

Reversible Gromov-Monge Sampler for Simulation-Based Inference

YoonHaeng Hur, Wenxuan Guo, and Tengyuan Liang

University of Chicago

September 22, 2022

Abstract

This paper introduces a new simulation-based inference procedure to model and sample from multi-dimensional probability distributions given access to i.i.d. samples, circumventing the usual approaches of explicitly modeling the density function or designing Markov chain Monte Carlo. Motivated by the seminal work on distance and isomorphism between metric measure spaces, we propose a new notion called the Reversible Gromov-Monge (RGM) distance and study how RGM can be used to design new transform samplers to perform simulation-based inference. Our RGM sampler can also estimate optimal alignments between two heterogeneous metric measure spaces $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ from empirical data sets, with estimated maps that approximately push forward one measure μ to the other ν , and vice versa. We study the analytic properties of the RGM distance and derive that under mild conditions, RGM equals the classic Gromov-Wasserstein distance. Curiously, drawing a connection to Brenier’s polar factorization, we show that the RGM sampler induces bias towards strong isomorphism with proper choices of $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$. Statistical rate of convergence, representation, and optimization questions regarding the induced sampler are studied. Synthetic and real-world examples showcasing the effectiveness of the RGM sampler are also demonstrated.

Keywords— Gromov-Wasserstein metric, transform sampling, simulation-based inference, generative models, isomorphism, likelihood-free inference

1 Introduction

One of the central tasks in statistics is to model and sample from a multi-dimensional probability distribution. Classic statistics approaches this problem by fitting a model to the target distribution and then sampling from a fitted model via Markov Chain Monte Carlo (MCMC) techniques. Although such model-based methods are widely used, MCMC sampling often entails several technicalities. Beyond diagnosing whether the chain mixes, obtaining i.i.d. samples from MCMC methods is complex as one has to control correlations between successive samples or run parallel chains.

An alternative approach available in statistics, reserved for the one-dimensional case, is usually referred to as the (inverse) *transform sampling*. Such an approach circumvents the calling for a parametric or nonparametric density and directly designs a sampler by transforming a simple uniform distribution. The idea is simple: one can transform a uniform measure $\mu = \text{Unif}([0, 1])$ to any one-dimensional target probability measure ν leveraging the following monotonic transformation $T: [0, 1] \rightarrow \mathbb{R}$ called the inverse Cumulative Distribution Function (CDF),

$$T(x) = \inf\{y \in \mathbb{R} : \nu((-\infty, y]) \geq x\}. \quad (1.1)$$

Define the pushforward measure $T_{\#}\mu$ by $T_{\#}\mu(S) = \mu(\{x : T(x) \in S\})$ for any Borel set $S \subseteq \mathbb{R}$, then one can easily check that $T_{\#}\mu = \nu$; namely, with a draw from the one-dimensional uniform distribution $x \sim \mu$, the transformed sample $T(x)$ has the target probability distribution ν .

The transform sampling idea can be extended to the multi-dimensional setting: given a target probability measure ν supported on \mathcal{Y} , one can specify a probability measure μ on \mathcal{X} , which is easy to sample from such as a multivariate Gaussian, and then find a measurable map $T: \mathcal{X} \rightarrow \mathcal{Y}$ such that $T_{\#}\mu = \nu$, where the pushforward measure $T_{\#}\mu$ is defined analogously to the one-dimensional case above. Such a map T , named as *transport map* from μ to ν , transforms i.i.d. samples from μ into i.i.d. samples from ν . Over the past few years, the generative modeling literature has been actively employing such transform sampling ideas by identifying $T_{\#}\mu = \nu$ through the following minimization:

$$\min_{T \in \mathcal{F}} \mathcal{L}(T_{\#}\mu, \nu), \quad (1.2)$$

where \mathcal{F} is a class of maps from \mathcal{X} to \mathcal{Y} parametrized by neural networks and \mathcal{L} measures certain discrepancies between two distributions. Different choices of \mathcal{L} have led to various models such as the Jensen-Shannon divergence for Generative Adversarial Networks (GANs) [24], the Wasserstein-1 distance for Wasserstein-GAN [2], and the Maximum Mean Discrepancy (MMD) for MMD-GAN [18, 32]. One caveat is that there can be infinitely many transport maps from μ to ν ; for instance, when $\mu = \nu = \text{Unif}([0, 1])$, define $T: [0, 1] \rightarrow [0, 1]$ by $T(x) = |2x - 1|$, then the n -fold compositions of T are valid transport maps for all $n \in \mathbb{N}$. In other words, finding a map T satisfying $T_{\#}\mu = \nu$ is an over-identified problem, where (1.2) has infinitely many minimizers. Though all minimizers are equivalent in terms of transform sampling, not all are equally preferred in light of Occam’s razor principle: one wishes to select simple, desirable transport maps among the over-identified set $\{T : T_{\#}\mu = \nu\}$.

Inductive biases solve the aforementioned over-identified problem by restricting the search to transport maps with desirable properties. In this context, there have been meaningful progresses based on optimal transport (OT) theory [53, 36]. The OT theory aims to identify an optimal transformation T , quantified by the transportation cost of moving mass from μ to ν ; for instance, when μ and ν lie in the same space \mathbb{R}^d , each transport map T is associated with the transport cost $C(T) := \int_{\mathbb{R}^d} \|x - T(x)\|^2 d\mu(x)$. Brenier [9] proved that, under mild regularity conditions, there exists a unique minimizer T^* of C among all transport maps, namely,

$$T^* = \arg \min_{T_{\#}\mu = \nu} C(T). \quad (1.3)$$

More importantly, T^* is the gradient of some convex function. On the one hand, Brenier’s result extends the one-dimensional (inverse) transform sampling to the multi-dimensional case. When $d = 1$ and $\mu = \text{Unif}([0, 1])$, the inverse CDF map in (1.1) turns out to be exactly T^* ; for $d > 1$, the multi-dimensional map $T^*: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the gradient of a convex function, generalizing monotonic functions on the real line to multi-dimensions. On the other hand, Brenier’s result naturally initiates an inductive bias in transform sampling: instead of searching any transport map, one may find T^* , the optimal one with the smallest cost. To contrast this with the plain transform sampling (1.2), let us rewrite (1.3) using a suitable Lagrangian multiplier $\lambda > 0$ to enforce the equality constraint $T_{\#}\mu = \nu$:

$$\min_{T \in \mathcal{F}} C(T) + \lambda \cdot \mathcal{L}(T_{\#}\mu, \nu). \quad (1.4)$$

Now, we can see that (1.4) incorporates an additional objective function of T —transport cost—in (1.2), thereby introducing an inductive bias towards the optimal transport map, the minimum of C .

Such an OT-based approach, however, can be inefficient in practice if the target ν is a high-dimensional embedding of some low-dimensional distribution. For instance, let ν be the distribution of handwritten digit images from the MNIST data set on \mathbb{R}^{784} .¹ To use the above OT-based approach, one must choose μ on \mathbb{R}^{784} and find a map $T: \mathbb{R}^{784} \rightarrow \mathbb{R}^{784}$. However, the support of ν is intrinsically low-dimensional (roughly

¹Images are normalized and fit into a 28×28 pixel bounding box, hence defined on $\mathbb{R}^{28 \times 28} \equiv \mathbb{R}^{784}$ [30].

\mathbb{R}^{15} as in [19]), hence other transform samplers with $\mathcal{X} = \mathbb{R}^{15}$ yielding $T: \mathbb{R}^{15} \rightarrow \mathbb{R}^{784}$ are more efficient than the OT-based method in terms of estimating T and computing $T(X)$ for $X \sim \mu$.

In this paper, we propose and study a transform sampler, combining the best of both worlds: it introduces beneficial inductive biases like the OT approach, while operating when \mathcal{X} and \mathcal{Y} are heterogeneous spaces. The key to our approach is to utilize a notion of isomorphism and the Gromov-Wasserstein (GW) distance between μ and ν . Given two cost functions $c_{\mathcal{X}}: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and $c_{\mathcal{Y}}: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, the GW distance [38, 13] is

$$\text{GW}(\mu, \nu) := \inf_{\gamma \in \Pi(\mu, \nu)} \left(\int_{\mathcal{X} \times \mathcal{Y}} \int_{\mathcal{X} \times \mathcal{Y}} (c_{\mathcal{X}}(x, x') - c_{\mathcal{Y}}(y, y'))^2 d\gamma(x, y) d\gamma(x', y') \right)^{1/2}, \quad (1.5)$$

where $\Pi(\mu, \nu)$ is the set of couplings between μ and ν . GW aims to match the cost functions defined on two heterogeneous spaces, intending to identify an isomorphism, namely, a transport map T such that $c_{\mathcal{X}}(x, x') = c_{\mathcal{Y}}(T(x), T(x'))$ for all $x, x' \in \mathcal{X}$. Inspired by these, one can define the following objective function of T to replace the transport cost C :

$$Q(T) := \int_{\mathcal{X}} \int_{\mathcal{X}} (c_{\mathcal{X}}(x, x') - c_{\mathcal{Y}}(T(x), T(x')))^2 d\mu(x) d\mu(x').$$

Roughly speaking, the new sampler takes a form of (1.4), but with Q instead of C , thereby introducing an inductive bias towards isomorphisms (the target when $\min_T Q(T) = 0$). Remark that GW is a Quadratic Program (QP) in γ , which is known to be computationally hard [11]; similarly, the objective function Q is quadratic in μ . More importantly, designing and analyzing the plug-in estimation of the quadratic objects $\text{GW}(\mu, \nu)$ and Q based on finite i.i.d. samples from μ and ν is not obvious.

Main contributions This paper considers computational and statistical questions regarding Gromov-Wasserstein outlined above, and aims to design a new transform sampler as an approach to model and sample from multi-dimensional probability distributions given access to i.i.d. samples, circumventing the usual ways of modeling the density function or MCMC. Our transform sampler can also estimate good alignments between two heterogeneous metric measure spaces $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ from empirical data sets, with estimated maps that approximately pushforward one measure μ to the other ν , and vice versa. Towards reaching these goals, we made the following specific contributions.

- We introduce a new notion, Reversible Gromov-Monge (RGM) distance, on metric measure spaces that majorizes the usual Gromov-Wasserstein distance. Moreover, we show several analytic properties possessed by GW naturally carry over to RGM; in particular, RGM induces a valid metric between metric measure spaces up to an isomorphism. Furthermore, under mild assumptions, we derive that RGM equals GW. Finally, we illustrate how RGM induces an inductive bias favoring strong isomorphisms through a new insight from Brenier’s polar factorization [9].
- Our RGM formulation induces a transform sampler, as a relaxation of the usual GW formulation. Rather than solving a QP which is quadratic in the coupling $\gamma \in \Pi(\mu, \nu)$, we decouple the pair as $(\text{Id}, F)_{\#}\mu$ and $(B, \text{Id})_{\#}\nu$ with $F: \mathcal{X} \rightarrow \mathcal{Y}$ and $B: \mathcal{Y} \rightarrow \mathcal{X}$, respectively, and then bind them later via the constraint $(\text{Id}, F)_{\#}\mu \approx (B, \text{Id})_{\#}\nu$. Such a decoupling and binding idea will prove suitable for the statistical estimation problem based on finite i.i.d. samples. We will also show, from an operator viewpoint, such a decoupling and binding idea can relax our RGM to an infinite-dimensional convex program in F, B that admits a simple representation theorem, as opposed to the otherwise intractable infinite-dimensional QP in GW.
- We derive non-asymptotic rates of convergence for the proposed RGM sampler using tools from empirical processes, for generic classes modeling the measurable maps F and B . Based on our non-asymptotic results, concrete upper bounds can be easily spelled out in the cases where F and B are parametrized by deep neural networks. As mentioned earlier, the RGM sampler also promises to identify good alignments between metric measure spaces, and to learn approximate isomorphism when possible. We demonstrate such a point using numerical experiments on MNIST.

Organization The rest of the paper is organized as follows. First, we briefly review other related studies omitted in the discussion above. Then, in Section 2, preliminary background on optimal transport and Gromov-Wasserstein distance is outlined. Next, Section 3 summarizes the primary methodology and theory regarding our proposed Reversible Gromov-Monge sampler. Synthetic and real-world examples showcasing the effectiveness of the RGM sampler are demonstrated in Section 4 as a proof of concept. The supplementary material collects details of the results in Sections 3 and 4 along with extensive discussions.

1.1 Related Literature

Inferring the underlying probability distributions from data has been a central problem in statistics and unsupervised machine learning since the invention of histograms by Pearson a century ago. Classic mathematical statistics explicitly models the density function in a parametric or a nonparametric way [47], and studies the minimax optimality of directly estimating such density functions [51]. It is also unclear how to proceed to sample from a possibly improper² density estimator, even with an optimal estimator at hand. One may employ Markov Chain Monte Carlo (MCMC) techniques for sampling from specific models. However, on the computational front, it is highly non-trivial how to ensure the mixing properties of MCMC for a designed sampler [44].

A recent trend in unsupervised machine learning is to learn complex, high-dimensional distributions via (deep) generative models, either explicitly by parametrizing the sufficient statistics of the exponential families [17, 28], or implicitly by parametrizing the pushforward map transporting distributions [18, 24], with a focus on tractability in computation. Surprisingly, though lacking theoretical underpinning and optimality, the generative models’ approach performs well empirically in large-scale applications where classical statistical procedures are destined to fail. There has been a growing literature on understanding distribution estimation with the implicit framework, with more general metrics and target distribution classes, to name a few, [41, 32, 18] on MMDs, [50, 33] on integral probability metrics, and [40, 3, 34, 48, 4, 56, 31, 12] on generative adversarial networks. Last but not least, we emphasize that an alternative implicit distribution estimation approach using the simulated method of moments has been formulated in the econometrics literature since [37, 42] and [25].

Originally introduced as a tool for comparing objects in computer graphics, analytic properties of the Gromov-Wasserstein distance have been studied extensively [38, 52]; the most important one is that it defines a distance between metric measure spaces, namely, metric spaces endowed with probability measures. Since many real-world data sets can be modeled as metric measure spaces, the GW distance has been utilized in various problems such as shape correspondence [49], graph matching [58], and protein comparison [22]. Certain statistical aspects of comparing metric measure spaces have been studied in [8, 57].

Computation of the GW distance amounts to a relaxation of the quadratic assignment problem [29]; both are known to be NP-hard [11] in the worst case. Several approaches have been proposed for the approximate computation of the GW distance. [38] studies lower bounds on the GW distance that are easier to compute. [43] adds an entropic regularization term to the GW distance, which leads to a fast iterative algorithm; [46] further modifies this by imposing a low-rank constraint on couplings. [54] proposes the Sliced Gromov-Wasserstein distance defined by integrating GW distances over one-dimensional projections. Last but not least, recent papers [58, 6, 14] study scalable partitioning schemes to approximately compute GW distances.

2 Background

In this section, we provide background on the Optimal Transport (OT) theory and the Gromov-Wasserstein distance. First, we start with some notations. Let $\|A\|$ denote the Frobenius norm of a matrix A and $\|x\|$ denote the Euclidean norm of a vector x . Given a set \mathcal{X} and a function $f: \mathcal{X} \rightarrow \mathbb{R}$, let $\|f\|_\infty = \sup_{x \in \mathcal{X}} |f(x)|$ denote the sup norm. For an integer $n \in \mathbb{N}$, we define $[n] = \{1, \dots, n\}$. For a metric space \mathcal{X} , we denote its metric as $d_{\mathcal{X}}$ and write $\mathcal{P}(\mathcal{X})$ to denote the collection of all Borel probability measures on \mathcal{X} ; we call \mathcal{X} a

²Here we mean that the estimated density is not always non-negative and integrates to one.

Polish space if it is complete and separable. We call a pair (\mathcal{X}, μ) a Polish probability space if \mathcal{X} is a Polish space and $\mu \in \mathcal{P}(\mathcal{X})$. Given two Polish probability spaces (\mathcal{X}, μ) and (\mathcal{Y}, ν) , the collection of all transport maps from μ to ν is denoted as $\mathcal{T}(\mu, \nu) := \{T: \mathcal{X} \rightarrow \mathcal{Y} \mid T_{\#}\mu = \nu\}$; we call $\gamma \in \mathcal{P}(\mathcal{X} \times \mathcal{Y})$ a coupling between μ and ν if $\gamma(A \times \mathcal{Y}) = \mu(A)$ and $\gamma(\mathcal{X} \times B) = \nu(B)$ for all Borel subsets $A \subset \mathcal{X}$ and $B \subset \mathcal{Y}$, and we denote the collection of all such couplings as $\Pi(\mu, \nu)$. For a sequence of numbers $a(n), b(n) \in \mathbb{R}$, we use $a(n) \lesssim b(n)$ to denote the relationship that $a(n)/b(n) \leq C, \forall n$ with some universal constant $C > 0$.

2.1 A Brief Overview of Optimal Transport Theory

A major goal of OT is minimizing the cost associated with the transport map between two Polish probability spaces, say (\mathcal{X}, μ) and (\mathcal{Y}, ν) . Consider a measurable function $c: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$; we view $c(x, y)$ as the cost associated with $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. For each transport map $T \in \mathcal{T}(\mu, \nu)$, we interpret $c(x, T(x))$ as a unit cost incurred by mapping each $x \in \mathcal{X}$ to $T(x) \in \mathcal{Y}$. We define the average cost incurred by the transport map T as the integration of all the unit costs with respect to μ , that is, $\int_{\mathcal{X}} c(x, T(x)) d\mu(x)$. Minimizing the cost over $\mathcal{T}(\mu, \nu)$ is referred to as the Monge problem named after Gaspard Monge. We call T^* an optimal transport map if T^* is minimizer, that is,

$$T^* \in \arg \min_{T \in \mathcal{T}(\mu, \nu)} \int_{\mathcal{X}} c(x, T(x)) d\mu(x),$$

Another important OT problem is minimizing the cost given by couplings. We define the average cost incurred by a coupling $\gamma \in \Pi(\mu, \nu)$ as the integration of the cost c with respect to γ , namely, $\int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\gamma(x, y)$. Minimizing this cost over $\Pi(\mu, \nu)$ is called the Kantorovich problem credited to Leonid Kantorovich. We call γ^* an optimal coupling if

$$\gamma^* \in \arg \min_{\gamma \in \Pi(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\gamma(x, y),$$

The two OT problems are closely related: the Kantorovich problem is a relaxation of the Monge problem. To see this, for each $T \in \mathcal{T}(\mu, \nu)$, define a map $(\text{Id}, T): \mathcal{X} \rightarrow \mathcal{X} \times \mathcal{Y}$ by $(\text{Id}, T)(x) = (x, T(x))$. One can verify $(\text{Id}, T)_{\#}\mu \in \Pi(\mu, \nu)$. Therefore, if we define $\Pi_{\mathcal{T}} := \{(\text{Id}, T)_{\#}\mu : T \in \mathcal{T}(\mu, \nu)\}$, then $\Pi_{\mathcal{T}} \subset \Pi(\mu, \nu)$ and thus

$$\inf_{T \in \mathcal{T}(\mu, \nu)} \int_{\mathcal{X}} c(x, T(x)) d\mu(x) = \inf_{\gamma \in \Pi_{\mathcal{T}}} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\gamma(x, y) \geq \inf_{\gamma \in \Pi(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\gamma(x, y),$$

where the first equality follows from change-of-variables. In other words, two OT problems share the same objective function as a function of couplings; however, the Kantorovich problem has a larger constraint set.

Unlike the Monge problem, the Kantorovich problem has favorable properties. First, the objective function is linear in γ . Moreover, $\Pi(\mu, \nu)$ is compact in the weak topology of Borel probability measures defined on $\mathcal{X} \times \mathcal{Y}$. This suggests that we can view the Kantorovich problem as an infinite-dimensional linear program.

Besides seeking optimal transport maps or couplings, another interesting aspect of OT problems is that the least possible cost can endow a metric structure among Polish probability spaces. If $\mathcal{X} = \mathcal{Y}$ and $c = d_{\mathcal{X}}^2$, the square root of the solution of the Kantorovich problem defines a distance between μ and ν , known as the Wasserstein distance.

Definition 1. Given a Polish space \mathcal{X} , we call

$$W_2(\mu, \nu) = \inf_{\gamma \in \Pi(\mu, \nu)} \left(\int_{\mathcal{X} \times \mathcal{X}} d_{\mathcal{X}}^2(x, y) d\gamma(x, y) \right)^{1/2}$$

the Wasserstein-2 distance³ between $\mu, \nu \in \mathcal{P}(\mathcal{X})$.

³One can define the Wasserstein- p distance by replacing the exponent 2 above with $p \in [1, \infty]$.

2.2 Gromov-Wasserstein and Gromov-Monge Distances

Although OT problems can be defined between arbitrary Polish probability spaces, in practice, it is unclear how to design a function $c: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ to represent meaningful cost associated with $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ in two heterogeneous spaces. For instance, if $\mathcal{X} = \mathbb{R}^p$ and $\mathcal{Y} = \mathbb{R}^q$ with $p \neq q$, there is no simple choice for a cost function c over $\mathbb{R}^p \times \mathbb{R}^q$. As a result, classic OT theory (including Brenier’s result) cannot be directly used for comparing heterogeneous Polish probability spaces.

Mémoli’s pioneering work [38] resolved this issue by considering a quadratic objective function of γ :

$$\int_{\mathcal{X} \times \mathcal{Y}} c(x, y) \, d\gamma(x, y) \Rightarrow \int_{\mathcal{X} \times \mathcal{Y}} \int_{\mathcal{X} \times \mathcal{Y}} (c_{\mathcal{X}}(x, x') - c_{\mathcal{Y}}(y, y'))^2 \, d\gamma(x, y) \, d\gamma(x', y'),$$

where $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$ are defined over $\mathcal{X} \times \mathcal{X}$ and $\mathcal{Y} \times \mathcal{Y}$, respectively. For instance, one can specify $c_{\mathcal{X}} = d_{\mathcal{X}}$ and $c_{\mathcal{Y}} = d_{\mathcal{Y}}$. Rather than considering a unit cost corresponding to each pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$, we associate two pairs $(x, y), (x', y') \in \mathcal{X} \times \mathcal{Y}$ with the discrepancy of intra-space quantities $c_{\mathcal{X}}(x, x')$ and $c_{\mathcal{Y}}(y, y')$. In summary, by switching from the integration $d\gamma$ to the double integration $d\gamma \, d\gamma$, we no longer need an otherwise inter-space quantity $c: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$. Therefore, we can always define this objective function whenever we have proper $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$ in each individual space, leading to the following definition.

Definition 2. A triple $(\mathcal{X}, \mu, c_{\mathcal{X}})$ is called a network space if (\mathcal{X}, μ) is a Polish probability space such that $\text{supp}(\mu) = \mathcal{X}$ and $c_{\mathcal{X}}: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is continuous. The Gromov-Wasserstein distance between network spaces $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ is defined as

$$\text{GW}(\mu, \nu) = \inf_{\gamma \in \Pi(\mu, \nu)} \left(\int_{\mathcal{X} \times \mathcal{Y}} \int_{\mathcal{X} \times \mathcal{Y}} (c_{\mathcal{X}}(x, x') - c_{\mathcal{Y}}(y, y'))^2 \, d\gamma(x, y) \, d\gamma(x', y') \right)^{1/2}.$$

Remark. On top of the network space definition introduced in [13], we impose continuity of $c_{\mathcal{X}}$ for a cleaner analysis. A network space $(\mathcal{X}, \mu, c_{\mathcal{X}})$ is called a metric measure space if $c_{\mathcal{X}} = d_{\mathcal{X}}$ as introduced in [38] and [52]. In short, a network space is a generalization of a metric measure space.

Like the Wasserstein distance, the GW distance has metric properties; it satisfies symmetry and the triangle inequality, and $\text{GW}(\mu, \nu) = 0$ if $(\mathcal{X}, \mu, c_{\mathcal{X}}) = (\mathcal{Y}, \nu, c_{\mathcal{Y}})$. However, the converse of this last statement does not hold in general: for its validity, a suitable equivalence relation needs to be defined on the collection of network spaces.

Definition 3. Network spaces $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ are strongly isomorphic if there exists $T \in \mathcal{T}(\mu, \nu)$ such that $T: \mathcal{X} \rightarrow \mathcal{Y}$ is bijective and $c_{\mathcal{X}}(x, x') = c_{\mathcal{Y}}(T(x), T(x'))$ for all $x, x' \in \mathcal{X}$. In this case, we write $(\mathcal{X}, \mu, c_{\mathcal{X}}) \cong (\mathcal{Y}, \nu, c_{\mathcal{Y}})$ and such a transport map T is called a strong isomorphism.

One can easily check that \cong is indeed an equivalence relation on the collection of network spaces. The following theorem states that the GW distance satisfies all metric axioms on the quotient space—under the equivalence relation \cong —of metric measure spaces.

Theorem 1 (Lemma 1.10 of [52]). Let \mathcal{M} be the collection of all network spaces $(\mathcal{X}, \mu, c_{\mathcal{X}})$ such that $c_{\mathcal{X}} = d_{\mathcal{X}}$. Also, let \mathcal{M}/\cong be the collection of all equivalence classes of \mathcal{M} induced by \cong . Then, GW satisfies the three metric axioms on \mathcal{M}/\cong .

Recall that the Monge problem is a restricted version of the Kantorovich problem with an additional constraint that couplings are given by a transport map; replacing $\Pi(\mu, \nu)$ in the Kantorovich problem with $\Pi_{\mathcal{T}}$ yields the Monge problem. Imposing the same constraint on the definition of GW leads to the Gromov-Monge distance.

Definition 4. The Gromov-Monge distance between network spaces $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ is defined as

$$\text{GM}(\mu, \nu) = \inf_{T \in \mathcal{T}(\mu, \nu)} \left(\int_{\mathcal{X}} \int_{\mathcal{X}} (c_{\mathcal{X}}(x, x') - c_{\mathcal{Y}}(T(x), T(x')))^2 \, d\mu(x) \, d\mu(x') \right)^{1/2}.$$

Loosely speaking, computing GM amounts to finding a transport map T such that $c_{\mathcal{X}}(x, x')$ best matches $c_{\mathcal{Y}}(T(x), T(x'))$ on average; we can view such a map T as a surrogate for an isomorphism. See Section 2.4 of [39] for more details of GM.

3 Summary of Results

Inspired by the Gromov-Wasserstein and Gromov-Monge distances, we propose a new metric—the reversible Gromov-Monge distance—between network spaces in this paper. Our formulation seeks a pair of transport maps $F \in \mathcal{T}(\mu, \nu)$ and $B \in \mathcal{T}(\nu, \mu)$ best approximating isomorphic relations between network spaces. We propose a novel transform sampling method that uses F as a push-forward map to obtain i.i.d. samples from a target distribution ν . We present two optimization formulations solving for such a pair (F, B) in order: a potentially non-convex formulation that employs the standard gradient descent method to optimize, and an infinite-dimensional convex formulation where global optima can be found efficiently. For the former, we analyze the statistical rate of convergence for generic classes $\mathcal{F} \times \mathcal{B}$ parametrizing (F, B) . For the latter, we derive a new representer theorem on a suitable reproducing kernel Hilbert space (RKHS).

3.1 Metric Properties of Reversible Gromov-Monge

Our formulation is based on the following observation: for a coupling γ such that $\gamma = (\text{Id}, F)_{\#}\mu = (B, \text{Id})_{\#}\nu$, which presents a binding constraint, we can simplify the objective function of GW as

$$\int_{\mathcal{X} \times \mathcal{Y}} (c_{\mathcal{X}}(x, B(y)) - c_{\mathcal{Y}}(F(x), y))^2 d\mu \otimes \nu ,$$

where $d\mu \otimes \nu := d\mu(x) d\nu(y)$ denotes the product measure of μ and ν . Imposing the binding constraint on the definition of GW leads to the following definition.

Definition 5. For network spaces $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$, we write $(F, B) \in \mathcal{I}(\mu, \nu)$ if measurable maps $F: \mathcal{X} \rightarrow \mathcal{Y}$ and $B: \mathcal{Y} \rightarrow \mathcal{X}$ satisfy the binding constraint $(\text{Id}, F)_{\#}\mu = (B, \text{Id})_{\#}\nu$. We define the reversible Gromov-Monge (RGM) distance between $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ as

$$\text{RGM}(\mu, \nu) := \inf_{(F, B) \in \mathcal{I}(\mu, \nu)} \left(\int_{\mathcal{X} \times \mathcal{Y}} (c_{\mathcal{X}}(x, B(y)) - c_{\mathcal{Y}}(F(x), y))^2 d\mu \otimes \nu \right)^{1/2} . \quad (3.1)$$

Remark. A few remarks are in place for the binding constraint. If $(\text{Id}, F)_{\#}\mu = (B, \text{Id})_{\#}\nu$, then $F_{\#}\mu = \nu$ and $B_{\#}\nu = \mu$ follow due to marginal conditions. However, the converse is not true in general. To see this, let $\mu = \nu = \text{Unif}([0, 1])$, then $F_{\#}\mu = \nu$ and $B_{\#}\nu = \mu$ hold for $F(x) = B(x) = |2x - 1|$. However, $(\text{Id}, F)_{\#}\mu \neq (B, \text{Id})_{\#}\nu$ because $(\text{Id}, F)_{\#}\mu$ is a uniform measure on $\{(x, |2x - 1|) : x \in [0, 1]\}$, whereas $(B, \text{Id})_{\#}\nu$ is a uniform measure on $\{(|2y - 1|, y) : y \in [0, 1]\}$.

