



HAL
open science

Theoretical Guarantees for Bridging Metric Measure Embedding and Optimal Transport

Mokhtar Z. Alaya, Maxime Berar, Gilles Gasso, Alain Rakotomamonjy

► **To cite this version:**

Mokhtar Z. Alaya, Maxime Berar, Gilles Gasso, Alain Rakotomamonjy. Theoretical Guarantees for Bridging Metric Measure Embedding and Optimal Transport. 2021. hal-02485039v3

HAL Id: hal-02485039

<https://hal.science/hal-02485039v3>

Preprint submitted on 22 Apr 2021

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.

Theoretical Guarantees for Bridging Metric Measure Embedding and Optimal Transport

Mokhtar Z. Alaya
LMAC EA 2222
Université de Technologie de Compiègne
elmokhtar.alaya@utc.fr

Maxime Bérar
LITIS EA 4108
Université de Rouen Normandie
maxime.berar@univ-rouen.fr

Gilles Gasso
LITIS EA 4108
Université de Rouen Normandie
& INSA Rouen
gilles.gasso@insa-rouen.fr

Alain Rakotomamonjy
LITIS EA 4108
Université de Rouen Normandie
& Criteo AI Lab
alain.rakoto@insa-rouen.fr

April 22, 2021

Abstract

We propose a novel approach for comparing distributions whose supports do not necessarily lie on the same metric space. Unlike Gromov-Wasserstein (GW) distance which compares pairwise distances of elements from each distribution, we consider a method allowing to embed the metric measure spaces in a common Euclidean space and compute an optimal transport (OT) on the embedded distributions. This leads to what we call a *sub-embedding robust Wasserstein* (SERW) distance. Under some conditions, SERW is a distance that considers an OT distance of the (low-distorted) embedded distributions using a common metric. In addition to this novel proposal that generalizes several recent OT works, our contributions stand on several theoretical analyses: (i) we characterize the embedding spaces to define SERW distance for distribution alignment; (ii) we prove that SERW mimics almost the same properties of GW distance, and we give a cost relation between GW and SERW. The paper also provides some numerical illustrations of how SERW behaves on matching problems.

1 Introduction

Many central tasks in machine learning often attempt to align or match real-world entities, based on computing distance (or dissimilarity) between pairs of corresponding probability distributions. Recently, optimal transport (OT) based data analysis has proven a significant usefulness to achieve such tasks, arising from designing loss functions (Frogner et al., 2015), unsupervised learning (Arjovsky et al., 2017), clustering (Ho et al., 2017), text classification (Kusner et al., 2015), domain adaptation (Courty et al., 2017), computer vision (Bonneel et al., 2011; Solomon et al., 2015), among many more applications (Kolouri et al., 2017; Peyré and Cuturi, 2019). Distances based on OT are referred to as the Monge-Kantorovich or Wasserstein distance (Monge, 1781; Kantorovich, 1942; Villani, 2009). OT tools allow for a natural geometric comparison of distributions, that takes into account the metric of the underlying space to find the most cost-efficiency way to transport mass from a set of sources to a set of targets. The success of machine learning algorithms based on Wasserstein distance is due to its nice properties (Villani, 2009) and to recent development of efficient computations using entropic regularization (Cuturi, 2013; Genevay et al., 2016; Altschuler et al., 2017; Alaya et al., 2019).

Distribution alignment using Wasserstein distance relies on the assumption that the two sets of entities in question belong to the same ground space, or at least pairwise distance between them can be computed. To overcome such limitations, one seeks to compute Gromov-Wasserstein (GW) distance (Sturm, 2006; Mémoli, 2011), which is a relaxation of Gromov-Hausdorff distance (Mémoli, 2008; Bronstein et al., 2010). GW distance allows to learn an optimal transport-like plan by measuring how the distances between pairs of samples within each ground space are similar. The GW framework has been used for solving alignment problems in several applications, for instance shape matching (Mémoli, 2011), graph partitioning and matching (Xu et al., 2019), matching of vocabulary sets between different languages (Alvarez-Melis and Jaakkola, 2018), generative models (Bunne et al., 2019), or matching weighted networks (Chowdhury and Mémoli, 2018), to name a few. However, computing GW distance induces a heavy computation burden as the underlying problem is a non-convex quadratic program and NP-hard (Peyré and Cuturi, 2019). Peyré et al. (2016) propose an entropic version called entropic GW discrepancy, that leads to approximate GW distance.

In this paper, we develop a distance, that is similarly to GW distance applied on sets of entities from different spaces. Our proposal builds upon metric embedding that allows for an approximation of some “hard” problem with complex metric spaces into another one involving “simpler” metric space (Matoušek, 2002) and upon Wasserstein OT cost on the embedding space. Hence, unlike GW distance that compares pairwise distances of elements from each distribution, we consider a method that embeds the metric measure spaces into a common Euclidean space and computes a Wasserstein OT distance between the embedded distributions. In this context, we introduce a distance, robust to isometry in the embedding space, that generalizes the “min-max” robust OT problem recently introduced in Paty and Cuturi (2019), where the authors consider orthogonal projections as embedding functions. Main contributions of this work are summarized in the following three points:

- We propose a framework for distribution alignment from different spaces using a sub-embedding robust Wasserstein (SERW) distance. As central contribution, we develop the theoretical analysis characterizing the embedding spaces so that SERW be a distance.
- We provide mathematical evidence on the relation between GW and our SERW distances. We show for instance, that one key point for approximating GW is that the embeddings be distance-preserving.
- We sketch a potential algorithm describing how our distance can be computed in practice and present numerical illustrations on simulated and real datasets that support our theoretical results.

The remainder of the paper is organized as follows. In Section 2 we introduce the definitions of Wasserstein and GW distances, and set up the embedding spaces. In Section 3 we investigate metric measure embedding for non-aligned distributions through an OT via SERW distance. Section 4 is dedicated to numerical experiments on matching tasks based on simulated and real data. The proofs of the main results are postponed to the appendices.

2 Preliminaries

We start here by reviewing basic definitions of the materials needed to introduce the main results. We consider two metric measure spaces (mm-space for short) (Gromov et al., 1999) (X, d_X, μ) and (Y, d_Y, ν) , where (X, d_X) is a compact metric space and μ is a probability measure with full support, i.e. $\mu(X) = 1$ and $\text{supp}[\mu] = X$. We recall that the support of a measure $\text{supp}[\mu]$ is the minimal closed subset $X_0 \subset X$ such that $\mu(X \setminus X_0) = 0$. Similarly, we define the mm-space (Y, d_Y, ν) . Let $\mathcal{P}(X)$ be the set of probability measures in X and $p \in \{1, 2\}$. We define $\mathcal{P}_p(X)$ as its subset consisting of measures with finite p -moment, i.e.,

$$\mathcal{P}_p(X) = \{\eta \in \mathcal{P}(X) : M_p(\eta) < \infty\},$$

where $M_p(\eta) := \int_X d_X^p(x, 0) d\eta(x)$. For $\mu \in \mathcal{P}(X)$ and $\nu \in \mathcal{P}(Y)$, we write $\Pi(\mu, \nu) \subset \mathcal{P}(X \times Y)$ for the collection of probability measures (couplings) on $X \times Y$ as

$$\Pi(\mu, \nu) = \{ \pi \in \mathcal{P}(X \times Y) : \forall A \subset X, B \subset Y, \pi(A \times Y) = \mu(A) \text{ and } \pi(X \times B) = \nu(B) \}.$$

Wasserstein distance. The Monge-Kantorovich or the 2-Wasserstein distance aims at finding an optimal mass transportation plan $\pi \in \mathcal{P}(X \times Y)$ such that the marginals of π are respectively μ and ν , and these two distributions are supposed to be defined over the same ground space, i.e., $X = Y$. It reads as

$$\mathcal{W}_2^2(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \int_{X \times X} d_X^2(x, x') d\pi(x, x'). \quad (1)$$

The infimum in (1) is attained, and any probability π which realizes the minimum is called an *optimal transport plan*.

Gromov-Wasserstein distance. In contrast to Wasserstein distance, GW one deals with measures that do not necessarily belong to the same ground space. It learns an optimal transport-like plan which transports samples from a source metric space X into a target metric space Y by measuring how the distances between pairs of samples within each space are similar. Following the pioneering work of Mémoli (2011), GW distance is defined as

$$\mathcal{GW}_2^2(\mu, \nu) = \frac{1}{2} \inf_{\pi \in \Pi(\mu, \nu)} J_\pi(\mu, \nu)$$

where

$$J_\pi(\mu, \nu) = \iint_{X^2 \times Y^2} \ell(d_X(x, x'), d_Y(y, y')) d\pi(x, y) d\pi(x', y')$$

with a quadratic loss function $\ell(a, b) = |a - b|^2$. Peyré et al. (2016) propose an entropic version called entropic GW discrepancy, allowing to tackle more flexible losses ℓ , such as mean-square-error or Kullback-Leibler divergence. The latter version includes an entropic regularization of π in the GW distance computation problem.

Metric embedding. Metric embedding consists in characterizing a new representation of the samples based on the concept of distance preserving.

Definition 1. A mapping $\phi : (X, d_X) \rightarrow (Z, d_Z)$ is an embedding with distortion τ , denoted as τ -embedding, if the following holds: there exists a constant $\kappa > 0$ (“scaling factor”) such that for all $x, x' \in X$,

$$\kappa d_X(x, x') \leq d_Z(\phi(x), \phi(x')) \leq \tau \kappa d_X(x, x'). \quad (2)$$

The approximation factor in metric embedding depends on a distortion parameter of the ϕ -embedding. This distortion is defined as the infimum of all $\tau \geq 1$ such that the above condition (2) holds. If no such τ exists, then the distortion of ϕ is infinity.

In this work we will focus on target spaces Z that are normed spaces endowed with Euclidean distance. Specifically, for some integer d to be specified later, we will consider the metric space $(Z = \mathbb{R}^d, d_Z = \|\cdot\|)$. Hence, one can always take the scaling factor κ to be equal to 1 (by replacing ϕ by $\frac{1}{\kappa}\phi$). Note that an embedding ϕ with distortion at most $\tau < \infty$ is necessarily *one-to-one* (injective). Isometric embeddings are for instance embeddings with distortion 1. For more details about embeddings, we invite the reader to look at the technical report of Matoušek (2013).

We suppose hereafter $\kappa = 1$ in (2) and we denote by $\mathcal{F}_d(X)$ and $\mathcal{F}_d(Y)$ the set of τ_ϕ -embedding $\phi : X \rightarrow \mathbb{R}^d$ and τ_ψ -embedding $\psi : Y \rightarrow \mathbb{R}^d$, respectively. We further assume that $\phi(0) = \psi(0) = 0$. It is worth noting that when X and Y are finite spaces, then the embedding spaces $\mathcal{F}_d(X)$ and $\mathcal{F}_d(Y)$ are non empty. Indeed, suppose we are given a set of n data points $\{x_1, x_2, \dots, x_n\} = X$, then *Bourgain's embedding theorem* (Bourgain, 1985) guarantees the existence of an embedding $\phi : X \rightarrow (\mathbb{R}^d, \|\cdot\|)$ with tight distortion at most $\mathcal{O}(\log n)$, i.e., $\tau_\phi = \mathcal{O}(\log n)$ ¹, and the target dimension $d = \mathcal{O}(\log^2 n)$. We stress that d is independent of the original dimensions of X and Y and depends only on the number of the given data points n and m and the accuracy-embedding parameters τ_ϕ and τ_ψ . Hence for data points $X = \{x_1, x_2, \dots, x_n\}$ and $Y = \{y_1, y_2, \dots, x_m\}$ underlying the distributions of interest, one has

$$d = \mathcal{O}(\log^2(\max(n, m))). \quad (3)$$

Remark 1. *For our setting, we cannot use the Johnson-Lindenstrauss flattening lemma (Johnson and Lindenstrauss, 1984) (JL-lemma in short) to guarantee the existence of the embeddings. To see that, let us first state this lemma: the JL-lemma asserts that if X is a Hilbert space, $\varepsilon > 0$, $n \in \mathbb{N}$, and $x_1, \dots, x_n \in X$ then there exists a linear mapping $\phi_{JL} : X \rightarrow \mathbb{R}^d$ with $d = \mathcal{O}(\log n)$ such that for all $i, i' \in \{1, \dots, n\}$, one has*

$$\|x_i - x_{i'}\|^2 \leq \|\phi_{JL}(x_i) - \phi_{JL}(x_{i'})\|^2 \leq (1 + \varepsilon)\|x_i - x_{i'}\|^2.$$

Note that the origin space X is supposed to be a Hilbert space in the JL-lemma and the embedding ϕ_{JL} is a linear mapping with a target dimension $d \leq C_\varepsilon \log n$. However, X is only supposed to be a metric space and the target dimension $d \leq C \log^2 n$ in Bourgain's theorem. Furthermore, the JL-lemma fails to hold true in certain non-Hilbertian settings (Matoušek, 1996; Brinkman and Charikar, 2005). In our approach, we consider a general setup that relies on non-linear embeddings that act on different data structures like text, graphs, images, etc. Therefore, we used Bourgain's theorem to guarantee the non-emptiness of the embedding spaces $\mathcal{F}_d(X)$ and $\mathcal{F}_d(Y)$.

Let's now highlight all the above criteria characterizing the metric embeddings we consider to define our novel distance and that help us shape some of its properties.

