , and (φ_n) be a sequence of continuous functions that converges uniformly towards φ . Then, for $\varepsilon > 0$, the SoftMin of the values of φ_n on α_n converges towards the SoftMin of the values of φ on α , i.e.*

$$\begin{aligned} (\alpha_n \rightharpoonup \alpha, \varphi_n \xrightarrow{\|\cdot\|_\infty} \varphi) \\ \implies \min_{x \sim \alpha_n} \varphi_n(x) \rightarrow \min_{x \sim \alpha} \varphi(x). \end{aligned}$$

B Proofs

B.1 Dual Potentials

We first state some important properties of solutions (f, g) to the dual problem (8). Please note that these results hold under the assumption that (\mathcal{X}, d) is a compact metric space, endowed with a *ground cost* function $C : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ that is κ -Lipschitz with respect to both of its input variables.

The existence of an optimal pair (f, g) of potentials that reaches the maximal value of the dual objective is proved using the contractance of the Sinkhorn map T , defined in (11), for the Hilbert projective metric (Franklin and Lorenz, 1989).

While optimal potentials are only defined (α, β) -a.e., as highlighted in Proposition 1, they are extended to the whole domain \mathcal{X} by imposing, similarly to the classical theory of OT (Santambrogio, 2015, Remark 1.13), that they satisfy

$$f = T(\beta, g) \quad \text{and} \quad g = T(\alpha, f), \quad (34)$$

with T defined in (11). We thus assume in the following that this condition holds. The following propositions studies the uniqueness and the smoothness (with respect to the spacial position and with respect to the input measures) of these functions (f, g) defined on the whole space.

Proposition 11 (Uniqueness of the dual potentials up to an additive constant). *Let (f_0, g_0) and (f_1, g_1) be two optimal pairs of dual potentials for a problem $\text{OT}_\varepsilon(\alpha, \beta)$ that satisfy (34). Then, there exists a constant $K \in \mathbb{R}$ such that*

$$f_0 = f_1 + K \quad \text{and} \quad g_0 = g_1 - K. \quad (35)$$

Proof. For $t \in [0, 1]$, let us define $f_t = f_0 + t(f_1 - f_0)$, $g_t = g_0 + t(g_1 - g_0)$ and

$$\begin{aligned} \varphi(t) &= \langle \alpha, f_t \rangle + \langle \beta, g_t \rangle \\ &\quad - \varepsilon \langle \alpha \otimes \beta, \exp\left(\frac{1}{\varepsilon}(f_t \oplus g_t - C)\right) - 1 \rangle, \end{aligned}$$

the value of the dual objective between the two optimal pairs. As φ is a concave function bounded above by $\varphi(0) = \varphi(1) = \text{OT}_\varepsilon(\alpha, \beta)$, it is *constant* with respect to t . Hence, for all t in $[0, 1]$,

$$\begin{aligned} 0 &= \varphi''(t) \\ &= -\frac{1}{\varepsilon} \langle \alpha \otimes \beta, e^{(f_t \oplus g_t - C)/\varepsilon} ((f_1 - f_0) \oplus (g_1 - g_0))^2 \rangle. \end{aligned}$$

This is only possible if, $\alpha \otimes \beta$ -a.e. in (x, y) ,

$$(f_1(x) - f_0(x) + g_1(y) - g_0(y))^2 = 0,$$

i.e. there exists a constant $K \in \mathbb{R}$ such that

$$\begin{aligned} f_1(x) - f_0(x) &= +K \quad \alpha\text{-a.e.} \\ g_1(y) - g_0(y) &= -K \quad \beta\text{-a.e.} \end{aligned}$$

As we extend the potentials through (34), the SoftMin operator commutes with the addition of K (33) and lets our result hold on the whole feature space. \square

Proposition 12 (Lipschitz property). *The optimal potentials (f, g) of the dual problem (8) are both κ -Lipschitz functions on the feature space (\mathcal{X}, d) .*