Roughly speaking, computing RGM consists in finding a pair $(F, B) \in \mathcal{I}(\mu, \nu)$ such that $c_{\mathcal{X}}(x, B(y))$ best matches $c_{\mathcal{Y}}(F(x), y)$ on average. Like a strong isomorphism, we can view such a pair as jointly capturing an isomorphic relation of $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$. We will use this observation later to build a transform sampling method.

We will prove that RGM possesses metric properties similar to the Gromov-Wasserstein. Motivated by Theorem 1, we derive the following result.

Theorem 2. Let $h: \mathbb{R}_+ \rightarrow \mathbb{R}$ be a continuous and strictly monotone function and \mathcal{N}^h be a collection of all network spaces $(\mathcal{X}, \mu, c_{\mathcal{X}})$ such that $c_{\mathcal{X}} = h(d_{\mathcal{X}})$. Then RGM satisfies the three metric axioms on \mathcal{N}^h / \cong , the collection of all equivalence classes of \mathcal{N}^h induced by \cong .

Remark. Suppose \mathcal{X} is a Euclidean space and $d_{\mathcal{X}}$ is the standard Euclidean distance. If $h(x) = \exp(-\alpha x^2)$ with $\alpha > 0$, then $h(d_{\mathcal{X}})$ is the radial basis function (RBF) kernel on \mathcal{X} ; we will use this in numerical experiments.

Readers may wonder about the generic relations among three distances GW, GM, and RGM, which will be established in the next proposition.

Proposition 1. *For network spaces $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ as in Definition 2,*

$$\text{GW}(\mu, \nu) \leq \text{GM}(\mu, \nu) \leq \text{RGM}(\mu, \nu) . \quad (3.2)$$

Interestingly, under mild conditions, the above inequalities in Proposition 1 hold as equality, thus showing that RGM provides the exact metric as GW. The proof is inspired by a construction in [10, 9].

Theorem 3. *Let $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ be two network spaces, where \mathcal{X} and \mathcal{Y} are closed subsets of \mathbb{R}^d with $d \geq 2$. Assume that $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$ are bounded and $\mu(\{x\}) = \nu(\{y\}) = 0$ for any $(x, y) \in \mathcal{X} \times \mathcal{Y}$. Then, $\text{GW}(\mu, \nu) = \text{GM}(\mu, \nu) = \text{RGM}(\mu, \nu)$.*

We defer the proof details of Theorem 2 and Proposition 1 to Appendix A.1, and Theorem 3 to Appendix A.2.

Finally, we conclude this section by pointing out a connection between inductive biases in RGMs and Brenier’s polar factorization [9]. Given two Polish probability spaces (\mathcal{X}, μ) and (\mathcal{Y}, ν) , there exist cost functions $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$ (that depend on μ, ν) such that the resulting network spaces $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ are strongly isomorphic. More importantly, among the infinitely many $F: \mathcal{X} \rightarrow \mathcal{Y}$, $B: \mathcal{Y} \rightarrow \mathcal{X}$ ’s such that $(\text{Id}, F)_{\#}\mu = (B, \text{Id})_{\#}\nu$ hold valid for transform sampling, the optimal pair (F^*, B^*) minimizing the RGM term (3.1) achieves the strong isomorphism

$$\text{RGM}(\mu, \nu) = \int_{\mathcal{X} \times \mathcal{Y}} (c_{\mathcal{X}}(x, B^*(y)) - c_{\mathcal{Y}}(F^*(x), y))^2 d\mu \otimes \nu = 0 , \quad (3.3)$$

and thus $c_{\mathcal{X}}(x, B^*(y)) = c_{\mathcal{Y}}(F^*(x), y)$ almost surely. In plain language, the RGM introduces an inductive bias favoring strong isomorphisms, in the same spirit as the Wasserstein-2 metric favors the transport map with the optimal cost seen in the introduction. The detailed discussions are deferred to Appendix A.3.

3.2 Transform Sampling via RGM

With the proposed notion of RGM, we design a transform sampling method in this section. The transform sampler is based on finding a minimizing pair (F, B) of RGM, which can capture isomorphic relations between network spaces. To implement this method, we need to estimate (F, B) using only i.i.d. samples from μ and ν . Leveraging the Lagrangian form, we derive a minimization problem that can be implemented based on finite samples.

First, we rewrite the population minimization problem with the binding constraint as follows,

$$\begin{aligned} \min_{\substack{F: \mathcal{X} \rightarrow \mathcal{Y} \\ B: \mathcal{Y} \rightarrow \mathcal{X}}} & \int_{\mathcal{X} \times \mathcal{Y}} (c_{\mathcal{X}}(x, B(y)) - c_{\mathcal{Y}}(F(x), y))^2 d\mu \otimes \nu \\ \text{s.t.} & \mathcal{L}_{\mathcal{X} \times \mathcal{Y}}((\text{Id}, F)_{\#}\mu, (B, \text{Id})_{\#}\nu) = 0 . \end{aligned} \quad (3.4)$$

Here, $\mathcal{L}_{\mathcal{X} \times \mathcal{Y}}$ is a suitable discrepancy measure on $\mathcal{P}(\mathcal{X} \times \mathcal{Y})$ so that the constraint of (3.4) is a surrogate for the original constraint $(\text{Id}, F)_{\#}\mu = (B, \text{Id})_{\#}\nu$. In practice, we do not require that $\mathcal{L}_{\mathcal{X} \times \mathcal{Y}} = 0$ implies $(\text{Id}, F)_{\#}\mu = (B, \text{Id})_{\#}\nu$; in fact, the former constraint can be a relaxation of the latter. The choice of $\mathcal{L}_{\mathcal{X} \times \mathcal{Y}}$ will be specified later. To solve this minimization problem, we propose utilizing the Lagrangian:

$$\min_{\substack{F: \mathcal{X} \rightarrow \mathcal{Y} \\ B: \mathcal{Y} \rightarrow \mathcal{X}}} \int_{\mathcal{X} \times \mathcal{Y}} (c_{\mathcal{X}}(x, B(y)) - c_{\mathcal{Y}}(F(x), y))^2 d\mu \otimes \nu + \lambda \cdot \mathcal{L}_{\mathcal{X} \times \mathcal{Y}}((\text{Id}, F)_{\#}\mu, (B, \text{Id})_{\#}\nu) .$$

Given i.i.d. samples $\{x_i\}_{i=1}^m$ and $\{y_j\}_{j=1}^n$ from μ and ν , respectively, we replace the population objective with its empirical estimates:

$$\min_{\substack{F: \mathcal{X} \rightarrow \mathcal{Y} \\ B: \mathcal{Y} \rightarrow \mathcal{X}}} \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n (c_{\mathcal{X}}(x_i, B(y_j)) - c_{\mathcal{Y}}(F(x_i), y_j))^2 + \lambda \cdot \mathcal{L}_{\mathcal{X} \times \mathcal{Y}}((\text{Id}, F)_{\#}\hat{\mu}_m, (B, \text{Id})_{\#}\hat{\nu}_n) ,$$

where $\hat{\mu}_m$ and $\hat{\nu}_n$ are the empirical measures based on $\{x_i\}_{i=1}^m$ and $\{y_j\}_{j=1}^n$, respectively. Empirically, we find that adding the following extra terms often enhances empirical results:

$$\begin{aligned} \min_{\substack{F: \mathcal{X} \rightarrow \mathcal{Y} \\ B: \mathcal{Y} \rightarrow \mathcal{X}}} & \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n (c_{\mathcal{X}}(x_i, B(y_j)) - c_{\mathcal{Y}}(F(x_i), y_j))^2 + \lambda_1 \cdot \mathcal{L}_{\mathcal{X} \times \mathcal{Y}}((\text{Id}, F)_{\#} \hat{\mu}_m, (B, \text{Id})_{\#} \hat{\nu}_n) \\ & + \lambda_2 \cdot \mathcal{L}_{\mathcal{X}}(\hat{\mu}_m, B_{\#} \hat{\nu}_n) + \lambda_3 \cdot \mathcal{L}_{\mathcal{Y}}(F_{\#} \hat{\mu}_m, \hat{\nu}_n) . \end{aligned}$$

Like $\mathcal{L}_{\mathcal{X} \times \mathcal{Y}}$, we utilize suitable discrepancy measures $\mathcal{L}_{\mathcal{X}}$ and $\mathcal{L}_{\mathcal{Y}}$ so that these additional terms help matching the marginals of $(\text{Id}, F)_{\#} \hat{\mu}_m$ and $(B, \text{Id})_{\#} \hat{\nu}_n$.

Lastly, we discuss the choice of $\mathcal{L}_{\mathcal{X}}$, $\mathcal{L}_{\mathcal{Y}}$, and $\mathcal{L}_{\mathcal{X} \times \mathcal{Y}}$. We use the square of Maximum Mean Discrepancy (MMD) as the leading example.⁴ MMD between two measures is a distance between their embeddings in some reproducing kernel Hilbert space (RKHS), which is indeed a metric under mild conditions [41]. Also, MMD is representable via the reproducing kernel of the RKHS, hence one may simply choose a kernel function to define it. Concretely, for any kernel $K_{\mathcal{X}}$ on \mathcal{X} , the square of MMD between $\hat{\mu}_m$ and $B_{\#} \hat{\nu}_n$ is

$$\frac{1}{m^2} \sum_{i, i'} K_{\mathcal{X}}(x_i, x_{i'}) + \frac{1}{n^2} \sum_{j, j'} K_{\mathcal{X}}(B(y_j), B(y_{j'})) - \frac{2}{mn} \sum_{i, j} K_{\mathcal{X}}(x_i, B(y_j)) .$$

To utilize such a convenient closed form, we specify $\mathcal{L}_{\mathcal{X}}$, $\mathcal{L}_{\mathcal{Y}}$, $\mathcal{L}_{\mathcal{X} \times \mathcal{Y}}$ as the square of corresponding MMDs by choosing kernels $K_{\mathcal{X}}$, $K_{\mathcal{Y}}$, $K_{\mathcal{X} \times \mathcal{Y}}$ on \mathcal{X} , \mathcal{Y} , $\mathcal{X} \times \mathcal{Y}$. For the kernel $K_{\mathcal{X} \times \mathcal{Y}}$ on the product space, we use the tensor product kernel $K_{\mathcal{X}} \otimes K_{\mathcal{Y}}$ given as

$$K_{\mathcal{X}} \otimes K_{\mathcal{Y}}((x, y), (x', y')) = K_{\mathcal{X}}(x, x') K_{\mathcal{Y}}(y, y') .$$

The tensor product notation is employed since the kernel on the product space inherits the feature map as the tensor product of two individual feature maps w.r.t. $K_{\mathcal{X}}$ and $K_{\mathcal{Y}}$.

Denoting the MMD associated with a kernel K as MMD_K , we obtain the following minimization problem:

$$\begin{aligned} \min_{\substack{F: \mathcal{X} \rightarrow \mathcal{Y} \\ B: \mathcal{Y} \rightarrow \mathcal{X}}} & \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n (c_{\mathcal{X}}(x_i, B(y_j)) - c_{\mathcal{Y}}(F(x_i), y_j))^2 \\ & + \lambda_1 \cdot \text{MMD}_{K_{\mathcal{X}} \otimes K_{\mathcal{Y}}}^2((\text{Id}, F)_{\#} \hat{\mu}_m, (B, \text{Id})_{\#} \hat{\nu}_n) \\ & + \lambda_2 \cdot \text{MMD}_{K_{\mathcal{X}}}^2(\hat{\mu}_m, B_{\#} \hat{\nu}_n) + \lambda_3 \cdot \text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#} \hat{\mu}_m, \hat{\nu}_n) . \end{aligned} \tag{3.5}$$

Once we solve the problem above, the solution $\hat{F}: \mathcal{X} \rightarrow \mathcal{Y}$ will serve as an approximate isomorphism and facilitate transform sampling of the target ν from a known distribution μ . The map \hat{B} possesses similar properties as \hat{F} , whereas the map \hat{F} is of our primary interest for sampling purposes. The reverse map $\hat{B}: \mathcal{Y} \rightarrow \mathcal{X}$ also embeds point clouds in \mathcal{Y} into \mathcal{X} , with approximate isomorphism properties in the sense of Gromov-Monge.

3.3 Statistical Rate of Convergence

Like other transform sampling approaches for generative models, we consider (3.5) using vector-valued function classes \mathcal{F} and \mathcal{B} parametrized by neural networks, and then optimize using a gradient descent algorithm. We emphasize this minimization problem is much simpler than adversarial formulations as in GANs: variational problems of GANs consist of minimization over a class of generators and maximization over a class of discriminators, which requires complex saddle-point dynamics [15, 35]. In contrast, our RGM only solves a single minimization problem in network parameters. Although generally non-convex in nature, the parameter minimization problem in neural networks can often be efficiently optimized by stochastic gradient descent, and can even provably achieve the global optima if the loss satisfies certain Polyak-Lojasiewicz conditions [5].

⁴This is merely a proof of concept. One may use other quantities in practice, described in Section 4.

We investigate the statistical rate of convergence for this minimization problem, assuming the empirical problem (3.5) can be solved accurately. First, define

$$\begin{aligned} C(\mu, \nu, F, B) &:= \int (c_{\mathcal{X}}(x, B(y)) - c_{\mathcal{Y}}(F(x), y))^2 d\mu \otimes \nu \\ &+ \lambda_1 \cdot \text{MMD}_{K_{\mathcal{X}} \otimes K_{\mathcal{Y}}}^2((\text{Id}, F)_{\#}\mu, (B, \text{Id})_{\#}\nu) \\ &+ \lambda_2 \cdot \text{MMD}_{K_{\mathcal{X}}}^2(\mu, B_{\#}\nu) + \lambda_3 \cdot \text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#}\mu, \nu). \end{aligned} \quad (3.6)$$

Then, the objective function of (3.5) is a plug-in estimator $C(\hat{\mu}_m, \hat{\nu}_n, F, B)$. We consider solving (3.5) over the transformation class $\mathcal{F} \times \mathcal{B}$ given as follows, for which we will state our non-asymptotic results in full generality. From now on, let \mathcal{X} and \mathcal{Y} be subsets of Euclidean spaces of dimensions $\dim(\mathcal{X})$ and $\dim(\mathcal{Y})$, respectively. \mathcal{F} (resp. \mathcal{B}) is a collection of vector-valued measurable functions from \mathcal{X} to \mathcal{Y} (resp. from \mathcal{Y} to \mathcal{X}). For each $F \in \mathcal{F}$ and $k \in [\dim(\mathcal{Y})]$, we write $F_k(x)$ to denote the k -th coordinate of $F(x)$. Accordingly, we define $\mathcal{F}_k = \{F_k : \mathcal{X} \rightarrow \mathbb{R} \mid F \in \mathcal{F}\}$, namely, a collection of real-valued measurable functions defined on \mathcal{X} that are given as the k -th coordinate of $F \in \mathcal{F}$. For $\ell \in [\dim(\mathcal{X})]$, we define B_ℓ and $\mathcal{B}_\ell = \{B_\ell : \mathcal{Y} \rightarrow \mathbb{R} \mid B \in \mathcal{B}\}$ analogously.

Then, solving (3.5) over $\mathcal{F} \times \mathcal{B}$ is written as $\min_{(F, B) \in \mathcal{F} \times \mathcal{B}} C(\hat{\mu}_m, \hat{\nu}_n, F, B)$. We prove that the empirical solution leads to an approximate infimum of $(F, B) \mapsto C(\mu, \nu, F, B)$ evaluated with the population measures μ, ν , with sufficiently large sample sizes m and n .

Overview of assumptions Before stating the next theorem, we present an overview of the assumptions. The complete statement of the assumptions and key definitions are designated to Sections B-C in the supplementary material due to space constraints. Assumptions 1 and 4 require the boundedness and Lipschitzness of the cost functions $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$. Similarly, boundedness and Lipschitzness of the kernel functions $K_{\mathcal{X}}, K_{\mathcal{Y}}$ corresponding to the MMD term are stated in Assumptions 2 and 5, respectively. The last two assumptions are imposed on the set of transformations $F: \mathcal{X} \rightarrow \mathcal{Y}$ and $B: \mathcal{Y} \rightarrow \mathcal{X}$: Assumption 3 requires the transformation class is bounded, and Assumption 6 states that the classes should contain non-trivial maps. We shall employ a notion of combinatorial dimension to measure the complexity of real-valued function classes—the pseudo-dimension—formally stated in Definition 7.

Theorem 4. *Let (\hat{F}, \hat{B}) be a solution to the empirical RGM problem*

$$(\hat{F}, \hat{B}) \in \arg \min_{(F, B) \in \mathcal{F} \times \mathcal{B}} C(\hat{\mu}_m, \hat{\nu}_n, F, B),$$

with $C: \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y}) \times \mathcal{F} \times \mathcal{B} \rightarrow \mathbb{R}$ defined in (3.6). Under Assumptions 1-6, the following inequality holds with probability $1 - \delta$ on $\{x_i\}_{i=1}^m$ and $\{y_j\}_{j=1}^n$

$$C(\mu, \nu, \hat{F}, \hat{B}) - \inf_{(F, B) \in \mathcal{F} \times \mathcal{B}} C(\mu, \nu, F, B) \lesssim \mathcal{M}(\mathcal{F}, \mathcal{B}, m, n, \delta). \quad (3.7)$$

Here, $\mathcal{M}(\mathcal{F}, \mathcal{B}, m, n, \delta)$ denotes a complexity measure of $(\mathcal{F}, \mathcal{B})$ given in terms of pseudo-dimensions (Pdim) of \mathcal{F}_k and \mathcal{B}_ℓ defined in Definition 7:

$$\mathcal{M}(\mathcal{F}, \mathcal{B}, m, n, \delta) := \sqrt{\frac{\log\left(\frac{m \vee n}{\delta}\right)}{m \wedge n}} + \sqrt{\frac{\log(m \vee n)}{m \wedge n} \left(\sum_{k=1}^{\dim(\mathcal{Y})} \text{Pdim}(\mathcal{F}_k) + \sum_{\ell=1}^{\dim(\mathcal{X})} \text{Pdim}(\mathcal{B}_\ell) \right)}.$$

We provide required assumptions and the full proof of Theorem 4 in Section B along with the definition of the pseudo-dimension (Definition 7). When \mathcal{F} and \mathcal{B} are parametrized by neural network classes (the ones we will use for numerical demonstrations in Section 4), tight pseudo-dimension bounds established in [1, 26] can be plugged in Theorem 4 for concrete non-asymptotic rates.

3.4 Convex Formulation and Representer Theorem

As the last bit of our contributions, we study a convex formulation of solving (3.5) by relaxing and lifting it to an infinite-dimensional space. There are two reasons behind our convex formulation: first, as a computational alternative to the possibly non-convex optimization; second, to point out a connection with the Nadaraya-Watson estimator in classic nonparametric statistics. The crux lies in relaxing optimizing over the map $F: \mathcal{X} \rightarrow \mathcal{Y}$ to optimizing over its induced (dual) linear operator $\mathbf{F}: L_{\mathcal{Y}}^2 \rightarrow L_{\mathcal{X}}^2$ that maps functions on \mathcal{Y} to functions on \mathcal{X} , where $L_{\mathcal{X}}^2$ is the collection of real-valued measurable functions f defined on \mathcal{X} such that $\int_{\mathcal{X}} f^2 d\pi_{\mathcal{X}} < \infty$ given a Borel measure $\pi_{\mathcal{X}}$ on \mathcal{X} ; similarly, define $L_{\mathcal{Y}}^2$ given a Borel measure $\pi_{\mathcal{Y}}$ on \mathcal{Y} . Then, for a measurable map $F: \mathcal{X} \rightarrow \mathcal{Y}$, we can define $\mathbf{F}: L_{\mathcal{Y}}^2 \rightarrow L_{\mathcal{X}}^2$ by letting $\mathbf{F}(g) = g \circ F$ for all $g \in L_{\mathcal{Y}}^2$. Similarly, we define $\mathbf{B}: L_{\mathcal{X}}^2 \rightarrow L_{\mathcal{Y}}^2$ for each measurable map $B: \mathcal{Y} \rightarrow \mathcal{X}$. We will see \mathbf{F} and \mathbf{B} are well-defined bounded linear operators in Section C under a mild assumption.

To state the representer theorem, consider (3.5) with $c_{\mathcal{X}} = K_{\mathcal{X}}$ and $c_{\mathcal{Y}} = K_{\mathcal{Y}}$, same as kernel functions specified in MMD terms. We show that this problem can be reduced to a finite-dimensional convex optimization by proving a representer theorem. Since finite-dimensional convex optimization can be optimized globally with provable guarantees, such a formulation can be solved numerically in an efficient way.

Let us lay out more details to state the result. Due to Mercer's theorem, let $\{\phi_k \in L_{\mathcal{X}}^2\}_{k \in \mathbb{N}}$ and $\{\psi_\ell \in L_{\mathcal{Y}}^2\}_{\ell \in \mathbb{N}}$ be countable orthonormal bases of $L_{\mathcal{X}}^2$ and $L_{\mathcal{Y}}^2$ where the kernels admit the following spectral decompositions:

$$K_{\mathcal{X}}(x, x') = \sum_k \lambda_k \phi_k(x) \phi_k(x'), \quad K_{\mathcal{Y}}(y, y') = \sum_\ell \gamma_\ell \psi_\ell(y) \psi_\ell(y'), \quad (3.8)$$

with positive eigenvalues $\lambda_k, \gamma_\ell > 0$. Since $\mathbf{F}: L_{\mathcal{Y}}^2 \rightarrow L_{\mathcal{X}}^2$ defines a bounded linear operator, one can represent \mathbf{F} (correspondingly \mathbf{B}) under the orthonormal bases

$$\mathbf{F}[\psi_\ell] = \sum_{k=1}^{\infty} \mathbf{F}_{k\ell} \phi_k, \quad \mathbf{B}[\phi_k] = \sum_{\ell=1}^{\infty} \mathbf{B}_{\ell k} \psi_\ell. \quad (3.9)$$

Here, $[\mathbf{F}_{k\ell}]$ is a semi-infinite matrix with each column describing the $L_{\mathcal{X}}^2$ representation of $\mathbf{F}[\psi_\ell]$ under the basis $\{\phi_k \in L_{\mathcal{X}}^2\}_{k \in \mathbb{N}}$. With a slight abuse of notation, we will write \mathbf{F} and \mathbf{B} to denote these matrices $[\mathbf{F}_{k\ell}]$ and $[\mathbf{B}_{\ell k}]$. Then, we will prove in Section C that the objective function in (3.5) with $c_{\mathcal{X}} = K_{\mathcal{X}}$ and $c_{\mathcal{Y}} = K_{\mathcal{Y}}$ is

$$\begin{aligned} \Omega(\mathbf{F}, \mathbf{B}) &:= \frac{1}{mn} \sum_{i,j} (\Psi_{y_j}^\top \mathbf{B} \Lambda \Phi_{x_i} - \Phi_{x_i}^\top \mathbf{F} \Gamma \Psi_{y_j})^2 \\ &+ \lambda_1 \cdot \left(\frac{1}{m^2} \sum_{i,i'} \Phi_{x_i}^\top \Lambda \Phi_{x_i'} \Phi_{x_i}^\top \mathbf{F} \Gamma \mathbf{F}^\top \Phi_{x_i'} + \frac{1}{n^2} \sum_{j,j'} \Psi_{y_j}^\top \Gamma \Psi_{y_j'} \Psi_{y_j}^\top \mathbf{B} \Lambda \mathbf{B}^\top \Psi_{y_j'} \right. \\ &\quad \left. - \frac{2}{mn} \sum_{i,j} \Psi_{y_j}^\top \mathbf{B} \Lambda \Phi_{x_i} \Phi_{x_i}^\top \mathbf{F} \Gamma \Psi_{y_j} \right) \\ &+ \lambda_2 \cdot \left(\frac{1}{m^2} \sum_{i,i'} \Phi_{x_i}^\top \Lambda \Phi_{x_i'} + \frac{1}{n^2} \sum_{j,j'} \Psi_{y_j}^\top \mathbf{B} \Lambda \mathbf{B}^\top \Psi_{y_j'} - \frac{2}{mn} \sum_{i,j} \Psi_{y_j}^\top \mathbf{B} \Lambda \Phi_{x_i} \right) \\ &+ \lambda_3 \cdot \left(\frac{1}{m^2} \sum_{i,i'} \Phi_{x_i}^\top \mathbf{F} \Gamma \mathbf{F}^\top \Phi_{x_i'} + \frac{1}{n^2} \sum_{j,j'} \Psi_{y_j}^\top \Gamma \Psi_{y_j'} - \frac{2}{mn} \sum_{i,j} \Phi_{x_i}^\top \mathbf{F} \Gamma \Psi_{y_j} \right). \end{aligned}$$

Here, \mathbf{F} and \mathbf{B} are the matrices denoting the operators induced by F and B , respectively, $\Phi_x = [\dots, \phi_k(x), \dots]^\top \in \mathbb{R}^\infty$ and $\Psi_y = [\dots, \psi_\ell(y), \dots]^\top \in \mathbb{R}^\infty$ for any $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, and $\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots)$ and $\Gamma =$

$\text{diag}(\gamma_1, \gamma_2, \dots)$ are diagonal matrices. Hence, (3.5) can be lifted to an infinite-dimensional optimization problem

$$\min_{(\mathbf{F}, \mathbf{B}) \in \mathcal{C}} \Omega(\mathbf{F}, \mathbf{B}), \quad (3.10)$$

where \mathcal{C} denotes the constraint set implying that \mathbf{F} and \mathbf{B} are matrices corresponding to bounded linear operators induced by some maps $F: \mathcal{X} \rightarrow \mathcal{Y}$ and $B: \mathcal{Y} \rightarrow \mathcal{X}$.

We will relax this problem by removing the constraint set \mathcal{C} , namely, by considering all matrices in $\mathbb{R}^{\infty \times \infty}$ as the decision variables,

$$\min_{\mathbf{F}, \mathbf{B} \in \mathbb{R}^{\infty \times \infty}} \Omega(\mathbf{F}, \mathbf{B}). \quad (3.11)$$

In other words, this relaxed problem minimizes Ω over any pair of infinite-dimensional matrices. The next result, which we refer to as the representer theorem, shows that (3.11) boils down to a finite-dimensional convex program.

Theorem 5. *Consider the optimization (3.10) under the assumptions in Proposition 10. Then, for any minimizer $(\mathbf{F}^*, \mathbf{B}^*)$ to the relaxed problem (3.11), we can find finite-dimensional matrices $\mathbf{F}_{m,n}^* \in \mathbb{R}^{m \times n}$ and $\mathbf{B}_{n,m}^* \in \mathbb{R}^{n \times m}$ such that*

$$\mathbf{F}^* = \Lambda \Phi_m \mathbf{F}_{m,n}^* \Psi_n^\top, \quad \mathbf{B}^* = \Gamma \Psi_n \mathbf{B}_{n,m}^* \Phi_m^\top,$$

where $\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots)$, $\Gamma = \text{diag}(\gamma_1, \gamma_2, \dots)$, and $\Phi_m \in \mathbb{R}^{\infty \times m}$ and $\Psi_n \in \mathbb{R}^{\infty \times n}$ are matrices whose elements are $\phi_k(x_i)$ and $\psi_\ell(y_j)$, as defined in (3.8). In this case, $\Omega(\mathbf{F}^*, \mathbf{B}^*)$ can be rewritten as $\omega(\mathbf{F}_{m,n}^*, \mathbf{B}_{n,m}^*)$ for some convex function ω defined over $\mathbb{R}^{m \times n} \times \mathbb{R}^{n \times m}$. Hence, by minimizing ω over $\mathbb{R}^{m \times n} \times \mathbb{R}^{n \times m}$, we obtain a relaxation of (3.11), that is,

$$\min_{\mathbf{F}, \mathbf{B} \in \mathbb{R}^{\infty \times \infty}} \Omega(\mathbf{F}, \mathbf{B}) \geq \min_{\substack{\mathbf{F}_{m,n} \in \mathbb{R}^{m \times n} \\ \mathbf{B}_{n,m} \in \mathbb{R}^{n \times m}}} \omega(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}).$$

In particular, the RHS is a finite-dimensional convex optimization. Lastly, this relaxation is tight, that is,

$$\min_{\mathbf{F}, \mathbf{B} \in \mathbb{R}^{\infty \times \infty}} \Omega(\mathbf{F}, \mathbf{B}) = \min_{\substack{\mathbf{F}_{m,n} \in \mathbb{R}^{m \times n} \\ \mathbf{B}_{n,m} \in \mathbb{R}^{n \times m}}} \omega(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}),$$

if kernel matrices $\mathbf{K}_{\mathcal{X}}$ and $\mathbf{K}_{\mathcal{Y}}$ whose elements are $K_{\mathcal{X}}(x_i, x_{i'})$ and $K_{\mathcal{Y}}(y_j, y_{j'})$, are positive definite.

Remark. Looking inside the proof of Theorem 5, we know the solution to the infinite-dimensional optimization is an operator taking form of $\mathbf{F}^* = \Lambda \Phi_m \mathbf{F}_{m,n}^* \Psi_n^\top$, with a finite-dimensional matrix $\mathbf{F}_{m,n}^* \in \mathbb{R}^{m \times n}$. Therefore, for any $g \in L_{\mathcal{Y}}^2$, we can deduce

$$\mathbf{F}^*[g](x) = \underbrace{K_{\mathcal{X}}(x, X_m)}_{1 \times m} \underbrace{\mathbf{F}_{m,n}^*}_{m \times n} \underbrace{g(Y_n)}_{n \times 1}, \quad (3.12)$$

where $K_{\mathcal{X}}(x, X_m)$ maps each $x \in \mathcal{X}$ to a row vector whose i -th element is $K_{\mathcal{X}}(x, x_i)$ and $g(Y_n)$ denotes a column vector whose j -th element is $g(y_j)$.