Assumption 1. *Assume that (X, d_X, μ) and (Y, d_Y, ν) are finite mm-spaces containing the origin 0, and endowed with measures μ and ν . Assume also that X and Y are of cardinalities n and m , the target dimension d satisfies (3), $\mathcal{F}_d(X) = \{\phi : X \rightarrow \mathbb{R}^d, \tau_\phi\text{-embedding, with } \phi(0) = 0\}$ and $\mathcal{F}_d(Y) = \{\psi : Y \rightarrow \mathbb{R}^d, \tau_\psi\text{-embedding, with } \psi(0) = 0\}$. The distortions parameters $\tau_\phi \in \mathcal{D}_{emb}(X)$, $\tau_\psi \in \mathcal{D}_{emb}(Y)$ where $\mathcal{D}_{emb}(X) = [1, \mathcal{O}(\log(n))]$ and $\mathcal{D}_{emb}(Y) = [1, \mathcal{O}(\log(m))]$.*

3 Metric measure embedding and OT for distribution alignment

Let us give first the overall structure of our approach of non-aligned distributions, which generalizes recent works (Alvarez-Melis and Jaakkola, 2018; Paty and Cuturi, 2019). The generalization stands on the fact that the two distributions lie on two different metric spaces and this fact raises several challenging questions about the characterization of the embeddings for yielding a distance. In our approach, we consider a general setup that relies on non-linear embeddings before aligning the measures. Note, if the metric spaces coincide for both distributions and the embeddings are restricted to be linear (subspace projection) then, our distance reduces to the one proposed by Paty and Cuturi (2019).

In this work, we aim at proposing a novel distance between two measures defined on different mm-spaces. This distance will be defined as the optimal objective of some optimization problem, we provide technical details and conditions ensuring its existence in the first part of this section. We then present formally our novel distance and its properties including its cost relation with GW distance.

¹There exists an absolute constant $C > 0$ such that $\tau_\phi \leq C \log n$.

In a nutshell, our distribution alignment distance between μ and ν is obtained as a Wasserstein distance between pushforwards (see Definition 2) of μ and ν w.r.t. some appropriate couple of embeddings $(\phi, \psi) \in \mathcal{F}_d(X) \times \mathcal{F}_d(Y)$. Towards this end, we need to exhibit some topological properties of the embeddings spaces, leading at first to the existence of the constructed OT approximate distances.

3.1 Topological properties of the embedding spaces

We may consider the function $\Gamma_X : \mathcal{F}_d(X) \times \mathcal{F}_d(X) \rightarrow \mathbb{R}_+$ such that $\Gamma_X(\phi, \phi') = \sup_{x \in X} \|\phi(x) - \phi'(x)\|$, for each pair of embeddings $\phi, \phi' \in \mathcal{F}_d(X)$. This function defines a proper metric on the space of embeddings $\mathcal{F}_d(X)$ and it is referred to as the *supremum metric* on $\mathcal{F}_d(X)$. Indeed, Γ_X satisfies all the conditions that define a general metric. We analogously define the metric Γ_Y on $\mathcal{F}_d(Y)$. With the aforementioned preparations, the embeddings spaces satisfy the following topological property.

Proposition 1. $(\mathcal{F}_d(X), \Gamma_X)$ and $(\mathcal{F}_d(Y), \Gamma_Y)$ are both compact metric spaces.

Endowing the embedding spaces with the supremum metrics is fruitful, since we get benefits from some existing topological results, based on this functional space metric, to prove the statement in Proposition 1.

To let it more readable, the proof of Proposition 1 is divided into 5 steps summarized as follows: first step is for metric property of $\mathcal{F}_d(X)$; second one shows completeness of $\mathcal{F}_d(X)$; third establishes the totally boundedness of $\mathcal{F}_d(X)$, namely that one can recover this space using balls centered on a finite number of embedding points; the last is a conclusion using Arzela-Ascoli’s Theorem for characterizing compactness of subsets of functional continuous space, see Appendix A.1 for all these details and their proofs.

We now give a definition of pushforward measures.

Definition 2. (*Pushforward measure*). Let (S, \mathcal{S}) and (T, \mathcal{T}) be two measurable spaces, $f : S \rightarrow T$ be a mapping, and η be a measure on S . The pushforward of η by f , written $f_{\#}\eta$, is the measure on T defined by $f_{\#}\eta(A) = \eta(f^{-1}(A))$ for $A \in \mathcal{T}$. If η is a measure and f is a measurable function, then $f_{\#}\eta$ is a measure.

3.2 Sub-Embedding OT

Assume Assumption 1 holds. Following Paty and Cuturi (2019), we define an embedding robust version of Wasserstein distance between pushforwards $\phi_{\#}\mu \in \mathcal{P}_p(\mathbb{R}^d)$ and $\psi_{\#}\nu \in \mathcal{P}_p(\mathbb{R}^d)$ for some appropriate couple of embeddings $(\phi, \psi) \in \mathcal{F}_d(X) \times \mathcal{F}_d(Y)$. We then consider the worst possible OT cost over all possible low-distortion embeddings. The notion of “robustness” in our distance stands from the fact that we look for this worst embedding.

Definition 3. The d -dimensional embedding robust 2-Wasserstein distance (ERW) between μ and ν reads as

$$\mathcal{E}_d^2(\mu, \nu) = \frac{1}{2} \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \mathcal{W}_2^2(\phi_{\#}\mu, (r \circ \psi)_{\#}\nu),$$

where \mathcal{R}_d stands for the set of orthogonal mappings on \mathbb{R}^d and \circ denotes the composition operator between functions.

The infimum over the orthogonal mappings on \mathbb{R}^d corresponds to a classical orthogonal procrustes problem (Gower, 1975; Grave et al., 2019). It learns the best rotation between the embedded points, allowing to an accurate alignment. The orthogonality constraint ensures that the distances between points are preserved by the transformation.

Note that $\mathcal{E}_d^2(\mu, \nu)$ is finite since the considered embeddings are Lipschitz and both of the distributions μ and ν have finite 2-moment due to Assumption 1. Next, using results of pushforward measures, for instance see Lemmas 7 and 8 in the Appendices, we explicit ERW in Lemma 1, whereas Lemmas 2 and 4 establish the existence of embeddings that achieve the suprema defining both ERW and SERW.

Lemma 1. *For any $(\phi, \psi) \in \mathcal{F}_d(X) \times \mathcal{F}_d(Y)$ and $r \in \mathcal{R}_d$, let $J_{\phi, \psi, r, \pi}(\mu, \nu) = \int_{X \times Y} \|\phi(x) - r(\psi(y))\|^2 d\pi(x, y)$. One has*

$$\mathcal{E}_d^2(\mu, \nu) = \frac{1}{2} \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \inf_{\pi \in \Pi(\mu, \nu)} J_{\phi, \psi, r, \pi}(\mu, \nu).$$

By the compactness property of the embedding spaces (see Proposition 1), the set of optima defining $\mathcal{E}_d^2(\mu, \nu)$ is not empty.

Lemma 2. *There exist a couple of embeddings $(\phi^*, \psi^*) \in \mathcal{F}_d(X) \times \mathcal{F}_d(Y)$ and $r^* \in \mathcal{R}_d$ such that*

$$\mathcal{E}_d^2(\mu, \nu) = \frac{1}{2} \mathcal{W}_2^2(\phi^*_{\#}\mu, (r^* \circ \psi^*)_{\#}\nu).$$

Clearly, the quantity $\mathcal{E}_d^2(\mu, \nu)$ is difficult to compute, since an OT is a linear programming problem that requires generally super cubic arithmetic operations (Peyré and Cuturi, 2019). Based on this observation, we focus on the corresponding “min-max” problem to define the d -dimensional sub-embedding robust 2-Wasserstein distance (SERW). For the sake, we make the next definition.

Definition 4. *The d -dimensional sub-embedding robust 2-Wasserstein distance (SERW) between μ and ν is defined as*

$$\mathcal{S}_d^2(\mu, \nu) = \frac{1}{2} \inf_{\pi \in \Pi(\mu, \nu, \pi)} \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} J_{\phi, \psi, r, \pi}(\mu, \nu).$$

Thanks to the minimax inequality, the following holds.

Lemma 3. $\mathcal{E}_d^2(\mu, \nu) \leq \mathcal{S}_d^2(\mu, \nu)$.

We emphasize that ERW and SERW quantities play a crucial role in our approach to match distributions in the common space \mathbb{R}^d regarding pushforwards of the measures μ and ν realized by a couple of optimal embeddings and a rotation. Optimal solutions for $\mathcal{S}_d^2(\mu, \nu)$ exist. Namely:

Lemma 4. *There exist a couple of embeddings $(\phi^*, \psi^*) \in \mathcal{F}_d(X) \times \mathcal{F}_d(Y)$ and $r^* \in \mathcal{R}_d$ such that*

$$\mathcal{S}_d^2(\mu, \nu) = \frac{1}{2} \inf_{\pi \in \Pi(\mu, \nu)} J_{\phi^*, \psi^*, r^*, \pi}(\mu, \nu).$$

The proofs of Lemmas 2 and 4 rely on the continuity under integral sign Theorem (Schilling, 2005), and the compactness property of the embedding spaces, the orthogonal mappings on \mathbb{R}^d space and the coupling transport plan $\Pi(\mu, \nu)$, see Appendices A.5 and A.3 for more details.

Recall that we are interested in distribution alignment for measures coming from different mm-spaces. One hence expects that SERW mimics some metric properties of GW distance. To proceed in this direction, we first prove that SERW defines a proper metric on the set of all *weakly isomorphism* classes of mm-spaces. In our setting the terminology of weakly isomorphism means that there exists a pushforward embedding between mm-spaces. If such a pushforward is 1-embedding the class is called *strongly isomorphism*.

Proposition 2. *Let Assumption 1 holds and assume $X \subseteq \mathbb{R}^D$ and $Y \subseteq \mathbb{R}^{D'}$ with $D \neq D'$. Then, $\mathcal{S}_d^2(\mu, \nu) = 0$ happens if and only if the couple of embeddings (ϕ^*, ψ^*) and $r^* \in \mathcal{R}_d$ optima of $\mathcal{S}_d^2(\mu, \nu)$ verify $\mu = (\phi^{*-1} \circ r^* \circ \psi^*)_{\#}\nu$ and $\nu = ((r^* \circ \psi^*)^{-1} \circ \phi^*)_{\#}\mu$.*




Figure 1: Illustration of the preserving measure mappings between the mm-spaces (X, d_X, μ) and (Y, d_Y, ν) given in Proposition 2. The embedding ϕ^* maps from X to \mathbb{R}^d while $r^* \circ \psi^*$ maps from Y to \mathbb{R}^d . Our distance $\mathcal{S}_d^2(\mu, \nu) = 0$ vanishes if and only if μ and ν are mapped through the embeddings ϕ^{*-1} and $(r^* \circ \psi^*)^{-1}$.

Figure 1 illustrates the mappings between the embedding spaces and how they are assumed to interact in order to satisfy condition in Proposition 2. In Mémoli (2011) (Theorem 5, property (a)), it is shown that $\mathcal{GW}_2^2(\mu, \nu) = 0$ if and only if (X, d_X, μ) and (Y, d_Y, ν) are strongly isomorphic. This means that there exists a Borel measurable bijection $\varphi : X \rightarrow Y$ (with Borel measurable inverse φ^{-1}) such that φ is 1-embedding and $\varphi\#\mu = \nu$. The statement in Proposition 2 is a weak version of the aforementioned result, because neither $\phi^{*-1} \circ r^* \circ \psi^*$ nor $(r^* \circ \psi^*)^{-1} \circ \phi^*$ are isometric embeddings. However, we succeed to find a measure-preserving mapping relating μ and ν to each other via the given pushforwards in Proposition 2. Note that $r^* \circ \psi^*$ maps from the Y space to \mathbb{R}^d while ϕ^* maps from X to \mathbb{R}^d . Our distance $\mathcal{S}_d^2(\mu, \nu)$ vanishes if and only if μ and ν are mapped through the embeddings ϕ^{*-1} and $(r^* \circ \psi^*)^{-1}$. With these elements, we can now prove that both ERW and SERW are further distances.

Proposition 3. *Assume that statement of Proposition 2 holds. Then, ERW and SERW define a proper distances between weakly isomorphism mm-spaces.*

3.3 Cost relation between GW and SERW

In addition to the afore-mentioned theoretical properties of SERW, we establish a *cost relation* metric between GW and SERW distances. The obtained upper and lower bounds depend on approximation constants that are linked to the distortions of the embeddings.

Proposition 4. *Let Assumption 1 holds. Then,*

$$\frac{1}{2}\mathcal{GW}_2^2(\mu, \nu) \leq \mathcal{S}_d^2(\mu, \nu) + \alpha \underline{M}_{\mu, \nu}$$

where $\alpha = 2 \inf_{\tau_\phi \in \mathcal{D}_{emb}(X), \tau_\psi \in \mathcal{D}_{emb}(Y)} (\tau_\phi \tau_\psi - 1)$ and $\underline{M}_{\mu, \nu} = 2(M_1(\mu) + M_1(\nu))$.

Proposition 5. *Let Assumption 1 holds. Then,*

$$\mathcal{S}_d^2(\mu, \nu) \leq \beta \mathcal{GW}_2^2(\mu, \nu) + 4\beta \overline{M}_{\mu, \nu}$$

where $\beta = 2 \sup_{\tau_\phi \in \mathcal{D}_{emb}(X), \tau_\psi \in \mathcal{D}_{emb}(Y)} (\tau_\phi^2 + \tau_\psi^2)$ and $\overline{M}_{\mu, \nu} = (\sqrt{M_2(\mu)} + \sqrt{M_1(\mu)})(\sqrt{M_2(\nu)} + \sqrt{M_1(\nu)}) + M_2(\mu) + M_2(\nu)$.

Proofs of Propositions 4 and 5 are presented in Appendices A.9 and A.8. We use upper and lower bounds of GW distance as provided in Mémoli (2008). The cost relation between SERW and GW distances obtained in Propositions 4 and 5 are up to the constants α, β which are depending on the distortion parameters of the embeddings, and up to an additive constant through the p -moments M_p of the measures μ and ν . In the following we highlight some particular cases leading to closed form of the upper and lower bounds for the cost relation between GW and SERW distances.