Proof. According to (34), f is a SoftMin combination of κ -Lipschitz functions of the variable x ; using the algebraic properties of the SoftMin operator detailed in (32-33), one can thus show that f is a κ -Lipschitz function on the feature space. The same argument holds for g . \square

Proposition 13 (The dual potentials vary continuously with the input measures). *Let $\alpha_n \rightharpoonup \alpha$ and $\beta_n \rightharpoonup \beta$ be weakly converging sequences of measures in $\mathcal{M}_1^+(\mathcal{X})$. Given some arbitrary anchor point $x_o \in \mathcal{X}$, let us denote by (f_n, g_n) the (unique) sequence of optimal potentials for $\text{OT}_\varepsilon(\alpha_n, \beta_n)$ such that $f_n(x_o) = 0$.*

Then, f_n and g_n converge uniformly towards the unique pair of optimal potentials (f, g) for $\text{OT}_\varepsilon(\alpha, \beta)$ such that $f(x_o) = 0$. Up to the value at the anchor point x_o , we thus have that

$$(\alpha_n \rightharpoonup \alpha, \beta_n \rightharpoonup \beta) \implies (f_n \xrightarrow{\|\cdot\|_\infty} f, g_n \xrightarrow{\|\cdot\|_\infty} g).$$

Proof. For all n in \mathbb{N} , the potentials f_n and g_n are κ -Lipschitz functions on the compact, bounded set \mathcal{X} . As $f_n(x_o)$ is set to zero, we can bound $|f_n|$ on \mathcal{X} by κ times the diameter of \mathcal{X} ; combining this with (34), we can then produce a uniform bound on both f_n and g_n : there exists a constant $M \in \mathbb{R}$ such that

$$\forall n \in \mathbb{N}, \forall x \in \mathcal{X}, -M \leq f_n(x), g_n(x) \leq +M.$$

Being equicontinuous and uniformly bounded on the compact set \mathcal{X} , the sequence $(f_n, g_n)_n$ satisfies the hypotheses of the Ascoli-Arzelà theorem: there exists a subsequence $(f_{n_k}, g_{n_k})_k$ that converges uniformly towards a pair (f, g) of continuous functions. k tend to infinity, we see that $f(x_o) = 0$ and, using the continuity of the SoftMin operator (Proposition 10) on the optimality equations (10), we show that (f, g) is an optimal pair for $\text{OT}_\varepsilon(\alpha, \beta)$.

Now, according to Proposition 11, such a limit pair of optimal potentials (f, g) is *unique*. $(f_n, g_n)_n$ is thus a *compact* sequence with a *single* possible adherence value: it has to converge, uniformly, towards (f, g) . \square

B.2 Proof of Proposition 2

The proof is mainly inspired from (Santambrogio, 2015, Proposition 7.17). Let us consider $\alpha, \delta\alpha, \beta, \delta\beta$ and times t in a neighborhood of 0, as in the statement above. We define $\alpha_t = \alpha + t\delta\alpha, \beta_t = \beta + t\delta\beta$ and the variation ratio Δ_t given by

$$\Delta_t \stackrel{\text{def.}}{=} \frac{\text{OT}_\varepsilon(\alpha_t, \beta_t) - \text{OT}_\varepsilon(\alpha, \beta)}{t}.$$

Using the very definition of OT_ε and the continuity property of Proposition 13, we now provide lower and upper bounds on Δ_t as t goes to 0.

Weak* continuity. As written in (13), $\text{OT}_\varepsilon(\alpha, \beta)$ can be computed through a straightforward, *continuous* expression that does not depend on the value of the optimal dual potentials (f, g) at the anchor point x_o :

$$\text{OT}_\varepsilon(\alpha, \beta) = \langle \alpha, f \rangle + \langle \beta, g \rangle.$$

Combining this equation with Proposition 13 (that guarantees the *uniform* convergence of potentials for weakly converging sequences of probability measures) allows us to conclude.