Now let's draw a connection between the classic Nadaraya-Watson estimator and (3.12). For now consider a special case: (x_i, y_i) 's are paired with $m = n$. In such a case, Nadaraya-Watson estimator takes the form

$$\sum_{i,j} K_{\mathcal{X}}(x, x_i) \cdot \frac{1}{m} \delta_{i=j} \cdot g(y_j); \quad (3.13)$$

Namely, for a new point x , the corresponding function value $g(y)$ evaluated on its coupled $y = F(x)$ is a weighted average of $g(y_j)$'s according to the affinity $K_{\mathcal{X}}(x, x_i)$. Our solution (3.12) extends the above

nonparametric smoothing idea to the decoupled data case, where the coupling weights $F_{m,n}^*$ is based on a solution to a convex program, with

$$(3.12) = \sum_{i,j} K_{\mathcal{X}}(x, x_i) \cdot F_{m,n}^*[i, j] \cdot g(y_j). \quad (3.14)$$

Lastly, we draw another connection to the Monte-Carlo integration. One downstream task after learning the distribution ν is to perform numerical integration of $g \in L^2_{\mathcal{Y}}$, under the measure $\nu \in \mathcal{P}(\mathcal{Y})$. In our transform sampling framework, this amounts to evaluating $\mathbb{E}_{y \sim F_{\#}^* \mu}[g(y)] = \mathbb{E}_{x \sim \mu}[g \circ F^*(x)]$. The integration, casted in the induced operator form, has the expression

$$\mathbb{E}_{x \sim \mu} [\mathbf{F}^*[g](x)] = \mathbb{E}_{x \sim \mu} \left[\underbrace{K_{\mathcal{X}}(x, X_m) F_{m,n}^*}_{=: W(x) \in \mathbb{R}^n} g(Y_n) \right] = \mathbb{E}_{x \sim \mu} \left[\sum_{j=1}^n W_j(x) g(y_j) \right] \quad (3.15)$$

where $W(x)$ can be interpreted as the importance weights in the Monte-Carlo integration. We conclude with one more remark: if plug in instead $x \sim \hat{\mu}_m$ in (3.15), one can verify that under mild conditions,

$$\mathbb{E}_{x \sim \hat{\mu}_m} [\mathbf{F}^*[g](x)] = \frac{1}{n} \sum_{j=1}^n g(y_j). \quad (3.16)$$

That is, with the empirical measure as input, (3.15) outputs the simple sample average.

4 Experiments

This section examines the empirical performance of the reversible Gromov-Monge sampler. Following Section 3.3, we find a minimum (\hat{F}, \hat{B}) of (3.5) over a suitable class $\mathcal{F} \times \mathcal{B}$ via gradient descent; we inspect the quality of transform sampling ($\hat{F}_{\#} \mu \approx \nu$) and space isomorphism. Complete technical details of the experiments are deferred to Section F.

Gaussian distributions Consider two strongly isomorphic Gaussian distributions on $\mathcal{X} = \mathcal{Y} = \mathbb{R}^2$: the base measure $\mu = N(0, I_2)$ and the target distribution $\nu = N(0, \Sigma)$, where I_2 is the identity matrix and the entries of Σ are $\Sigma_{11} = \Sigma_{22} = 1$ and $\Sigma_{12} = \Sigma_{21} = 0.7$. We let $c_{\mathcal{X}}(x, x') = x^{\top} x'$ and $c_{\mathcal{Y}}(y, y') = y^{\top} \Sigma^{-1} y'$, then two network spaces are strongly isomorphic by design; indeed, any pair (F, B) given by $F(x) = \Sigma^{1/2} Qx$ and $B(y) = Q^{\top} \Sigma^{-1/2} y$ for $Q \in O(2)$, where $O(2)$ is the orthogonal group, yields $c_{\mathcal{X}}(x, B(y)) = c_{\mathcal{Y}}(F(x), y)$ for all $x, y \in \mathbb{R}^2$, hence F and B are strong isomorphisms. We aim at obtaining such a pair of (linear) isomorphisms by letting $\mathcal{F} = \mathcal{B} = \{x \mapsto Wx : W \in \mathbb{R}^{2 \times 2}\}$, that is, the collection of all linear maps from \mathbb{R}^2 to \mathbb{R}^2 . We set $K_{\mathcal{X}} = K_{\mathcal{Y}}$ as a degree-2 polynomial kernel that maps (x, y) to $(x^{\top} y + 1)^2$; the resulting MMD compares distributions by matching the first two moments, which is sufficient to distinguish Gaussian distributions. The resulting linear maps are given by $\hat{F}(x) = \mathbf{F}x$ and $\hat{B}(y) = \mathbf{B}y$ for some $\mathbf{F}, \mathbf{B} \in \mathbb{R}^{2 \times 2}$ satisfying

$$\mathbf{F}\mathbf{F}^{\top} = \begin{pmatrix} 1.035 & 0.751 \\ 0.751 & 1.094 \end{pmatrix}, \quad \mathbf{B}\Sigma\mathbf{B}^{\top} = \begin{pmatrix} 0.940 & 0.001 \\ 0.001 & 0.944 \end{pmatrix}, \quad \mathbf{F}\mathbf{B} = \begin{pmatrix} 0.966 & 0.029 \\ -0.020 & 1.029 \end{pmatrix}.$$

Since $\mathbf{F}\mathbf{F}^{\top} \approx \Sigma$, $\mathbf{B}\Sigma\mathbf{B}^{\top} \approx I_2$, and $\mathbf{F}\mathbf{B} \approx I_2$, the pair (\hat{F}, \hat{B}) can be seen as an instance of the pair of strong isomorphisms described above. Figure 1 illustrates that \hat{F} is a strong isomorphism (Definition 3): (a) shows that $\hat{F}_{\#} \mu \approx \nu$, that is, \hat{F} is roughly a transport map, and (b) implies that $c_{\mathcal{X}}(x, x') \approx c_{\mathcal{Y}}(\hat{F}(x), \hat{F}(x'))$ holds.

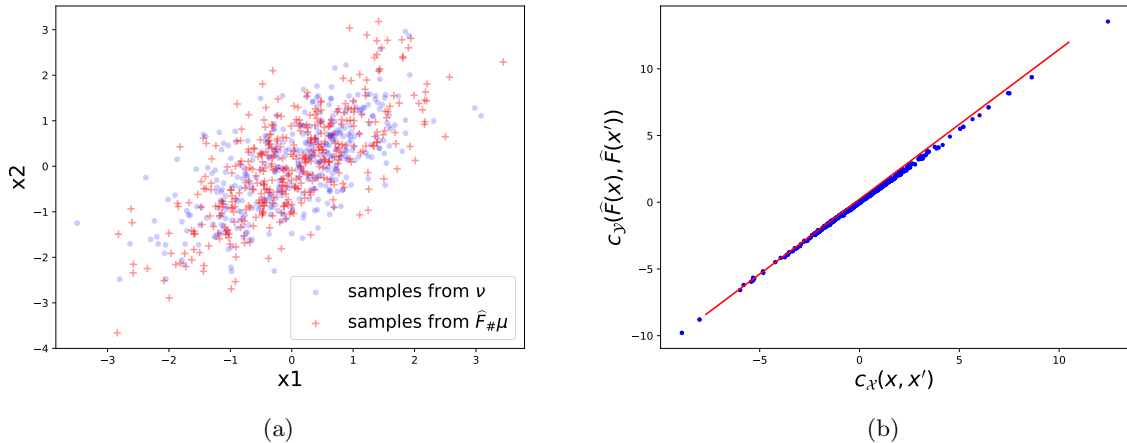


Figure 1: Gaussian experiment: $m = n = 1000$ and $\lambda_1 = \lambda_2 = \lambda_3 = 1$. (a) shows $\{\tilde{y}_j\}_{j=1}^{400}$ versus $\{\widehat{F}(\tilde{x}_i)\}_{i=1}^{400}$, where $\{\tilde{y}_j\}_{j=1}^{400}$ and $\{\tilde{x}_i\}_{i=1}^{400}$ are i.i.d. from $\nu = N(0, \Sigma)$ and $\mu = N(0, I_2)$, respectively; they are new samples independent from $\{y_j\}_{j=1}^{1000}$ and $\{x_i\}_{i=1}^{1000}$ used in (3.5). (b) shows the points $\{(c_{\mathcal{X}}(\tilde{x}_i, \tilde{x}_{i'}), c_{\mathcal{Y}}(\widehat{F}(\tilde{x}_i), \widehat{F}(\tilde{x}_{i'})))\}_{i, i'=1}^{40}$ and a straight line $y = x$.

MNIST Next, let ν be a distribution of images corresponding to four digits (2, 4, 6, 7) from the MNIST data set, which is supported on \mathbb{R}^{784} . Recall from Section 1 that the support \mathcal{Y} of ν is low-dimensional [19], hence choosing $\mathcal{X} = \mathbb{R}^d$ with $d \ll 784$ is reasonable. Here, for visualization, we try an extreme embedding task with $d = 2$ and $\mu = N(0, I_2)$, that is, generate MNIST images by transforming two-dimensional Gaussian samples.

Unlike the Gaussian example where we design the cost functions in advance to make the two spaces strongly isomorphic, specifying them can be more complicated in general cases, which might affect the quality of the RGM sampler. Here, we briefly discuss some of the most commonly used cost functions: given a fixed exponent $p \in \mathbb{N}$ or constant $\alpha > 0$,

$$(x, y) \mapsto \underbrace{\|x - y\|^p}_{\text{distance-based}} \quad \text{or} \quad \underbrace{\exp(-\alpha\|x - y\|^2)}_{\text{RBF kernel}}.$$

Clearly, $\|x - y\|^p$ is the most straightforward choice in Euclidean cases; $p = 1$ and $p = 2$ are indeed widely used in the literature [43]. The RBF kernel, also referred to as the heat kernel, is a common choice in the object matching literature [49]. In this MNIST example, we have found that these cost functions provide reasonable performance once they are scaled properly. Here, we will present the results based on the RBF kernel. Concretely, first define the RBF kernel $K_d(x, y) = \exp(-\|x - y\|^2/d)$ for $d \in \mathbb{N}$ and $x, y \in \mathbb{R}^d$; here, the constant $(1/d)$ serves as a scaling factor. Then, we define the cost functions as $c_{\mathcal{X}} = (K_2 - m_{\mathcal{X}})/\text{sd}_{\mathcal{X}}$ and $c_{\mathcal{Y}} = (K_{784} - m_{\mathcal{Y}})/\text{sd}_{\mathcal{Y}}$, where $m_{\mathcal{X}}$ and $\text{sd}_{\mathcal{X}}$ are the median and the standard error of $\{K_{\mathcal{X}}(x_i, x_{i'})\}_{i, i'=1}^m$, respectively; $m_{\mathcal{Y}}$ and $\text{sd}_{\mathcal{Y}}$ are defined analogously. This additional standardization process helps aligning the cost functions.

In the same vein, $K_{\mathcal{X}}$ and $K_{\mathcal{Y}}$ must be properly specified; comparing the first two moments using the degree-2 polynomial kernel is no longer sufficient as the target distribution is non-Gaussian. We suggest using RBF kernels for the MMD terms as well; let $K_{\mathcal{X}} = K_2$ and $K_{\mathcal{Y}} = K_{784}$. The MMD induced by the RBF kernel indeed defines a metric between distributions under mild assumptions [41], which allows the resulting MMD terms to represent the original constraint of the RGM distance as mentioned in Section 3.2.

For the function classes \mathcal{F} and \mathcal{B} , we need richer classes instead of the linear maps used in the Gaussian

case. To this end, we will use fully connected neural networks with three hidden layers, each of which consists of 50 neurons. Lastly, we let $m = n = 20000$ and $\lambda_1 = \lambda_2 = \lambda_3 = 100$. Figure 2(a) shows the images generated by applying the resulting map \widehat{F} to new i.i.d. samples from $\mu = N(0, I_2)$. Though not perfect, we see that recognizable images can be generated by transforming two-dimensional Gaussian samples, efficient in computation.⁵ Meanwhile, the map \widehat{B} shows how the MNIST images can be embedded in \mathbb{R}^2 . Figure

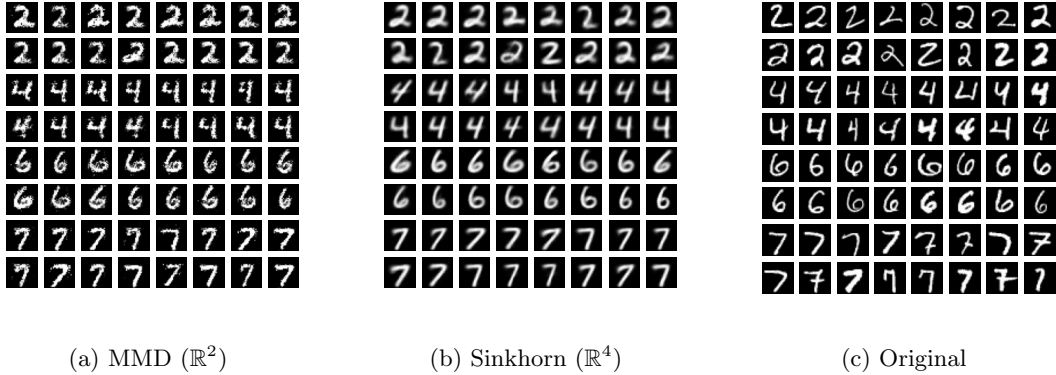


Figure 2: (a) and (b) are generated by transforming new i.i.d. samples from μ using \widehat{F} : (a) from $\mu = N(0, I_2)$ with MMDs and (b) from $\mu = N(0, I_4)$ with Sinkhorn divergences. (c) shows real MNIST images.

3(a) shows $\{\widehat{B}(\tilde{y}_j)\}_{j=1}^{500}$, where $\{\tilde{y}_j\}_{j=1}^{500}$ are i.i.d. from ν (125×4 digits), independent from $\{y_j\}_{j=1}^{20000}$ used in (3.5). We see that each digit forms a local cluster in \mathbb{R}^2 , each of which is roughly representable according to the range of the angular coordinate. Lastly, though not perfect as in Figure 1(b) (strongly isomorphic case), Figure 3(b) shows that \widehat{B} leads to a reasonable alignment of $c_{\mathcal{X}}(\widehat{B}(y), \widehat{B}(y'))$ versus $c_{\mathcal{Y}}(y, y')$.




Figure 3: (a) is generated by applying \widehat{B} to 500 out-of-sample MNIST images, i.i.d. $\{\tilde{y}_j\}_{j=1}^{500}$ from ν . (b) shows the points $\{(c_{\mathcal{X}}(\widehat{B}(\tilde{y}_j), \widehat{B}(\tilde{y}_{j'})), c_{\mathcal{Y}}(\tilde{y}_j, \tilde{y}_{j'}))\}_{j, j'=1}^{50}$ and a straight line $y = x$.

⁵Computational cost for obtaining $\widehat{F}: \mathbb{R}^2 \rightarrow \mathbb{R}^{784}$ and computing $\widehat{F}(X)$ from $X \sim \mu$ is far less than that of the OT-based sampler as explained in Section 1.

We clarify that the current experiment with $\mu = N(0, I_2)$ is a proof of concept. Suppose one aims to obtain images comparable to those from dedicated MNIST generators. In that case, exhaustive tests should be done for tuning each component of the RGM sampler, which is beyond the scope of this paper. Instead, we highlight that the RGM sampler with a simple modification can indeed generate significantly improved images seen in Figure 2(b). These images are generated from the following settings that are fully introduced in Section F: $(\mathcal{X}, \mu, c_{\mathcal{X}}) = (\mathbb{R}^4, N(0, I_4), K_4)$ and MMD terms in (3.5) are replaced by Sinkhorn divergences [23]. As such, the RGM sampler is amendable to other more general choices of $\mathcal{L}_{\mathcal{X} \times \mathcal{Y}}$ in its practical implementation.

5 Discussions

In this work, we proposed the Reversible Gromov-Monge (RGM) sampler, a new variant of transform sampling based on the RGM distance operating between distributions on heterogeneous spaces. We discuss the following important aspects of the RGM sampler.

Inductive bias Inductive bias alleviates the over-identified issue in transform sampling, as seen in Section 1. RGM sampler induces a bias towards finding approximate isomorphisms: namely among all possible $(F, B) \in \mathcal{I}(\mu, \nu)$, RGM sampler favors the ones that $c_{\mathcal{X}}(x, B(y))$ best matches $c_{\mathcal{Y}}(F(x), y)$ on average. In particular, if the user of the RGM sampler specifies suitable $c_{\mathcal{X}}, c_{\mathcal{Y}}$ such that the resulting network spaces are strongly isomorphic, the RGM sampler will identify a strong isomorphism, providing a beneficial inductive bias to transform sampling. Intriguingly, borrowing insights from Brenier’s polar factorization [9], we can always find $c_{\mathcal{X}}, c_{\mathcal{Y}}$ making two Polish spaces strongly isomorphic under mild conditions. Details are provided in Section A.3.

Optimization We have introduced two optimization methods for computing RGM given finite samples and proved their properties. As explained in Section 3.3, minimization via the gradient descent algorithm is easy to implement and less complicated than the min-max optimization of GAN-type samplers. Nevertheless, it is unclear whether one can derive global convergence results from this minimization. In general, different optimization algorithms of transform samplers might output maps that favor distinctive inductive biases. Such a phenomenon is an important research direction in the optimization literature.

Acknowledgments

Liang acknowledges the generous support from the NSF Career award (DMS-2042473), and the William Ladany faculty fellowship at the University of Chicago Booth School of Business. Liang wishes to thank Maxim Raginsky and Chris Hansen for discussions on simulation-based inference.

References

- [1] Anthony, M. and Bartlett, P. L. (1999). *Neural Network Learning: Theoretical Foundations*. Cambridge University Press.
- [2] Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223. PMLR.
- [3] Arora, S., Ge, R., Liang, Y., Ma, T., and Zhang, Y. (2017). Generalization and equilibrium in generative adversarial nets (gans). In *International Conference on Machine Learning*, pages 224–232. PMLR.
- [4] Bai, Y., Ma, T., and Risteski, A. (2018). Approximability of discriminators implies diversity in gans. *arXiv preprint arXiv:1806.10586*.

- [5] Bassily, R., Belkin, M., and Ma, S. (2018). On exponential convergence of sgd in non-convex over-parametrized learning. *arXiv preprint arXiv:1811.02564*.
- [6] Blumberg, A. J., Carriere, M., Mandell, M. A., Rabadan, R., and Villar, S. (2020). Mrec: a fast and versatile framework for aligning and matching point clouds with applications to single cell molecular data. *arXiv preprint arXiv:2001.01666*.
- [7] Boucheron, S., Lugosi, G., and Massart, P. (2013). *Concentration inequalities. A nonasymptotic theory of independence*. Oxford University Press.
- [8] Bréchet, C. (2019). A statistical test of isomorphism between metric-measure spaces using the distance-to-a-measure signature. *Electron. J. Stat.*, 13(1):795–849.
- [9] Brenier, Y. (1991). Polar factorization and monotone rearrangement of vector-valued functions. *Communications on Pure and Applied Mathematics*, 44(4):375–417.
- [10] Brenier, Y. and Gangbo, W. (2003). L^p approximation of maps by diffeomorphisms. *Calc. Var. Partial Differential Equations*, 16(2):147–164.
- [11] Çela, E. (1998). *The quadratic assignment problem: theory and algorithms*. Springer.
- [12] Chen, M., Liao, W., Zha, H., and Zhao, T. (2020). Statistical guarantees of generative adversarial networks for distribution estimation. *arXiv preprint arXiv:2002.03938*.
- [13] Chowdhury, S. and Mémoli, F. (2019). The Gromov–Wasserstein distance between networks and stable network invariants. *Information and Inference: A Journal of the IMA*, 8(4):757–787.
- [14] Chowdhury, S., Miller, D. W., and Needham, T. (2021). Quantized gromov-wasserstein. In *ECML/PKDD*.
- [15] Daskalakis, C., Ilyas, A., Syrgkanis, V., and Zeng, H. (2017). Training gans with optimism. *arXiv preprint arXiv:1711.00141*.
- [16] Diamond, S. and Boyd, S. (2016). CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5.
- [17] Doersch, C. (2016). Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*.
- [18] Dziugaite, G. K., Roy, D. M., and Ghahramani, Z. (2015). Training generative neural networks via Maximum Mean Discrepancy optimization. In *Proceedings of the 31st Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 258–267.
- [19] Facco, E., d’Errico, M., Rodriguez, A., and Laio, A. (2017). Estimating the intrinsic dimension of datasets by a minimal neighborhood information. *Sci. Rep.*, 7(1):1–8.
- [20] Feydy, J., Séjourné, T., Vialard, F.-X., Amari, S.-i., Trounev, A., and Peyré, G. (2019). Interpolating between optimal transport and mmd using sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 2681–2690.
- [21] Flamary, R., Courty, N., Gramfort, A., Alaya, M. Z., Boisbunon, A., Chambon, S., Chapel, L., Corenflos, A., Fatras, K., Fournier, N., et al. (2021). Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8.
- [22] Gellert, M., Hossain, M. F., Berens, F. J. F., Bruhn, L. W., Urbainsky, C., Liebscher, V., and Lillig, C. H. (2019). Substrate specificity of thioredoxins and glutaredoxins—towards a functional classification. *Heliyon*, 5(12):e02943.

- [23] Genevay, A., Peyré, G., and Cuturi, M. (2018). Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617. PMLR.
- [24] 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*, volume 27.
- [25] Gouriéroux, C. and Monfort, A. (1997). *Simulation-Based Econometric Methods*. Oxford University Press.
- [26] Harvey, N., Liaw, C., and Mehrabian, A. (2017). Nearly-tight VC-dimension bounds for piecewise linear neural networks. In *Conference on Learning Theory*, pages 1064–1068. PMLR.
- [27] Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [28] Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- [29] Koopmans, T. C. and Beckmann, M. (1957). Assignment Problems and the Location of Economic Activities. *Econometrica*, 25(1):53–76.
- [30] LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324.
- [31] Lei, Q., Lee, J. D., Dimakis, A. G., and Daskalakis, C. (2019). Sgd learns one-layer networks in wgens. *arXiv preprint arXiv:1910.07030*.
- [32] Li, Y., Swersky, K., and Zemel, R. (2015). Generative moment matching networks. In *International Conference on Machine Learning*, pages 1718–1727. PMLR.
- [33] Liang, T. (2019). Estimating certain integral probability metric (IPM) is as hard as estimating under the IPM. *arXiv preprint arXiv:1911.00730*.
- [34] Liang, T. (2021). How well generative adversarial networks learn distributions. *Journal of Machine Learning Research*, 22(228):1–41.
- [35] Liang, T. and Stokes, J. (2019). Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. In *International Conference on Artificial Intelligence and Statistics*, volume 89, pages 907–915. PMLR.
- [36] Makuva, A., Taghvaei, A., Oh, S., and Lee, J. (2020). Optimal transport mapping via input convex neural networks. In *International Conference on Machine Learning*, pages 6672–6681. PMLR.
- [37] McFadden, D. (1989). A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration. *Econometrica*, 57(5):995–1026.
- [38] Mémoli, F. (2011). Gromov–Wasserstein Distances and the Metric Approach to Object Matching. *Foundations of Computational Mathematics*, 11(4):417–487.
- [39] Mémoli, F. and Needham, T. (2018). Distance distributions and inverse problems for metric measure spaces. *arXiv preprint arXiv:1810.09646*.
- [40] Mroueh, Y., Li, C.-L., Sercu, T., Raj, A., and Cheng, Y. (2017). Sobolev gan. *arXiv preprint arXiv:1711.04894*.

- [41] Muandet, K., Fukumizu, K., Sriperumbudur, B., and Schölkopf, B. (2017). Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning*, 10(1-2):1–141.
- [42] Pakes, A. and Pollard, D. (1989). Simulation and the Asymptotics of Optimization Estimators. *Econometrica*, 57(5):1027–1057.
- [43] Peyré, G., Cuturi, M., and Solomon, J. (2016). Gromov-wasserstein averaging of kernel and distance matrices. In *International Conference on Machine Learning*, pages 2664–2672. PMLR.
- [44] Robert, C. P. and Casella, G. (1999). *Monte Carlo statistical methods*. Springer.
- [45] Royden, H. L. (1988). *Real analysis*. Macmillan, third edition.
- [46] Scetbon, M., Peyré, G., and Cuturi, M. (2021). Linear-time gromov wasserstein distances using low rank couplings and costs. *arXiv preprint arXiv:2106.01128*.
- [47] Silverman, B. W. (1986). *Density estimation for statistics and data analysis*. Chapman & Hall.
- [48] Singh, S. and Póczos, B. (2018). Minimax distribution estimation in wasserstein distance. *arXiv preprint arXiv:1802.08855*.
- [49] Solomon, J., Peyré, G., Kim, V. G., and Sra, S. (2016). Entropic metric alignment for correspondence problems. *ACM Trans. Graph.*, 35(4).
- [50] Sriperumbudur, B. K., Fukumizu, K., Gretton, A., Schölkopf, B., and Lanckriet, G. R. (2012). On the empirical estimation of integral probability metrics. *Electron. J. Stat.*, 6:1550–1599.
- [51] Stone, C. J. (1982). Optimal global rates of convergence for nonparametric regression. *Ann. Statist.*, 10(4):1040–1053.
- [52] Sturm, K.-T. (2012). The space of spaces: curvature bounds and gradient flows on the space of metric measure spaces. *arXiv preprint arXiv:1208.0434*.
- [53] Taghvaei, A. and Jalali, A. (2019). 2-Wasserstein Approximation via Restricted Convex Potentials with Application to Improved Training for GANs. *arXiv preprint arXiv:1902.07197*.
- [54] Titouan, V., Flamary, R., Courty, N., Tavenard, R., and Chapel, L. (2019). Sliced gromov-wasserstein. In *Advances in Neural Information Processing Systems*, volume 32.
- [55] Wainwright, M. J. (2019). *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge University Press.
- [56] Weed, J. and Berthet, Q. (2019). Estimation of smooth densities in wasserstein distance. *arXiv preprint arXiv:1902.01778*.
- [57] Weitkamp, C. A., Proksch, K., Tameling, C., and Munk, A. (2020). Gromov-wasserstein distance based object matching: asymptotic inference. *arXiv preprint arXiv:2006.12287*.
- [58] Xu, H., Luo, D., and Carin, L. (2019). Scalable gromov-wasserstein learning for graph partitioning and matching. In *Advances in Neural Information Processing Systems*, volume 32.

SUPPLEMENTARY MATERIAL

This supplementary material collects details of Sections 3 and 4 along with relevant discussions and technicalities.

- Section A studies analytic properties of the RGM distance along with the details of the inductive bias of the RGM sampler based on Brenier’s polar factorization.
- Section B derives the non-asymptotic rate of convergence by analyzing the statistical properties of the RGM sampler.
- Section C discusses a further relaxation of the RGM into an infinite-dimensional convex program that relies on a new representer theorem.
- Section D contains the proofs of the results in B and auxiliary lemmas.
- Section E discusses computational aspects of the RGM distance.
- Section F provides the implementation details of the experiments in Section 4.

A Analytic Properties and Inductive Biases Based on Polar Factorization

A.1 Metric Properties and Some Basic Properties

In this section, we derive metric properties of the proposed RGM distance. First, observe that our RGM is symmetric while the original GM is not. Next, we prove a triangle inequality using a gluing technique, as in OT.

Proposition 2. *RGM satisfies the triangle inequality, that is,*

$$\text{RGM}(\mu_{\mathcal{X}}, \mu_{\mathcal{Z}}) \leq \text{RGM}(\mu_{\mathcal{X}}, \mu_{\mathcal{Y}}) + \text{RGM}(\mu_{\mathcal{Y}}, \mu_{\mathcal{Z}})$$

holds for three network spaces $(\mathcal{X}, \mu_{\mathcal{X}}, c_{\mathcal{X}})$, $(\mathcal{Y}, \mu_{\mathcal{Y}}, c_{\mathcal{Y}})$, and $(\mathcal{Z}, \mu_{\mathcal{Z}}, c_{\mathcal{Z}})$.