3.4 Fixed sub-embedding for distribution alignment

From the computational point of view, computing SERW distance seems a daunting task, since one would have to optimize over the product of two huge embedding spaces $\mathcal{F}_d(X) \times \mathcal{F}_d(Y)$. However in some applications we may not require solving over $\mathcal{F}_d(X) \times \mathcal{F}_d(Y)$ and rather have at disposal known embeddings in advance. For instance, for image-text alignment we may leverage on features extracted from pre-trained deep architectures (VGG (Simonyan and Zisserman, 2015) for image embedding, and Word2vec Mikolov et al. (2013) for projection the text). Roughly speaking, our SERW procedure with respect to these fixed embeddings can be viewed as an *embedding-dependent distribution alignment* for matching. More precisely, the alignment quality is strongly dependent on the given embeddings; the lower distorted embeddings, the more accurate alignment.

Definition 5. For a fixed couple of embeddings $(\phi_f, \psi_f) \in \mathcal{F}_d(X) \times \mathcal{F}_d(Y)$, we define the fixed sub-embedding robust Wasserstein (FSERW) as

$$\widetilde{\mathcal{S}}_d^2 = \frac{1}{2} \inf_{\pi \in \Pi(\mu, \nu)} \inf_{r \in \mathcal{R}_d} J_{\phi_f, \psi_f, r, \pi}(\mu, \nu).$$

Lemma 5. $\widetilde{\mathcal{S}}_d^2$ defines a proper distance if and only if $\mu = (\phi_f^{-1} \circ (r_f \circ \psi_f))_{\#} \nu$ and $\nu = ((r_f \circ \psi_f)^{-1} \circ \phi_f)_{\#} \mu$, where $r_f = \inf_{r \in \mathcal{R}_d} J_{\phi_f, \psi_f, r, \pi}(\mu, \nu)$.

The cost relation guarantees given in Propositions 3 and 4 are dependent on the distortions of the fixed embeddings, i.e., the constants α and β become: $\alpha_f = 2(\tau_{\phi_f} \tau_{\psi_f} - 1)$ and $\beta_f = 2(\tau_{\phi_f}^2 + \tau_{\psi_f}^2)$. Then the following holds

Lemma 6. One has $\frac{1}{2} \mathcal{GW}_2^2(\mu, \nu) \leq \mathcal{S}_d^2(\mu, \nu) + \alpha_f \underline{M}_{\mu, \nu}$ and $\mathcal{S}_d^2(\mu, \nu) \leq \beta_f \mathcal{GW}_2^2(\mu, \nu) + 4\beta_f \overline{M}_{\mu, \nu}$.

In a particular case of isometric embeddings, our procedure gives the following cost relation

$$\frac{1}{2} \mathcal{GW}_2^2(\mu, \nu) \leq \mathcal{S}_d^2(\mu, \nu) \leq 4 \mathcal{GW}_2^2(\mu, \nu) + 16 \overline{M}_{\mu, \nu}.$$

The additive constants $\underline{M}_{\mu, \nu}$ and $\overline{M}_{\mu, \nu}$ can be upper bounded in a setting of data preprocessing, for instance in the case of a ℓ_2 -normalization preprocessing we have $\underline{M}_{\mu, \nu} \leq 4$ and $\overline{M}_{\mu, \nu} \leq 6$.

4 Numerical experiments

Here we illustrate how SERW distance behaves on numerical problems. We apply it on some toy problems as well as on some problems usually addressed using GW distance.

Sketch of practical implementation. Based on the above presented theory, we have several options for computing the distance between non-aligned measures and they all come with some guarantees compared to a GW distance. In the simpler case of fixed embedding, if the original spaces are subspaces of \mathbb{R}^d , any distance preserving embedding can be a good option for having an embedding with low distortion. Typically, methods like multidimensional scaling (MDS) (Kruskal and Wish, 1978),

Isomap (Tenenbaum et al., 2000) or Local linear embedding (LLE) (Roweis and Saul, 2000) can be good candidates. One of the key advantages of SERW is that it considers non-linear embedding before measure alignment. Hence, it has the ability of leveraging over the large zoo of recent embedding methods that act on different data structures like text (Grave et al., 2018), graphs (Grover and Leskovec, 2016; Narayanan et al., 2017), images (Simonyan and Zisserman, 2015), or even histograms (Courtney et al., 2018).

In the general setting our theoretical results require computing $\mathcal{S}_d^2(\mu, \nu)$ to solve the problem given in the Definition (4). We sketch in Section 4.1 a practical procedure to learn from samples $X = \{x_i\}_{i=1}^n$ and $Y = \{y_j\}_{j=1}^m$, non-linear neural-network based embedding functions ϕ and ψ that maximize the Wasserstein distance between the embedded samples while minimizing the embedding distortion.

4.1 Implementation details on learning the embeddings

In practice, for computing $\mathcal{S}_d^2(\mu, \nu)$, we need to solve the problem given in Equation (4). As stated above in some practical situations, we leverage on existing embeddings and consider the problem without the maximization over the embeddings as the space is restricted to an unique singleton. In some other cases, it is possible to learn the embedding that maximizes the Wasserstein distance between embedded examples and that minimizes the distance distortion of the embedding. In what follows, we detail how we have numerically implemented the computation of $\mathcal{S}_d^2(\mu, \nu)$ from samples $\{x_i\}$ and $\{y_j\}$ respectively sampled from X and Y according to μ and ν .

For all $i, i' \in \{1, \dots, n\}, j, j' \in \{1, \dots, m\}$, we denote the pairwise distances $C_{ii'}^X = d_X(x_i, x_{i'})$, $C_{jj'}^Y = d_Y(y_j, y_{j'})$ in the origin space X, Y and $C_{ii'}^\phi = \|\phi(x_i) - \phi(x_{i'})\|^2$, $C_{jj'}^\psi = \|\psi(y_j) - \psi(y_{j'})\|^2$ in the embedded space \mathbb{R}^d . The problem we want to solve is

$$\min_{\pi \in \Pi(\mu, \nu)} \min_{r \in \mathcal{R}_d} \max_{\phi, \psi} \left\{ \frac{1}{2} \sum_{i,j} \|\phi(x_i) - r(\psi(y_j))\|^2 \pi_{i,j} - \sum_{i,i:i \neq i'} D_X^\phi(C_{i,j}^X, C_{i,j}^\phi) - \sum_{j,j':j \neq j'} D_Y^\psi(C_{i,j}^Y, C_{i,j}^\psi) \right\} \quad (4)$$

with $\Pi(\mu, \nu) = \{\pi \in \mathbb{R}^{n \times m} | \pi \mathbf{1}_m = \mu, \pi^\top \mathbf{1}_n = \nu\}$. In this equation, the first sum corresponds to the optimal transport cost function and the other two sums compute the distortion between pairwise distances in the input space and embedded space respectively for the x and the y samples. In the notation, $D_X^\phi(\cdot, \cdot)$ is a loss function that penalizes the discrepancy between the input $C_{ii'}^X$ and embedded $C_{ii'}^\phi$ distances. This distance loss D_X^ϕ has been designed so as to encourage the embedding to preserve pairwise distance up to a $\tilde{\tau}_\phi$ factor. Hence

$$D_X^\phi(C_{ii'}^X, C_{ii'}^\phi) = \mathbb{1} \left(\frac{C_{ii'}^\phi}{C_{ii'}^X} > \tilde{\tau}_\phi \right) \frac{C_{ii'}^\phi}{C_{ii'}^X} \quad D_Y^\psi(C_{jj'}^Y, C_{jj'}^\psi) = \mathbb{1} \left(\frac{C_{jj'}^\psi}{C_{jj'}^Y} > \tilde{\tau}_\psi \right) \frac{C_{jj'}^\psi}{C_{jj'}^Y}$$

with

$$\min_{i \neq i'} \frac{C_{ii'}^\phi}{C_{ii'}^X} \leq \tilde{\tau}_\phi \leq \max_{i \neq i'} \frac{C_{ii'}^\phi}{C_{ii'}^X} \quad \min_{j \neq j'} \frac{C_{jj'}^\psi}{C_{jj'}^Y} \leq \tilde{\tau}_\psi \leq \max_{j \neq j'} \frac{C_{jj'}^\psi}{C_{jj'}^Y}$$

and $\mathbb{1}(\cdot)$ denotes the indicator function. In the experiments $\tilde{\tau}_\phi$ is fixed as $\max(0.9 \max_{i \neq i'} \frac{C_{ii'}^\phi}{C_{ii'}^X}, \min_{i \neq i'} \frac{C_{ii'}^\phi}{C_{ii'}^X})$ (similarly for $\tilde{\tau}_\psi$). It penalizes the embedded couples of inputs whose embedded pairwise distances are the most dissimilar to the input pairwise distances. As these specific discrepancies impact the estimation of the distortion rate of the embedding, the designed loss has been tailored to reduce the distortion rate comparatively to those of the initial embeddings.

In practice, the embedding functions ϕ and ψ have been implemented in the following way

$$\phi = (\mathbf{I} + g_X) \circ h_{\theta_X} \quad \psi = (\mathbf{I} + g_Y) \circ h_{\theta_Y} \quad (5)$$

where \mathbf{I} is the identity matrix, $g_X : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $g_Y : \mathbb{R}^d \rightarrow \mathbb{R}^d$ are trainable neural networks based embeddings and $h_{\theta_X} : X \rightarrow \mathbb{R}^d$ and $h_{\theta_Y} : Y \rightarrow \mathbb{R}^d$ are data-dependent low-dimensional projections that preserves (local) distances. Typically, for the h functions, we have considered in our experiments algorithms like MDS or LLE. Our choice of the embedding dimension can be then principally giving by the d -largest number of eigenvalues of the spectral decomposition arising within these algorithms.

So the learning problem described in Equation (4) involves a max-min problem over the Wasserstein distance of the mapped samples. For solving the problem, we have adopted an alternate optimization strategy where for each mini-batch of samples from $\{x_i\}$ and $\{y_j\}$, we first optimize r and π at fixed ϕ and ψ and then optimize the embeddings for fixed optimal r and π . In practice, the sub-problem with respects to r and π is an invariant OT problem and can be solved using the algorithm proposed by Alvarez-Melis et al. (2019). g_X and g_Y is implemented as two fully connected neural networks with leaky ReLU activation functions and no bias. They are optimized using stochastic gradient descent using Adam as optimizer. Some details of the algorithms is provided in Algorithm 1.

Algorithm 1: COMPUTING SERW WITH TRAINED ϕ AND ψ

Input: Source and target samples: (X, μ) and (Y, ν) , Epochs (E), batch sizes (B);

Output: the embeddings $\phi = (\mathbf{I} + g_X) \circ h_{\theta_X}$ and $\psi = (\mathbf{I} + g_Y) \circ h_{\theta_Y}$;

for $e=1$ **to** E **do**

for $b = 1$ **to** B **do**

 sample batches x and y from the two input spaces;

 get $x_{\text{emb}} \leftarrow \phi(x)$ and $y_{\text{emb}} \leftarrow \psi(y)$;

 get r^* and π^* by minimizing Equation (4) with using x_{emb} and y_{emb} ;

 update θ_X and θ_Y by maximizing Equation (4) using r^* and π^* ;

end for

end for

Return: $\pi^{(e)}$;

4.2 Toy example

In this example, we extracted randomly $n = m = 1000$ samples from MNIST, USPS and Fashion MNIST data sets, denoted by X , Y and Z . We compare GW distances between three possible matchings with the assorted SERW distances. We pre-process the data in order to fix the parameter $\underline{M}_{\mu,\nu}$ and $\overline{M}_{\mu,\nu}$ as discussed previously. We then vary the dimension of the embedded points from $\log^2 n$ up to the smallest dimension of the original samples. We perform the embeddings by using LLE followed by a non-linear embedding scheme aiming at minimizing the distance distortion as described in Section 4.1.