Lower bound. First, let us remark that (f, g) is a *suboptimal* pair of dual potentials for $\text{OT}_\varepsilon(\alpha_t, \beta_t)$. Hence,

$$\begin{aligned} \text{OT}_\varepsilon(\alpha_t, \beta_t) &\geq \langle \alpha_t, f \rangle + \langle \beta_t, g \rangle \\ &\quad - \varepsilon \langle \alpha_t \otimes \beta_t, \exp\left(\frac{1}{\varepsilon}(f \oplus g - C)\right) - 1 \rangle \end{aligned}$$

and thus, since

$$\begin{aligned} \text{OT}_\varepsilon(\alpha, \beta) &= \langle \alpha, f \rangle + \langle \beta, g \rangle \\ &\quad - \varepsilon \langle \alpha \otimes \beta, \exp\left(\frac{1}{\varepsilon}(f \oplus g - C)\right) - 1 \rangle, \end{aligned}$$

one has

$$\begin{aligned} \Delta_t &\geq \langle \delta\alpha, f \rangle + \langle \delta\beta, g \rangle \\ &\quad - \varepsilon \langle \delta\alpha \otimes \beta + \alpha \otimes \delta\beta, \exp\left(\frac{1}{\varepsilon}(f \oplus g - C)\right) \rangle + o(1) \\ &\geq \langle \delta\alpha, f - \varepsilon \rangle + \langle \delta\beta, g - \varepsilon \rangle + o(1), \end{aligned}$$

since g and f satisfy the optimality equations (10).

Upper bound. Conversely, let us denote by (g_t, f_t) the optimal pair of potentials for $\text{OT}_\varepsilon(\alpha_t, \beta_t)$ satisfying $g_t(x_o) = 0$ for some arbitrary anchor point $x_o \in \mathcal{X}$. As (f_t, g_t) are suboptimal potentials for $\text{OT}_\varepsilon(\alpha, \beta)$, we get that

$$\begin{aligned} \text{OT}_\varepsilon(\alpha, \beta) &\geq \langle \alpha, f_t \rangle + \langle \beta, g_t \rangle \\ &\quad - \varepsilon \langle \alpha \otimes \beta, \exp\left(\frac{1}{\varepsilon}(f_t \oplus g_t - C)\right) - 1 \rangle \end{aligned}$$

and thus, since

$$\begin{aligned} \text{OT}_\varepsilon(\alpha_t, \beta_t) &= \langle \alpha_t, f_t \rangle + \langle \beta_t, g_t \rangle \\ &\quad - \varepsilon \langle \alpha_t \otimes \beta_t, \exp\left(\frac{1}{\varepsilon}(f_t \oplus g_t - C)\right) - 1 \rangle, \end{aligned}$$

$$\begin{aligned} \Delta_t &\leq \langle \delta\alpha, f_t \rangle + \langle \delta\beta, g_t \rangle \\ &\quad - \varepsilon \langle \delta\alpha \otimes \beta_t + \alpha_t \otimes \delta\beta, \exp\left(\frac{1}{\varepsilon}(f_t \oplus g_t - C)\right) \rangle + o(1) \\ &\leq \langle \delta\alpha, f_t - \varepsilon \rangle + \langle \delta\beta, g_t - \varepsilon \rangle + o(1). \end{aligned}$$

Conclusion. Now, let us remark that as t goes to 0

$$\alpha + t\delta\alpha \rightharpoonup \alpha \quad \text{and} \quad \beta + t\delta\beta \rightharpoonup \beta.$$

Thanks to Proposition 13, we thus know that f_t and g_t converge uniformly towards f and g . Combining the lower and upper bound, we get

$$\begin{aligned} \Delta_t &\xrightarrow{t \rightarrow 0} \langle \delta\alpha, f - \varepsilon \rangle + \langle \delta\beta, g - \varepsilon \rangle = \langle \delta\alpha, f \rangle + \langle \delta\beta, g \rangle, \end{aligned}$$

since $\delta\alpha$ and $\delta\beta$ both have an overall mass that sums up to zero.