Proof of Proposition 2. Recall that


$$\text{RGM}(\mu_{\mathcal{X}}, \mu_{\mathcal{Z}}) = \inf_{(F,B) \in \mathcal{I}(\mu_{\mathcal{X}}, \mu_{\mathcal{Z}})} C_{\mathcal{X}\mathcal{Z}}(F, B),$$

where

$$C_{\mathcal{X}\mathcal{Z}}(F, B) = \left(\int (c_{\mathcal{X}}(x, B(z)) - c_{\mathcal{Z}}(F(x), z))^2 d\mu_{\mathcal{X}} \otimes \mu_{\mathcal{Z}}(x, z) \right)^{1/2}.$$

First, $(F_{\mathcal{Z}} \circ F_{\mathcal{Y}}, B_{\mathcal{X}} \circ B_{\mathcal{Y}}) \in \mathcal{I}(\mu_{\mathcal{X}}, \mu_{\mathcal{Z}})$ holds for $(F_{\mathcal{Y}}, B_{\mathcal{X}}) \in \mathcal{I}(\mu_{\mathcal{X}}, \mu_{\mathcal{Y}})$ and $(F_{\mathcal{Z}}, B_{\mathcal{Y}}) \in \mathcal{I}(\mu_{\mathcal{Y}}, \mu_{\mathcal{Z}})$ since

$$\begin{aligned} (\text{Id}, F_{\mathcal{Z}} \circ F_{\mathcal{Y}})_{\#} \mu_{\mathcal{X}} &= (\text{Id}, F_{\mathcal{Z}})_{\#} (\text{Id}, F_{\mathcal{Y}})_{\#} \mu_{\mathcal{X}} \\ &= (\text{Id}, F_{\mathcal{Z}})_{\#} (B_{\mathcal{X}}, \text{Id})_{\#} \mu_{\mathcal{Y}} \quad (\because (F_{\mathcal{Y}}, B_{\mathcal{X}}) \in \mathcal{I}(\mu_{\mathcal{X}}, \mu_{\mathcal{Y}})) \\ &= (B_{\mathcal{X}}, \text{Id})_{\#} (\text{Id}, F_{\mathcal{Z}})_{\#} \mu_{\mathcal{Y}} \\ &= (B_{\mathcal{X}}, \text{Id})_{\#} (B_{\mathcal{Y}}, \text{Id})_{\#} \mu_{\mathcal{Z}} \quad (\because (F_{\mathcal{Z}}, B_{\mathcal{Y}}) \in \mathcal{I}(\mu_{\mathcal{Y}}, \mu_{\mathcal{Z}})) \\ &= (B_{\mathcal{X}} \circ B_{\mathcal{Y}}, \text{Id})_{\#} \mu_{\mathcal{Z}}. \end{aligned}$$



Moreover, in this case, we have

$$C_{\mathcal{X}\mathcal{Z}}(F_{\mathcal{Z}} \circ F_{\mathcal{Y}}, B_{\mathcal{X}} \circ B_{\mathcal{Y}}) \leq C_{\mathcal{X}\mathcal{Y}}(F_{\mathcal{Y}}, B_{\mathcal{X}}) + C_{\mathcal{Y}\mathcal{Z}}(F_{\mathcal{Z}}, B_{\mathcal{Y}})$$

since

$$\begin{aligned} & C_{\mathcal{X}\mathcal{Z}}(F_{\mathcal{Z}} \circ F_{\mathcal{Y}}, B_{\mathcal{X}} \circ B_{\mathcal{Y}}) \\ &= \left(\int [c_{\mathcal{X}}(x, B_{\mathcal{X}} \circ B_{\mathcal{Y}}(z)) - c_{\mathcal{Z}}(F_{\mathcal{Z}} \circ F_{\mathcal{Y}}(x), z)]^2 d\mu_{\mathcal{X}} \otimes \mu_{\mathcal{Z}}(x, z) \right)^{1/2} \\ &\leq \left(\int \int [c_{\mathcal{Y}}(x, B_{\mathcal{X}} \circ B_{\mathcal{Y}}(z)) - c_{\mathcal{Y}}(F_{\mathcal{Y}}(x), B_{\mathcal{Y}}(z))]^2 d\mu_{\mathcal{Z}}(z) d\mu_{\mathcal{X}}(x) \right)^{1/2} \\ &\quad + \left(\int \int [c_{\mathcal{Y}}(F_{\mathcal{Y}}(x), B_{\mathcal{Y}}(z)) - c_{\mathcal{Z}}(F_{\mathcal{Z}} \circ F_{\mathcal{Y}}(x), z)]^2 d\mu_{\mathcal{X}}(x) d\mu_{\mathcal{Z}}(z) \right)^{1/2} \\ &= \left(\int \int [c_{\mathcal{Y}}(x, B_{\mathcal{X}}(y)) - c_{\mathcal{Y}}(F_{\mathcal{Y}}(x), y)]^2 d\mu_{\mathcal{Y}}(y) d\mu_{\mathcal{X}}(x) \right)^{1/2} \quad (\because (B_{\mathcal{Y}})_{\#}\mu_{\mathcal{Z}} = \mu_{\mathcal{Y}}) \\ &\quad + \left(\int \int [c_{\mathcal{Y}}(y, B_{\mathcal{Y}}(z)) - c_{\mathcal{Z}}(F_{\mathcal{Z}}(y), z)]^2 d\mu_{\mathcal{Y}}(y) d\mu_{\mathcal{Z}}(z) \right)^{1/2} \quad (\because (F_{\mathcal{Y}})_{\#}\mu_{\mathcal{X}} = \mu_{\mathcal{Y}}) \\ &= C_{\mathcal{X}\mathcal{Y}}(F_{\mathcal{Y}}, B_{\mathcal{X}}) + C_{\mathcal{Y}\mathcal{Z}}(F_{\mathcal{Z}}, B_{\mathcal{Y}}). \end{aligned}$$

Hence,

$$\begin{aligned} \text{RGM}(\mu_{\mathcal{X}}, \mu_{\mathcal{Z}}) &= \inf_{(F, B) \in \mathcal{I}(\mu_{\mathcal{X}}, \mu_{\mathcal{Z}})} C_{\mathcal{X}\mathcal{Z}}(F, B) \\ &\leq \inf_{\substack{(F_{\mathcal{Y}}, B_{\mathcal{X}}) \in \mathcal{I}(\mu_{\mathcal{X}}, \mu_{\mathcal{Y}}) \\ (F_{\mathcal{Z}}, B_{\mathcal{Y}}) \in \mathcal{I}(\mu_{\mathcal{Y}}, \mu_{\mathcal{Z}})}} C_{\mathcal{X}\mathcal{Z}}(F_{\mathcal{Z}} \circ F_{\mathcal{Y}}, B_{\mathcal{X}} \circ B_{\mathcal{Y}}) \\ &\leq \inf_{(F_{\mathcal{Y}}, B_{\mathcal{X}}) \in \mathcal{I}(\mu_{\mathcal{X}}, \mu_{\mathcal{Y}})} C_{\mathcal{X}\mathcal{Y}}(F_{\mathcal{Y}}, B_{\mathcal{X}}) + \inf_{(F_{\mathcal{Z}}, B_{\mathcal{Y}}) \in \mathcal{I}(\mu_{\mathcal{Y}}, \mu_{\mathcal{Z}})} C_{\mathcal{Y}\mathcal{Z}}(F_{\mathcal{Z}}, B_{\mathcal{Y}}) \\ &= \text{RGM}(\mu_{\mathcal{X}}, \mu_{\mathcal{Y}}) + \text{RGM}(\mu_{\mathcal{Y}}, \mu_{\mathcal{Z}}). \end{aligned}$$

□

Next, we study whether $\text{RGM}(\mu, \nu) = 0$ holds if and only if $(\mathcal{X}, \mu, c_{\mathcal{X}}) \cong (\mathcal{Y}, \nu, c_{\mathcal{Y}})$. Here the equivalence relation induced by \cong can be read from Definition 3. As in the Gromov-Wasserstein distance, in general, we can only assert the if part without further conditions. The following proposition states that $\text{RGM}(\mu, \nu) = 0$ if and only if $(\mathcal{X}, \mu, c_{\mathcal{X}}) \cong (\mathcal{Y}, \nu, c_{\mathcal{Y}})$ under some additional conditions on $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$, thereby implying Theorem 2.

Proposition 3. *Let $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ be two network spaces. If $(\mathcal{X}, \mu, c_{\mathcal{X}}) \cong (\mathcal{Y}, \nu, c_{\mathcal{Y}})$, then $\text{RGM}(\mu, \nu) = 0$. The converse is true if there exists a continuous and strictly monotone function $h: \mathbb{R}_+ \rightarrow \mathbb{R}$ such that $c_{\mathcal{X}} = h(d_{\mathcal{X}})$ and $c_{\mathcal{Y}} = h(d_{\mathcal{Y}})$.*

Proof of Proposition 3. Suppose $\text{RGM}(\mu, \nu) = 0$. Due to the inequality $\text{GW}(\mu, \nu) \leq \text{RGM}(\mu, \nu)$, we have $\text{GW}(\mu, \nu) = 0$, that is,

$$\inf_{\gamma \in \Pi(\mu, \nu)} \left(\int_{\mathcal{X} \times \mathcal{Y}} \int_{\mathcal{X} \times \mathcal{Y}} (h(d_{\mathcal{X}}(x, x')) - h(d_{\mathcal{Y}}(y, y')))^2 d\gamma(x, y) d\gamma(x', y') \right)^{1/2} = 0.$$

Since there exists a coupling γ^* that achieves the minimum of GW due to Theorem 2.2 of [13], we conclude

$$h(d_{\mathcal{X}}(x, x')) = h(d_{\mathcal{Y}}(y, y'))$$

holds $\gamma^* \otimes \gamma^*$ almost surely on $(\mathcal{X} \times \mathcal{Y})^2$. Since h is strictly monotone, this means

$$d_{\mathcal{X}}(x, x') = d_{\mathcal{Y}}(y, y')$$

holds $\gamma^* \otimes \gamma^*$ almost surely on $(\mathcal{X} \times \mathcal{Y})^2$. Therefore,

$$\begin{aligned} & \inf_{\gamma \in \Pi(\mu, \nu)} \left(\int_{\mathcal{X} \times \mathcal{Y}} \int_{\mathcal{X} \times \mathcal{Y}} (d_{\mathcal{X}}(x, x') - d_{\mathcal{Y}}(y, y'))^2 d\gamma(x, y) d\gamma(x', y') \right)^{1/2} \\ & \geq \left(\int_{\mathcal{X} \times \mathcal{Y}} \int_{\mathcal{X} \times \mathcal{Y}} (d_{\mathcal{X}}(x, x') - d_{\mathcal{Y}}(y, y'))^2 d\gamma^*(x, y) d\gamma^*(x', y') \right)^{1/2} \\ & = 0. \end{aligned}$$

Theorem 1 implies that metric measure spaces $(\mathcal{X}, \mu, d_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, d_{\mathcal{Y}})$ are strongly isomorphic. Since $c_{\mathcal{X}} = h(d_{\mathcal{X}})$ and $c_{\mathcal{Y}} = h(d_{\mathcal{Y}})$, it follows easily that $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ are strongly isomorphic as well.

To prove the if part, suppose $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ are strongly isomorphic and consider a strong isomorphism T . Then, $(T, T^{-1}) \in \mathcal{I}(\mu, \nu)$ holds since $(\text{Id}, T)_{\#}\mu = (T^{-1}, \text{Id})_{\#}T_{\#}\mu = (T^{-1}, \text{Id})_{\#}\nu$. Also, by definition of T , we have $c_{\mathcal{X}}(x, T^{-1}(y)) = c_{\mathcal{Y}}(T(x), T \circ T^{-1}(y)) = c_{\mathcal{Y}}(T(x), y)$ for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$, thus

$$\text{RGM}(\mu, \nu) \leq \int_{\mathcal{X} \times \mathcal{Y}} (c_{\mathcal{X}}(x, T^{-1}(y)) - c_{\mathcal{Y}}(T(x), y))^2 d\mu \otimes \nu = 0.$$

□

We conclude this section with a few more properties and examples. We first complete the proof of Proposition 1, which characterizes the relations among three distances: GW, GM, and RGM.

Proof of Proposition 1. Define

$$Q(\gamma) = \int_{\mathcal{X} \times \mathcal{Y}} \int_{\mathcal{X} \times \mathcal{Y}} (c_{\mathcal{X}}(x, x') - c_{\mathcal{Y}}(y, y'))^2 d\gamma(x, y) d\gamma(x', y') \quad (\text{A.1})$$

for all $\gamma \in \Pi(\mu, \nu)$ so that

$$\text{GW}(\mu, \nu)^2 = \inf_{\gamma \in \Pi(\mu, \nu)} Q(\gamma).$$

Recall that $\Pi_{\mathcal{T}} = \{(\text{Id}, T)_{\#}\mu : T \in \mathcal{T}(\mu, \nu)\} \subset \Pi(\mu, \nu)$. Hence, as noted in Section 2,

$$\text{GM}(\mu, \nu)^2 = \inf_{\gamma \in \Pi_{\mathcal{T}}} Q(\gamma).$$

Define $\Pi' = \{\gamma \in \Pi(\mu, \nu) : \gamma = (\text{Id}, F)_{\#}\mu = (B, \text{Id})_{\#}\nu \exists (F, B) \in \mathcal{I}(\nu, \mu)\}$, then one can check

$$\text{RGM}(\mu, \nu)^2 = \inf_{\gamma \in \Pi'} Q(\gamma)$$

using change-of-variables. Note that Π' may be rewritten as $\{\gamma \in \Pi_{\mathcal{T}} : \gamma = (B, \text{Id})_{\#}\nu \exists B \in \mathcal{T}(\nu, \mu)\}$. Hence, $\Pi' \subseteq \Pi_{\mathcal{T}} \subseteq \Pi(\mu, \nu)$, thus we conclude $\text{GW}(\mu, \nu) \leq \text{GM}(\mu, \nu) \leq \text{RGM}(\mu, \nu)$. □

In short, Proposition 1 shows that RGM, GM, and GW are minimization problems of a common objective function Q over different constraint sets of couplings, namely, $\Pi' \subseteq \Pi_{\mathcal{T}} \subseteq \Pi(\mu, \nu)$, respectively.

Next, we discuss the binding constraint $(\text{Id}, F)_{\#}\mu = (B, \text{Id})_{\#}\nu$, or equivalently, the feasible set $\mathcal{I}(\mu, \nu)$. One should notice that $\mathcal{I}(\mu, \nu)$ might be empty, for instance, if μ and ν are discrete and their supports have different cardinality; say, $\mu = \delta_x$ and $\nu = (\delta_{y_1} + \delta_{y_2})/2$, namely, Dirac measures supported on $x \in \mathcal{X}$ and $y_1, y_2 \in \mathcal{Y}$, then even $\mathcal{T}(\mu, \nu)$ is empty, meaning that the Monge problem is infeasible and so is RGM. On the flip side, the following lemma gives a sufficient condition for $(F, B) \in \mathcal{I}(\mu, \nu)$ which can be useful in practice.

Lemma 1. Let $(F, B) \in \mathcal{T}(\mu, \nu) \times \mathcal{T}(\nu, \mu)$. If $F \circ B = \text{Id}$ or $B \circ F = \text{Id}$ holds, then $(F, B) \in \mathcal{I}(\mu, \nu)$.

Proof. Without loss of generality, assume $B \circ F = \text{Id}$. Then,

$$(\text{Id}, F)_{\#}\mu = (B \circ F, F)_{\#}\mu = (B, \text{Id})_{\#}(F_{\#}\mu) = (B, \text{Id})_{\#}\nu.$$

Hence, $(F, B) \in \mathcal{I}(\mu, \nu)$. □

The following example illustrates that this condition can be used to find a pair $(F, B) \in \mathcal{I}(\mu, \nu)$ when μ and ν are Gaussian distributions.

Example 1. Given $p < q$, suppose $\mu = N(0, I_p)$ and $\nu = N(0, \Sigma)$, where $I_p \in \mathbb{R}^{p \times p}$ is the identity matrix and $\Sigma \in \mathbb{R}^{q \times q}$ is of rank p . Then, we can find a rank- p matrix $A \in \mathbb{R}^{q \times p}$ such that $\Sigma = AA^\top$. Let $F(x) = Ax$ and $B(y) = A^\dagger y$, then one can easily check $F_{\#}\mu = \nu$, $B_{\#}\nu = \mu$, and $B \circ F = \text{Id}$. Hence, $(F, B) \in \mathcal{I}(\mu, \nu)$.

Lastly, we provide a simple example that shows that properly chosen cost functions give a strong isomorphism between two Gaussian distributions.

Example 2. Consider two Gaussian distributions on \mathbb{R}^d , say $\mu = N(0, \Sigma_1)$ and $\nu = N(0, \Sigma_2)$. Assume Σ_1 and Σ_2 are invertible. Then two network spaces (\mathbb{R}^d, μ, c_X) and (\mathbb{R}^d, ν, c_Y) are strongly isomorphic if c_X and c_Y are Mahalanobis distances, that is,

$$c_X(x, x') = \sqrt{(x - x')^\top \Sigma_1^{-1} (x - x')}, \quad c_Y(y, y') = \sqrt{(y - y')^\top \Sigma_2^{-1} (y - y')}.$$

To see this, let $T = \Sigma_2^{1/2} \Sigma_1^{-1/2}$, where $\Sigma_1^{1/2}$ and $\Sigma_2^{1/2}$ are the square roots of Σ_1 and Σ_2 , respectively. Obviously, a linear map T satisfies $T \in \mathcal{T}(\mu, \nu)$ and $c_X(x, x') = c_Y(Tx, Tx')$ for all $x, x' \in \mathbb{R}^d$. According to Definition 3, a linear map T is a strong isomorphism. Proposition 3 implies $\text{RGM}(\mu, \nu) = 0$. Notice that the same results hold for $c_X(x, x') = x^\top \Sigma_1^{-1} x'$ and $c_Y(y, y') = y^\top \Sigma_2^{-1} y'$ as well.

A.2 Analysis of $\text{GW} = \text{RGM}$ in Non-atomic Cases

We have seen in Proposition 1 that $\text{GW} \leq \text{GM} \leq \text{RGM}$ holds in general. This section shows that these inequalities become equalities under mild conditions. The key idea is that Π' is dense in $\Pi(\mu, \nu)$ with respect to the weak topology if μ and ν have no atoms. It is already studied in the optimal transport literature that $\Pi_{\mathcal{T}}$ is dense in $\Pi(\mu, \nu)$ under certain conditions, which has been used to show that the Monge and the Kantorovich problems have the same infimum, for instance, see Theorem 1.1 of [10]. The following lemma shows the denseness of Π' in $\Pi(\mu, \nu)$ in a similar fashion.

Lemma 2. Let (\mathcal{X}, μ) and (\mathcal{Y}, ν) be two Polish probability spaces, where \mathcal{X} and \mathcal{Y} are closed subsets of \mathbb{R}^d with $d \geq 2$. Suppose that $\mu(\{x\}) = \nu(\{y\}) = 0$ for any $(x, y) \in \mathcal{X} \times \mathcal{Y}$. Then, for any $\gamma \in \Pi(\mu, \nu)$, we can find a sequence $(\gamma_n)_{n \in \mathbb{N}}$ in Π' converging weakly to γ .

Proof. Fix a $\gamma \in \Pi(\mu, \nu)$. Our goal is to construct a sequence $(F_n, B_n)_{n \in \mathbb{N}}$ in $\mathcal{I}(\mu, \nu)$ such that a sequence $(\gamma_n)_{n \in \mathbb{N}}$, defined by $\gamma_n = (\text{Id}, F_n)_{\#}\mu = (B_n, \text{Id})_{\#}\nu$, converges weakly to γ .

The first step is to establish equivalence of the two Polish probability spaces $(\gamma, \mathcal{X} \times \mathcal{Y})$ and $([0, 1]^d, \lambda)$, where λ is the Lebesgue measure on \mathbb{R}^d restricted to $[0, 1]^d$. To this end, we invoke Theorem 16 in Chapter 15 of [45], which states that any Polish probability space that has no atoms is equivalent to $([0, 1], \lambda_1)$, where λ_1 is the Lebesgue measure on \mathbb{R} restricted to $[0, 1]$. Note that γ has no atoms, namely, $\gamma(\{(x, y)\}) = 0$ for any $(x, y) \in \mathcal{X} \times \mathcal{Y}$; also, λ has no atoms. Therefore, $(\gamma, \mathcal{X} \times \mathcal{Y})$ and $([0, 1]^d, \lambda)$ are equivalent as both of them are equivalent to a common space $([0, 1], \lambda_1)$.⁶ More precisely, this equivalence implies that there exist maps $T_1: [0, 1]^d \rightarrow \mathbb{R}^d$ and $T_2: [0, 1]^d \rightarrow \mathbb{R}^d$ such that $(T_1, T_2)_{\#}\lambda = \gamma$. Accordingly, $(T_1)_{\#}\lambda = \mu$ and $(T_2)_{\#}\lambda = \nu$ as $\gamma \in \Pi(\mu, \nu)$.

⁶For the precise definition of this equivalence and related details, see Chapter 15 of [45]. Here, we only provide and paraphrase some results from this chapter that are relevant to the proof.

Next, we approximate T_1 and T_2 by bijective transport maps. Theorem 2.1 of [10] shows that we can find a sequence $(T_1^n, T_2^n)_{n \in \mathbb{N}}$ such that $T_1^n: [0, 1]^d \rightarrow \mathcal{X}$ and $T_2^n: [0, 1]^d \rightarrow \mathcal{Y}$ are bijective diffeomorphisms satisfying $(T_1^n)_\# \lambda = \mu$ and $(T_2^n)_\# \lambda = \nu$, respectively, and

$$\lim_{n \rightarrow \infty} \int_{[0,1]^d} \|T_1^n(z) - T_1(z)\| d\lambda(z) = \lim_{n \rightarrow \infty} \int_{[0,1]^d} \|T_2^n(z) - T_2(z)\| d\lambda(z) = 0. \quad (\text{A.2})$$

For each $n \in \mathbb{N}$, define $F_n = T_2^n \circ (T_1^n)^{-1}$ and $B_n = T_1^n \circ (T_2^n)^{-1}$, then $(F_n, B_n) \in \mathcal{I}(\mu, \nu)$ by Lemma 1; hence $\gamma_n := (\text{Id}, F_n)_\# \mu = (B_n, \text{Id})_\# \nu \in \Pi'$.

Lastly, we prove that $(\gamma_n)_{n \in \mathbb{N}}$ converges weakly to γ . It suffices to show that

$$\int_{\mathcal{X} \times \mathcal{Y}} f(x, y) d\gamma(x, y) = \lim_{n \rightarrow \infty} \int_{\mathcal{X} \times \mathcal{Y}} f(x, y) d\gamma_n(x, y) \quad (\text{A.3})$$

for all bounded and Lipschitz functions f on $\mathcal{X} \times \mathcal{Y}$. As $(T_1, T_2)_\# \lambda = \gamma$,

$$\int_{\mathcal{X} \times \mathcal{Y}} f(x, y) d\gamma(x, y) = \int_{\mathcal{X} \times \mathcal{Y}} f(x, y) d(T_1, T_2)_\# \lambda(x, y) = \int_{[0,1]^d} f(T_1(z), T_2(z)) d\lambda(z).$$

By construction, $\gamma_n = (\text{Id}, F_n)_\# (T_1^n)_\# \lambda = (T_1^n, T_2^n)_\# \lambda$, hence

$$\int_{\mathcal{X} \times \mathcal{Y}} f(x, y) d\gamma_n(x, y) = \int_{\mathcal{X} \times \mathcal{Y}} f(x, y) d(T_1^n, T_2^n)_\# \lambda(x, y) = \int_{[0,1]^d} f(T_1^n(z), T_2^n(z)) d\lambda(z).$$

If f is Lipschitz with constant L ,

$$\begin{aligned} & \left| \int_{\mathcal{X} \times \mathcal{Y}} f(x, y) d\gamma_n(x, y) - \int_{\mathcal{X} \times \mathcal{Y}} f(x, y) d\gamma(x, y) \right| \\ & \leq \int_{[0,1]^d} |f(T_1^n(z), T_2^n(z)) - f(T_1(z), T_2(z))| d\lambda(z) \\ & \leq L \cdot \int_{[0,1]^d} (\|T_1^n(z) - T_1(z)\| + \|T_2^n(z) - T_2(z)\|) d\lambda(z), \end{aligned}$$

where the last inequality is due to Lipschitzness. Due to (A.2), we have (A.3). \square

Using Lemma 2, we are now ready to establish the following result showing $\text{GW} = \text{RGM}$ under mild conditions.

Proof of Theorem 3. Recall that Q is the objective function of couplings defined in (A.1). As $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$ are bounded, one can verify that Q is continuous on $\Pi(\mu, \nu)$ with respect to the weak topology, for instance, see Lemma 2.3 of [13]. By Theorem 2.2 of [13], there exists an optimal coupling $\gamma^* \in \Pi(\mu, \nu)$ such that $\text{GW}(\mu, \nu) = Q(\gamma^*)$. By Lemma 2, we can find a sequence $(\gamma_n)_{n \in \mathbb{N}}$ in $\Pi'(\mu, \nu)$ converging weakly to γ^* . Accordingly,

$$\text{GW}(\mu, \nu) = Q(\gamma^*) = \lim_{n \rightarrow \infty} Q(\gamma_n) \geq \text{RGM}(\mu, \nu).$$

As $\text{GW}(\mu, \nu) \leq \text{RGM}(\mu, \nu)$ by Proposition 1, we have $\text{GW}(\mu, \nu) = \text{RGM}(\mu, \nu)$. \square

Remark. Roughly speaking, the fact that Π' is dense in $\Pi(\mu, \nu)$ implies that the binding constraint of the RGM distance is in fact, a mild restriction on the set of couplings $\Pi(\mu, \nu)$. In other words, restricting to couplings Π' does not incur a loss of information in terms of minimizing the quadratic objective function Q .

A.3 Inductive Bias of RGM and Brenier’s Polar Factorization

As mentioned in Section 5, given two Polish probability spaces (\mathcal{X}, μ) and (\mathcal{Y}, ν) , we prove that there is a pair of cost functions $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$ such that the resulting network spaces $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$ are strongly isomorphic. More importantly, we discuss how these costs give rise to a specific strong isomorphism based on Brenier’s polar factorization, providing a deeper insight into the inductive bias of the RGM sampler.

Preliminaries Throughout this section, we focus on Borel probability measures on \mathbb{R}^d that are absolutely continuous with respect to the d -dimensional Lebesgue measure and have finite second moments; $\mathcal{P}_2^{ac}(\mathbb{R}^d)$ denotes the collection of such measures. We can always find a unique optimal transport map—given as the gradient of a convex function—between any two elements in $\mathcal{P}_2^{ac}(\mathbb{R}^d)$; this is Brenier’s theorem briefly mentioned in Section 1, which we formally state as follows.

Theorem 6 (Brenier’s Theorem). *Given $\lambda, \rho \in \mathcal{P}_2^{ac}(\mathbb{R}^d)$, we can find a convex function $\psi: \mathbb{R}^d \rightarrow \mathbb{R}$ such that $\nabla\psi$ and $(\nabla\psi)^{-1}$ are unique optimal transport maps between them:*

$$\nabla\psi = \arg \min_{T_{\#}\lambda = \rho} \int_{\mathbb{R}^d} \|x - T(x)\|^2 d\lambda(x) \quad \text{and} \quad (\nabla\psi)^{-1} = \arg \min_{T_{\#}\rho = \lambda} \int_{\mathbb{R}^d} \|x - T(x)\|^2 d\rho(x) .$$

Remark. *The proof of Brenier’s theorem requires intricate convex analysis techniques. It is helpful to consider a simple case—where both λ and ρ are Gaussian—to better understand the main message of Brenier’s theorem. Suppose $\lambda = \mathcal{N}(0, I_d)$ and $\rho = \mathcal{N}(0, \Sigma)$ —assuming Σ is invertible—and focus on linear transport maps, that is, $T \in \mathbb{R}^{d \times d}$ such that $TT^{\top} = \Sigma$ the transport cost boils down to*

$$\mathbb{E}_{x \sim \mathcal{N}(0, I_d)} \|(I_d - T)x\|^2 = \text{tr}((I_d - T)(I_d - T)^{\top}) = \text{tr}(I_d) - 2\text{tr}(T) + \text{tr}(\Sigma) .$$

Therefore, the transport cost is minimized by a linear transport map T that maximizes its trace under the constraint $TT^{\top} = \Sigma$. Using linear algebra techniques, one can verify that $\Sigma^{1/2}$, the unique square root of Σ , is the optimal linear transport map;⁷ we can indeed see that it is the gradient of a convex function $x \mapsto \frac{1}{2}x^{\top}\Sigma^{1/2}x$. Last but not least, notice that any admissible linear transport map $T \in \mathbb{R}^{d \times d}$, that satisfies $TT^{\top} = \Sigma$, can be decomposed as $T = \Sigma^{1/2}S$ for some $S \in O(d)$, where $O(d)$ is the orthogonal group in dimension d . The last fact is called the matrix factorization theorem, which will be elaborated in A.3.2 to shed light on Brenier’s polar factorization, as done here to example Brenier’s theorem.

A.3.1 Designing Cost Functions for Strong Isomorphism

We have already seen in Example 2 that Mahalanobis distances make two Gaussian distributions strongly isomorphic. Though this constructive example seems to be a special case, it indicates a fundamental principle that carries over to general cases. We will elaborate on the general constructive principle in this section by referring to a common space Polish probability space (\mathcal{Z}, λ) to design cost functions. Figure 4(a) visualizes Example 2 along with an additional base measure $\lambda := \mathcal{N}(0, I_d)$; as we have remarked after Theorem 6, linear maps $\Sigma_1^{1/2}$ and $\Sigma_2^{1/2}$ are the optimal transport maps from λ to μ and ν , respectively.

Letting $c_{\mathcal{Z}}$ be the Euclidean distance, notice that we may rewrite the Mahalanobis distances in Example 2 as

$$\begin{aligned} c_{\mathcal{X}}(x, x') &= c_{\mathcal{Z}}(\Sigma_1^{-1/2}x, \Sigma_1^{-1/2}x') , \\ c_{\mathcal{Y}}(y, y') &= c_{\mathcal{Z}}(\Sigma_2^{-1/2}y, \Sigma_2^{-1/2}y') . \end{aligned}$$

This shows that $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$ are derived by properly combining the base distance $c_{\mathcal{Z}}$ with the optimal transport maps $\Sigma_1^{-1/2}$ and $\Sigma_2^{-1/2}$, respectively. Also, in this case, $\Sigma_2^{1/2}\Sigma_1^{-1/2}$, a composition of the two optimal transport maps, gives rise to a strong isomorphism under these cost functions.