In Figure 2 we report plots of the distortion rate, the additive constant $\beta \overline{M}_{\mu,\nu}$ in the upper bound in Proposition 5, and the distance ratio of SERW for the three data sets X , Y and Z . As can be seen the rates decrease as the embedding dimension increase. Note that to determine the distortion coefficient for each given embedded dimension, we compute the quotient of the pairwise distances both in the original and embedding spaces. Thus, this high magnitudes of the upper bounds are due to a ‘‘crude’’ estimation of the distortion rate. One may explore a better estimation to reach a tighter upper bound. For this toy set, we investigate a useful property in our approach called *proximity preservation*, a property stating that:

$$\mathcal{GW}_2(\mu, \nu) \leq \mathcal{GW}_2(\mu, \eta) \Rightarrow \mathcal{S}_d(\mu, \nu) \leq \mathcal{S}_d(\mu, \eta).$$

In order to confirm this property, we compute the ratio between $\mathcal{S}_d(X, Y)/\mathcal{S}_d(Y, Z)$ and $\mathcal{S}_d(X, Z)/\mathcal{S}_d(Y, Z)$ for various embeddings and compare the resulting order with the quantities $\mathcal{GW}_2(X, Y)/\mathcal{GW}_2(Y, Z)$

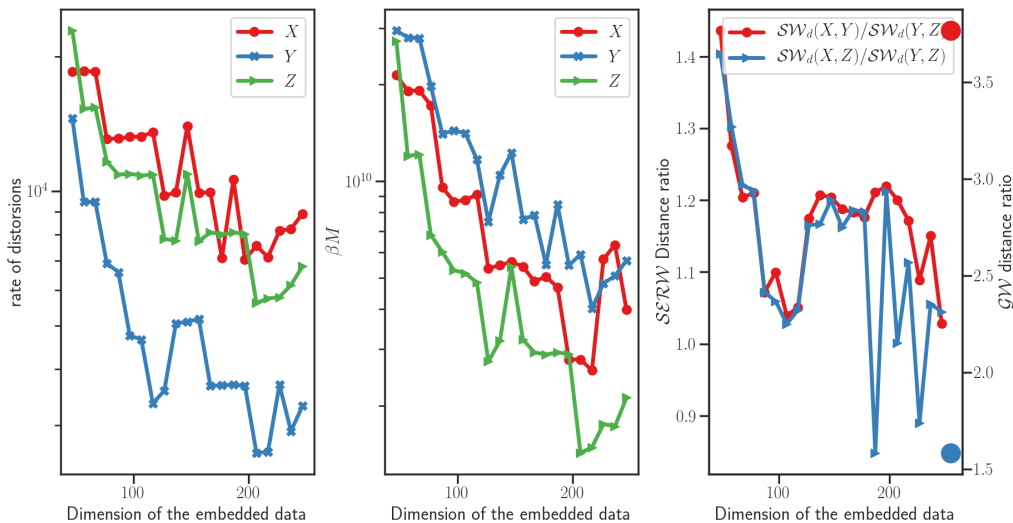


Figure 2: Plots of the distortion rates (left); various bounds in Proposition 5 (middle); GW cost and the distance ratio of SERW between the data points (right); as a function of the target dimension of embedded data d . The bold points in the right panel correspond to $\mathcal{GW}_2(X, Y)/\mathcal{GW}_2(Y, Z)$ (red) and $\mathcal{GW}_2(X, Z)/\mathcal{GW}_2(Y, Z)$ (blue). Note that the distance ratio (red curve) is upper than the blue one which highlights that the proximity preservation for SERW distance is fulfilled as the embeded dimension d goes up.

and $\mathcal{GW}_2(X, Z)/\mathcal{GW}_2(Y, Z)$. As seen in Figure 2, while the ratios vary their order is often preserved for large embedding dimensions.

4.3 Meshes comparison

GW distance is frequently used in computer graphics for computing correspondence between meshes. Those distances are then exploited for organizing shape collection, for instance for shape retrieval or search. One of the useful key property of GW distance for those applications is that it is isometry-invariant. In order to show that our proposed approach approximately satisfies this property, we reproduce an experiment already considered by Solomon et al. (2016) and Vayer et al. (2019).

We have at our disposal a time-series of 45 meshes of galloping horses. When comparing all meshes with the first one, the goal is to show that the distance presents a cyclic nature related to gallop cycle. Each mesh is composed of 8400 samples and in our case, we have embedded them into a 2-dimensional space using a multi-dimensional scaling algorithm followed by a non-linear embedding aiming at minimizing distortion as described in Section 4.1.

Figure 3 shows the (centered and max-normalized) distances between meshes we obtain with SERW and with a Sliced Gromov-Wasserstein (SGW) (Vayer et al., 2019) and a genuine GW distance. In both cases, due to the random aspect of the first two algorithms, distances are averaged over 10 runs. We note that our approach is able to recover the cyclic nature of the galloping motion as described by GW distance.

To quantify the cyclic nature of the galloping motion as described by GW, SERW and SGW, we compute the dynamic time warping measures (DTW) (Berndt and Clifford, 1994). DTW aims to find an optimal alignment between two given signals. Figure 4 illustrates that the couple distances (GW, SERW) and (GW, SGW) behave well according to the recovery of galloping motion. Indeed, the alignment suggested by DTW in both cases is almost diagonal showing that the motion curves provided by SERW and SGW are in accordance with the obtained motion curve via the plain GW. We further compute the mean and the standard deviation of DTW measures between GW and SERW one




Figure 3: GW, SGW, and SERW distances between 3D meshes of galloping horses. We can note that both SGW and SERW distances are able to retrieve the cyclic nature of the movement (see also Figure 4).

the one hand, and between GW and SGW distances on the other hand, over the 10 runs. We get the following: $DTW(GW, SERW) = 10.62 \pm 1.35$ and $DTW(GW, SGW) = 10.57 \pm 1.89$. This emphasizes the fact that SERW retrieves the cyclic nature of the movement as does by GW.

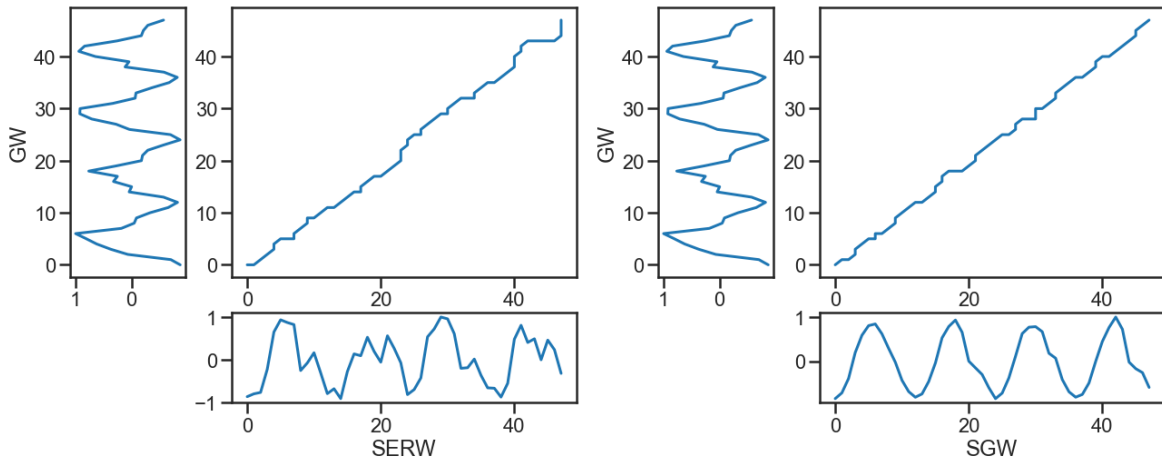


Figure 4: Alignment of the couple distances (GW, SERW) and (GW, SGW) between 3D meshes images as yielded by DTW measure. Optimal alignment is diagonal. On the right and bottom of each panel are the compared curves.

4.4 Text-Image alignment

To show that our proposed also provides relevant coupling when considering out-of-the-shelves embeddings, we present here results on aligning text and images distributions. The problem we address is related to identifying different states of objects, scene and materials (Isola et al., 2015). We have images labeled by some nouns modified by some adjectives describing state of the objects. In our

experiment, we want to show that our approach provides coupling between labels and images semantically meaningful as those obtained by a Gromov-Wasserstein approach. As for proof of concept, from the 115 available adjectives, we have considered only three of them *ruffled*, *weathered*, *engraved* and extracted all the classes associated with those adjectives. In total, we obtain 109 different classes of objects and about 525 images in total (as each class contains at most 5 objects).

The composed name (adjective + noun) of each label is embedded into \mathbb{R}^{100} using a word vector representation issued by `fasttext` model (Grave et al., 2018) trained on the first 1 billion bytes of English Wikipedia according to Mikolov et al. (2018). The 256×256 images have been embedded into a vector of dimension 4096 using a pre-trained VGG-16 model (Simonyan and Zisserman, 2015). These embeddings are extracted from the first dense layer of a VGG-16. The Gromov-Wasserstein distance of

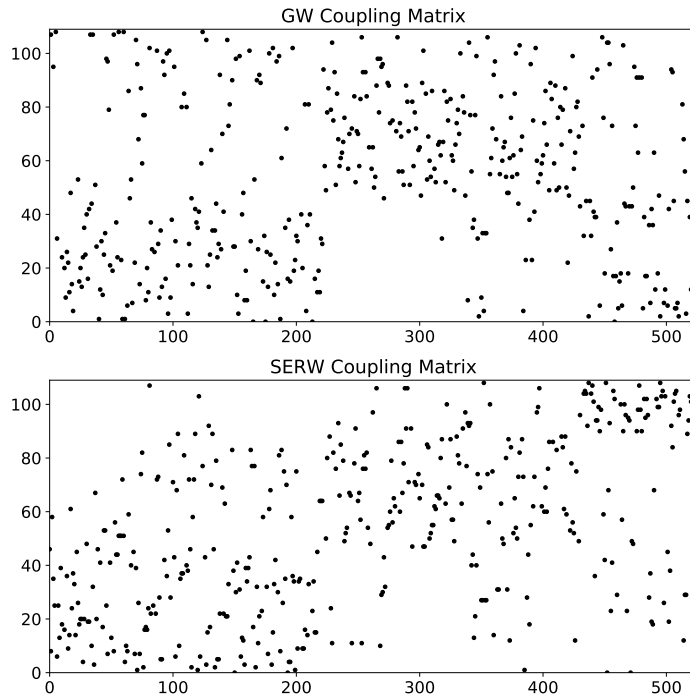


Figure 5: Coupling matrices between text and image embeddings. (top) Gromov-Wasserstein coupling matrix obtained in the original embedding spaces (bottom) our SERW coupling matrix after projecting embeddings into same dimension space.

those embeddings has been computed for coupling labels and images in the two different embedding spaces. For our SERW approach, we have further reduce the dimension of the image embeddings using Isomap with 100 dimensions. When computing the distance matrix, objects have been organized by class of adjectives for an easy visual inspection.

Figure 5 presents coupling matrices obtained using GW and our SERW. Since in both cases, the distance is not approximated by the Sinkhorn algorithm, the obtained matching is not smooth. Our results show that both GW and SERW distances are able to retrieve the 3 classes of adjectives and matches appropriate images with the relevant labels. Figure 6 illustrates the best matched images by GW and SERW (according to the transportation map) to the texts *Engraved Copper* and *Engraved Metal*. We can remark that in both cases GW and SERW do not suggest the same images. However, the retrieved images are meaningful according to the text queries. We shall notice that the embeddings

used by SERW do not distort the discriminative information, leading to interesting matched images as shown by the last row of Figure 6.

5 Conclusion

In this paper we introduced the SERW distance for distribution alignment lying in different mm-spaces. It is based on metric measure embedding of the original mm-spaces into a common Euclidean space and computes an optimal transport on the (low-distorted) embedded distributions. We prove that SERW defines a proper distance behaving like GW distance and we further show a cost relation between SERW and GW. Some of numerical experiments are tailored using a fully connected neural network to learn the maximization problem defining SERW, while other ones are conducted with fixed embeddings. In particular, SERW can be viewed as an embedding-dependent alignment for distributions coming from different mm-spaces, that its quality is strongly dependent on the given embeddings.

Acknowledgments

The works of Maxime Bérar, Gilles Gasso and Alain Rakotomamonjy have been supported by the OATMIL ANR-17-CE23-0012 Project of the French National Research Agency (ANR).



Figure 6: Best matched images obtained through GW transportation plan, and our SERW distance. The first block of images correspond to the class *Engraved Copper* and the second one to *Engraved Metal*. Within each block, the top row shows the results of GW and the bottom row illustrates the matching proposed by SERW.

A Proofs

In the proofs, we frequently use the two following lemmas. Lemma 7 writes an integration result using push-forward measures; it relates integrals with respect to a measure η and its push-forward under a measurable map $f : X \rightarrow Y$. Lemma 8 proves that the admissible set of couplings between the embedded measures are exactly the embedded of the admissible couplings between the original measures.

Lemma 7. *Let $f : S \rightarrow T$ be a measurable mapping, let η be a measurable measure on S , and let g be a measurable function on T . Then $\int_T g d f_{\#} \eta = \int_S (g \circ f) d \eta$.*

Lemma 8. *For all $\phi \in \mathcal{F}_d(X)$, $\psi \in \mathcal{F}_d(Y)$, $r \in \mathcal{R}_d$, and $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$, one has*

$$\Pi(\phi_{\#} \mu, (r \circ \psi)_{\#} \nu) = \{(\phi \otimes (r \circ \psi))_{\#} \pi \text{ s.t. } \pi \in \Pi(\mu, \nu)\}$$

where $\phi \otimes (r \circ \psi) : X \times Y \rightarrow X \times Y$ such that $(\phi \otimes (r \circ \psi))(x, y) = (\phi(x), r(\psi(y)))$ for all $x, y \in X \times Y$.

Proof of Lemma 8. See Paty and Cuturi (2019).

A.1 Proof of Proposition 1

To let it more readable, the proof is divided into 5 steps summarized as follows: first step is for metric property of $\mathcal{F}_d(X)$; second one shows completeness of $\mathcal{F}_d(X)$; third establishes the totally boundedness of $\mathcal{F}_d(X)$, namely that one can recover this space using balls centred on a finite number of embedding points; the last is a conclusion using Arzela-Ascoli's Theorem for characterizing compactness of subsets of functional continuous space.

Since the arguments of the proof are similar for the two spaces, we only focus on proving the topological property of $\mathcal{F}_d(X)$. Let us refresh the memories by some results in topology: we denote $\mathcal{C}(X, \mathbb{R}^d)$ the set of all continuous mappings of X into $(\mathbb{R}^d, \|\cdot\|)$ and recall the notions of totally boundedness in order to characterize the compactness of $(\mathcal{F}_d(X), \Gamma_X)$. The material here is taken from Kubrusly (2011) and O'Searcoid (2006).