B.3 Proof of Proposition 3

The definition of $\text{OT}_\varepsilon(\alpha, \alpha)$ is that

$$\text{OT}_\varepsilon(\alpha, \alpha) = \max_{(f,g) \in \mathcal{C}(\mathcal{X})^2} \langle \alpha, f + g \rangle - \varepsilon \langle \alpha \otimes \alpha, e^{(f \oplus g - C)/\varepsilon} - 1 \rangle.$$

Reduction of the problem. Thanks to the symmetry of this concave problem with respect to the variables f and g , we know that there exists a pair $(f, g = f)$ of optimal potentials on the diagonal, and

$$\text{OT}_\varepsilon(\alpha, \alpha) = \max_{f \in \mathcal{C}(\mathcal{X})} 2\langle \alpha, f \rangle - \varepsilon \langle \alpha \otimes \alpha, e^{(f \oplus f - C)/\varepsilon} - 1 \rangle.$$

Thanks to the density of continuous functions in the set of simple measurable functions, just as in the proof of Proposition 7, we show that this maximization can be done in the full set of measurable functions $\mathcal{F}_b(\mathcal{X}, \mathbb{R})$:

$$\begin{aligned} \text{OT}_\varepsilon(\alpha, \alpha) &= \max_{f \in \mathcal{F}_b(\mathcal{X}, \mathbb{R})} 2\langle \alpha, f \rangle \\ &\quad - \varepsilon \langle \alpha \otimes \alpha, e^{(f \oplus f - C)/\varepsilon} - 1 \rangle \\ &= \max_{f \in \mathcal{F}_b(\mathcal{X}, \mathbb{R})} 2\langle \alpha, f \rangle \\ &\quad - \varepsilon \langle \exp(f/\varepsilon)\alpha, k_\varepsilon \star \exp(f/\varepsilon)\alpha \rangle + \varepsilon, \end{aligned}$$

where \star denotes the smoothing (convolution) operator defined through

$$[k \star \mu](x) = \int_{\mathcal{X}} k(x, y) d\mu(y)$$

for $k \in \mathcal{C}(\mathcal{X} \times \mathcal{X})$ and $\mu \in \mathcal{M}^+(\mathcal{X})$.

Optimizing on measures. Through a change of variables

$$\mu = \exp(f/\varepsilon)\alpha \quad \text{i.e.} \quad f = \varepsilon \log \frac{d\mu}{d\alpha},$$

keeping in mind that α is a probability measure, we then get that

$$\begin{aligned} \text{OT}_\varepsilon(\alpha, \alpha) &= \varepsilon \max_{\mu \in \mathcal{M}^+(\mathcal{X}), \alpha \ll \mu \ll \alpha} 2\langle \alpha, \log \frac{d\mu}{d\alpha} \rangle \\ &\quad - \langle \mu, k_\varepsilon \star \mu \rangle + 1 \\ -\frac{1}{2}\text{OT}_\varepsilon(\alpha, \alpha) &= \varepsilon \min_{\mu \in \mathcal{M}^+(\mathcal{X}), \alpha \ll \mu \ll \alpha} \langle \alpha, \log \frac{d\mu}{d\alpha} \rangle \\ &\quad + \frac{1}{2}\langle \mu, k_\varepsilon \star \mu \rangle - \frac{1}{2}, \end{aligned}$$

where we optimize on positive measures $\mu \in \mathcal{M}^+(\mathcal{X})$ such that $\alpha \ll \mu$ and $\mu \ll \alpha$.