⁷Though this argument is designed to show that $\Sigma^{1/2}$ is optimal among linear transport maps to get insights, it can be shown that $\Sigma^{1/2}$ is, in fact, optimal among all admissible transport maps.




Figure 4: Optimal transport maps by Brenier’s theorem

The aforementioned procedure is indeed applicable to any $\mu, \nu \in \mathcal{P}_2^{ac}(\mathbb{R}^d)$. As visualized in Figure 4(b), simply replace the linear maps $\Sigma_1^{1/2}$ and $\Sigma_2^{1/2}$ with the optimal transport maps $\nabla\psi_1$ and $\nabla\psi_2$ from any suitable base measure $\lambda \in \mathcal{P}_2^{ac}(\mathbb{R}^d)$, respectively, by Brenier’s theorem. Then, for any suitable base distance $c_{\mathcal{Z}}$, define

$$\begin{aligned} c_{\mathcal{X}}(x, x') &= c_{\mathcal{Z}}((\nabla\psi_1)^{-1}(x), (\nabla\psi_1)^{-1}(x')), \\ c_{\mathcal{Y}}(y, y') &= c_{\mathcal{Z}}((\nabla\psi_2)^{-1}(y), (\nabla\psi_2)^{-1}(y')). \end{aligned} \tag{A.4}$$

Then, the network spaces $(\mathcal{X}, \mu, c_{\mathcal{X}})$ and $(\mathcal{Y}, \nu, c_{\mathcal{Y}})$,⁸ where $\mathcal{X} = \text{supp}(\mu)$ and $\mathcal{Y} = \text{supp}(\nu)$, are strongly isomorphic; also, $\nabla\psi_2 \circ (\nabla\psi_1)^{-1}$ is a strong isomorphism. Notice that we have derived this fundamental result by simply rethinking Example 2 via Brenier’s theorem and generalizing the diagram in Figure 4. We reiterate that the principle behind the isomorphic Gaussian example is fundamental and generalizable to generic measures in $\mathcal{P}_2^{ac}(\mathbb{R}^d)$; it is certainly not just a toy example.

A.3.2 Identifying Strong Isomorphism

In the previous section, we have seen how to define suitable cost functions $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$ that make $\mu, \nu \in \mathcal{P}_2^{ac}(\mathbb{R}^d)$ strongly isomorphic. As pointed out in Section 5, the RGM sampler brings in an inductive bias towards a strong isomorphism; indeed, we have shown that $\nabla\psi_2 \circ (\nabla\psi_1)^{-1}$ is a strong isomorphism under the cost functions in (A.4), leading to $\text{RGM}(\mu, \nu) = 0$. In this section, we look at these results from a different angle using an insight from Brenier’s polar factorization, highlighting unseen aspects of the inductive bias.

We start from Figure 4(b): fix $\mu, \nu \in \mathcal{P}_2^{ac}(\mathbb{R}^d)$ and let $\mathcal{X} = \text{supp}(\mu)$ and $\mathcal{Y} = \text{supp}(\nu)$. Also, let λ be the Lebesgue measure on $\mathcal{Z} = [0, 1]^d$. Recall that the key ingredients of the diagram were optimal transport maps $\nabla\psi_1$ and $\nabla\psi_2$ from the base measure. It turns out that we may replace them with other transport maps—possibly not optimal—from the base measure. This observation comes from Brenier’s polar factorization [9], which we paraphrase as follows:

Theorem 7 (Brenier’s Polar Factorization). *For any transport map T_1 from λ to μ , we can find a unique map $s_1: \mathcal{Z} \rightarrow \mathcal{Z}$ such that $(s_1)_{\#}\lambda = \mu$ and $T_1 = \nabla\psi_1 \circ s_1$.*

In other words, any transport map from λ to μ is factorized into a composition of the unique optimal transport map $\nabla\psi_1$ and a (Lebesgue) measure-preserving map $s_1: \mathcal{Z} \rightarrow \mathcal{Z}$. Let $S(\mathcal{Z})$ be the collection of all measure-preserving maps from \mathcal{Z} to \mathcal{Z} , then Brenier’s polar factorization essentially shows a one-to-one correspondence between $\mathcal{T}(\lambda, \mu)$ and $S(\mathcal{Z})$. Analogously, this implies a one-to-one correspondence between $\mathcal{T}(\lambda, \nu)$ and $S(\mathcal{Z})$.

Remark (Matrix polar factorization). *To get insights into this sophisticated result, let us pause and go back to the remark below Theorem 6, where we have mentioned that in the Gaussian case, any linear transport*

⁸Technically, one needs to check continuity of $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$. If μ and ν have smooth densities, it is known that $\nabla\psi_1$ and $\nabla\psi_2$ are sufficiently smooth, hence continuity is guaranteed.

map is decomposed as $T = \Sigma^{1/2}S$, the multiplication of the optimal transport map $\Sigma^{1/2}$ and an orthogonal matrix $S \in O(d)$,⁹ which exactly correspond to $\nabla\psi_1$ and s_1 in Theorem 7, respectively. In other words, Brenier’s Polar Factorization is a generalization of the matrix factorization that we have discussed earlier by restricting to the Gaussian case and linear transport maps.

As mentioned earlier, we now replace the two arrows—the optimal transport maps $\nabla\psi_1$ and $\nabla\psi_2$ from the base measure—in Figure 4(b) with any transport maps, namely, $\nabla\psi_1 \circ s_1$ and $\nabla\psi_2 \circ s_2$, respectively, for any $s_1, s_2 \in S(\mathcal{Z})$. One technicality here is that we restrict our focus to bijective s_1 and s_2 to use invertibility to reverse the arrows.¹⁰ Then, we can also define a transport map $\nabla\psi_2 \circ s_2 \circ (\nabla\psi_1 \circ s_1)^{-1}$ from \mathcal{X} to \mathcal{Y} by chaining the transport maps $\mu \rightarrow \lambda$ and $\lambda \rightarrow \nu$ as in Figure 4(b). These changes are shown in the new diagram Figure 5.




Figure 5: Generalization of Figure 4(b) via Brenier’s polar factorization.

Now, let us go back to transform sampling. The arrow from μ to ν in Figure 5 in fact shows that there are many transport maps from μ to ν , that is,

$$\mathcal{F} = \{\nabla\psi_2 \circ s_2 \circ (\nabla\psi_1 \circ s_1)^{-1} : \text{bijective } s_1, s_2 \in S(\mathcal{Z})\} \subset \mathcal{T}(\mu, \nu).$$

Though \mathcal{F} is a collection of transport maps constructed in a certain way, that is, chaining $\mu \rightarrow \lambda$ and $\lambda \rightarrow \nu$, it still consists of infinitely many transport maps, reiterating that transform sampling is indeed over-identified. We show how the RGM sampler induces an inductive bias in this case, thereby choosing a strong isomorphism from the collection \mathcal{F} .

Consider the cost functions $c_{\mathcal{X}}$ and $c_{\mathcal{Y}}$ defined by (A.4) with a suitable distance $c_{\mathcal{Z}}$. As transport maps in \mathcal{F} are invertible, by Lemma 1,

$$\mathcal{I}' := \{(F, F^{-1}) : F \in \mathcal{F}\} \subset \mathcal{I}(\mu, \nu).$$

Recall that computing the RGM distance amounts to finding (F, B) such that $c_{\mathcal{X}}(x, B(y))$ best matches $c_{\mathcal{Y}}(F(x), y)$ on average. For $(F, B) \in \mathcal{I}'$, by (A.4), we have

$$\begin{aligned} c_{\mathcal{X}}(x, B(y)) &= c_{\mathcal{Z}}((\nabla\psi_1)^{-1}(x), s_1 \circ s_2^{-1} \circ (\nabla\psi_2)^{-1}(y)), \\ c_{\mathcal{Y}}(F(x), y) &= c_{\mathcal{Z}}(s_2 \circ s_1^{-1} \circ (\nabla\psi_1)^{-1}(x), (\nabla\psi_2)^{-1}(y)). \end{aligned}$$

Therefore, $c_{\mathcal{X}}(x, B(y)) = c_{\mathcal{Y}}(F(x), y)$ indeed holds provided $s_1 = s_2$, which amounts to a pair $(F, B) = (\nabla\psi_2 \circ (\nabla\psi_1)^{-1}, \nabla\psi_1 \circ (\nabla\psi_2)^{-1}) \in \mathcal{I}'$. In other words, among infinitely many possible elements in \mathcal{F} characterized by s_1 and s_2 , the RGM sampler with cost functions in (A.4) favors $\nabla\psi_2 \circ (\nabla\psi_1)^{-1}$, the strong isomorphism under these cost functions as visualized in Figure 4(b). In summary,

- we can characterize a collection \mathcal{F} of transport maps by pairs of bijective measure-preserving maps (s_1, s_2) via Brenier’s polar factorization,
- the RGM sampler with cost functions in (A.4) finds the one in \mathcal{F} satisfying the rearrangement correspondance $s_1 = s_2$, which is exactly the strong isomorphism (inductive bias).

⁹Note that $S \in O(d)$ transports $\mathcal{N}(0, I_d)$ to itself.

¹⁰Elements of $S(\mathcal{Z})$ may not be invertible in general. That said, a subset of $S(\mathcal{Z})$ consisting of bijective measure-preserving maps is dense in $S(\mathcal{Z})$ as discussed in [10].

B Statistical Theory

This section serves to prove Theorem 4. Without loss of generality, we assume $\lambda_1 = \lambda_2 = \lambda_3 = 1$ in $C(\mu, \nu, F, B)$ since the proof is essentially identical with any constants $\lambda_i, 1 \leq i \leq 3$. For convenience, we denote

$$\begin{aligned} C_0(F, B) &= \int (c_{\mathcal{X}}(x, B(y)) - c_{\mathcal{Y}}(F(x), y))^2 d\mu \otimes \nu, \\ M(F, B) &= \text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#}\mu, \nu) + \text{MMD}_{K_{\mathcal{X}}}^2(\mu, B_{\#}\nu) + \text{MMD}_{K_{\mathcal{X}} \otimes K_{\mathcal{Y}}}^2((\text{Id}, F)_{\#}\mu, (B, \text{Id})_{\#}\nu) \end{aligned}$$

and therefore $C(\mu, \nu, F, B) = C_0(F, B) + M(F, B)$. Similarly, define the empirical counterparts as

$$\begin{aligned} \widehat{C}_0(F, B) &= \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n (c_{\mathcal{X}}(x_i, B(y_j)) - c_{\mathcal{Y}}(F(x_i), y_j))^2, \\ \widehat{M}(F, B) &= \text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#}\widehat{\mu}_m, \widehat{\nu}_n) + \text{MMD}_{K_{\mathcal{X}}}^2(\widehat{\mu}_m, B_{\#}\widehat{\nu}_n) + \text{MMD}_{K_{\mathcal{X}} \otimes K_{\mathcal{Y}}}^2((\text{Id}, F)_{\#}\widehat{\mu}_m, (B, \text{Id})_{\#}\widehat{\nu}_n) \end{aligned}$$

and thus $C(\widehat{\mu}_m, \widehat{\nu}_n, F, B) = \widehat{C}_0(F, B) + \widehat{M}(F, B)$.

Our goal is to give an upper bound on $C(\mu, \nu, \widehat{F}, \widehat{B}) - \inf_{(F, B) \in \mathcal{F} \times \mathcal{B}} C(\mu, \nu, F, B)$. To this end, first recall that

$$C(\widehat{\mu}_m, \widehat{\nu}_n, \widehat{F}, \widehat{B}) \leq C(\widehat{\mu}_m, \widehat{\nu}_n, F, B)$$

holds for any $F \in \mathcal{F}$ and $B \in \mathcal{B}$ by definition of \widehat{F} and \widehat{B} given in Theorem 4. Therefore,

$$C(\mu, \nu, \widehat{F}, \widehat{B}) - C(\mu, \nu, F, B) \leq C(\mu, \nu, \widehat{F}, \widehat{B}) - C(\widehat{\mu}_m, \widehat{\nu}_n, \widehat{F}, \widehat{B}) + C(\widehat{\mu}_m, \widehat{\nu}_n, F, B) - C(\mu, \nu, F, B).$$

The RHS can be decomposed as

$$C_0(\widehat{F}, \widehat{B}) - \widehat{C}_0(\widehat{F}, \widehat{B}) + M(\widehat{F}, \widehat{B}) - \widehat{M}(\widehat{F}, \widehat{B}) + \widehat{C}_0(F, B) - C_0(F, B) + \widehat{M}(F, B) - M(F, B).$$

To further control the expression, we will first derive probabilistic bounds on $|\widehat{C}_0(F, B) - C_0(F, B)|$ and $|\widehat{M}(F, B) - M(F, B)|$ that hold for a fixed $(F, B) \in \mathcal{F} \times \mathcal{B}$ via standard concentration inequalities. Later, we will establish uniform probabilistic bounds on $\sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} |\widehat{C}_0(F, B) - C_0(F, B)|$ and $\sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} |\widehat{M}(F, B) - M(F, B)|$, using tools from empirical process theory.

B.1 Concentration Inequalities

We utilize the McDiarmid's inequality to derive bounds on $|\widehat{C}_0(F, B) - C_0(F, B)|$ and $|\widehat{M}(F, B) - M(F, B)|$. To give a bound on the former, we make the following boundedness assumption.

Assumption 1. $c_{\mathcal{X}}(\cdot, \cdot), c_{\mathcal{Y}}(\cdot, \cdot)$ is uniformly bounded, that is, there exists a constant $H > 0$ such that

$$\sup_{(x, x') \in \mathcal{X} \times \mathcal{X}} c_{\mathcal{X}}(x, x'), \quad \sup_{(y, y') \in \mathcal{Y} \times \mathcal{Y}} c_{\mathcal{Y}}(y, y') \leq \sqrt{\frac{H}{4}}.$$

Proposition 4. Under Assumption 1, for any pair $(F, B) \in \mathcal{F} \times \mathcal{B}$ and $\delta > 0$,

$$|\widehat{C}_0(F, B) - C_0(F, B)| \lesssim \sqrt{\frac{\log(\frac{m \vee n}{\delta})}{m \wedge n}}$$

holds with probability at least $1 - 4\delta$.

To derive a similar bound on $|\widehat{M}(F, B) - M(F, B)|$, we assume that kernels are bounded.

Assumption 2. *There exists $K > 0$ such that*

$$\sup_{x \in \mathcal{X}} |K_{\mathcal{X}}(x, x)|, \sup_{y \in \mathcal{Y}} |K_{\mathcal{Y}}(y, y)| \leq K.$$

Proposition 5. *Under Assumption 2, for any pair $(F, B) \in \mathcal{F} \times \mathcal{B}$ and $\delta > 0$,*

$$|\widehat{M}(F, B) - M(F, B)| \lesssim \sqrt{\frac{\log(1/\delta)}{m}} + \sqrt{\frac{\log(1/\delta)}{n}}$$

holds with probability at least $1 - 6\delta$.

B.2 Uniform Deviations

We now derive uniform deviation bounds for

$$\sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} |\widehat{C}_0(F, B) - C_0(F, B)|, \sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} |\widehat{M}(F, B) - M(F, B)|.$$

For the former, we use the notion of uniform covering numbers defined below.

Definition 6 (Uniform Covering Number). *Let \mathcal{G} be a collection of real-valued functions defined on a set \mathcal{Z} . Given m points $z_1, \dots, z_m \in \mathcal{Z}$ and any $\delta > 0$, we define $N_{\infty}(\delta, \mathcal{G}, \{z_i\}_{i=1}^m)$ to be the δ -covering number of \mathcal{G} under the pseudometric d induced by points z_1, \dots, z_m :*

$$d(g, g') := \max_{i \in [m]} |g(z_i) - g'(z_i)|.$$

Also, we define the uniform δ -covering number of \mathcal{G} as follows:

$$N_{\infty}(\delta, \mathcal{G}, m) := \sup \{N_{\infty}(\delta, \mathcal{G}, \{z_i\}_{i=1}^m) : z_1, \dots, z_m \in \mathcal{Z}\}.$$

Here, the supremum is taken over all possible combinations of m points in \mathcal{Z} .

Also, we make the following assumptions.

Assumption 3. \mathcal{F}_k and \mathcal{B}_{ℓ} (see Section 3.3) consist of uniformly bounded functions, that is, there exists a constant $b > 0$ such that

$$\max_{k \in [\dim(\mathcal{Y})]} \sup_{F_k \in \mathcal{F}_k} \|F_k\|_{\infty}, \max_{\ell \in [\dim(\mathcal{X})]} \sup_{B_{\ell} \in \mathcal{B}_{\ell}} \|B_{\ell}\|_{\infty} \leq b.$$

Assumption 4. *There exists a constant $L > 0$ such that*

$$|c_{\mathcal{X}}(x, x_1) - c_{\mathcal{X}}(x, x_2)| \leq L\|x_1 - x_2\|, |c_{\mathcal{Y}}(y_1, y) - c_{\mathcal{Y}}(y_2, y)| \leq L\|y_1 - y_2\|.$$

This Lipschitzness assumption ensures the smoothness of a map $(F, B) \mapsto |\widehat{C}_0(F, B) - C_0(F, B)|$ over $\mathcal{F} \times \mathcal{B}$, which allows us to utilize the uniform covering numbers.

Proposition 6. *Under Assumptions 1, 3, 4, for any $\epsilon > 0$ and $\delta > 0$,*

$$\begin{aligned} & \sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} |\widehat{C}_0(F, B) - C_0(F, B)| \\ & \lesssim \sqrt{\frac{\log\left(\frac{m\sqrt{n}}{\delta}\right)}{m \wedge n}} + \epsilon + \sqrt{\frac{\sum_{k=1}^{\dim(\mathcal{Y})} \log N_{\infty}(\epsilon, \mathcal{F}_k, m) + \sum_{\ell=1}^{\dim(\mathcal{X})} \log N_{\infty}(\epsilon, \mathcal{B}_{\ell}, n)}{m \wedge n}} \end{aligned}$$

holds with probability at least $1 - 2\delta$.

Now, the remaining task is to choose ϵ carefully in Proposition 6 for a concrete upper bound. To this end, we utilize the pseudo-dimension defined below.

Definition 7 (Pseudo-Dimension). *Let \mathcal{G} be a collection of real-valued functions defined on a set \mathcal{Z} . Given a subset $S := \{z_1, \dots, z_m\} \subset \mathcal{Z}$, we say S is pseudo-shattered by \mathcal{G} if there are $r_1, \dots, r_m \in \mathbb{R}$ such that for each $b \in \{0, 1\}^m$ we can find $g_b \in \mathcal{G}$ satisfying $\text{sign}(g_b(z_i) - r_i) = b_i$ for all $i \in [m]$. We define the pseudo-dimension of \mathcal{G} , denoted as $\text{Pdim}(\mathcal{G})$, as the maximum cardinality of a subset $S \subset \mathcal{Z}$ that is pseudo-shattered by \mathcal{G} .*

Using a well-established relation of the uniform covering number and the pseudo-dimension (Lemma 5), we can simplify Proposition 6 as follows.

Corollary 1. *Under Assumptions 1, 3, 4, for any $\delta > 0$,*

$$\sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} |\widehat{C}_0(F, B) - C_0(F, B)| \\ \lesssim \sqrt{\frac{\log\left(\frac{m \vee n}{\delta}\right)}{m \wedge n}} + \sqrt{\frac{\log(m \vee n)}{m \wedge n} \left(\sum_{k=1}^{\dim(\mathcal{Y})} \text{Pdim}(\mathcal{F}_k) + \sum_{\ell=1}^{\dim(\mathcal{X})} \text{Pdim}(\mathcal{B}_\ell) \right)}$$

holds with probability at least $1 - 2\delta$.

To derive an upper bound on $\sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} |\widehat{M}(F, B) - M(F, B)|$, we first introduce Rademacher complexities defined below.

Definition 8 (Rademacher Complexity). *Let (\mathcal{Z}, ρ) be a probability space and \mathcal{G} be a collection of measurable functions defined on \mathcal{Z} . We define the Rademacher complexity of \mathcal{G} with respect to m samples from ρ as follows:*

$$R_m(\mathcal{G}, \rho) = \mathbb{E} \mathbb{E} \sup_{z_i \stackrel{\text{iid}}{\sim} \rho} \left| \frac{1}{m} \sum_{i=1}^m \epsilon_i g(z_i) \right|,$$

Here, z_1, \dots, z_m are i.i.d. samples from ρ and $\epsilon_1, \dots, \epsilon_m$ are i.i.d. Rademacher random variables such that (z_1, \dots, z_m) and $(\epsilon_1, \dots, \epsilon_m)$ are independent.

Proposition 7. *Denote a closed unit ball of any RKHS \mathcal{H} as $\mathcal{H}(1)$. Also, let $(\text{Id}, \mathcal{F}) := \{(\text{Id}, F) : F \in \mathcal{F}\}$ and $(\mathcal{B}, \text{Id}) := \{(B, \text{Id}) : B \in \mathcal{B}\}$; hence, they are classes of maps from \mathcal{X} to $\mathcal{X} \times \mathcal{Y}$ and from \mathcal{Y} to $\mathcal{X} \times \mathcal{Y}$, respectively. Under Assumption 2, for any $\delta > 0$,*

$$\sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} |\widehat{M}(F, B) - M(F, B)| \lesssim \sqrt{\frac{\log(1/\delta)}{m}} + \sqrt{\frac{\log(1/\delta)}{n}} + R_m(\mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}, \mu) + R_n(\mathcal{H}_{\mathcal{X}}(1) \circ \mathcal{B}, \nu) \\ + R_m(\mathcal{H}_{\mathcal{X} \times \mathcal{Y}}(1) \circ (\text{Id}, \mathcal{F}), \mu) + R_n(\mathcal{H}_{\mathcal{X} \times \mathcal{Y}}(1) \circ (\mathcal{B}, \text{Id}), \nu)$$

holds with probability at least $1 - 6\delta$. Here, $\mathcal{F} \circ \mathcal{G} = \{f \circ g : f \in \mathcal{F}, g \in \mathcal{G}\}$ for any function classes \mathcal{F} and \mathcal{G} with matching input and output space.

Now, the only remaining task is to bound four Rademacher complexities. We will derive upper bounds using the chaining technique. To illustrate the main idea, let us consider $\mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}$. Recall that

$$R_m(\mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}, \mu) = \mathbb{E} R_m(\mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}, \{x_i\}_{i=1}^m),$$

where $R_m(\mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}, \{x_i\}_{i=1}^m)$ is the empirical Rademacher complexity of $\mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}$ associated with $\{x_i\}_{i=1}^m$:

$$R_m(\mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}, \{x_i\}_{i=1}^m) = \mathbb{E} \sup_{\epsilon_i} \sup_{h \in \mathcal{H}_{\mathcal{Y}}(1), F \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \epsilon_i h(F(x_i)) \right| = \mathbb{E} \sup_{\epsilon_i} \sup_{h \in \mathcal{H}_{\mathcal{Y}}(1), F \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \epsilon_i h(F(x_i)).$$

Notice that we may remove the absolute value since $\mathcal{H}_Y(1) = -\mathcal{H}_Y(1)$. Now, considering $\{x_i\}_{i=1}^m$ as fixed, we will first bound the empirical Rademacher complexity by replacing the Rademacher random variables with Gaussian random variables. Let g_i be i.i.d. standard Gaussian random variables, then it is well known that

$$R_m(\mathcal{H}_Y(1) \circ \mathcal{F}, \{x_i\}_{i=1}^m) \leq \sqrt{\frac{\pi}{2}} \mathbb{E} \sup_{g_i, h \in \mathcal{H}_Y(1), F \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m g_i h(F(x_i)) =: \sqrt{\frac{\pi}{2}} \mathcal{G}_m(\mathcal{H}_Y(1) \circ \mathcal{F}, \{x_i\}_{i=1}^m).$$

Also, under the assumption that K_Y is bounded by K , the reproducing property and the Cauchy-Schwarz inequality imply

$$\begin{aligned} & \sup_{h \in \mathcal{H}_Y(1), F \in \mathcal{F}} \sum_{i=1}^m g_i h(F(x_i)) \\ &= \sup_{h \in \mathcal{H}_Y(1), F \in \mathcal{F}} \left\langle h, \sum_{i=1}^m g_i K_Y(\cdot, F(x_i)) \right\rangle_{\mathcal{H}_Y} \\ &\leq \sup_{h \in \mathcal{H}_Y(1), F \in \mathcal{F}} \|h\|_{\mathcal{H}_Y} \left[\sum_{i=1}^m g_i^2 K_Y(F(x_i), F(x_i)) + \sum_{i \neq j} g_i g_j K_Y(F(x_i), F(x_j)) \right]^{1/2} \\ &\leq \sup_{F \in \mathcal{F}} \left[\sum_{i=1}^m g_i^2 K + \sum_{i \neq j} g_i g_j K_Y(F(x_i), F(x_j)) \right]^{1/2} \\ &\leq \left[\sum_{i=1}^m g_i^2 K + \sup_{F \in \mathcal{F}} \sum_{i \neq j} g_i g_j K_Y(F(x_i), F(x_j)) \right]^{1/2}. \end{aligned}$$

Here, $\langle \cdot, \cdot \rangle_{\mathcal{H}_Y}$ denotes the inner product on \mathcal{H}_Y . Hence,

$$\begin{aligned} \mathcal{G}_m(\mathcal{H}_Y(1) \circ \mathcal{F}, \{x_i\}_{i=1}^m) &\leq \frac{1}{m} \mathbb{E} \left[\sum_{i=1}^m g_i^2 K + \sup_{F \in \mathcal{F}} \sum_{i \neq j} g_i g_j K_Y(F(x_i), F(x_j)) \right]^{1/2} \\ &\leq \frac{1}{m} \left[mK + \mathbb{E} \sup_{g_i, F \in \mathcal{F}} \sum_{i \neq j} g_i g_j K_Y(F(x_i), F(x_j)) \right]^{1/2}, \end{aligned}$$

where the second inequality follows from the Jensen's inequality and $\mathbb{E} g_i^2 = 1$.

For any $F: \mathcal{X} \rightarrow \mathcal{Y}$, let $A_F \in \mathbb{R}^{m \times m}$ be a matrix whose diagonal elements are zero and (i, j) -th element is $K_Y(F(x_i), F(x_j))$ for $i \neq j$. Then, the last term amounts to the supremum of a quadratic process

$$\mathbb{E} \sup_{g, F \in \mathcal{F}} g^\top A_F g,$$

where $g := [g_1, \dots, g_m]^\top \sim N(0, I_m)$.

We rely on the following chaining bound for the quadratic processes, derived in Section D.

Lemma 3 (Chaining Bound). *Let $\mathbb{S}_0^{m \times m}$ be the collection of all symmetric matrices A whose diagonal elements are zero. Endow $\mathbb{S}_0^{m \times m}$ with a metric d given by $d(A, A') := \|A - A'\|$. Given $\mathcal{T} \subset \mathbb{S}_0^{m \times m}$ and a fixed $A_0 \in \mathcal{T}$, define $\Delta = \sup_{A \in \mathcal{T}} d(A, A_0)$. Let $N(\delta, \mathcal{T})$ be the covering number of \mathcal{T} under the metric $d(\cdot, \cdot)$, then*

$$\mathbb{E} \sup_{g, A \in \mathcal{T}} g^\top A g \leq \inf_{J \in \mathbb{N}} \left\{ m\delta_J + 12 \int_{\delta_J/2}^{\Delta/2} \sqrt{2 \log N(\delta, \mathcal{T})} d\delta + 24 \int_{\delta_J/2}^{\Delta/2} \log N(\delta, \mathcal{T}) d\delta \right\}, \quad (\text{B.1})$$

where for any integer $J \geq 0$, we define $\delta_J = 2^{-J} \Delta$.

With the above chaining bound, we can directly upper bound the Rademacher complexities of the compositional classes such as $R_m(\mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}, \mu)$ and $R_m(\mathcal{H}_{\mathcal{X} \times \mathcal{Y}}(1) \circ (\text{Id}, \mathcal{F}), \mu)$. More specifically, for the former class, we will apply this chaining bound to $\mathcal{T} := \{A_F : F \in \mathcal{F}\}$. Then, to further bound the RHS of (B.1), we make the following assumptions.

Assumption 5. *Suppose $K_{\mathcal{X}}$ and $K_{\mathcal{Y}}$ are Lipschitz: there exists $L > 0$ such that*

$$|K_{\mathcal{X}}(x_1, x') - K_{\mathcal{X}}(x_2, x')| \leq L\|x_1 - x_2\|, \quad |K_{\mathcal{Y}}(y_1, y') - K_{\mathcal{Y}}(y_2, y')| \leq L\|y_1 - y_2\|.$$

This plays a similar role as Assumption 4: we can derive an upper bound on $d(A_F, A_{F'})$ via closeness of F and F' in \mathcal{F} . As a result, we will see that the covering number $N(\delta, \mathcal{T})$ can be bounded by the complexity of \mathcal{F} .

Assumption 6. *There exist y_0 and y'_0 in \mathcal{Y} with $K_{\mathcal{Y}}(y_0, y'_0) \neq K_{\mathcal{Y}}(y_0, y_0)$ such that*

- \mathcal{F} contains a constant map F satisfying $F(x) = y_0$ for all $x \in \mathcal{X}$,
- whenever we have $x \neq x' \in \mathcal{X}$, we can find a non-constant map $F \in \mathcal{F}$ such that $F(x) = y_0$ and $F(x') = y'_0$.