Definition 6. *i) (Totally bounded) Let A be a subset of a metric space (S, d_S) . A subset A_ε of A is an ε -net for A if for every point s of A there exists a point t in A_ε such that $d(s, t) < \varepsilon$. A subset A of S is totally bounded (precompact) in (S, d_S) if for every real number $\varepsilon > 0$ there exists a finite ε -net for A .*
ii) (Pointwise totally bounded) A subset \mathcal{S} of $\mathcal{C}((S, d_S), (T, d_T))$ is pointwise totally bounded if for each s in S the set $\mathcal{S}(s) = \{f(s) \in T : f \in \mathcal{S}\}$ is totally bounded in T .
iii) (Equicontinuous) A subset \mathcal{S} of $\mathcal{C}(S, T)$ is equicontinuous at a point $s_0 \in S$ if for each $\varepsilon > 0$ there exists a $\delta > 0$ such that $d_T(f(s), f(s_0)) < \varepsilon$ whenever $d_S(s, s_0) < \delta$ for every $f \in \mathcal{S}$

Proposition 6. *If S is a metric space, then S is compact if and only if S is complete and totally bounded.*

Let $\mathcal{C}((S, d_S), (T, d_T))$ consisting of all continuous, bounded mappings of S into (T, d_T) , endowed with the supremum metric $d_\infty(f, g) = \sup_{s \in S} d_T(f(s), g(s))$. Proving the totally boundedness of some topological spaces may need more technical tricks. Fortunately, in our case we use Arzelà-Ascoli Theorem that gives compactness criteria for subspaces of $\mathcal{C}((S, d_S), (T, d_T))$ in terms of pointwise totally bounded and equicontinuous, namely they are a necessary and sufficient condition to guarantee that the totally boundedness of a subset \mathcal{S} in $(\mathcal{C}(S, T), d_\infty)$.

Theorem 1. *(Arzelà-Ascoli Theorem) If S is compact, then a subset of the metric space $\mathcal{C}((S, d_S), (T, d_T))$ is totally bounded if and only if it is pointwise totally bounded and equicontinuous.*

The proof is divided on 5 Steps:

• **Step 1.** $(\mathcal{F}_d(X), \Gamma_X)$ is a metric space. It is clear that for all $\phi, \phi' \in \mathcal{F}_d(X)$, $\Gamma_X(\phi, \phi') \geq 0$ (nonnegativeness) and $\Gamma_X(\phi, \phi') = 0$ if and only if $\phi = \phi'$. To verify the triangle inequality, we proceed as follows. Take an arbitrary $x \in X$ and note that, if ϕ, ϕ' , and ϕ'' are embeddings in $\mathcal{F}_d(X)$ then by triangle inequality in the Euclidean space \mathbb{R}^d .

$$\|\phi(x) - \phi'(x)\| \leq \|\phi(x) - \phi''(x)\| + \|\phi''(x) - \phi'(x)\| \leq \Gamma_X(\phi, \phi'') + \Gamma_X(\phi'', \phi'),$$

hence $\Gamma_X(\phi, \phi') \leq \Gamma_X(\phi, \phi'') + \Gamma_X(\phi'', \phi')$, and therefore $(\mathcal{F}_d(X), \Gamma_X)$ is a metric space.

• **Step 2.** $\mathcal{F}_d(X) \subset \mathcal{C}(X, \mathbb{R}^d)$. First recall that for each $\phi \in \mathcal{F}_d(X)$ is a τ_ϕ -embedding then it is Lipschitzian mapping. It is readily verified that every Lipschitzian mapping is uniformly continuous, that is for each real number $\varepsilon > 0$ there exists a real number $\delta > 0$ such that $d_X(x, x') < \delta$ implies $\|\phi(x) - \phi(x')\| < \varepsilon$ for all $x, x' \in X$. So it is sufficient to take $\delta = \frac{\varepsilon}{\tau_\phi}$.

• **Step 3.** $(\mathcal{F}_d(X), \Gamma_X)$ is complete. The proof of this step is classic in the topology literature of the continuous space endowed with the supremum metric. For the sake of completeness, we adapt it in our case. Let $\{\phi_k\}_{k \leq 1}$ be a Cauchy sequence in $(\mathcal{F}_d(X), \Gamma_X)$. Thus $\{\phi_k(x)\}_{k \leq 1}$ is a Cauchy sequence in $(\mathbb{R}^d, \|\cdot\|)$ for every $x \in X$. This can be as follows: $\|\phi_k(x) - \phi_{k'}(x)\| \leq \sup_{x \in X} \|\phi_k(x) - \phi_{k'}(x)\| = \Gamma_X(\phi_k, \phi_{k'})$ for each pair of integers k, k' and every $x \in X$, and hence $\{\phi_k(x)\}_{k \leq 1}$ converges in \mathbb{R}^d for every $x \in X$ (since \mathbb{R}^d is complete). Let $\phi(x) = \lim_{k \rightarrow \infty} \phi_k(x)$ for each $x \in X$ (i.e., $\phi_k(x) \rightarrow \phi(x)$ in \mathbb{R}^d , which defines a mapping ϕ of X into \mathbb{R}^d). We shall show that $\phi \in \mathcal{F}_d(X)$ and that $\{\phi_k\}$ converges to ϕ in $\mathcal{F}_d(X)$, thus proving that $(\mathcal{F}_d(X), \Gamma_X)$ is complete. Note that for any integer n and every pair of points x, x' in $\mathcal{F}_d(X)$, we have $\|\phi(x) - \phi(x')\| \leq \|\phi(x) - \phi_k(x)\| + \|\phi_k(x) - \phi_k(x')\| + \|\phi_k(x') - \phi(x')\|$ by the triangle inequality. Now take an arbitrary real number $\varepsilon > 0$. Since $\{\phi_k(x)\}_k$ is a Cauchy sequence in $(\mathcal{F}_d(X), \Gamma_X)$, it follows that there exists a positive integer $k_\varepsilon \in \mathbb{N}$ such that $\Gamma(\phi_k, \phi_{k'}) < \varepsilon$, and hence $\|\phi_k(x) - \phi_{k'}(x)\| < \varepsilon$ for all $x \in X$, whenever $k, k' \geq k_\varepsilon$. Moreover, since $\phi_k(x) \rightarrow \phi(x)$ in \mathbb{R}^d for every $x \in X$, and the Euclidean distance is a continuous function from the metric space \mathbb{R}^d to the metric space \mathbb{R} for each $y \in \mathbb{R}^d$, it also follows that $\|\phi(x) - \phi_k(x)\| = \|\lim_{k' \rightarrow \infty} \phi_{k'}(x) - \phi_k(x)\|$ for each positive integer k and every $x \in X$. Thus $\|\phi(x) - \phi_k(x)\| \leq \varepsilon$ for all $x \in X$ whenever $k \geq k_\varepsilon$. Furthermore, since each ϕ_k lies in $(\mathcal{F}_d(X), \Gamma_X)$, it follows that there exists a real number $\gamma(k_\varepsilon)$ such that $\sup d(\phi_{k_\varepsilon}(s), \phi_{k_\varepsilon}(x')) \leq \gamma(k_\varepsilon)$, $x, x' \in X$. Therefore, for any $\varepsilon > 0$ there exists a positive integer k_ε such that $\|\phi(x) - \phi(x')\| \leq 2\varepsilon + \gamma(k_\varepsilon)$ for all $x, x' \in X$ so that $\phi \in (\mathcal{F}_d(X), \Gamma_X)$, and $\Gamma_X(\phi, \phi') = \sup_{x \in X} \|\phi(x) - \phi'(x)\| \leq \varepsilon$, $x \in X$ whenever $k \geq k_\varepsilon$, so that ϕ_k converges to ϕ in $(\mathcal{F}_d(X), \Gamma_X)$.

• **Step 4.** $\mathcal{F}_d(X)$ is pointwise totally bounded and equicontinuous. From (iii) in Definition 6 and the details in **Step 3**, $\mathcal{F}_d(X)$ is readily equicontinuous. Next we shall prove that the subset $\{\hat{x}\} = \{\phi(x) \in \mathbb{R}^d : \phi \in \mathcal{F}_d(X)\}$ is totally bounded in \mathbb{R}^d . To proceed we use another result characterizing total boundedness that reads as:

(S, d_S) is totally bounded metric space if and only if every sequence in S has a Cauchy subsequence.

Since for any $\phi \in \mathcal{F}_d(X)$ is Lipschitzian then it is uniformly continuous as explained above. Furthermore uniformly continuous functions have some very nice conserving properties. They map totally bounded sets onto totally bounded sets and Cauchy sequences onto Cauchy sequences.

Now suppose that Suppose $\{y_l\}_{l \geq 1}$ is any sequence in $\{\hat{x}\} \subset \phi(X)$. For each $l \in \mathbb{N}$, the subset $X \cap \phi^{-1}(\{y_l\}) \subset X$ is non empty for each $l \in \mathbb{N}$ (Axiom of Countable Choice see O'Searcoid (2006)). Then $\phi(x_l) = y_l$ for each $l \in \mathbb{N}$. By the Cauchy criterion for total boundedness of X , the sequence $\{x_l\}_l$ has a Cauchy subsequence $\{x_{l_j}\}_{l_j}$. Then, by what we have just proved, $\{\phi(x_{l_j})\}_{l_j} = \{y_{l_j}\}_{l_j}$ is a Cauchy subsequence of $\{y_l\}$. Since $\{y_l\}$ is an arbitrary sequence in $\{\hat{x}\}$, $\{\hat{x}\}$ satisfies the Cauchy criterion for total boundedness and so is totally bounded.

• **Step 5.** $\mathcal{F}_d(X)$ is compact. Using Arzela-Ascoli Theorem 1 and **Step 2** we conclude that $\mathcal{F}_d(X)$ is totally bounded. Together with **Step 3** $\mathcal{F}_d(X)$ is compact.

A.2 Proof of Lemma 1

Notice that for $\mu \in \mathcal{P}_p(X)$, $\nu \in \mathcal{P}_p(Y)$, and $(\phi, \psi) \in \mathcal{F}_d(X) \times \mathcal{F}_d(Y)$, $r \in \mathcal{R}_d$ one has $W_2^2(\frac{1}{\sqrt{2}}\phi\#\mu, \frac{1}{\sqrt{2}}(r \circ \psi)\#\nu) < \infty$. It can be seen easily using the facts that

$$\int_{\mathbb{R}^d} \|u\|^2 d\frac{1}{\sqrt{2}}\phi\#\mu(u) = \frac{1}{2} \int_X \|\phi(x)\|^2 d\mu(x) \leq \frac{\tau_\phi^2}{2} M_2(\mu)$$

and

$$\int_{\mathbb{R}^d} \|v\|^2 d\frac{1}{\sqrt{2}}(r \circ \psi)\#\nu(v) = \frac{1}{2} \int_Y \|r(\psi(y))\|^2 d\nu(y) = \frac{1}{2} \int_Y \|\psi(y)\|^2 d\nu(y) \leq \frac{\tau_\psi^2}{2} M_2(\nu)$$

where $M_2(\mu) = \int_X \|x\|_X^2 d\mu(x) < \infty$ and $M_2(\nu) = \int_Y \|y\|_Y^2 d\nu(y) < \infty$, by Assumption 1. Now, thanks to Lemmas 7 and 8, we have

$$\begin{aligned} \mathcal{E}_d^2(\mu, \nu) &= \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \inf_{\gamma \in \Pi(\frac{1}{\sqrt{2}}\phi\#\mu, \frac{1}{\sqrt{2}}(r \circ \psi)\#\nu)} \int_{X \times Y} \|u - v\|^2 d\gamma(u, v) \\ &= \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \inf_{\pi \in \Pi(\mu, \nu)} \int_{X \times X} \|u - v\|^2 d(\frac{1}{\sqrt{2}}\phi \otimes \frac{1}{\sqrt{2}}(r \circ \psi))\#\pi(u, v) \\ &= \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \inf_{\pi \in \Pi(\mu, \nu)} \frac{1}{2} \int_{X \times Y} \|\phi(x) - r(\psi(y))\|^2 d\pi(x, y) \\ &= \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \frac{1}{2} \inf_{\pi \in \Pi(\mu, \nu)} J_{\phi, \psi, r, \pi}(\mu, \nu). \end{aligned}$$

A.3 Proof of Lemma 2

In one hand, for any fixed $\pi \in \Pi(\mu, \nu)$ the application $h_\pi : (\phi, \psi, r) \mapsto \int_{X \times Y} \|\phi(x) - r(\psi(y))\|^2 d\pi(x, y)$ is continuous. To show that, we use the continuity under integral sign Theorem. Indeed,

- for π -almost (x, y) , the mapping $(\phi, \psi, r) \mapsto \|\phi(x) - r(\psi(y))\|^2$ is continuous. To show that fix $\varepsilon > 0$, and $\phi, \psi, r, \phi_0, \psi_0, r_0 \in \mathcal{F}_d(X) \times \mathcal{F}_d(Y) \times \mathcal{R}_d$. We endow the product space $\mathcal{F}_d(X) \times \mathcal{F}_d(Y) \times \mathcal{R}_d$ by the metric $\Gamma_{X, Y}$ defined as follows:

$$\Gamma_{X, Y}((\phi, \psi), (\phi', \psi'), (r, r')) = \Gamma_X(\phi, \phi') + \Gamma_Y(r \circ \psi, r' \circ \psi')$$

We have

$$\begin{aligned} \left| \|\phi(x) - (r \circ \psi)(y)\|^2 - \|\phi_0(x) - (r_0 \circ \psi_0)(y)\|^2 \right| &\leq \|(\phi(x) - \phi_0(x)) - (r(\psi(y)) - r_0(\psi_0(y)))\|^2 \\ &\leq 2(\|\phi(x) - \phi_0(x)\|^2 + \|r(\psi(y)) - r_0(\psi_0(y))\|^2) \\ &\leq 2(\Gamma_X^2(\phi, \phi_0) + \Gamma_Y^2(r \circ \psi, r_0 \circ \psi_0)) \\ &\leq 2\Gamma_{X, Y}^2((\phi, \psi), (\phi_0, \psi_0), (r, r_0)). \end{aligned}$$

Letting $\delta_\varepsilon = \sqrt{\varepsilon/2}$, then if $\Gamma_{X, Y}((\phi, \psi), (\phi_0, \psi_0), (r, r_0)) < \delta_\varepsilon$, one has $|\|\phi(x) - r(\psi(y))\|^2 - \|\phi_0(x) - r_0(\psi_0(y))\|^2| < \varepsilon$. This yields that $\lim_{(\phi, \psi, r) \rightarrow (\phi_0, \psi_0, r_0)} \|\phi(x) - r(\psi(y))\|^2 = \|\phi_0(x) - r_0(\psi_0(y))\|^2$.