Expansion of the problem. As $k_\varepsilon(x, y) = \exp(-C(x, y)/\varepsilon)$ is positive for all x and y in \mathcal{X} , we can remove the $\mu \ll \alpha$ constraint from the optimization problem:

$$-\frac{1}{2}\text{OT}_\varepsilon(\alpha, \alpha) = \varepsilon \min_{\mu \in \mathcal{M}^+(\mathcal{X}), \alpha \ll \mu} \langle \alpha, \log \frac{d\mu}{d\alpha} \rangle + \frac{1}{2}\langle \mu, k_\varepsilon \star \mu \rangle - \frac{1}{2}.$$

Indeed, restricting a positive measure μ to the support of α lowers the right-hand term $\langle \mu, k_\varepsilon \star \mu \rangle$ without having any influence on the density of α with respect to μ . Finally, let us remark that the $\alpha \ll \mu$ constraint is already encoded in the $\log \frac{d\mu}{d\alpha}$ operator, which blows up to infinity if α has no density with respect to μ ; in all, we thus have:

$$\begin{aligned} F_\varepsilon(\alpha) &= -\frac{1}{2}\text{OT}_\varepsilon(\alpha, \alpha) \\ &= \varepsilon \min_{\mu \in \mathcal{M}^+(\mathcal{X})} \langle \alpha, \log \frac{d\mu}{d\alpha} \rangle + \frac{1}{2}\langle \mu, k_\varepsilon \star \mu \rangle - \frac{1}{2}, \end{aligned}$$

which is the desired result.

Existence of the optimal measure μ . In the expression above, the existence of an optimal μ is given as a consequence of the well-known fact from OT theory that optimal dual potentials f and g exist, so that the dual OT problem (8) is a max and not a mere supremum. Nevertheless, since this property of F_ε is key to the metrization of the convergence in law by Sinkhorn divergences, let us endow it with a direct, alternate proof:

Proposition 14. *For any $\alpha \in \mathcal{M}_1^+(\mathcal{X})$, assuming that \mathcal{X} is compact, there exists a unique $\mu_\alpha \in \mathcal{M}^+(\mathcal{X})$ such that*

$$F_\varepsilon(\alpha) = \varepsilon \left[\langle \alpha, \log \frac{d\mu_\alpha}{d\alpha} \rangle + \frac{1}{2}\langle \mu_\alpha, k_\varepsilon \star \mu_\alpha \rangle - \frac{1}{2} \right].$$

Moreover, $\alpha \ll \mu_\alpha \ll \alpha$.

Proof. Notice that for $(\alpha, \mu) \in \mathcal{M}_1^+(\mathcal{X}) \times \mathcal{M}^+(\mathcal{X})$,

$$\begin{aligned} E_\varepsilon(\alpha, \mu) &\stackrel{\text{def.}}{=} \langle \alpha, \log \frac{d\mu}{d\alpha} \rangle + \frac{1}{2}\langle \mu, k_\varepsilon \star \mu \rangle \\ &= \text{KL}(\alpha, \mu) + \langle \alpha - \mu, 1 \rangle + \frac{1}{2}\|\mu\|_{k_\varepsilon}^2 - \frac{1}{2}. \end{aligned}$$

Since C is bounded on the compact set $\mathcal{X} \times \mathcal{X}$ and α is a probability measure, we can already say that

$$\frac{1}{\varepsilon}F_\varepsilon(\alpha) \leq E_\varepsilon(\alpha, \alpha) - \frac{1}{2} = \frac{1}{2}\langle \alpha \otimes \alpha, e^{-C/\varepsilon} \rangle - \frac{1}{2} < +\infty.$$

Upper bound on the mass of μ . Since $\mathcal{X} \times \mathcal{X}$ is compact and $k_\varepsilon(x, y) > 0$, there exists $\eta > 0$ such that $k(x, y) > \eta$ for all x and y in \mathcal{X} . We thus get

$$\|\mu\|_{k_\varepsilon}^2 \geq \langle \mu, 1 \rangle^2 \eta$$

and show that

$$\begin{aligned} E_\varepsilon(\alpha, \mu) &\geq \langle \alpha - \mu, 1 \rangle + \frac{1}{2} \|\mu\|_{k_\varepsilon}^2 - \frac{1}{2} \\ &\geq \langle \mu, 1 \rangle (\langle \mu, 1 \rangle \eta - 1) - \frac{1}{2}. \end{aligned}$$

As we build a minimizing sequence (μ_n) for $F_\varepsilon(\alpha)$, we can thus assume that $\langle \mu_n, 1 \rangle$ is uniformly bounded by some constant $M > 0$.