Similarly, there exist x_0 and x'_0 in \mathcal{X} with $K_{\mathcal{X}}(x_0, x'_0) \neq K_{\mathcal{X}}(x_0, x_0)$ such that

- \mathcal{B} contains a constant map B such that $B(y) = x_0$ for all $y \in \mathcal{Y}$,
- whenever we have $y \neq y' \in \mathcal{Y}$, we can find a non-constant map $B \in \mathcal{B}$ such that $B(y) = x_0$ and $B(y') = x'_0$.

The main purpose of this assumption is to exclude overly restrictive \mathcal{F} and \mathcal{B} , and is minimal: \mathcal{F} and \mathcal{B} should contain constant maps, as well as non-constant maps. With these assumptions, we can derive the following result.

Proposition 8. *Under Assumptions 2, 3, 5, 6,*

$$R_m(\mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}, \mu), R_m(\mathcal{H}_{\mathcal{X} \times \mathcal{Y}}(1) \circ (\text{Id}, \mathcal{F}), \mu) \lesssim \sqrt{\frac{\log m}{m} \sum_{k=1}^{\dim(\mathcal{Y})} \text{Pdim}(\mathcal{F}_k)},$$

$$R_n(\mathcal{H}_{\mathcal{X}}(1) \circ \mathcal{B}, \mu), R_n(\mathcal{H}_{\mathcal{X} \times \mathcal{Y}}(1) \circ (\mathcal{B}, \text{Id}), \nu) \lesssim \sqrt{\frac{\log n}{n} \sum_{k=1}^{\dim(\mathcal{X})} \text{Pdim}(\mathcal{B}_k)}.$$

In summary, Propositions 4, 5, 7, 8 and Corollary 1 directly imply Theorem 4.

C Representer Theorem and Convex Formulation

This section provides details of the results presented in Section 3.4. Again, without loss of generality, we only consider $\lambda_1 = \lambda_2 = \lambda_3 = 1$ in (3.5).

First, we clarify how measurable maps correspond to bounded linear operators between L^2 spaces.

Proposition 9. *Let $F: \mathcal{X} \rightarrow \mathcal{Y}$ be a measurable map such that $\|dF_{\#}\pi_{\mathcal{X}} / d\pi_{\mathcal{Y}}\|_{\infty} < \infty$. If we define*

$$\mathbf{F}(g) = g \circ F$$

for all $g \in L^2_{\mathcal{Y}}$, then $\mathbf{F}: L^2_{\mathcal{Y}} \rightarrow L^2_{\mathcal{X}}$ is a bounded linear operator. Similarly, a measurable map $B: \mathcal{Y} \rightarrow \mathcal{X}$ satisfying $\|dB_{\#}\pi_{\mathcal{Y}} / d\pi_{\mathcal{X}}\|_{\infty} < \infty$ induces a bounded linear operator $\mathbf{B}: L^2_{\mathcal{X}} \rightarrow L^2_{\mathcal{Y}}$ such that $\mathbf{B}(g) = g \circ B$ for all $g \in L^2_{\mathcal{X}}$.

Proof. Linearity of \mathbf{F} is obvious. Since

$$\int_{\mathcal{X}} g(F(x))^2 d\pi_{\mathcal{X}} = \int_{\mathcal{Y}} g(y)^2 dF_{\#}\pi_{\mathcal{X}}(y) = \int_{\mathcal{Y}} g(y)^2 \frac{dF_{\#}\pi_{\mathcal{X}}}{d\pi_{\mathcal{Y}}}(y) d\pi_{\mathcal{Y}}(y) \leq \|g\|_{L^2(\pi_{\mathcal{Y}})}^2 \left\| \frac{dF_{\#}\pi_{\mathcal{X}}}{d\pi_{\mathcal{Y}}} \right\|_{\infty},$$

we can see $\mathbf{F}(g) = g \circ F \in L^2_{\mathcal{X}}$ and thus $\mathbf{F}: L^2_{\mathcal{Y}} \rightarrow L^2_{\mathcal{X}}$. From this inequality, the operator norm of \mathbf{F} is bounded by $\|dF_{\#}\pi_{\mathcal{X}}/d\pi_{\mathcal{Y}}\|_{\infty}^{1/2}$; hence boundedness of \mathbf{F} follows. The same argument applies to \mathbf{B} . \square

Next, we prove that (3.5) can be written in terms of \mathbf{F} and \mathbf{B} if $K_{\mathcal{X}}$ and $K_{\mathcal{Y}}$ are given by the Mercer's representation:

$$K_{\mathcal{X}}(x, x') = \sum_{k=1}^{\infty} \lambda_k \phi_k(x) \phi_k(x'), \quad (\text{C.1})$$

$$K_{\mathcal{Y}}(y, y') = \sum_{\ell=1}^{\infty} \gamma_{\ell} \psi_{\ell}(y) \psi_{\ell}(y'). \quad (\text{C.2})$$

Let $\Phi_x = [\dots, \phi_k(x), \dots]^{\top} \in \mathbb{R}^{\infty}$ and $\Psi_y = [\dots, \psi_{\ell}(y), \dots]^{\top} \in \mathbb{R}^{\infty}$. Then, $K_{\mathcal{X}}(x, x') = \Phi_x^{\top} \Lambda \Phi_{x'}$ and $K_{\mathcal{Y}}(y, y') = \Psi_y^{\top} \Gamma \Psi_{y'}$. Also,

$$\begin{aligned} K_{\mathcal{X}}(x, B(y)) &= \sum_k \lambda_k \phi_k(x) [\phi_k \circ B](y) \\ &= \sum_k \lambda_k \phi_k(x) \mathbf{B}[\phi_k](y) \\ &= \sum_{k, \ell} \lambda_k \phi_k(x) \mathbf{B}_{\ell k} \psi_{\ell}(y) \\ &= \Psi_y^{\top} \mathbf{B} \Lambda \Phi_x. \end{aligned}$$

Analogously, we can obtain

$$\begin{aligned} K_{\mathcal{Y}}(F(x), y) &= \Phi_x^{\top} \mathbf{F} \Gamma \Psi_y, \\ K_{\mathcal{X}}(B(y), B(y')) &= \Psi_y^{\top} \mathbf{B} \Lambda \mathbf{B}^{\top} \Psi_{y'}, \\ K_{\mathcal{Y}}(F(x), F(x')) &= \Phi_x^{\top} \mathbf{F} \Gamma \mathbf{F}^{\top} \Phi_{x'}. \end{aligned}$$

Using this, we have

$$\frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n (K_{\mathcal{X}}(x_i, B(y_j)) - K_{\mathcal{Y}}(F(x_i), y_j))^2 = \frac{1}{mn} \sum_{i, j} (\Psi_{y_j}^{\top} \mathbf{B} \Lambda \Phi_{x_i} - \Phi_{x_i}^{\top} \mathbf{F} \Gamma \Psi_{y_j})^2. \quad (\text{C.3})$$

Also,

$$\begin{aligned} \text{MMD}_{K_{\mathcal{X}}}^2(\hat{\mu}_m, B_{\#}\hat{\nu}_n) &= \frac{1}{m^2} \sum_{i, i'} K_{\mathcal{X}}(x_i, x_{i'}) + \frac{1}{n^2} \sum_{j, j'} K_{\mathcal{X}}(B(y_j), B(y_{j'})) - \frac{2}{mn} \sum_{i, j} K_{\mathcal{X}}(x_i, B(y_j)) \\ &= \frac{1}{m^2} \sum_{i, i'} \Phi_{x_i}^{\top} \Lambda \Phi_{x_{i'}} + \frac{1}{n^2} \sum_{j, j'} \Psi_{y_j}^{\top} \mathbf{B} \Lambda \mathbf{B}^{\top} \Psi_{y_{j'}} - \frac{2}{mn} \sum_{i, j} \Psi_{y_j}^{\top} \mathbf{B} \Lambda \Phi_{x_i}. \end{aligned} \quad (\text{C.4})$$

Similarly, we have

$$\text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#}\hat{\mu}_m, \hat{\nu}_n) = \frac{1}{m^2} \sum_{i, i'} \Phi_{x_i}^{\top} \mathbf{F} \Gamma \mathbf{F}^{\top} \Phi_{x_{i'}} + \frac{1}{n^2} \sum_{j, j'} \Psi_{y_j}^{\top} \Gamma \Psi_{y_{j'}} - \frac{2}{mn} \sum_{i, j} \Phi_{x_i}^{\top} \mathbf{F} \Gamma \Psi_{y_j} \quad (\text{C.5})$$

and

$$\begin{aligned}
& \text{MMD}_{K_{\mathcal{X}} \otimes K_{\mathcal{Y}}}^2((\text{Id}, F)_{\#} \widehat{\mu}_m, (B, \text{Id})_{\#} \widehat{\nu}_n) \\
&= \frac{1}{m^2} \sum_{i, i'} \Phi_{x_i}^{\top} \Lambda \Phi_{x_{i'}} \Phi_{x_i}^{\top} \mathbf{F} \Gamma \mathbf{F}^{\top} \Phi_{x_{i'}} + \frac{1}{n^2} \sum_{j, j'} \Psi_{y_j}^{\top} \Gamma \Psi_{y_{j'}} \Psi_{y_j}^{\top} \mathbf{B} \Lambda \mathbf{B}^{\top} \Psi_{y_{j'}} \\
&\quad - \frac{2}{mn} \sum_{i, j} \Psi_{y_j}^{\top} \mathbf{B} \Lambda \Phi_{x_i} \Phi_{x_i}^{\top} \mathbf{F} \Gamma \Psi_{y_j}.
\end{aligned} \tag{C.6}$$

The following proposition summarizes the discussion so far.

Proposition 10. *Given Borel measures $\pi_{\mathcal{X}}$ and $\pi_{\mathcal{Y}}$ over \mathcal{X} and \mathcal{Y} , respectively, suppose their corresponding L^2 spaces $L_{\mathcal{X}}^2$ and $L_{\mathcal{Y}}^2$ have countable orthonormal bases: $\{\phi_k\}_{k \in \mathbb{N}}$ and $\{\psi_{\ell}\}_{\ell \in \mathbb{N}}$. Also, assume $K_{\mathcal{X}}$ and $K_{\mathcal{Y}}$ are given by the Mercer's representation (C.1) and (C.2). Let \mathcal{F}_o and \mathcal{B}_o be collections of all $F: \mathcal{X} \rightarrow \mathcal{Y}$ and $B: \mathcal{Y} \rightarrow \mathcal{X}$ such that $\|dF_{\#} \pi_{\mathcal{X}} / d\pi_{\mathcal{Y}}\|_{\infty} < \infty$ and $\|dB_{\#} \pi_{\mathcal{Y}} / d\pi_{\mathcal{X}}\|_{\infty} < \infty$, respectively. Then, solving (3.5) over $\mathcal{F}_o \times \mathcal{B}_o$ is equivalent to (3.10), where \mathcal{C} denotes the collection of all pairs of matrices (\mathbf{F}, \mathbf{B}) that correspond to a pair of bounded linear operators induced by $(F, B) \in \mathcal{F}_o \times \mathcal{B}_o$. Also, Ω is defined as*

$$\begin{aligned}
\Omega(\mathbf{F}, \mathbf{B}) &:= \frac{1}{mn} \sum_{i, j} (\Psi_{y_j}^{\top} \mathbf{B} \Lambda \Phi_{x_i} - \Phi_{x_i}^{\top} \mathbf{F} \Gamma \Psi_{y_j})^2 \\
&\quad + \frac{1}{m^2} \sum_{i, i'} \Phi_{x_i}^{\top} \Lambda \Phi_{x_{i'}} + \frac{1}{n^2} \sum_{j, j'} \Psi_{y_j}^{\top} \mathbf{B} \Lambda \mathbf{B}^{\top} \Psi_{y_{j'}} - \frac{2}{mn} \sum_{i, j} \Psi_{y_j}^{\top} \mathbf{B} \Lambda \Phi_{x_i} \\
&\quad + \frac{1}{m^2} \sum_{i, i'} \Phi_{x_i}^{\top} \mathbf{F} \Gamma \mathbf{F}^{\top} \Phi_{x_{i'}} + \frac{1}{n^2} \sum_{j, j'} \Psi_{y_j}^{\top} \Gamma \Psi_{y_{j'}} - \frac{2}{mn} \sum_{i, j} \Phi_{x_i}^{\top} \mathbf{F} \Gamma \Psi_{y_j} \\
&\quad + \frac{1}{m^2} \sum_{i, i'} \Phi_{x_i}^{\top} \Lambda \Phi_{x_{i'}} \Phi_{x_i}^{\top} \mathbf{F} \Gamma \mathbf{F}^{\top} \Phi_{x_{i'}} + \frac{1}{n^2} \sum_{j, j'} \Psi_{y_j}^{\top} \Gamma \Psi_{y_{j'}} \Psi_{y_j}^{\top} \mathbf{B} \Lambda \mathbf{B}^{\top} \Psi_{y_{j'}} \\
&\quad - \frac{2}{mn} \sum_{i, j} \Psi_{y_j}^{\top} \mathbf{B} \Lambda \Phi_{x_i} \Phi_{x_i}^{\top} \mathbf{F} \Gamma \Psi_{y_j}.
\end{aligned}$$

Based on this, we now prove Theorem 5.

Proof of Theorem 5. Let $\mathbf{x}_i = \Lambda^{1/2} \Phi_{x_i}$ and $\mathbf{y}_j = \Gamma^{1/2} \Psi_{y_j}$. Notice that we can view them as elements of a Hilbert space $\ell_{\mathbb{N}}^2$, that is, the space of square-summable sequences:

$$\ell_{\mathbb{N}}^2 = \{(a_k)_{k \in \mathbb{N}} : \sum_k a_k^2 < \infty\}.$$

Also, define $\bar{\mathbf{F}} = \Lambda^{-1/2} \mathbf{F} \Gamma^{1/2}$ and $\bar{\mathbf{B}} = \Gamma^{-1/2} \mathbf{B} \Lambda^{1/2}$ where $\Lambda, \Gamma \succ 0$. By rewriting (C.3)-(C.6) using $\mathbf{x}_i, \mathbf{y}_j, \bar{\mathbf{F}}$, and $\bar{\mathbf{B}}$, we have

$$\begin{aligned}
& \Omega(\mathbf{F}, \mathbf{B}) \\
&= \frac{1}{mn} \sum_{i, j} (\mathbf{y}_j^{\top} \bar{\mathbf{B}} \mathbf{x}_i - \mathbf{x}_i^{\top} \bar{\mathbf{F}} \mathbf{y}_j)^2 \tag{i} \\
&\quad + \frac{1}{m^2} \sum_{i, i'} \mathbf{x}_i^{\top} \mathbf{x}_{i'} + \frac{1}{n^2} \sum_{j, j'} \mathbf{y}_j^{\top} \bar{\mathbf{B}} \bar{\mathbf{B}}^{\top} \mathbf{y}_{j'} - \frac{2}{mn} \sum_{i, j} \mathbf{y}_j^{\top} \bar{\mathbf{B}} \mathbf{x}_i \tag{ii} \\
&\quad + \frac{1}{m^2} \sum_{i, i'} \mathbf{x}_i^{\top} \bar{\mathbf{F}} \bar{\mathbf{F}}^{\top} \mathbf{x}_{i'} + \frac{1}{n^2} \sum_{j, j'} \mathbf{y}_j^{\top} \mathbf{y}_{j'} - \frac{2}{mn} \sum_{i, j} \mathbf{x}_i^{\top} \bar{\mathbf{F}} \mathbf{y}_j \tag{iii} \\
&\quad + \frac{1}{m^2} \sum_{i, i'} (\mathbf{x}_i^{\top} \mathbf{x}_{i'}) (\mathbf{x}_i^{\top} \bar{\mathbf{F}} \bar{\mathbf{F}}^{\top} \mathbf{x}_{i'}) + \frac{1}{n^2} \sum_{j, j'} (\mathbf{y}_j^{\top} \mathbf{y}_{j'}) (\mathbf{y}_j^{\top} \bar{\mathbf{B}} \bar{\mathbf{B}}^{\top} \mathbf{y}_{j'}) - \frac{2}{mn} \sum_{i, j} (\mathbf{y}_j^{\top} \bar{\mathbf{B}} \mathbf{x}_i) (\mathbf{x}_i^{\top} \bar{\mathbf{F}} \mathbf{y}_j) \tag{iv} \\
&=: \bar{\Omega}(\bar{\mathbf{F}}, \bar{\mathbf{B}}).
\end{aligned}$$

As a result, (3.11) reduces to $\min_{\bar{\mathbf{F}}, \bar{\mathbf{B}} \in \mathbb{R}^{\infty \times \infty}} \bar{\Omega}(\bar{\mathbf{F}}, \bar{\mathbf{B}})$. Now, we define two finite-dimensional subspaces of $\ell_{\mathbb{N}}^2$ spanned by $(\mathbf{x}_1, \dots, \mathbf{x}_m)$ and $(\mathbf{y}_1, \dots, \mathbf{y}_n)$, respectively:

$$U_m := \text{span}\{\mathbf{x}_1, \dots, \mathbf{x}_m\}, \quad V_n := \text{span}\{\mathbf{y}_1, \dots, \mathbf{y}_n\}.$$

Also, we define P_{U_m} and P_{V_n} to be matrices that correspond to the orthogonal projection operators from $\ell_{\mathbb{N}}^2$ to U_m and to V_n , respectively. Recall that P_{U_m} and P_{V_n} are symmetric and idempotent by definition.

Our goal is to prove

$$\bar{\Omega}(\bar{\mathbf{F}}, \bar{\mathbf{B}}) \geq \bar{\Omega}(P_{U_m} \bar{\mathbf{F}} P_{V_n}, P_{V_n} \bar{\mathbf{B}} P_{U_m}).$$

More precisely, we show that four terms (i)-(iv) decrease if we replace $\bar{\mathbf{F}}$ and $\bar{\mathbf{B}}$ with $P_{U_m} \bar{\mathbf{F}} P_{V_n}$ and $P_{V_n} \bar{\mathbf{B}} P_{U_m}$, respectively. First, observe that (i) remains the same. By definition, $P_{U_m} \mathbf{x}_i = \mathbf{x}_i$ and $P_{V_n} \mathbf{y}_j = \mathbf{y}_j$, thus $\mathbf{y}_j^\top \bar{\mathbf{B}} \mathbf{x}_i = \mathbf{y}_j^\top P_{V_n} \bar{\mathbf{B}} P_{U_m} \mathbf{x}_i$ and $\mathbf{x}_i^\top \bar{\mathbf{F}} \mathbf{y}_j = \mathbf{x}_i^\top P_{U_m} \bar{\mathbf{F}} P_{V_n} \mathbf{y}_j$. Hence, (i) does not change.

To prove that (ii) decrease, it suffices to prove

$$\sum_{j, j'} \mathbf{y}_j^\top \bar{\mathbf{B}} \bar{\mathbf{B}}^\top \mathbf{y}_{j'} \geq \sum_{j, j'} \mathbf{y}_j^\top (P_{V_n} \bar{\mathbf{B}} P_{U_m}) (P_{V_n} \bar{\mathbf{B}} P_{U_m})^\top \mathbf{y}_{j'} = \sum_{j, j'} \mathbf{y}_j^\top \bar{\mathbf{B}} P_{U_m} \bar{\mathbf{B}}^\top \mathbf{y}_{j'}.$$

To this end, define $P_{U_m}^\perp$ to be a matrix that corresponds to the orthogonal projection from $\ell_{\mathbb{N}}^2$ to U_m^\perp , the orthogonal complement of U_m . By definition, $P_{U_m} + P_{U_m}^\perp$ is the identity matrix and $P_{U_m} P_{U_m}^\perp = 0$. Hence,

$$\sum_{j, j'} \mathbf{y}_j^\top \bar{\mathbf{B}} \bar{\mathbf{B}}^\top \mathbf{y}_{j'} = \left\| \bar{\mathbf{B}}^\top \sum_j \mathbf{y}_j \right\|^2 = \left\| P_{U_m} \bar{\mathbf{B}}^\top \sum_j \mathbf{y}_j \right\|^2 + \left\| P_{U_m}^\perp \bar{\mathbf{B}}^\top \sum_j \mathbf{y}_j \right\|^2 \geq \sum_{j, j'} \mathbf{y}_j^\top \bar{\mathbf{B}} P_{U_m} \bar{\mathbf{B}}^\top \mathbf{y}_{j'}$$

Here, the second equality is the Pythagorean theorem. Therefore, we can see (ii) decreases if we replace $\bar{\mathbf{B}}$ with $P_{V_n} \bar{\mathbf{B}} P_{U_m}$. Similarly, (iii) decreases.

For (iv), it suffices to prove

$$\begin{aligned} \sum_{j, j'} (\mathbf{y}_j^\top \mathbf{y}_{j'}) (\mathbf{y}_j^\top \bar{\mathbf{B}} \bar{\mathbf{B}}^\top \mathbf{y}_{j'}) &\geq \sum_{j, j'} (\mathbf{y}_j^\top \mathbf{y}_{j'}) (\mathbf{y}_j^\top (P_{V_n} \bar{\mathbf{B}} P_{U_m}) (P_{V_n} \bar{\mathbf{B}} P_{U_m})^\top \mathbf{y}_{j'}) \\ &= \sum_{j, j'} (\mathbf{y}_j^\top \mathbf{y}_{j'}) (\mathbf{y}_j^\top \bar{\mathbf{B}} P_{U_m} \bar{\mathbf{B}}^\top \mathbf{y}_{j'}). \end{aligned}$$

To see this,

$$\begin{aligned} \sum_{j, j'} (\mathbf{y}_j^\top \mathbf{y}_{j'}) (\mathbf{y}_j^\top \bar{\mathbf{B}} \bar{\mathbf{B}}^\top \mathbf{y}_{j'}) &= \sum_{j, j'} (\mathbf{y}_j^\top \mathbf{y}_{j'}) (\mathbf{y}_j^\top \bar{\mathbf{B}} P_{U_m} \bar{\mathbf{B}}^\top \mathbf{y}_{j'}) + \sum_{j, j'} (\mathbf{y}_j^\top \mathbf{y}_{j'}) (\mathbf{y}_j^\top \bar{\mathbf{B}} P_{U_m}^\perp \bar{\mathbf{B}}^\top \mathbf{y}_{j'}) \\ &\geq \sum_{j, j'} (\mathbf{y}_j^\top \mathbf{y}_{j'}) (\mathbf{y}_j^\top \bar{\mathbf{B}} P_{U_m} \bar{\mathbf{B}}^\top \mathbf{y}_{j'}), \end{aligned}$$

where the inequality holds since

$$\sum_{j, j'} (\mathbf{y}_j^\top \mathbf{y}_{j'}) (\mathbf{y}_j^\top \bar{\mathbf{B}} P_{U_m}^\perp \bar{\mathbf{B}}^\top \mathbf{y}_{j'}) = \text{Tr} \left[\left(\sum_j P_{U_m}^\perp \bar{\mathbf{B}}^\top \mathbf{y}_j \mathbf{y}_j^\top \right) \left(\sum_j P_{U_m}^\perp \bar{\mathbf{B}}^\top \mathbf{y}_j \mathbf{y}_j^\top \right)^\top \right] \geq 0.$$

Similarly, we can obtain

$$\sum_{i, i'} (\mathbf{x}_i^\top \mathbf{x}_{i'}) (\mathbf{x}_i^\top \bar{\mathbf{F}} \bar{\mathbf{F}}^\top \mathbf{x}_{i'}) \geq \sum_{i, i'} (\mathbf{x}_i^\top \mathbf{x}_{i'}) (\mathbf{x}_i^\top \bar{\mathbf{F}} P_{V_n} \bar{\mathbf{F}}^\top \mathbf{x}_{i'}).$$

Hence, (iv) decreases.

Consequently, we have

$$(3.11) = \min_{\bar{\mathbf{F}}, \bar{\mathbf{B}} \in \mathbb{R}^{\infty \times \infty}} \bar{\Omega}(\bar{\mathbf{F}}, \bar{\mathbf{B}}) = \min_{\bar{\mathbf{F}}, \bar{\mathbf{B}} \in \mathbb{R}^{\infty \times \infty}} \bar{\Omega}(P_{U_m} \bar{\mathbf{F}} P_{V_n}, P_{V_n} \bar{\mathbf{B}} P_{U_m}).$$

By definition of a projection operator, we can find $\mathbf{U}_m \in \mathbb{R}^{m \times \infty}$ and $\mathbf{V}_n \in \mathbb{R}^{n \times \infty}$ such that

$$P_{U_m} = \Lambda^{1/2} \Phi_m \mathbf{U}_m, \quad P_{V_n} = \Gamma^{1/2} \Psi_n \mathbf{V}_n.$$

By letting $\mathbf{U}_m \bar{\mathbf{F}} \mathbf{V}_n^\top = \mathbf{F}_{m,n} \in \mathbb{R}^{m \times n}$ and $\mathbf{V}_n \bar{\mathbf{B}} \mathbf{U}_m^\top = \mathbf{B}_{n,m} \in \mathbb{R}^{n \times m}$, we have

$$\begin{aligned} P_{U_m} \bar{\mathbf{F}} P_{V_n} &= \Lambda^{1/2} \Phi_m \mathbf{F}_{m,n} \Psi_n^\top \Gamma^{1/2}, \\ P_{V_n} \bar{\mathbf{B}} P_{U_m} &= \Gamma^{1/2} \Psi_n \mathbf{B}_{n,m} \Phi_m^\top \Lambda^{1/2}. \end{aligned}$$

Hence,

$$\min_{\bar{\mathbf{F}}, \bar{\mathbf{B}} \in \mathbb{R}^{\infty \times \infty}} \bar{\Omega}(P_{U_m} \bar{\mathbf{F}} P_{V_n}, P_{V_n} \bar{\mathbf{B}} P_{U_m}) = \min_{(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}) \in \mathbf{C}} \omega(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}),$$

where

$$\omega(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}) := \bar{\Omega}(\Lambda^{1/2} \Phi_m \mathbf{F}_{m,n} \Psi_n^\top \Gamma^{1/2}, \Gamma^{1/2} \Psi_n \mathbf{B}_{n,m} \Phi_m^\top \Lambda^{1/2}).$$

Here, \mathbf{C} is a constraint set implying that $\mathbf{F}_{m,n}$ and $\mathbf{B}_{n,m}$ are associated with $\bar{\mathbf{F}}$ and $\bar{\mathbf{B}}$, respectively, namely,

$$\mathbf{C} = \{(\mathbf{U}_m \bar{\mathbf{F}} \mathbf{V}_n^\top, \mathbf{V}_n \bar{\mathbf{B}} \mathbf{U}_m^\top) : \bar{\mathbf{F}}, \bar{\mathbf{B}} \in \mathbb{R}^{\infty \times \infty}\} \subset \mathbb{R}^{m \times n} \times \mathbb{R}^{n \times m}.$$

Therefore,

$$(3.11) = \min_{(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}) \in \mathbf{C}} \omega(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}) \geq \min_{\substack{\mathbf{F}_{m,n} \in \mathbb{R}^{m \times n} \\ \mathbf{B}_{n,m} \in \mathbb{R}^{n \times m}}} \omega(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}).$$

Finally, note that $\mathbf{C} = \mathbb{R}^{m \times n} \times \mathbb{R}^{n \times m}$ if \mathbf{U}_m and \mathbf{V}_n are full rank, that is, row spaces of \mathbf{U}_m and \mathbf{V}_n are rank- m and rank- n , respectively. This is true if kernel matrices

$$\mathbf{K}_{\mathcal{X}} = (\Lambda^{1/2} \Phi_m)^\top (\Lambda^{1/2} \Phi_m), \quad \mathbf{K}_{\mathcal{Y}} = (\Gamma^{1/2} \Psi_n)^\top (\Gamma^{1/2} \Psi_n)$$

are invertible. This is equivalent to say that they are positive definite. In this case,

$$\mathbf{U}_m = \mathbf{K}_{\mathcal{X}}^{-1} (\Lambda^{1/2} \Phi_m)^\top, \quad \mathbf{V}_n = \mathbf{K}_{\mathcal{Y}}^{-1} (\Gamma^{1/2} \Psi_n)^\top,$$

which are indeed full rank. Accordingly, we have

$$(3.11) = \min_{(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}) \in \mathbf{C}} \omega(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}) = \min_{\substack{\mathbf{F}_{m,n} \in \mathbb{R}^{m \times n} \\ \mathbf{B}_{n,m} \in \mathbb{R}^{n \times m}}} \omega(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}).$$

Finally, we prove ω is convex. To see this, verify

$$\begin{aligned} \omega(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}) &= \frac{1}{mn} \|\mathbf{K}_{\mathcal{Y}} \mathbf{B}_{n,m} \mathbf{K}_{\mathcal{X}} - \mathbf{K}_{\mathcal{Y}} \mathbf{F}_{m,n}^\top \mathbf{K}_{\mathcal{X}}\|^2 \\ &+ \left\| \mathbf{K}_{\mathcal{X}}^{1/2} \cdot \left(\frac{1}{m} \mathbf{1}_m - \mathbf{B}_{n,m}^\top \mathbf{K}_{\mathcal{Y}} \frac{1}{n} \mathbf{1}_n \right) \right\|^2 + \left\| \mathbf{K}_{\mathcal{Y}}^{1/2} \cdot \left(\frac{1}{n} \mathbf{1}_n - \mathbf{F}_{m,n}^\top \mathbf{K}_{\mathcal{X}} \frac{1}{m} \mathbf{1}_m \right) \right\|^2 \\ &+ \left\| \frac{1}{m} \mathbf{K}_{\mathcal{X}}^{3/2} \mathbf{F}_{m,n} \mathbf{K}_{\mathcal{Y}}^{1/2} - \frac{1}{n} \mathbf{K}_{\mathcal{X}}^{1/2} \mathbf{B}_{n,m}^\top \mathbf{K}_{\mathcal{Y}}^{3/2} \right\|^2, \end{aligned}$$

where $\mathbf{K}_{\mathcal{X}}^{1/2}$ and $\mathbf{K}_{\mathcal{Y}}^{1/2}$ are the square root matrices of $\mathbf{K}_{\mathcal{X}}$ and $\mathbf{K}_{\mathcal{Y}}$, respectively, and $\mathbf{1}_m \in \mathbb{R}^m$ and $\mathbf{1}_n \in \mathbb{R}^n$ are all-ones vectors. \square