- for a fixed (ϕ, ψ, r) and $(x, y) \in X \times Y$, we have $\|\phi(x) - r(\psi(y))\|^2 \leq \|\phi(x)\|^2 + \|r(\psi(y))\|^2 \leq g(x, y) := \tau_\phi^2 \|x\|^2 + \tau_\psi^2 \|y\|^2$ with $\int_{X \times Y} g(x, y) d\pi(x, y) = \tau_\phi^2 \int_X \|x\|^2 d\mu(x) + \tau_\psi^2 \int_Y \|y\|^2 d\nu(y) < \infty$.

Therefore, the family $(h_\pi)_{\pi \in \Pi(\mu, \nu)}$ is continuous then it is upper semicontinuous. We know that the pointwise infimum of a family of upper semicontinuous functions is upper semicontinuous (see Lemma 2.41 in Aliprantis and Border (2006)). This entails $\inf_{\pi \in \Pi(\mu, \nu)} h_\pi$ is upper semicontinuous. Since the product of compact sets is a compact set (Tychonoff Theorem), then $\mathcal{F}_d(X) \times \mathcal{F}_d(Y)$ is compact, hence $\sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \inf_{\pi \in \Pi(\mu, \nu)} h_\pi(\phi, \psi, r)$ attains a maximum value (see Theorem 2.44 in Aliprantis and Border (2006)). So, there exists a couple of embeddings $(\phi^*, \psi^*) \in \mathcal{F}_d(X) \times \mathcal{F}_d(Y)$ and $\pi^* \in \Pi(\mu, \nu)$ such that $\sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \inf_{\pi \in \Pi(\mu, \nu)} h_\pi(\phi, \psi, r) = h_{\pi^*}(\phi^*, \psi^*, r)$ for all $r \in \mathcal{R}_d$. Finally, it is easy to show that $r \mapsto h_{\pi^*}(\phi^*, \psi^*, r)$ is continuous, hence the infimum over the orthogonal mappings \mathcal{R}_d (compact) exists.

A.4 Proof of Lemma 3

Let us recall first the minimax inequality:

Lemma 9. (*Minimax inequality*) Let $\Xi : \mathcal{U} \times \mathcal{V} \rightarrow \mathbb{R} \cup \{\pm\infty\}$ be a function. Then

$$\sup_{v \in \mathcal{V}} \inf_{u \in \mathcal{U}} \Xi(u, v) \leq \inf_{u \in \mathcal{U}} \sup_{v \in \mathcal{V}} \Xi(u, v).$$

Using minimax inequality, one has

$$\mathcal{E}_d^2(\mu, \nu) \leq \frac{1}{2} \inf_{r \in \mathcal{R}_d} \inf_{\pi \in \Pi(\mu, \nu)} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} J_{\phi, \psi, r, \pi}(\mu, \nu).$$

Note that for a fixed $\pi \in \Pi(\mu, \nu)$ and $r \in \mathcal{R}_d$ one has $\sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} J_{\phi, \psi, r, \pi}(\mu, \nu)$ exists (continuity of $\pi, r \mapsto \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} J_{\phi, \psi, r, \pi}(\mu, \nu)$ + compact set as shown in Proof of Lemma 2). Then

$$\inf_{r \in \mathcal{R}_d} \inf_{\pi \in \Pi(\mu, \nu)} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} J_{\phi, \psi, r, \pi}(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} J_{\phi, \psi, r, \pi}(\mu, \nu).$$

Thus $\mathcal{E}_d^2(\mu, \nu) \leq \mathcal{S}_d^2(\mu, \nu)$.

A.5 Proof of Lemma 4

As we proved in Lemma 2 that for any fixed $\pi \in \Pi(\mu, \nu)$, $h_\pi : (\phi, \psi, r) \mapsto \int_{X \times Y} \|\phi(x) - r(\psi(y))\|^2 d\pi(x, y)$ is continuous, then it is lower semicontinuous. The pointwise supremum of a family of lower semicontinuous functions is lower semicontinuous (Lemma 2.41 in Aliprantis and Border (2006)). Moreover, the pointwise infimum of a compact family of lower semicontinuous functions is lower semicontinuous (here \mathcal{R}_d is compact) then $\pi \mapsto \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \int_{X \times Y} \|\phi(x) - r(\psi(y))\|^2 d\pi(x, y)$ is lower semicontinuous. Furthermore $\Pi(\mu, \nu)$ is compact set with respect to the topology of narrow convergence (Villani, 2003), then $\inf_{\pi \in \Pi(\mu, \nu)} \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \int_{X \times Y} \|\phi(x) - r(\psi(y))\|^2 d\pi(x, y)$ exists (see Theorem 2.44 in Aliprantis and Border (2006)).

A.6 Proof of Proposition 2

• “ \Rightarrow ” Suppose that $\mathcal{S}_d(\mu, \nu) = 0$ then $\mathcal{E}_d(\mu, \nu) = 0$, that gives the Wasserstein distance $\mathcal{W}_2(\frac{1}{\sqrt{2}}\phi_{\#}\mu, \frac{1}{\sqrt{2}}(r \circ \psi)_{\#}\nu) = 0$ and hence $\phi_{\#}\mu = (r \circ \psi)_{\#}\nu$ for any $\phi, \psi \in \mathcal{F}_d(X) \times \mathcal{F}_d(Y)$ and $r \in \mathcal{R}_d$. Then for any $K \subseteq \mathbb{R}^d$ Borel, we have $\mu(\phi^{-1}(K)) = \nu((r \circ \psi)^{-1}(K))$. Recall that $X \subseteq \mathbb{R}^D$ and $Y \subseteq \mathbb{R}^{D'}$, then through the proof lines we regard to μ and ν as probability measures on \mathbb{R}^D and $\mathbb{R}^{D'}$, allowing us to use the following key result of Cramér and Wold (1936).

Theorem 2. (*Cramér and Wold, 1936*) Let γ, β be Borel probability measures on \mathbb{R}^D and agree at every open half-space of X . Then $\gamma = \beta$. In other words if, for $\omega \in \mathbb{S}^D = \{x \in \mathbb{R}^D : \|x\| = 1\}$ and $\alpha \in \mathbb{R}$ we write $H_{\omega, \alpha} = \{x \in \mathbb{R}^D : \langle \omega, x \rangle < \alpha\}$ and if $\gamma(H_{\omega, \alpha}) = \beta(H_{\omega, \alpha})$, for all $\omega \in \mathbb{S}^D$ and $\alpha \in \mathbb{R}$ then one has $\gamma = \beta$.

The fundamental Cramér-Wold theorem states that a Borel probability measure μ on \mathbb{R}^D is uniquely determined by the values it gives to halfspaces $H_{\omega,\alpha} = \{x \in \mathbb{R}^D : \langle \omega, x \rangle < \alpha\}$ for $\omega \in \mathbb{S}^D$ and $\alpha \in \mathbb{R}$. Equivalently, γ is uniquely determined by its one-dimensional projections $(\Delta_\omega)_\# \mu$, where Δ_ω is the projection $x \in \mathbb{R}^D \mapsto \langle x, \omega \rangle \in \mathbb{R}$ for $\omega \in \mathbb{S}^D$.

Straightforwardly, we have

$$\begin{aligned}
\phi_\#^{-1}((r \circ \psi)_\# \nu)(H_{\omega,\alpha}) &= (r \circ \psi)_\# \nu((\phi^{-1})^{-1}(H_{\omega,\alpha})) \\
&= (r \circ \psi)_\# \nu(\{u \in X : \phi^{-1}(u) \in H_{\omega,\alpha}\}) \\
&= (r \circ \psi)_\# \nu(\{u \in X : \langle w, \phi^{-1}(u) \rangle < \alpha\}) \\
&= \phi_\# \mu(\{u \in X : \langle w, \phi^{-1}(u) \rangle < \alpha\}) \text{ (by hypothesis)} \\
&= \mu(\phi^{-1}(\{u \in X : \langle w, \phi^{-1}(u) \rangle < \alpha\})) \\
&= \mu(\{x \in X : \phi(x) \in \{u \in X : \langle w, \phi^{-1}(u) \rangle < \alpha\}\}) \\
&= \mu(\{x \in X : \langle w, \phi^{-1}(\phi(x)) \rangle < \alpha\}) \\
&= \mu(\{x \in X : \langle w, x \rangle < \alpha\}) \text{ (since } \phi \text{ is one-to-one)} \\
&= \mu(H_{\omega,\alpha}).
\end{aligned}$$

Analogously, we prove that $(r \circ \psi)_\#^{-1}(\phi_\# \mu(H_{\omega,\alpha})) = \nu(H_{\omega,\alpha})$. Therefore, for all $A \subseteq X$ and $B \subseteq Y$ Borels, we have $\mu(A) = \phi_\#^{-1}((r \circ \psi)_\# \nu)(A)$ and $\nu(B) = (r \circ \psi)_\#^{-1}(\phi_\# \mu)(B)$.

• “ \Leftarrow ” Thanks to Lemma 4 in the core of the paper, there exists a couple of embeddings (ϕ^*, ψ^*) and $r^* \in \mathcal{R}_d$ optimum for $\mathcal{S}_d^2(\mu, \nu)$. We assume now that $\nu = ((r^* \circ \psi^*)^{-1} \circ \phi^*)_\# \mu$, then

$$\begin{aligned}
\mathcal{S}_d^2(\mu, \nu) &= \frac{1}{2} \inf_{\pi \in \Pi(\mu, (r^* \circ \psi^*)_\#^{-1}(\phi^*_\# \mu))} \int_{X \times Y} \|\phi^*(x) - r^*(\psi^*(y))\|^2 d\pi(x, y) \\
&= \frac{1}{2} \inf_{\pi \in \Pi(\mu, \phi^*_\# \mu)} \int_{X \times Y} \|\phi^*(x) - r^*(\psi^*(y))\|^2 d(I \otimes (r^* \circ \psi^*)^{-1})_\# \pi(x, y) \\
&= \frac{1}{2} \inf_{\pi \in \Pi(\mu, \mu)} \int_{X \times Y} \|\phi^*(x) - r^*(\psi^*(y))\|^2 d(I \otimes \phi^*)_\# ((I \otimes (r^* \circ \psi^*)^{-1})_\# \pi(x, y)).
\end{aligned}$$

On the other hand, it is clear that $(I \otimes \phi^*)_\# ((I \otimes (r^* \circ \psi^*)^{-1})_\# \pi)(\cdot) = (I \otimes \phi^* \circ (r^* \circ \psi^*)^{-1})_\# \pi(\cdot)$. Using the fact that ϕ^* is τ_{ϕ^*} -embedding then we get

$$\begin{aligned}
\mathcal{S}_d^2(\mu, \nu) &= \frac{1}{2} \inf_{\pi \in \Pi(\mu, \mu)} \int_{X \times Y} \|\phi^*(x) - r^*(\psi^*(y))\|^2 d(I \otimes \phi^* \circ (r^* \circ \psi^*)^{-1})_\# \pi(x, y) \\
&= \frac{1}{2} \inf_{\pi \in \Pi(\mu, \mu)} \int_{X \times X} \|\phi^*(x) - \phi^*(x')\|^2 d\pi(x, x') \\
&\leq \frac{\tau_{\phi^*}^2}{2} \inf_{\pi \in \Pi(\mu, \mu)} \int_{X \times X} d_X^2(x, x') d\pi(x, x') \\
&\leq \frac{\tau_{\phi^*}^2}{2} W_2^2(\mu, \mu) \\
&= 0.
\end{aligned}$$

A.7 Proof of Proposition 3

Symmetry is clear for both objects. In order to prove the triangle inequality, we use a classic lemma known as “gluing lemma” that allows to produce a sort of composition of two transport plans, as if they are maps.

Lemma 10. (Villani, 2003) Let X, Y, Z be three Polish spaces and let $\gamma^1 \in \mathcal{P}(X \times Y)$, $\gamma^2 \in \mathcal{P}(Y \times Z)$, be such that $\Delta_{\#}^Y \gamma^1 = \Delta_{\#}^Y \gamma^2$ where Δ^Y is the natural projection from $X \times Y$ (or $Y \times Z$) onto Y . Then there exists a measure $\gamma \in \mathcal{P}(X \times Y \times Z)$ such that $\Delta_{\#}^{X \times Y} \gamma = \gamma^1$ and $\Delta_{\#}^{Y \times Z} \gamma = \gamma^2$.