Weak continuity. Crucially, the Banach-Alaoglu theorem asserts that

$$\{\mu \in \mathcal{M}^+(\mathcal{X}) \mid \langle \mu, 1 \rangle \leq M\}$$

is weakly compact; we can thus extract a weakly converging subsequence $\mu_{n_k} \rightharpoonup \mu_\infty$ from the minimizing sequence (μ_n) . Using Proposition 8 and the fact that k_ε is continuous on $\mathcal{X} \times \mathcal{X}$, we show that $\mu \mapsto E_\varepsilon(\alpha, \mu)$ is a weakly lower semi-continuous function: $\mu_\infty = \mu_\alpha$ realizes the minimum of E_ε and we get our existence result.

Uniqueness. We assumed that our kernel k_ε is *positive universal*. The squared norm $\mu \mapsto \|\mu\|_{k_\varepsilon}^2$ is thus a strictly convex functional and using Proposition 6, we can show that $\mu \mapsto E_\varepsilon(\alpha, \mu)$ is *strictly* convex. This ensures that μ_α is uniquely defined. \square

B.4 Proof of Proposition 4

Let us take a pair of measures $\alpha_0 \neq \alpha_1$ in $\mathcal{M}_1^+(\mathcal{X})$, and $t \in (0, 1)$; according to Proposition 14, there exists a pair of measures μ_0, μ_1 in $\mathcal{M}^+(\mathcal{X})$ such that

$$\begin{aligned} (1-t)F_\varepsilon(\alpha_0) + tF_\varepsilon(\alpha_1) &= \varepsilon(1-t)E_\varepsilon(\alpha_0, \mu_0) + \varepsilon tE_\varepsilon(\alpha_1, \mu_1) \\ &> \varepsilon E_\varepsilon((1-t)\alpha_0 + t\alpha_1, (1-t)\mu_0 + t\mu_1) \\ &\geq F_\varepsilon((1-t)\alpha_0 + t\alpha_1), \end{aligned}$$

which is enough to conclude. To show the strict inequality, let us remark that

$$\begin{aligned} (1-t)E_\varepsilon(\alpha_0, \mu_0) + tE_\varepsilon(\alpha_1, \mu_1) \\ = E_\varepsilon((1-t)\alpha_0 + t\alpha_1, (1-t)\mu_0 + t\mu_1) \end{aligned}$$

would imply that $\mu_0 = \mu_1$, since $\mu \mapsto \|\mu\|_{k_\varepsilon}^2$ is strictly convex. As $\alpha \mapsto \text{KL}(\alpha, \beta)$ is strictly convex on the set of measures α that are absolutely continuous with respect to β , we would then have $\alpha_0 = \alpha_1$ and a contradiction with our first hypothesis.

B.5 Proof of the Metrization of the Convergence in Law

The regularized OT cost is weakly continuous, and the uniform convergence for dual potentials ensures that

H_ε and S_ε are both continuous too. Paired with (6), this property guarantees the convergence towards 0 of the Hausdorff and Sinkhorn divergences, as soon as $\alpha_n \rightharpoonup \alpha$.

Conversely, let us assume that $S_\varepsilon(\alpha_n, \alpha) \rightarrow 0$ (resp. $H_\varepsilon(\alpha_n, \alpha)$). Any weak limit α_{n_∞} of a subsequence $(\alpha_{n_k})_k$ is equal to α : since our divergence is weakly continuous, we have $S_\varepsilon(\alpha_{n_\infty}, \alpha) = 0$ (resp. $H_\varepsilon(\alpha_{n_\infty}, \alpha)$), and positive definiteness holds through (6).

In the meantime, since \mathcal{X} is compact, the set of probability Radon measures $\mathcal{M}_1^+(\mathcal{X})$ is sequentially compact for the weak- \star topology. α_n is thus a compact sequence with a unique adherence value: it converges, towards α .