D Supporting Proofs of Section B

D.1 Proofs in Section B

Proof of Proposition 4. Let $h_{F,B}(x, y) := (c_{\mathcal{X}}(x, B(y)) - c_{\mathcal{Y}}(F(x), y))^2$, then

$$\begin{aligned}\widehat{C}_0(F, B) - C_0(F, B) &= \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n h_{F,B}(x_i, y_j) - \mathbb{E}_{(x,y) \sim \mu \otimes \nu} h_{F,B}(x, y) \\ &= \frac{1}{m} \sum_{i=1}^m \left(\frac{1}{n} \sum_{j=1}^n h_{F,B}(x_i, y_j) - \mathbb{E}_{y \sim \nu} h_{F,B}(x_i, y) \right) \\ &\quad + \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{y \sim \nu} h_{F,B}(x_i, y) - \mathbb{E}_{(x,y) \sim \mu \otimes \nu} h_{F,B}(x, y).\end{aligned}$$

Assumption 1 implies that a function $x \mapsto \mathbb{E}_{y \sim \nu} h_{F,B}(x, y)$ is bounded in $[0, H]$. Thus, by the McDiarmid's inequality,

$$\frac{1}{m} \sum_{i=1}^m \mathbb{E}_{y \sim \nu} h_{F,B}(x_i, y) - \mathbb{E}_{(x,y) \sim \mu \otimes \nu} h_{F,B}(x, y) \leq \sqrt{\frac{H^2 \log(1/\delta)}{2m}}$$

holds with probability at least $1 - \delta$. By the same logic, for fixed x_i ,

$$\frac{1}{n} \sum_{j=1}^n h_{F,B}(x_i, y_j) - \mathbb{E}_{y \sim \nu} h_{F,B}(x_i, y) \leq \sqrt{\frac{H^2 \log(1/\delta)}{2n}}$$

holds with probability at least $1 - \delta$, where the probability is the conditional probability of y_1, \dots, y_n given x_1, \dots, x_m . Since this is true for all x_i , the union bound implies

$$\frac{1}{m} \sum_{i=1}^m \left(\frac{1}{n} \sum_{j=1}^n h_{F,B}(x_i, y_j) - \mathbb{E}_{y \sim \nu} h_{F,B}(x_i, y) \right) \leq \sqrt{\frac{H^2 \log(m/\delta)}{2n}}$$

holds with probability at least $1 - \delta$. Hence,

$$\widehat{C}_0(F, B) - C_0(F, B) \lesssim \sqrt{\frac{\log(m/\delta)}{m}} \leq \sqrt{\frac{\log(\frac{m \vee n}{\delta})}{m \wedge n}}$$

holds with probability at least $1 - 2\delta$. The same result holds for $C_0(F, B) - \widehat{C}_0(F, B)$, hence we complete the proof. \square

Proof of Proposition 5. By the triangle inequality, $|\widehat{M}(F, B) - M(F, B)|$ is bounded above by the sum of the following three terms:

$$\begin{aligned}& |\text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#} \widehat{\mu}_m, \widehat{\nu}_n) - \text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#} \mu, \nu)|, \\ & |\text{MMD}_{K_{\mathcal{X}}}^2(\widehat{\mu}_m, B_{\#} \widehat{\nu}_n) - \text{MMD}_{K_{\mathcal{X}}}^2(\mu, B_{\#} \nu)|, \\ & |\text{MMD}_{K_{\mathcal{X}} \otimes K_{\mathcal{Y}}}^2((\text{Id}, F)_{\#} \widehat{\mu}_m, (B, \text{Id})_{\#} \widehat{\nu}_n) - \text{MMD}_{K_{\mathcal{X}} \otimes K_{\mathcal{Y}}}^2((\text{Id}, F)_{\#} \mu, (B, \text{Id})_{\#} \nu)|.\end{aligned}$$

First, we give an upper bound on the first term. Boundedness of kernels (Assumption 2) implies

$$\text{MMD}_{K_{\mathcal{Y}}}(F_{\#} \widehat{\mu}_m, \widehat{\nu}_n), \text{MMD}_{K_{\mathcal{Y}}}(F_{\#} \mu, \nu) \leq 2\sqrt{K}.$$

Hence,

$$|\text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#} \widehat{\mu}_m, \widehat{\nu}_n) - \text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#} \mu, \nu)| \leq 4\sqrt{K} |\text{MMD}_{K_{\mathcal{Y}}}(F_{\#} \widehat{\mu}_m, \widehat{\nu}_n) - \text{MMD}_{K_{\mathcal{Y}}}(F_{\#} \mu, \nu)|.$$

Due to the triangle inequality of MMD, we have

$$|\text{MMD}_{K_Y}(F_{\#}\hat{\mu}_m, \hat{\nu}_n) - \text{MMD}_{K_Y}(F_{\#}\mu, \nu)| \leq \text{MMD}_{K_Y}(F_{\#}\hat{\mu}_m, F_{\#}\mu) + \text{MMD}_{K_Y}(\hat{\nu}_n, \nu).$$

By Theorem 3.4 of [41],

$$\text{MMD}_{K_Y}(\hat{\nu}_n, \nu) \leq \sqrt{\frac{K}{n}} + \sqrt{\frac{2K \log(1/\delta)}{n}}$$

holds with probability at least $1 - \delta$. Next, note that $F_{\#}\hat{\mu}_m = \frac{1}{m} \sum_i \delta_{F(x_i)}$ is the empirical measure constructed from $\{F(x_i)\}_{i=1}^m$. Since they are m many i.i.d. samples from $F_{\#}\mu$, by the same theorem,

$$\text{MMD}_{K_Y}(F_{\#}\hat{\mu}_m, F_{\#}\mu) \leq \sqrt{\frac{K}{m}} + \sqrt{\frac{2K \log(1/\delta)}{m}}$$

holds with probability at least $1 - \delta$. Hence,

$$|\text{MMD}_{K_Y}^2(F_{\#}\hat{\mu}_m, \hat{\nu}_n) - \text{MMD}_{K_Y}^2(F_{\#}\mu, \nu)| \lesssim \sqrt{\frac{\log(1/\delta)}{m}} + \sqrt{\frac{\log(1/\delta)}{n}}$$

holds with probability at least $1 - 2\delta$. Similarly, we have

$$\begin{aligned} |\text{MMD}_{K_X}^2(\hat{\mu}_m, B_{\#}\hat{\nu}_n) - \text{MMD}_{K_X}^2(\mu, B_{\#}\nu)| &\lesssim \sqrt{\frac{\log(1/\delta)}{m}} + \sqrt{\frac{\log(1/\delta)}{n}}, \\ |\text{MMD}_{K_X \otimes K_Y}^2((\text{Id}, F)_{\#}\hat{\mu}_m, (B, \text{Id})_{\#}\hat{\nu}_n) - \text{MMD}_{K_X \otimes K_Y}^2((\text{Id}, F)_{\#}\mu, (B, \text{Id})_{\#}\nu)| \\ &\lesssim \sqrt{\frac{\log(1/\delta)}{m}} + \sqrt{\frac{\log(1/\delta)}{n}}, \end{aligned}$$

each of which holds with probability at least $1 - 2\delta$. Combining these three probabilistic bounds, we obtain a bound for $|\widehat{M}(F, B) - M(F, B)|$. \square

Proof of Proposition 6. Without loss of generality, assume $n \geq m$. From the proof of Proposition 4,

$$\begin{aligned} \sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} |\widehat{C}_0(F, B) - C_0(F, B)| &\leq \frac{1}{m} \sum_{i=1}^m \sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{n} \sum_{j=1}^n h_{F, B}(x_i, y_j) - \mathbb{E}_{y \sim \nu} h_{F, B}(x_i, y) \right| \\ &\quad + \sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{y \sim \nu} h_{F, B}(x_i, y) - \mathbb{E}_{x \sim \mu} \mathbb{E}_{y \sim \nu} h_{F, B}(x, y) \right|. \end{aligned}$$

Since $x \mapsto \mathbb{E}_{y \sim \nu} h_{F, B}(x, y)$ is bounded in $[0, H]$, Lemma 4 implies

$$\begin{aligned} \sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{y \sim \nu} h_{F, B}(x_i, y) - \mathbb{E}_{x \sim \mu} \mathbb{E}_{y \sim \nu} h_{F, B}(x, y) \right| \\ \lesssim \sqrt{\frac{\log(1/\delta)}{m}} + \mathbb{E} \mathbb{E}_{x_i \epsilon_i} \sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{E}_{y \sim \nu} h_{F, B}(x_i, y) \right| \end{aligned}$$

holds with probability at least $1 - \delta$. Since $n \geq m$,

$$\begin{aligned} \mathbb{E} \mathbb{E}_{x_i \epsilon_i} \sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{E}_{y \sim \nu} h_{F, B}(x_i, y) \right| &= \mathbb{E} \mathbb{E}_{x_i \epsilon_i} \sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} \left| \mathbb{E}_{y_1, \dots, y_m} \frac{1}{m} \sum_{i=1}^m \epsilon_i h_{F, B}(x_i, y_i) \right| \\ &\leq \mathbb{E} \mathbb{E}_{x_i \epsilon_i} \mathbb{E}_{y_1, \dots, y_m} \sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{m} \sum_{i=1}^m \epsilon_i h_{F, B}(x_i, y_i) \right| \\ &= \mathbb{E} \mathbb{E}_{x_i y_i \epsilon_i} \sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{m} \sum_{i=1}^m \epsilon_i h_{F, B}(x_i, y_i) \right|. \end{aligned}$$

We first give an upper bound on

$$\mathbb{E} \sup_{\epsilon_i \substack{(F,B) \in \mathcal{F} \times \mathcal{B} \\ =: X_{F,B}}} \left| \frac{1}{m} \sum_{i=1}^m \epsilon_i h_{F,B}(x_i, y_i) \right|.$$

First, observe that Assumption 1 and Assumption 4 imply

$$\begin{aligned} |h_{F,B}(x, y) - h_{F',B'}(x, y)| &\leq \left| \sqrt{h_{F,B}(x, y)} + \sqrt{h_{F',B'}(x, y)} \right| \left| \sqrt{h_{F,B}(x, y)} - \sqrt{h_{F',B'}(x, y)} \right| \\ &\leq 2\sqrt{H} (|c_{\mathcal{X}}(x, B(y)) - c_{\mathcal{X}}(x, B'(y))| + |c_{\mathcal{Y}}(F(x), y) - c_{\mathcal{Y}}(F'(x), y)|) \\ &\leq 2\sqrt{H}L (\|F(x) - F'(x)\| + \|B(y) - B'(y)\|) \\ &= 2\sqrt{H}L \left[\sqrt{\sum_{k=1}^{\dim(\mathcal{Y})} |F_k(x) - F'_k(x)|^2} + \sqrt{\sum_{\ell=1}^{\dim(\mathcal{X})} |B_{\ell}(y) - B'_{\ell}(y)|^2} \right] \\ &\leq 2\sqrt{H}L \left(\sum_{k=1}^{\dim(\mathcal{Y})} |F_k(x) - F'_k(x)| + \sum_{\ell=1}^{\dim(\mathcal{X})} |B_{\ell}(y) - B'_{\ell}(y)| \right). \end{aligned}$$

Therefore,

$$\begin{aligned} |X_{F,B} - X_{F',B'}| &\leq \frac{1}{m} \sum_{i=1}^m |h_{F,B}(x_i, y_i) - h_{F',B'}(x_i, y_i)| \\ &\leq 2\sqrt{H}L \left(\sum_{k=1}^{\dim(\mathcal{Y})} \frac{1}{m} \sum_{i=1}^m |F_k(x_i) - F'_k(x_i)| + \sum_{\ell=1}^{\dim(\mathcal{X})} \frac{1}{m} \sum_{i=1}^m |B_{\ell}(y_i) - B'_{\ell}(y_i)| \right) \\ &\leq 2\sqrt{H}L \left(\sum_{k=1}^{\dim(\mathcal{Y})} \max_{i \in [m]} |F_k(x_i) - F'_k(x_i)| + \sum_{\ell=1}^{\dim(\mathcal{X})} \max_{i \in [m]} |B_{\ell}(y_i) - B'_{\ell}(y_i)| \right) \\ &=: \rho((F, B), (F', B')). \end{aligned}$$

For $\epsilon > 0$, let $\mathcal{N}_{\infty}(\epsilon, \mathcal{F}_k, \{x_i\}_{i=1}^m)$ be the minimal ϵ -covering net of \mathcal{F}_k under the pseudometric d induced by x_1, \dots, x_m :

$$d(F_k, F'_k) := \max_{i \in [m]} |F_k(x_i) - F'_k(x_i)|.$$

In other words, for any $F_k \in \mathcal{F}_k$, we can find $F'_k \in \mathcal{N}_{\infty}(\epsilon, \mathcal{F}_k, \{x_i\}_{i=1}^m)$ such that $d(F_k, F'_k) \leq \epsilon$. Also, $|\mathcal{N}_{\infty}(\epsilon, \mathcal{F}_k, \{x_i\}_{i=1}^m)| = N_{\infty}(\epsilon, \mathcal{F}_k, \{x_i\}_{i=1}^m)$. We define $\mathcal{N}_{\infty}(\epsilon, \mathcal{B}_{\ell}, \{y_i\}_{i=1}^m)$ in a similar fashion.

Given $\epsilon > 0$, let $T_{\epsilon} = \otimes_{k=1}^{\dim(\mathcal{Y})} \mathcal{N}_{\infty}(\epsilon, \mathcal{F}_k, \{x_i\}_{i=1}^m) \times \otimes_{\ell=1}^{\dim(\mathcal{X})} \mathcal{N}_{\infty}(\epsilon, \mathcal{B}_{\ell}, \{y_i\}_{i=1}^m)$. Then, for any $(F, B) \in \mathcal{F} \times \mathcal{B}$, we can find $(F', B') \in T_{\epsilon}$ such that

$$\rho((F, B), (F', B')) \leq \eta \epsilon,$$

where $\eta = 2\sqrt{H}L(\dim(\mathcal{X}) + \dim(\mathcal{Y}))$. As a result, one can easily check

$$\begin{aligned} \sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} X_{F,B} &\leq \sup_{\rho((F,B), (F',B')) \leq \eta \epsilon} |X_{F,B} - X_{F',B'}| + \sup_{(F',B') \in T_{\epsilon}} X_{F',B'} \\ &\leq \eta \epsilon + \sup_{(F',B') \in T_{\epsilon}} X_{F',B'}. \end{aligned}$$

Note that $X_{F,B}$ is the absolute value of a sub-Gaussian random variable with parameter $H^2/(4m)$. Hence, the maximal inequality yields

$$\mathbb{E} \sup_{\epsilon_i (F,B) \in \mathcal{F} \times \mathcal{B}} X_{F,B} \leq \eta\epsilon + \mathbb{E} \sup_{\epsilon_i (F,B) \in T_\epsilon} X_{F,B} \leq \eta\epsilon + \sqrt{\frac{H^2 \log(|T_\epsilon|)}{m}}.$$

Using $|T_\epsilon| = \prod_{k=1}^{\dim(\mathcal{Y})} N_\infty(\epsilon, \mathcal{F}_k, \{x_i\}_{i=1}^m) \times \prod_{\ell=1}^{\dim(\mathcal{X})} N_\infty(\epsilon, \mathcal{B}_\ell, \{y_i\}_{i=1}^m)$, we have

$$\begin{aligned} & \mathbb{E} \sup_{\epsilon_i (F,B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{m} \sum_{i=1}^m \epsilon_i h_{F,B}(x_i, y_i) \right| \\ & \leq \eta\epsilon + H \sqrt{\frac{\sum_{k=1}^{\dim(\mathcal{Y})} \log N_\infty(\epsilon, \mathcal{F}_k, \{x_i\}_{i=1}^m) + \sum_{\ell=1}^{\dim(\mathcal{X})} \log N_\infty(\epsilon, \mathcal{B}_\ell, \{y_i\}_{i=1}^m)}{m}} \\ & \leq \eta\epsilon + H \sqrt{\frac{\sum_{k=1}^{\dim(\mathcal{Y})} \log N_\infty(\epsilon, \mathcal{F}_k, m) + \sum_{\ell=1}^{\dim(\mathcal{X})} \log N_\infty(\epsilon, \mathcal{B}_\ell, m)}{m}}. \end{aligned}$$

The second inequality is obvious from the definition of the uniform covering number. Since the last equation is independent of x_i and y_i , we have

$$\begin{aligned} & \mathbb{E} \mathbb{E} \mathbb{E} \sup_{x_i y_i \epsilon_i (F,B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{m} \sum_{i=1}^m \epsilon_i h_{F,B}(x_i, y_i) \right| \\ & \leq \eta\epsilon + H \sqrt{\frac{\sum_{k=1}^{\dim(\mathcal{Y})} \log N_\infty(\epsilon, \mathcal{F}_k, m) + \sum_{\ell=1}^{\dim(\mathcal{X})} \log N_\infty(\epsilon, \mathcal{B}_\ell, m)}{m}}. \end{aligned}$$

As a result,

$$\begin{aligned} & \sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{y \sim \nu} h_{F,B}(x_i, y) - \mathbb{E}_{x \sim \mu} \mathbb{E}_{y \sim \nu} h_{F,B}(x, y) \right| \\ & \lesssim \sqrt{\frac{\log(1/\delta)}{m}} + \epsilon + \sqrt{\frac{\sum_{k=1}^{\dim(\mathcal{Y})} \log N_\infty(\epsilon, \mathcal{F}_k, m) + \sum_{\ell=1}^{\dim(\mathcal{X})} \log N_\infty(\epsilon, \mathcal{B}_\ell, m)}{m}} \quad (\text{D.1}) \\ & \leq \sqrt{\frac{\log(1/\delta)}{m}} + \epsilon + \sqrt{\frac{\sum_{k=1}^{\dim(\mathcal{Y})} \log N_\infty(\epsilon, \mathcal{F}_k, m) + \sum_{\ell=1}^{\dim(\mathcal{X})} \log N_\infty(\epsilon, \mathcal{B}_\ell, n)}{m}} \end{aligned}$$

holds with probability at least $1 - \delta$. Here, $\log N_\infty(\epsilon, \mathcal{B}_\ell, m) \leq \log N_\infty(\epsilon, \mathcal{B}_\ell, n)$ holds since $n \geq m$, which is obvious from the definition of the uniform covering number.

Next, we give a bound on

$$\frac{1}{m} \sum_{i=1}^m \sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{n} \sum_{j=1}^n h_{F,B}(x_i, y_j) - \mathbb{E}_{y \sim \nu} h_{F,B}(x_i, y) \right|.$$

Considering x_1, \dots, x_m are fixed, Lemma 4 implies

$$\begin{aligned} & \sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{n} \sum_{j=1}^n h_{F,B}(x_i, y_j) - \mathbb{E}_{y \sim \nu} h_{F,B}(x_i, y) \right| \\ & \lesssim \sqrt{\frac{\log(1/\delta)}{n}} + \mathbb{E} \mathbb{E} \sup_{y_j \epsilon_j (F,B) \in \mathcal{F} \times \mathcal{B}} \underbrace{\left| \frac{1}{n} \sum_{j=1}^n \epsilon_j h_{F,B}(x_i, y_j) \right|}_{Y_{F,B}} \end{aligned}$$

holds with probability at least $1 - \delta$. Here, the probability should be understood as a conditional probability of y_1, \dots, y_n given x_1, \dots, x_m . Again, we have

$$\begin{aligned} |Y_{F,B} - Y_{F',B'}| &\leq 2\sqrt{HL} \left(\sum_{k=1}^{\dim(\mathcal{Y})} |F_k(x_i) - F'_k(x_i)| + \sum_{\ell=1}^{\dim(\mathcal{X})} \frac{1}{n} \sum_{j=1}^n |B_\ell(y_j) - B'_\ell(y_j)| \right) \\ &\leq 2\sqrt{HL} \left(\sum_{k=1}^{\dim(\mathcal{Y})} \max_{i \in [m]} |F_k(x_i) - F'_k(x_i)| + \sum_{\ell=1}^{\dim(\mathcal{X})} \max_{j \in [n]} |B_\ell(y_j) - B'_\ell(y_j)| \right). \end{aligned}$$

Also, $Y_{F,B}$ is the absolute value of a sub-Gaussian random variable with parameter $H^2/(4n)$. By the same argument as before,

$$\begin{aligned} &\mathbb{E} \mathbb{E} \sup_{y_j \in \mathcal{B}} \left| \frac{1}{n} \sum_{j=1}^n \epsilon_j h_{F,B}(x_i, y_j) \right| \\ &\leq \eta\epsilon + H \sqrt{\frac{\sum_{k=1}^{\dim(\mathcal{Y})} \log N_\infty(\epsilon, \mathcal{F}_k, m) + \sum_{\ell=1}^{\dim(\mathcal{X})} \log N_\infty(\epsilon, \mathcal{B}_\ell, n)}{n}}. \end{aligned}$$

Hence,

$$\begin{aligned} &\sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{n} \sum_{j=1}^n h_{F,B}(x_i, y_j) - \mathbb{E}_{y \sim \nu} h_{F,B}(x_i, y) \right| \\ &\lesssim \sqrt{\frac{\log(1/\delta)}{n}} + \epsilon + \sqrt{\frac{\sum_{k=1}^{\dim(\mathcal{Y})} \log N_\infty(\epsilon, \mathcal{F}_k, m) + \sum_{\ell=1}^{\dim(\mathcal{X})} \log N_\infty(\epsilon, \mathcal{B}_\ell, n)}{n}} \end{aligned}$$

holds with probability (conditional probability as explained earlier) at least $1 - \delta$. Since this holds for all x_i , the union bound implies

$$\begin{aligned} &\frac{1}{m} \sum_{i=1}^m \sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} \left| \frac{1}{n} \sum_{j=1}^n h_{F,B}(x_i, y_j) - \mathbb{E}_{y \sim \nu} h_{F,B}(x_i, y) \right| \\ &\lesssim \sqrt{\frac{\log(m/\delta)}{n}} + \epsilon + \sqrt{\frac{\sum_{k=1}^{\dim(\mathcal{Y})} \log N_\infty(\epsilon, \mathcal{F}_k, m) + \sum_{\ell=1}^{\dim(\mathcal{X})} \log N_\infty(\epsilon, \mathcal{B}_\ell, n)}{n}} \quad (\text{D.2}) \\ &\leq \sqrt{\frac{\log(m/\delta)}{m}} + \epsilon + \sqrt{\frac{\sum_{k=1}^{\dim(\mathcal{Y})} \log N_\infty(\epsilon, \mathcal{F}_k, m) + \sum_{\ell=1}^{\dim(\mathcal{X})} \log N_\infty(\epsilon, \mathcal{B}_\ell, n)}{m}} \end{aligned}$$

holds with probability at least $1 - \delta$. Combining (D.1) and (D.2), for any $\epsilon > 0$, we have

$$\begin{aligned} &\sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} |\widehat{C}_0(F, B) - C_0(F, B)| \\ &\lesssim \sqrt{\frac{\log\left(\frac{m \vee n}{\delta}\right)}{m \wedge n}} + \epsilon + \sqrt{\frac{\sum_{k=1}^{\dim(\mathcal{Y})} \log N_\infty(\epsilon, \mathcal{F}_k, m) + \sum_{\ell=1}^{\dim(\mathcal{X})} \log N_\infty(\epsilon, \mathcal{B}_\ell, n)}{m \wedge n}} \end{aligned}$$

holds with probability at least $1 - 2\delta$. □

Proof of Corollary 1. Combining Assumption 3 and Lemma 5, we have

$$N_\infty(\epsilon, \mathcal{F}_k, m) \leq \left(\frac{2emb}{\epsilon \cdot \text{Pdim}(\mathcal{F}_k)} \right)^{\text{Pdim}(\mathcal{F}_k)}.$$

Hence,

$$\begin{aligned}
& \sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} |\widehat{C}_0(F,B) - C_0(F,B)| \\
& \lesssim \sqrt{\frac{\log\left(\frac{m \vee n}{\delta}\right)}{m \wedge n}} + \epsilon + \sqrt{\frac{\sum_{k=1}^{\dim(\mathcal{Y})} \log N_\infty(\epsilon, \mathcal{F}_k, m) + \sum_{\ell=1}^{\dim(\mathcal{X})} \log N_\infty(\epsilon, \mathcal{B}_\ell, n)}{m \wedge n}} \\
& \leq \sqrt{\frac{\log\left(\frac{m \vee n}{\delta}\right)}{m \wedge n}} + \epsilon + \sqrt{\frac{\log\left(\frac{2eb(m \vee n)}{\epsilon}\right)}{m \wedge n} \left(\sum_{k=1}^{\dim(\mathcal{Y})} \text{Pdim}(\mathcal{F}_k) + \sum_{\ell=1}^{\dim(\mathcal{X})} \text{Pdim}(\mathcal{B}_\ell) \right)} \\
& \lesssim \sqrt{\frac{\log\left(\frac{m \vee n}{\delta}\right)}{m \wedge n}} + \sqrt{\frac{\log(m \vee n)}{m \wedge n} \left(\sum_{k=1}^{\dim(\mathcal{Y})} \text{Pdim}(\mathcal{F}_k) + \sum_{\ell=1}^{\dim(\mathcal{X})} \text{Pdim}(\mathcal{B}_\ell) \right)}
\end{aligned}$$

holds with probability at least $1 - 2\delta$, where the last bound comes from choosing $\epsilon = (m \wedge n)^{-1/2}$. \square

Proof of Proposition 7. Using the triangle inequality, we bound $\sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} |\widehat{M}(F,B) - M(F,B)|$ by the sum of the following three terms:

$$\begin{aligned}
& \sup_{F \in \mathcal{F}} |\text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#}\widehat{\mu}_m, \widehat{\nu}_n) - \text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#}\mu, \nu)|, \\
& \sup_{B \in \mathcal{B}} |\text{MMD}_{K_{\mathcal{X}}}^2(\widehat{\mu}_m, B_{\#}\widehat{\nu}_n) - \text{MMD}_{K_{\mathcal{X}}}^2(\mu, B_{\#}\nu)|, \\
& \sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} |\text{MMD}_{K_{\mathcal{X}} \otimes K_{\mathcal{Y}}}^2((\text{Id}, F)_{\#}\widehat{\mu}_m, (B, \text{Id})_{\#}\widehat{\nu}_n) - \text{MMD}_{K_{\mathcal{X}} \otimes K_{\mathcal{Y}}}^2((\text{Id}, F)_{\#}\mu, (B, \text{Id})_{\#}\nu)|.
\end{aligned}$$

As in the proof of Proposition 5, we have

$$\begin{aligned}
& \sup_{F \in \mathcal{F}} |\text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#}\widehat{\mu}_m, \widehat{\nu}_n) - \text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#}\mu, \nu)| \\
& \leq 4\sqrt{K} \left[\sup_{F \in \mathcal{F}} \text{MMD}_{K_{\mathcal{Y}}}(F_{\#}\widehat{\mu}_m, F_{\#}\mu) + \text{MMD}_{K_{\mathcal{Y}}}(\widehat{\nu}_n, \nu) \right].
\end{aligned}$$

$\text{MMD}_{K_{\mathcal{Y}}}(\widehat{\nu}_n, \nu)$ has already been bounded in Proposition 5. For the first term on the RHS, observe that

$$\begin{aligned}
\sup_{F \in \mathcal{F}} \text{MMD}_{K_{\mathcal{Y}}}(F_{\#}\widehat{\mu}_m, F_{\#}\mu) &= \sup_{F \in \mathcal{F}} \sup_{f \in \mathcal{H}_{\mathcal{Y}}(1)} \left| \int f \, dF_{\#}\widehat{\mu}_m - \int f \, dF_{\#}\mu \right| \\
&= \sup_{F \in \mathcal{F}} \sup_{f \in \mathcal{H}_{\mathcal{Y}}(1)} \left| \int f \circ F \, d\widehat{\mu}_m - \int f \circ F \, d\mu \right| \\
&= \sup_{f \in \mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}} \left| \int f \, d\widehat{\mu}_m - \int f \, d\mu \right|,
\end{aligned}$$

where the second equality follows from change-of-variables.