Let $\eta \in \mathcal{P}_2(Z)$ and $\pi^1 \in \Pi(\mu, \nu)$ and $\pi^2 \in \Pi(\nu, \eta)$. By the gluing lemma we know that there exists $\gamma \in \mathcal{P}_2(X \times Y \times Z)$ such that $\Delta_{\#}^{X \times Y} \gamma = \pi^1$ and $\Delta_{\#}^{Y \times Z} \gamma = \pi^2$. Since $\Delta_{\#}^X \gamma = \mu$ and $\Delta_{\#}^Z \gamma = \eta$, we have $\pi = \Delta_{\#}^{X \times Z} \gamma \in \Pi(\mu, \eta)$. On the other hand

$$\begin{aligned} & \int_{X \times Z} \|\phi(x) - \vartheta(\zeta(z))\|^2 d\pi(x, z) \\ &= \int_{X \times Y \times Z} \|\phi(x) - \vartheta(\zeta(z))\|^2 d\gamma(x, y, z) \\ &\leq 2 \int_{X \times Y \times Z} (\|\phi(x) - r(\psi(y))\|^2 + \|r(\psi(y)) - \vartheta(\zeta(z))\|^2) d\gamma(x, y, z) \\ &\leq 2 \int_{X \times Y \times Z} \|\phi(x) - r(\psi(y))\|^2 d\gamma(x, y, z) + 2 \int_{X \times Y \times Z} \|r(\psi(y)) - \vartheta(\zeta(z))\|^2 d\gamma(x, y, z) \\ &= 2 \int_{X \times Y} \|\phi(x) - r(\psi(y))\|^2 d\pi^1(x, y) + 2 \int_{Y \times Z} \|\tilde{\psi}(y) - \vartheta(\zeta(z))\|^2 d\pi^2(y, z), \end{aligned}$$

where $\tilde{\psi} = r \circ \psi \in \mathcal{F}_d(Y)$ ($\|r(\psi(y))\|_2 = \|\psi(y)\|_2, \forall y$). Hence, we end up with the desired result, $\mathcal{S}_d^2(\mu, \eta) \leq \mathcal{S}_d^2(\mu, \nu) + \mathcal{S}_d^2(\nu, \eta)$.

A.8 Proof of Proposition 4

As the embedding ϕ is Lipschitzian then it is continuous. Since X is compact hence $\phi(X)$ is also compact. Consequently $\text{supp}[\phi_{\#}\mu] \subset \phi(X)$ is compact (closed subset of a compact). The same observation is fulfilled by $\text{supp}[\psi_{\#}\nu] \subset \psi(Y)$. Letting $Z = \{\text{supp}[\phi_{\#}\mu] \cup \text{supp}[(r \circ \psi)_{\#}\nu]\} \subseteq \mathbb{R}^d$. Hence, $(Z, \|\cdot\|)$ is compact metric space and $\phi_{\#}\mu$ and $(r \circ \psi)_{\#}\nu$ are Borel probability measures on Z . Thanks to Theorem 5 (property (c)) in Mémoli (2011), we have that

$$\mathcal{W}_2^2\left(\frac{1}{\sqrt{2}}\phi_{\#}\mu, \frac{1}{\sqrt{2}}(r \circ \psi)_{\#}\nu\right) \geq \mathcal{GW}_2^2\left(\frac{1}{\sqrt{2}}\phi_{\#}\mu, \frac{1}{\sqrt{2}}(r \circ \psi)_{\#}\nu\right), \text{ for any } \phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y).$$

So

$$\mathcal{E}_d^2(\mu, \nu) \geq \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \mathcal{GW}_2^2\left(\frac{1}{\sqrt{2}}\phi_{\#}\mu, \frac{1}{\sqrt{2}}(r \circ \psi)_{\#}\nu\right).$$

Together with the minimax inequality we arrive at

$$\begin{aligned}
& \mathcal{S}_d^2(\mu, \nu) \\
& \geq \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \mathcal{GW}_2^2\left(\frac{1}{\sqrt{2}}\phi_{\#}\mu, \frac{1}{\sqrt{2}}(r \circ \psi)_{\#}\nu\right) \\
& \geq \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \frac{1}{2} \inf_{\gamma \in \Pi\left(\frac{1}{\sqrt{2}}\phi_{\#}, \frac{1}{\sqrt{2}}(r \circ \psi)_{\#}\right)} \iint_{Z \times Z} (\|u - u'\|_2 - \|v - v'\|_2)^2 d\gamma(u, v) d\gamma(u', v') \\
& \geq \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \frac{1}{4} \inf_{\pi \in \Pi(\mu, \nu)} \iint_{Z \times Z} (\|u - u'\|_2 - \|v - v'\|_2)^2 d(\phi \otimes (r \circ \psi))_{\#}\pi(u, v) d(\phi \otimes (r \circ \psi))_{\#}\pi(u', v') \\
& \geq \inf_{r \in \mathcal{R}_d} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \frac{1}{4} \inf_{\pi \in \Pi(\mu, \nu)} \iint_{X \times Y} (\|\phi(x) - \phi(x')\|_2 - \|r(\psi(y)) - r(\psi(y'))\|_2)^2 d\pi(x, y) d\pi(x', y') \\
& = \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \frac{1}{4} \inf_{\pi \in \Pi(\mu, \nu)} \iint_{X \times Y} (\|\phi(x) - \phi(x')\|_2 - \|\psi(y) - \psi(y')\|_2)^2 d\pi(x, y) d\pi(x', y') \\
& \geq \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} \frac{1}{4} \inf_{\pi \in \Pi(\mu, \nu)} \iint_{X \times Y} (d_X^2(x, x') + d_Y^2(y, y') - 2\tau_\phi\tau_\psi d_X(x, x')d_Y(y, y')) d\pi(x, y) d\pi(x', y') \\
& \geq \frac{1}{2}\mathcal{GW}_2^2(\mu, \nu) + \frac{1}{2} \sup_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} (1 - \tau_\phi\tau_\psi) \inf_{\pi \in \Pi(\mu, \nu)} \iint_{X \times Y} d_X(x, x')d_Y(y, y') d\pi(x, y) d\pi(x', y').
\end{aligned}$$

Using the fact that $-\sup -x = \inf x$, we get

$$\mathcal{GW}_2^2(\mu, \nu) \leq 2\mathcal{S}_d^2(\mu, \nu) + \inf_{\phi \in \mathcal{F}_d(X), \psi \in \mathcal{F}_d(Y)} (\tau_\phi\tau_\psi - 1)\mathcal{I}(\mu, \nu),$$

where $\mathcal{I}_1(\mu, \nu) := \inf_{\pi \in \Pi(\mu, \nu)} \iint_{X \times Y} d_X(x, x')d_Y(y, y') d\pi(x, y) d\pi(x', y')$. Using Bourgain's embedding theorem Bourgain (1985), $\tau_\phi \in [1, \mathcal{O}(\log n)]$ and $\tau_\psi \in [1, \mathcal{O}(\log m)]$, then

$$\mathcal{GW}_2^2(\mu, \nu) \leq 2\mathcal{S}_d^2(\mu, \nu) + \inf_{\tau_\phi \in \mathcal{D}_{\text{emb}}(X), \tau_\psi \in \mathcal{D}_{\text{emb}}(Y)} (\tau_\phi\tau_\psi - 1)\mathcal{I}_1(\mu, \nu).$$

In another hand, we have

$$\begin{aligned}
\mathcal{I}_1(\mu, \nu) &= \inf_{\pi \in \Pi(\mu, \nu)} \iint_{X \times Y} d_X(x, x')d_Y(y, y') d\pi(x, y) d\pi(x', y') \\
&\leq \iint_{X \times Y} d_X(x, x')d_Y(y, y') d\mu(x) d\nu(y) d\mu(x') d\nu(y') \\
&\leq \int_{X \times X} d_X(x, x') d\mu(x) d\mu(x') \int_{Y \times Y} d_Y(y, y') d\nu(y) d\nu(y') \\
&\leq 4 \left(\int_X d_X(x, 0) d\mu(x) + \int_Y d_Y(y, 0) d\nu(y) \right) \\
&\leq 4 \left(\int_X \|x\|_X d\mu(x) + \int_Y \|y\|_Y d\nu(y) \right) \\
&\leq 4(M_1(\mu) + M_1(\nu)),
\end{aligned}$$

$M_1(\mu) = \int_X \|x\|_X d\mu(x) < \infty$. Hence,

$$\frac{1}{2}\mathcal{GW}_2^2(\mu, \nu) \leq \mathcal{S}_d^2(\mu, \nu) + 2 \inf_{\tau_\phi \in \mathcal{D}_{\text{emb}}(X), \tau_\psi \in \mathcal{D}_{\text{emb}}(Y)} (\tau_\phi\tau_\psi - 1)(M_1(\mu) + M_1(\nu)).$$

A.9 Proof of Proposition 5

The proof of this proposition is based on a lower bound for the Gromov-Wasserstein distance (Proposition 6.1 in Mémoli (2011)):

$$\mathcal{GW}_2^2(\mu, \nu) \geq \text{FLB}_2^2(\mu, \nu) := \frac{1}{2} \inf_{\pi \in \Pi(\mu, \nu)} \int_{X \times Y} |s_{X,2}(x) - s_{Y,2}(y)|^2 d\pi(x, y),$$

where $s_{X,2} : X \rightarrow \mathbb{R}_+$, $s_{X,2}(x) = \left(\int_X d_X^2(x, x') d\mu(x') \right)^{1/2}$ defines an eccentricity function. Note that FLB_2^2 leads to a mass transportation problem for the cost $c(x, y) := |s_{X,2}(x) - s_{Y,2}(y)|^2$.

Now, for any $x, y \in X \times Y$, and $\phi, \psi \in \mathcal{F}_d(X) \times \mathcal{F}_d(Y)$, $r \in \mathcal{R}_d$ we have (by triangle inequality)

$$\begin{aligned} & \|\phi(x) - r(\psi(y))\|_2^2 \\ &= \int_{X \times Y} \|\phi(x) - r(\psi(y))\|_2^2 d\mu(x') d\nu(y') \\ &\leq 4 \int_X \|\phi(x) - \phi(x')\|_2^2 d\mu(x') + 4 \int_Y \|r(\psi(y)) - r(\psi(y'))\|_2^2 d\nu(y') \\ &\quad + 2 \int_{X \times Y} \|\phi(x') - r(\psi(y'))\|_2^2 d\mu(x') d\nu(y') \\ &\leq 4\tau_\phi^2 \int_X d_X^2(x, x') d\mu(x') + 4\tau_\psi^2 \int_Y d_Y^2(y, y') d\nu(y') + 2 \int_{X \times Y} \|\phi(x') - r(\psi(y'))\|_2^2 d\mu(x') d\nu(y') \\ &\leq 4(\tau_\phi^2 + \tau_\psi^2) \left(\int_X d_X^2(x, x') d\mu(x') + \int_Y d_Y^2(y, y') d\nu(y') \right) \\ &\quad - 2 \left(\int_X d_X^2(x, x') d\mu(x') \right)^{1/2} \left(\int_Y d_Y^2(y, y') d\nu(y') \right)^{1/2} \\ &\quad + (\tau_\phi^2 + \tau_\psi^2) \left(\int_X d_X^2(x, x') d\mu(x') \right)^{1/2} \left(\int_Y d_Y^2(y, y')^2 d\nu(y') \right)^{1/2} \\ &\quad + 2 \int_{X \times Y} \|\phi(x') - r(\psi(y'))\|_2^2 d\mu(x') d\nu(y') \\ &\leq 4(\tau_\phi^2 + \tau_\psi^2) |s_{X,2}(x) - s_{Y,2}(y)|^2 + 8(\tau_\phi^2 + \tau_\psi^2) \sqrt{\mathcal{I}_{2,x,y}(\mu, \nu)} \\ &\quad + 2 \int_{X \times Y} \|\phi(x') - r(\psi(y'))\|_2^2 d\mu(x') d\nu(y'), \end{aligned}$$

where

$$\mathcal{I}_{2,x,y}(\mu, \nu) := \left(\int_X d_X^2(x, x') d\mu(x') \right) \left(\int_Y d_Y^2(y, y') d\nu(y') \right).$$

We observe that

$$\mathcal{I}_{2,x,y}(\mu, \nu) \leq 4(M_2(\mu) + d_X^2(x, 0))(M_2(\nu) + d_Y^2(y, 0)).$$

Moreover,

$$\int_{X \times Y} \|\phi(x') - r(\psi(y'))\|_2^2 d\mu(x') d\nu(y') \leq 2(\tau_\phi^2 + \tau_\psi^2)(M_2(\mu) + M_2(\nu)).$$

Therefore, for any $\pi \in \Pi(\mu, \nu)$

$$\begin{aligned}
& \int_{X \times Y} \|\phi(x) - r(\psi(y))\|_2^2 d\pi(x, y) \\
& \leq 2(\tau_\phi^2 + \tau_\psi^2) \int_{X, Y} |s_{X,2}(x) - s_{Y,2}(y)|^2 d\pi(x, y) \\
& \quad + 8(\tau_\phi^2 + \tau_\psi^2) \int_{X \times Y} \sqrt{4(M_2(\mu) + d_X^2(x, 0))(M_2(\nu) + d_Y^2(y, 0))} d\pi(x, y) \\
& \quad + 2(\tau_\phi^2 + \tau_\psi^2)(M_2(\mu) + M_2(\nu)) \\
& \leq 4(\tau_\phi^2 + \tau_\psi^2) \int_{X, Y} |s_{X,2}(x) - s_{Y,2}(y)|^2 d\pi(x, y) \\
& \quad + 16(\tau_\phi^2 + \tau_\psi^2) \int_X \sqrt{(M_2(\mu) + d_X^2(x, 0))} d\mu(x) \int_Y \sqrt{(M_2(\nu) + d_Y^2(y, 0))} d\nu(y) \\
& \quad + 2(\tau_\phi^2 + \tau_\psi^2)(M_2(\mu) + M_2(\nu)).
\end{aligned}$$

Note that

$$\int_X \sqrt{M_2(\mu) + d_X^2(x, 0)} d\mu(x) \leq \sqrt{M_2(\mu)} + \int_X d_X(x, 0) d\mu(x) \leq \sqrt{M_2(\mu)} + \sqrt{M_1(\mu)},$$

and

$$\int_Y \sqrt{M_2(\nu) + d_Y^2(y, 0)} d\nu(y) \leq \sqrt{M_2(\nu)} + \int_Y d_Y(y, 0) d\nu(y) \leq \sqrt{M_2(\nu)} + \sqrt{M_1(\nu)}.$$

So

$$\begin{aligned}
\int_{X \times Y} \|\phi(x) - r(\psi(y))\|_2^2 d\pi(x, y) & \leq 4(\tau_\phi^2 + \tau_\psi^2) \int_{X, Y} |s_{X,2}(x) - s_{Y,2}(y)|^2 d\pi(x, y) \\
& \quad + 16(\tau_\phi^2 + \tau_\psi^2)(\sqrt{M_2(\mu)} + \sqrt{M_1(\mu)})(\sqrt{M_2(\nu)} + \sqrt{M_1(\nu)}) \\
& \quad + 2(\tau_\phi^2 + \tau_\psi^2)(M_2(\mu) + M_2(\nu)).
\end{aligned}$$

Finally,

$$\mathcal{S}_d^2(\mu, \nu) \leq 2 \sup_{\tau_\phi \in \mathcal{D}_{\text{emb}}(X), \tau_\psi \in \mathcal{D}_{\text{emb}}(Y)} (\tau_\phi^2 + \tau_\psi^2)(\mathcal{GW}_2^2(\mu, \nu) + \bar{M}_{\mu, \nu}),$$

where

$$\bar{M}_{\mu, \nu} = 8(\sqrt{M_2(\mu)} + \sqrt{M_1(\mu)})(\sqrt{M_2(\nu)} + \sqrt{M_1(\nu)}) + (M_2(\mu) + M_2(\nu)).$$

This finishes the proof.

A.10 Proof of Lemma 5

Since \mathcal{R}_d is compact set and the mapping $r \mapsto J_{\phi_f, \psi_f, r, \pi}(\mu, \nu)$ is continuous, then there exists $r_f \in \mathcal{R}_d$ such that $\inf_{r \in \mathcal{R}_d} J_{\phi_f, \psi_f, r, \pi}(\mu, \nu) = J_{\phi_f, \psi_f, r_f, \pi}(\mu, \nu)$. Using Lemma 8, we then get

$$\begin{aligned}
\widetilde{\mathcal{S}}_d^2(\mu, \nu) & = \frac{1}{2} \inf_{\pi \in \Pi(\mu, \nu)} J_{\phi_f, \psi_f, r_f, \pi}(\mu, \nu) \\
& = \frac{1}{2} \inf_{\pi \in \Pi(\mu, \nu)} \int_{X \times Y} \|\phi_f(x) - r_f(\psi_f(y))\|^2 d\pi(x, y) \\
& = \inf_{\gamma \in \Pi(\frac{1}{\sqrt{2}}(\phi_f) \# \mu, \frac{1}{\sqrt{2}}(r_f \circ \psi_f) \# \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|u - v\|^2 d\gamma(u, v) \\
& = \mathcal{W}_2^2(\mu_f, \nu_f),
\end{aligned}$$

where $\mu_f = (\frac{1}{\sqrt{2}}\phi_f)_\# \mu$ and $\nu_f = (\frac{1}{\sqrt{2}}r_f \circ \psi_f)_\# \nu$. Therefore, $\widetilde{\mathcal{S}}_d^2(\mu, \nu)$ is the 2-Wasserstein distance between μ_f and ν_f . Hence $\widetilde{\mathcal{S}}_d^2(\mu, \nu) = 0$ if and only if $\mu_f = \nu_f$ that is $\phi_f \# \mu = (r_f \circ \psi_f)_\# \nu$. On the other hand, one has

$$\mu = (\phi_f^{-1} \circ \phi_f)_\# \mu = (\phi_f^{-1})_\# (\phi_f \# \mu) = (\phi_f^{-1})_\# ((r_f \circ \psi_f)_\# \nu) = (\phi_f^{-1} \circ (r_f \circ \psi_f))_\# \nu.$$

The triangle inequality follows the same lines as proof of Proposition 3.

References

- Alaya, M. Z., M. Bézar, G. Gasso, and A. Rakotomamonjy (2019). Screening Sinkhorn algorithm for regularized optimal transport. In H. Wallach, H. Larochelle, A. Beygelzimer, F. Alché-Buc, E. Fox, and R. Garnett (Eds.), *Advances in Neural Information Processing Systems 32*, pp. 12169–12179. Curran Associates, Inc.
- Aliprantis, C. and K. Border (2006). *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Springer Berlin Heidelberg.
- Altschuler, J., J. Weed, and P. Rigollet (2017). Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), *Advances in Neural Information Processing Systems 30*, pp. 1964–1974. Curran Associates, Inc.
- Alvarez-Melis, D. and T. Jaakkola (2018). Gromov–Wasserstein alignment of word embedding spaces. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 1881–1890. Association for Computational Linguistics.
- Alvarez-Melis, D., S. Jegelka, and T. S. Jaakkola (2019, 16–18 Apr). Towards optimal transport with global invariances. In K. Chaudhuri and M. Sugiyama (Eds.), *Proceedings of Machine Learning Research*, Volume 89 of *Proceedings of Machine Learning Research*, pp. 1870–1879. PMLR.
- Arjovsky, M., S. Chintala, and L. Bottou (2017). Wasserstein generative adversarial networks. In D. Precup and Y. W. Teh (Eds.), *Proceedings of the 34th International Conference on Machine Learning*, Volume 70 of *Proceedings of Machine Learning Research*, International Convention Centre, Sydney, Australia, pp. 214–223. PMLR.
- Berndt, D. J. and J. Clifford (1994). Using dynamic time warping to find patterns in time series. In *Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining, AAAIWS’94*, pp. 359–370. AAAI Press.
- Bonneel, N., M. van de Panne, S. Paris, and W. Heidrich (2011). Displacement interpolation using lagrangian mass transport. *ACM Trans. Graph.* 30(6), 158:1–158:12.
- Bourgain, J. (1985). On lipschitz embedding of finite metric spaces in hilbert space. *Israel Journal of Mathematics* 52.
- Brinkman, B. and M. Charikar (2005, September). On the impossibility of dimension reduction in l_1 . *J. ACM* 52(5), 766–788.
- Bronstein, A. M., M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro (2010). A Gromov–Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *International Journal of Computer Vision* 89(2-3), 266–286.

- Bunne, C., D. Alvarez-Melis, A. Krause, and S. Jegelka (2019, 09–15 Jun). Learning generative models across incomparable spaces. In K. Chaudhuri and R. Salakhutdinov (Eds.), *Proceedings of the 36th International Conference on Machine Learning*, Volume 97 of *Proceedings of Machine Learning Research*, Long Beach, California, USA, pp. 851–861. PMLR.
- Chowdhury, S. and F. Mémoli (2018). The Gromov–Wasserstein distance between networks and stable network invariants. *CoRR abs/1808.04337*.
- Courty, N., R. Flamary, and M. Ducoffe (2018). Learning Wasserstein embeddings. In *International Conference on Learning Representations (ICLR)*.
- Courty, N., R. Flamary, D. Tuia, and A. Rakotomamonjy (2017). Optimal transport for domain adaptation. *IEEE transactions on pattern analysis and machine intelligence* 39(9), 1853–1865.
- Cramér, H. and H. Wold (1936). Some theorems on distribution functions. *Journal of the London Mathematical Society s1-11* (4), 290–294.
- Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems 26*, pp. 2292–2300. Curran Associates, Inc.
- Frogner, C., C. Zhang, H. Mobahi, M. Araya, and T. A. Poggio (2015). Learning with a Wasserstein loss. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Eds.), *Advances in Neural Information Processing Systems 28*, pp. 2053–2061. Curran Associates, Inc.
- Genevay, A., M. Cuturi, G. Peyré, and F. Bach (2016). Stochastic optimization for large-scale optimal transport. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.), *Advances in Neural Information Processing Systems 29*, pp. 3440–3448. Curran Associates, Inc.
- Gower, J. C. (1975). Generalized procrustes analysis. *Psychometrika* 40(1), 33–51.
- Grave, E., P. Bojanowski, P. Gupta, A. Joulin, and T. Mikolov (2018). Learning word vectors for 157 languages. *arXiv preprint arXiv:1802.06893*.
- Grave, E., A. Joulin, and Q. Berthet (2019, 16–18 Apr). Unsupervised alignment of embeddings with wasserstein procrustes. In K. Chaudhuri and M. Sugiyama (Eds.), *Proceedings of Machine Learning Research*, Volume 89 of *Proceedings of Machine Learning Research*, pp. 1880–1890. PMLR.
- Gromov, M., J. Lafontaine, and P. Pansu (1999). *Metric Structures for Riemannian and Non-Riemannian Spaces*. Progress in Mathematics - Birkhäuser. Birkhäuser.
- Grover, A. and J. Leskovec (2016). Node2Vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 855–864.
- Ho, N., X. L. Nguyen, M. Yurochkin, H. H. Bui, V. Huynh, and D. Phung (2017). Multilevel clustering via Wasserstein means. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML’17, pp. 1501–1509. JMLR.org.
- Isola, P., J. J. Lim, and E. H. Adelson (2015). Discovering states and transformations in image collections. In *CVPR*.
- Johnson, W. B. and J. Lindenstrauss (1984). Extensions of lipschitz mappings into hilbert space. *Contemporary mathematics* 26, 189–206.

- Kantorovich, L. (1942). On the transfer of masses (in russian). *Doklady Akademii Nauk* 2, 227–229.
- Kolouri, S., S. R. Park, M. Thorpe, D. Slepcev, and G. K. Rohde (2017, July). Optimal mass transport: Signal processing and machine-learning applications. *IEEE Signal Processing Magazine* 34(4), 43–59.
- Kruskal, J. B. and M. Wish (1978). *Multidimensional scaling*. Number 11. Sage.
- Kubrusly, C. (2011). *Elements of Operator Theory*. Birkhäuser Boston.
- Kusner, M., Y. Sun, N. Kolkin, and K. Weinberger (2015, 07–09 Jul). From word embeddings to document distances. In F. Bach and D. Blei (Eds.), *Proceedings of the 32nd International Conference on Machine Learning*, Volume 37 of *Proceedings of Machine Learning Research*, Lille, France, pp. 957–966. PMLR.
- Matoušek, J. (1996). On the distortion required for embedding finite metric spaces into normed spaces. *Israel Journal of Mathematics* 93(1), 333–344.
- Matoušek, J. (2002). *Embedding Finite Metric Spaces into Normed Spaces*, pp. 355–400. New York, NY: Springer New York.
- Matoušek, J. (2013). Lecture notes on metric embeddings. *Technical Report*.
- Mémoli, F. (2008). Gromov–Hausdorff distances in Euclidean spaces. In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1–8.
- Mémoli, F. (2011). Gromov–Wasserstein distances and the metric approach to object matching. *Foundations of Computational Mathematics* 11(4), 417–487.
- Mikolov, T., E. Grave, P. Bojanowski, C. Puhersch, and A. Joulin (2018). Advances in pre-training distributed word representations. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- Mikolov, T., I. Sutskever, K. Chen, G. S. Corrado, and J. Dean (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119.
- Monge, G. (1781). Mémoire sur la théotie des déblais et des remblais. *Histoire de l’Académie Royale des Sciences*, 666–704.
- Narayanan, A., M. Chandramohan, R. Venkatesan, L. Chen, Y. Liu, and S. Jaiswal (2017). Graph2Vec: Learning distributed representations of graphs. *arXiv preprint arXiv:1707.05005*.
- O’Searcoid, M. (2006). *Metric Spaces*. Springer Undergraduate Mathematics Series. Springer London.
- Paty, F.-P. and M. Cuturi (2019). Subspace robust Wasserstein distances. In K. Chaudhuri and R. Salakhutdinov (Eds.), *Proceedings of the 36th International Conference on Machine Learning*, Volume 97 of *Proceedings of Machine Learning Research*, Long Beach, California, USA, pp. 5072–5081. PMLR.
- Peyré, G., M. Cuturi, and J. Solomon (2016). Gromov–Wasserstein averaging of kernel and distance matrices. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML’16*, pp. 2664–2672. JMLR.org.
- Peyré, G. and M. Cuturi (2019). Computational optimal transport. *Foundations and Trends® in Machine Learning* 11(5-6), 355–607.

- Roweis, S. T. and L. K. Saul (2000). Nonlinear dimensionality reduction by locally linear embedding. *science* 290(5500), 2323–2326.
- Schilling, R. L. (2005). *Measures, Integrals and Martingales*. Cambridge University Press.
- Simonyan, K. and A. Zisserman (2015). Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*.
- Solomon, J., F. de Goes, G. Peyré, M. Cuturi, A. Butscher, A. Nguyen, T. Du, and L. Guibas (2015). Convolutional Wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Trans. Graph.* 34(4), 66:1–66:11.
- Solomon, J., G. Peyré, V. G. Kim, and S. Sra (2016). Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics (TOG)* 35(4), 1–13.
- Sturm, K. T. (2006). On the geometry of metric measure spaces. ii. *Acta Math.* 196(1), 133–177.
- Tenenbaum, J. B., V. De Silva, and J. C. Langford (2000). A global geometric framework for nonlinear dimensionality reduction. *science* 290(5500), 2319–2323.
- Vayer, T., R. Flamary, N. Courty, R. Tavenard, and L. Chapel (2019). Sliced Gromov–Wasserstein. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alché-Buc, E. Fox, and R. Garnett (Eds.), *Advances in Neural Information Processing Systems 32*, pp. 14726–14736. Curran Associates, Inc.
- Villani, C. (2003). *Topics in Optimal Transportation*. Graduate studies in mathematics. American Mathematical Society.
- Villani, C. (2009). *Optimal Transport: Old and New*, Volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer Berlin Heidelberg.
- Xu, H., D. Luo, and L. Carin (2019). Scalable Gromov–Wasserstein learning for graph partitioning and matching. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pp. 3046–3056.