First, we show $\mathcal{H}_{\mathcal{Y}}(1)$ consists of \sqrt{K} -uniformly bounded functions. Let $\|\cdot\|_{\mathcal{H}_{\mathcal{Y}}}$ be the norm of $\mathcal{H}_{\mathcal{Y}}$ so that $f \in \mathcal{H}_{\mathcal{Y}}(1)$ is equivalent to $\|f\|_{\mathcal{H}_{\mathcal{Y}}} \leq 1$. Then, the reproducing property implies

$$|f(y)| \leq \|f\|_{\mathcal{H}_{\mathcal{Y}}} \sqrt{K_{\mathcal{Y}}(y,y)} \leq \sqrt{K}$$

for any $f \in \mathcal{H}_{\mathcal{Y}}(1)$. Accordingly, $\mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}$ also consists of \sqrt{K} -uniformly bounded functions. Hence, Lemma 4 implies that

$$\begin{aligned}
\sup_{F \in \mathcal{F}} \text{MMD}_{K_{\mathcal{Y}}}(F_{\#}\widehat{\mu}_m, F_{\#}\mu) &= \sup_{f \in \mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}} \left| \int f \, d\widehat{\mu}_m - \int f \, d\mu \right| \\
&\leq 2R_m(\mathcal{H}_{\mathcal{Y}}(1) \circ \mathcal{F}, \mu) + \sqrt{\frac{2K \log(1/\delta)}{m}}
\end{aligned}$$

holds with probability at least $1 - \delta$. Therefore, combining this with the upper bound on $\text{MMD}_{K_Y}(\widehat{\nu}_n, \nu)$ derived in Proposition 5,

$$\sup_{F \in \mathcal{F}} |\text{MMD}_{K_Y}^2(F \# \widehat{\mu}_m, \widehat{\nu}_n) - \text{MMD}_{K_Y}^2(F \# \mu, \nu)| \lesssim R_m(\mathcal{H}_Y(1) \circ \mathcal{F}, \mu) + \sqrt{\frac{\log(1/\delta)}{m}} + \sqrt{\frac{\log(1/\delta)}{n}} \quad (\text{D.3})$$

holds with probability at least $1 - 2\delta$. Similarly, we can prove that

$$\sup_{B \in \mathcal{B}} |\text{MMD}_{K_X}^2(\widehat{\mu}_m, B \# \widehat{\nu}_n) - \text{MMD}_{K_X}^2(\mu, B \# \nu)| \lesssim R_n(\mathcal{H}_X(1) \circ \mathcal{B}, \nu) + \sqrt{\frac{\log(1/\delta)}{m}} + \sqrt{\frac{\log(1/\delta)}{n}} \quad (\text{D.4})$$

holds with probability at least $1 - 2\delta$.

Lastly, since $K_X \otimes K_Y$ is bounded by K^2 , that is,

$$\sup_{(x,y),(x',y') \in \mathcal{X} \times \mathcal{Y}} K_X \otimes K_Y((x,y), (x',y')) \leq K^2,$$

by the same argument, we have

$$\begin{aligned} & \sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} |\text{MMD}_{K_X \otimes K_Y}^2((\text{Id}, F) \# \widehat{\mu}_m, (B, \text{Id}) \# \widehat{\nu}_n) - \text{MMD}_{K_X \otimes K_Y}^2((\text{Id}, F) \# \mu, (B, \text{Id}) \# \nu)| \\ & \leq 4K \left[\sup_{F \in \mathcal{F}} \text{MMD}_{K_X \otimes K_Y}((\text{Id}, F) \# \widehat{\mu}_m, (\text{Id}, F) \# \mu) + \sup_{B \in \mathcal{B}} \text{MMD}_{K_X \otimes K_Y}((B, \text{Id}) \# \widehat{\nu}_n, (B, \text{Id}) \# \nu) \right]. \end{aligned}$$

Analogously, $\mathcal{H}_{\mathcal{X} \times \mathcal{Y}}(1)$ consists of K -uniformly bounded functions, hence

$$\begin{aligned} \sup_{F \in \mathcal{F}} \text{MMD}_{K_X \otimes K_Y}((\text{Id}, F) \# \widehat{\mu}_m, (\text{Id}, F) \# \mu) &= \sup_{f \in \mathcal{H}_{\mathcal{X} \times \mathcal{Y}}(1) \circ (\text{Id}, \mathcal{F})} \left| \int f \, d\widehat{\mu}_m - \int f \, d\mu \right| \\ &\leq 2R_m(\mathcal{H}_{\mathcal{X} \times \mathcal{Y}}(1) \circ (\text{Id}, \mathcal{F}), \mu) + \sqrt{\frac{2K^2 \log(1/\delta)}{m}}, \\ \sup_{B \in \mathcal{B}} \text{MMD}_{K_X \otimes K_Y}((B, \text{Id}) \# \widehat{\nu}_n, (B, \text{Id}) \# \nu) &= \sup_{f \in \mathcal{H}_{\mathcal{X} \times \mathcal{Y}}(1) \circ (\mathcal{B}, \text{Id})} \left| \int f \, d\widehat{\nu}_n - \int f \, d\nu \right| \\ &\leq 2R_n(\mathcal{H}_{\mathcal{X} \times \mathcal{Y}}(1) \circ (\mathcal{B}, \text{Id}), \nu) + \sqrt{\frac{2K^2 \log(1/\delta)}{n}}, \end{aligned}$$

each of which holds with probability at least $1 - \delta$. Therefore,

$$\begin{aligned} & \sup_{(F,B) \in \mathcal{F} \times \mathcal{B}} |\text{MMD}_{K_X \otimes K_Y}^2((\text{Id}, F) \# \widehat{\mu}_m, (B, \text{Id}) \# \widehat{\nu}_n) - \text{MMD}_{K_X \otimes K_Y}^2((\text{Id}, F) \# \mu, (B, \text{Id}) \# \nu)| \\ & \lesssim R_m(\mathcal{H}_{\mathcal{X} \times \mathcal{Y}}(1) \circ (\text{Id}, \mathcal{F}), \mu) + R_n(\mathcal{H}_{\mathcal{X} \times \mathcal{Y}}(1) \circ (\mathcal{B}, \text{Id}), \nu) + \sqrt{\frac{\log(1/\delta)}{m}} + \sqrt{\frac{\log(1/\delta)}{n}} \end{aligned} \quad (\text{D.5})$$

holds with probability at least $1 - 2\delta$. We complete the proof by combining (D.3), (D.4), (D.5). \square

Proof of Lemma 3. For any integer $j \in \mathbb{N} \cup \{0\}$, define $\delta_j = 2^{-j}\Delta$, and let $\mathcal{T}_j \subset \mathbb{S}_0^{m \times m}$ be a minimal δ_j -covering net of \mathcal{T} ; clearly, $|\mathcal{T}_j| = N(\delta_j, \mathcal{T})$. For each j , the covering set induces a mapping $\Pi_j : \mathcal{T} \rightarrow \mathcal{T}_j$ such that

$$\sup_{A \in \mathcal{T}} d(A, \Pi_j(A)) \leq \delta_j.$$

By definition of Δ , we may assume $\mathcal{T}_0 = \{A_0\}$ so that $\Pi_0(A) = A_0$ for all $A \in \mathcal{T}$.

Note that $\mathbb{E} g^\top A_0 g = 0$ by definition. Using this, we write $\mathbb{E} \sup_{A \in \mathcal{T}} g^\top A g$ as a chaining sum:

$$\begin{aligned}
\mathbb{E} \sup_{A \in \mathcal{T}} g^\top A g &= \mathbb{E} \sup_g \sup_{A \in \mathcal{T}} g^\top A g - \mathbb{E} g^\top A_0 g \\
&= \mathbb{E} \sup_{A \in \mathcal{T}} (g^\top A g - g^\top A_0 g) \\
&= \mathbb{E} \sup_{A \in \mathcal{T}} \left(g^\top A g - g^\top \Pi_J(A) g + \sum_{j=0}^{J-1} g^\top \Pi_{j+1}(A) g - g^\top \Pi_j(A) g \right) \\
&\leq \mathbb{E} \sup_{A \in \mathcal{T}} (g^\top A g - g^\top \Pi_J(A) g) + \sum_{j=0}^{J-1} \mathbb{E} \sup_{A \in \mathcal{T}} (g^\top \Pi_{j+1}(A) g - g^\top \Pi_j(A) g) .
\end{aligned}$$

For the first term on RHS, using the Cauchy-Schwarz inequality and Jensen's inequality, we have

$$\mathbb{E} \sup_{A \in \mathcal{T}} (g^\top A g - g^\top \Pi_J(A) g) \leq \mathbb{E} \left[\left(\sum_{i \neq j} g_i^2 g_j^2 \right)^{1/2} \cdot \delta_J \right] \leq m \delta_J .$$

For each summand in the second term on RHS, use Lemma 7. Note that for any j , the maximal cardinality of

$$|\{(\Pi_{j+1}(A), \Pi_j(A)) : A \in \mathcal{T}\}| \leq N(\delta_{j+1}, \mathcal{T}) \times N(\delta_j, \mathcal{T}) \leq N(\delta_{j+1}, \mathcal{T})^2$$

and that

$$d(\Pi_{j+1}(A), \Pi_j(A)) \leq d(\Pi_{j+1}(A), A) + d(A, \Pi_j(A)) \leq 3\delta_{j+1} .$$

Since $\Pi_{j+1}(A) - \Pi_j(A) \in \mathbb{S}_0^{m \times m}$ and $\|\Pi_{j+1}(A) - \Pi_j(A)\| \leq 3\delta_{j+1}$, Lemma 7 asserts that for any j

$$\mathbb{E} \sup_{A \in \mathcal{T}} (g^\top \Pi_{j+1}(A) g - g^\top \Pi_j(A) g) \leq 6\delta_{j+1} \sqrt{2 \log N(\delta_{j+1}, \mathcal{T})} + 12\delta_{j+1} \log N(\delta_{j+1}, \mathcal{T}) .$$

Summing over j , we have for any J , the following inequality,

$$\mathbb{E} \sup_{A \in \mathcal{T}} (g^\top A g - g^\top A_0 g) \leq m \delta_J + 12 \int_{\delta_J/2}^{\Delta/2} \sqrt{2 \log N(\delta, \mathcal{T})} d\delta + 24 \int_{\delta_J/2}^{\Delta/2} \log N(\delta, \mathcal{T}) d\delta .$$

□

Proof of Proposition 8. To make use of the chaining inequality, we are only left to bound the covering number $N(\delta, \mathcal{T})$ with

$$\mathcal{T} := \{A_F : F \in \mathcal{F}\} \subset \mathbb{S}_0^{m \times m} .$$

Lipschitzness of $K_{\mathcal{Y}}$ (Assumption 5) implies

$$|K_{\mathcal{Y}}(F(x_i), F(x_j)) - K_{\mathcal{Y}}(F'(x_i), F'(x_j))| \leq \frac{L}{2} \|F(x_i) - F'(x_i)\| + \frac{L}{2} \|F(x_j) - F'(x_j)\| .$$

Hence,

$$\begin{aligned}
d(A_F, A_{F'}) &\leq m \max_{i \neq j} |K_{\mathcal{Y}}(F(x_i), F(x_j)) - K_{\mathcal{Y}}(F'(x_i), F'(x_j))| \\
&\leq \frac{mL}{2} \left(\max_{i \in [m]} \|F(x_i) - F'(x_i)\| + \max_{j \in [m]} \|F(x_j) - F'(x_j)\| \right) \\
&= mL \max_{i \in [m]} \|F(x_i) - F'(x_i)\| \\
&\leq mL \max_{i \in [m]} \left(\sum_{k=1}^{\dim(\mathcal{Y})} |F_k(x_i) - F'_k(x_i)|^2 \right)^{1/2} \\
&\leq mL \left[\sum_{k=1}^{\dim(\mathcal{Y})} \left(\max_{i \in [m]} |F_k(x_i) - F'_k(x_i)| \right)^2 \right]^{1/2}.
\end{aligned}$$

As in the proof of Proposition 6, for $\epsilon > 0$, let $\mathcal{N}_\infty(\epsilon, \mathcal{F}_k, \{x_i\}_{i=1}^m)$ be a minimal ϵ -covering net of \mathcal{F}_k . Then, one can easily see that

$$\left\{ A_F : F \in \otimes_{k=1}^{\dim(\mathcal{Y})} \mathcal{N}_\infty \left(\frac{\delta}{mL\sqrt{\dim(\mathcal{Y})}}, \mathcal{F}_k, \{x_i\}_{i=1}^m \right) \right\}$$

is a δ -covering of \mathcal{T} . Therefore, we conclude

$$N(\delta, \mathcal{T}) \leq \prod_{k=1}^{\dim(\mathcal{Y})} N_\infty \left(\frac{\delta}{mL\sqrt{\dim(\mathcal{Y})}}, \mathcal{F}_k, \{x_i\}_{i=1}^m \right) \leq \prod_{k=1}^{\dim(\mathcal{Y})} N_\infty \left(\frac{\delta}{mL\sqrt{\dim(\mathcal{Y})}}, \mathcal{F}_k, m \right).$$

Lastly, we bound $N_\infty(\delta, \mathcal{F}_k, m)$ by the pseudo-dimension of \mathcal{F}_k via Lemma 5. Combining Assumption 3 and Lemma 5, we have

$$N(\delta, \mathcal{T}) \leq \prod_{k=1}^{\dim(\mathcal{Y})} \left(\frac{2emb}{\frac{\delta}{mL\sqrt{\dim(\mathcal{Y})}} \cdot \text{Pdim}(\mathcal{F}_k)} \right)^{\text{Pdim}(\mathcal{F}_k)}.$$

Now, to apply Lemma 3, fix $F_0 \in \mathcal{F}$ and let $A_0 = A_{F_0}$. Then,

$$\int_{\delta_{J/2}}^{\Delta/2} \log N(\delta, \mathcal{T}) \, d\delta \leq \Delta/2 \cdot \sum_{k=1}^{\dim(\mathcal{Y})} \text{Pdim}(\mathcal{F}_k) \cdot \log \left(\frac{2emb}{\frac{1}{mL\sqrt{\dim(\mathcal{Y})}} \Delta 2^{-J/2} \cdot \text{Pdim}(\mathcal{F}_k)} \right).$$

First, we obtain an upper bound on Δ :

$$\Delta = \sup_{F \in \mathcal{F}} \|A_F - A_0\| \leq m \max_{i \neq j} |K_{\mathcal{Y}}(F(x_i), F(x_j)) - K_{\mathcal{Y}}(F_0(x_i), F_0(x_j))| \leq 2mK.$$

Next, we claim that Δ is bounded below by a universal constant; this is to upper bound Δ in the denominator. Consider y_0 and y'_0 given in Assumption 6. We may assume that F_0 is the constant map explained in Assumption 6: $F_0(x) = y_0$ for all $x \in \mathcal{X}$. Without loss of generality, we assume $x_1 \neq x_2$. Then, we can find $F \in \mathcal{F}$ such that $F(x_1) = y_0$ and $F(x_2) = y'_0$ according to Assumption 6. Hence,

$$\Delta \geq |K_{\mathcal{Y}}(F(x_1), F(x_2)) - K_{\mathcal{Y}}(F_0(x_1), F_0(x_2))| = |K_{\mathcal{Y}}(y_0, y'_0) - K_{\mathcal{Y}}(y_0, y_0)| > 0.$$

Therefore, with the choice of J such that $m2^{-J} \asymp \sum_k \text{Pdim}(\mathcal{F}_k)$,

$$\begin{aligned}
\int_{\delta_{J/2}}^{\Delta/2} \log N(\delta, \mathcal{T}) \, d\delta &\lesssim m \left[\sum_{k=1}^{\dim(\mathcal{Y})} \text{Pdim}(\mathcal{F}_k) \right] \cdot \log \left(\frac{m}{\min_{k \in [\dim(\mathcal{Y})]} \text{Pdim}(\mathcal{F}_k)} \right) \\
&\leq m \log(m) \left[\sum_{k=1}^{\dim(\mathcal{Y})} \text{Pdim}(\mathcal{F}_k) \right].
\end{aligned}$$

Analogously,

$$\int_{\delta_J/2}^{\Delta/2} \sqrt{2 \log N(\delta, \mathcal{T})} d\delta \lesssim m \log(m) \left[\sum_{k=1}^{\dim(\mathcal{Y})} \text{Pdim}(\mathcal{F}_k) \right].$$

Thus,

$$R_m(\mathcal{H}_y(1) \circ \mathcal{F}, \{x_i\}_{i=1}^m) \lesssim \frac{1}{m} \left[mK + \mathbb{E} \sup_{g \in \mathcal{F}} g^\top A_F g \right]^{1/2} \lesssim \sqrt{\frac{\log m}{m} \sum_{k=1}^{\dim(\mathcal{Y})} \text{Pdim}(\mathcal{F}_k)}$$

The same argument can be applied to the other three Rademacher complexities. Hence, we have proved the proposition. \square

D.2 Auxiliary Lemmas

Lemma 4 (Theorem 4.10 of [55]). *Let (\mathcal{Z}, ρ) be a probability space and \mathcal{G} be a class of b -uniformly bounded measurable functions defined on \mathcal{Z} , that is, $\sup_{g \in \mathcal{G}} \|g\|_\infty \leq b$. Let z_1, \dots, z_m are i.i.d. samples from ρ and let $\hat{\rho}_m$ be the empirical measure constructed from them. Then, for any $\delta > 0$,*

$$\sup_{g \in \mathcal{G}} \left| \int g d\hat{\rho}_m - \int g d\rho \right| \leq 2R_m(\mathcal{G}, \rho) + \sqrt{\frac{2b^2 \log(1/\delta)}{m}}$$

holds with probability at least $1 - \delta$.

Lemma 5 (Theorem 12.2 of [1]). *Let \mathcal{G} be a collection of real-valued functions defined on a set \mathcal{Z} . Suppose $\sup_{g \in \mathcal{G}} \|g\|_\infty = b < \infty$. For $\epsilon > 0$ and $m \geq \text{Pdim}(\mathcal{G})$,*

$$N_\infty(\epsilon, \mathcal{G}, m) \leq \left(\frac{2emb}{\epsilon \cdot \text{Pdim}(\mathcal{G})} \right)^{\text{Pdim}(\mathcal{G})}.$$

Lemma 6 (Example 2.12 of [7]). *For any $A \in \mathbb{S}_0^{m \times m}$ and $0 \leq \lambda < 1/(2\|A\|_{\text{op}})$,*

$$\log \mathbb{E}_g e^{\lambda g^\top A g} \leq \frac{\lambda^2 \|A\|}{1 - 2\lambda \|A\|_{\text{op}}}, \quad (\text{D.6})$$

where $g \sim N(0, I_m)$. Here, $\|\cdot\|_{\text{op}}$ denotes the operator norm of A .

This lemma tells that $g^\top A g$ is a sub-Gamma random variable with variance factor $2\|A\|^2$ and scale parameter $2\|A\|_{\text{op}}$ (see Chapter 2.4 of [7] for the definition). Using Corollary 2.6 of the same text, we can derive the following maximal inequality.

Lemma 7 (Maximal inequality). *For $A_1, \dots, A_N \in \mathbb{S}_0^{m \times m}$, suppose $\max_{i=1, \dots, N} \|A_i\| \leq \delta$. Then,*

$$\mathbb{E}_g \max_{i=1, \dots, N} g^\top A_i g \leq 2\delta \left(\sqrt{\log N} + \log N \right), \quad (\text{D.7})$$

where $g \sim N(0, I_m)$.

E Computational Aspects of the RGM Distance

Recall from Section 3.2 that we have utilized the Lagrangian form instead of the constrained form (3.1) of the RGM distance. Such a practical computation allows us to easily implement the RGM sampler and we have observed its good empirical performance in Section 4. This section shifts our focus to the exact

computation of the RGM distance; we discuss conditions under which minimizing the Lagrangian form leads to a close approximation to the RGM distance. Another important computational aspect is the comparison of the RGM distance and the GW distance. $\text{GW}(\mu, \nu) \leq \text{RGM}(\mu, \nu)$ holds in theory by Proposition 1; using a concrete example, we approximate both quantities numerically and see how large the gap between them is. Lastly, we examine the numerical performance of the convex formulation discussed in 3.4 and compare with the results from the Lagrangian form.

E.1 Approximation with the Lagrangian Form

First, we derive connections between the RGM distance and its Lagrangian formulation. As in Section B, let

$$C_0(F, B) = \int (c_{\mathcal{X}}(x, B(y)) - c_{\mathcal{Y}}(F(x), y))^2 d\mu \otimes \nu$$

for any $F: \mathcal{X} \rightarrow \mathcal{Y}$ and $B: \mathcal{Y} \rightarrow \mathcal{X}$ so that $\text{RGM}(\mu, \nu)^2 = \inf_{(F, B) \in \mathcal{I}(\mu, \nu)} C_0(F, B)$. Define the Lagrangian form as

$$L_{\lambda_1, \lambda_2, \lambda_3}(F, B) = C_0(F, B) + \lambda_1 \cdot \underbrace{\mathcal{L}_{\mathcal{X} \times \mathcal{Y}}((\text{Id}, F)_{\#}\mu, (B, \text{Id})_{\#}\nu)}_{\ell_1(F, B)} + \lambda_2 \cdot \underbrace{\mathcal{L}_{\mathcal{X}}(\mu, B_{\#}\nu)}_{\ell_2(F, B)} + \lambda_3 \cdot \underbrace{\mathcal{L}_{\mathcal{Y}}(F_{\#}\mu, \nu)}_{\ell_3(F, B)},$$

where $\mathcal{L}_{\mathcal{X} \times \mathcal{Y}}$, $\mathcal{L}_{\mathcal{X}}$, and $\mathcal{L}_{\mathcal{Y}}$ are suitable nonnegative discrepancy measures as in Section 3.2; in particular, we assume $\ell_1(F, B) = \ell_2(F, B) = \ell_3(F, B) = 0$ for $(F, B) \in \mathcal{I}(\mu, \nu)$.¹¹

Suppose $\lambda_1, \lambda_2, \lambda_3 \geq 0$, then

$$\inf_{\substack{F: \mathcal{X} \rightarrow \mathcal{Y} \\ B: \mathcal{Y} \rightarrow \mathcal{X}}} L_{\lambda_1, \lambda_2, \lambda_3}(F, B) \leq \inf_{(F, B) \in \mathcal{I}(\mu, \nu)} L_{\lambda_1, \lambda_2, \lambda_3}(F, B) = \text{RGM}(\mu, \nu)^2 = \inf_{(F, B) \in \mathcal{I}(\mu, \nu)} C_0(F, B).$$

We seek a minimizer of $L_{\lambda_1, \lambda_2, \lambda_3}$ over $\mathcal{F} \times \mathcal{B}$, namely, the product of suitable function classes as discussed in Section 3.3. Let

$$(F^*, B^*) \in \arg \min_{(F, B) \in \mathcal{F} \times \mathcal{B}} L_{\lambda_1, \lambda_2, \lambda_3}(F, B). \quad (\text{E.1})$$

If $\mathcal{I}(\mu, \nu) \subseteq \mathcal{F} \times \mathcal{B}$,

$$\begin{aligned} C_0(F^*, B^*) &\leq L_{\lambda_1, \lambda_2, \lambda_3}(F^*, B^*) \quad (\because \lambda_1, \lambda_2, \lambda_3 \geq 0) \\ &= \inf_{(F, B) \in \mathcal{F} \times \mathcal{B}} L_{\lambda_1, \lambda_2, \lambda_3}(F, B) \quad (\because (\text{E.1})) \\ &\leq \inf_{(F, B) \in \mathcal{I}(\mu, \nu)} L_{\lambda_1, \lambda_2, \lambda_3}(F, B) \quad (\because \mathcal{I}(\mu, \nu) \subseteq \mathcal{F} \times \mathcal{B}) \\ &= \text{RGM}(\mu, \nu)^2. \end{aligned}$$

Roughly speaking, if the function classes are rich enough to ensure $\mathcal{I}(\mu, \nu) \subseteq \mathcal{F} \times \mathcal{B}$, the minimizer (F^*, B^*) produces a lower bound $C_0(F^*, B^*)$ on $\text{RGM}(\mu, \nu)^2$. On the other hand, if the minimizer satisfies the constraint $(F^*, B^*) \in \mathcal{I}(\mu, \nu)$, then $C_0(F^*, B^*)$ is an upper bound on $\text{RGM}(\mu, \nu)^2$ by definition.

Therefore, a sufficient condition for $C_0(F^*, B^*) = \text{RGM}(\mu, \nu)^2$ is that the following two hold: $\mathcal{I}(\mu, \nu) \subseteq \mathcal{F} \times \mathcal{B}$ and $(F^*, B^*) \in \mathcal{I}(\mu, \nu)$.

E.2 Numerical Experiments

Using a concrete example, we compute the aforementioned quantities related to the RGM distance, approximate the GW distance, and compare them; we will also discuss the results from the convex formulation in Section 3.4. Throughout, we consider two point clouds on \mathbb{R}^2 as in Figure 6(a), that is, μ and ν are uniform distributions supported on 30 grid points of a segment and a circle, respectively. We fix the cost functions: $c_{\mathcal{X}} = c_{\mathcal{Y}}$ is the RBF kernel that maps (x, y) to $\exp(-\|x - y\|^2)$.

¹¹Though we may define the Lagrangian form without $\mathcal{L}_{\mathcal{X}}$ and $\mathcal{L}_{\mathcal{Y}}$, we include them for a seamless connection with the experiment results in Section E.2.




Figure 6: The supports of μ and ν are the grid points of a segment connecting $(-1, -1)$ and $(1, 1)$ and a circle $x^2 + y^2 = 1$, respectively.

First, we aim to compute the quantities discussed in Section E.1. To this end, we specify the discrepancy measures and the function classes as follows.

- $\mathcal{L}_{\mathcal{X} \times \mathcal{Y}} = \text{MMD}_{K_{\mathcal{X} \otimes K_{\mathcal{Y}}}}^2$, $\mathcal{L}_{\mathcal{X}} = \text{MMD}_{K_{\mathcal{X}}}^2$, $\mathcal{L}_{\mathcal{Y}} = \text{MMD}_{K_{\mathcal{Y}}}^2$, where $K_{\mathcal{X}} = K_{\mathcal{Y}} = c_{\mathcal{X}}$.
- $\mathcal{F} = \mathcal{B}$ is the class of neural networks with two hidden layers as follows:

$$\{x \mapsto \tanh(W_2 \tanh(W_1 x + b_1)) + b_2\} : W_1 \in \mathbb{R}^{2 \times 30}, b_1 \in \mathbb{R}^{30}, W_2 \in \mathbb{R}^{30 \times 2}, b_2 \in \mathbb{R}^2\},$$

where \tanh is the tangent hyperbolic function applied elementwise, that is, $\tanh(x) = (\tanh(x_1), \dots, \tanh(x_k)) \in \mathbb{R}^k$ for $x = (x_1, \dots, x_k) \in \mathbb{R}^k$.

We use Adam [27], a variant of stochastic gradient descent, to find a minimizer of $L_{\lambda_1, \lambda_2, \lambda_3}$ over $\mathcal{F} \times \mathcal{B}$ as in (E.1). After 10000 iterations, we can see that the loss converges as in Figure 7, indicating we have a local minimizer $(\hat{F}^*, \hat{B}^*) \in \mathcal{F} \times \mathcal{B}$. Since this optimization problem may be nonconvex, there is no guarantee that this is a global minimizer, hence

$$C_0(F^*, B^*) \leq C_0(\hat{F}^*, \hat{B}^*).$$

To examine the experiment results in light of Section E.1, we will assume

- \mathcal{F} and \mathcal{B} are rich enough to ensure $\mathcal{I}(\mu, \nu) \subset \mathcal{F} \times \mathcal{B}$,
- $(F, B) \in \mathcal{I}(\mu, \nu)$ if and only if $\ell_1(F, B) = \ell_2(F, B) = \ell_3(F, B) = 0$,¹²
- (\hat{F}^*, \hat{B}^*) is indeed a global minimizer: $(\hat{F}^*, \hat{B}^*) \in \arg \min_{(F, B) \in \mathcal{F} \times \mathcal{B}} L_{\lambda_1, \lambda_2, \lambda_3}(F, B)$.

Under these assumptions, $(\hat{F}^*, \hat{B}^*) \in \mathcal{I}(\mu, \nu)$ implies $C_0(\hat{F}^*, \hat{B}^*) = \text{RGM}(\mu, \nu)^2$ as discussed in Section E.1. To verify $(\hat{F}^*, \hat{B}^*) \in \mathcal{I}(\mu, \nu)$, we check the values of $\ell_1(F, B)$, $\ell_2(F, B)$, and $\ell_3(F, B)$ in Table 1. We observe that they get smaller as we increase the values of the Lagrangian multipliers. For the cases where $\lambda_1 = \lambda_2 = \lambda_3 = 10^2$ or $\lambda_1 = \lambda_2 = \lambda_3 = 10^4$, the values ℓ_1, ℓ_2, ℓ_3 are sufficiently small to conclude $(\hat{F}^*, \hat{B}^*) \in \mathcal{I}(\mu, \nu)$, hence we can roughly estimate $\text{RGM}(\mu, \nu)^2 \approx C_0(\hat{F}^*, \hat{B}^*) \approx 0.39$.

¹²This is true since we are using the RBF kernel as mentioned in Section 4.





Figure 7: Training curves (10000 iterations).

Comparison with GW As discussed earlier, exact computation of the GW distance (Definition 2) is impossible in general. Here, we estimate it using an off-the-shelf computational tool called Python Optimal Transport (POT) [21] widely used in literature, which yields $\text{GW}^2(\mu, \nu) \approx 0.171$.¹³ Combined with the previous computation, we can say

$$1 \leq \frac{\text{RGM}(\mu, \nu)}{\text{GW}(\mu, \nu)} \approx \sqrt{\frac{0.390}{0.171}} = 1.515,$$

indicating that the RGM distance is approximately the GW distance times 1.5. This rough computation is based on the aforementioned assumptions regarding the RGM computation and the accuracy of POT in computing the GW distance.

Instead, we may give an upper bound on the ratio of the two distances using the well-known lower bounds on the GW distance: the First Lower Bound (FLB) and the Second Lower Bound (SLB) on GW

¹³Technically, this should be an upper bound on the exact value of $\text{GW}^2(\mu, \nu)$ because the result of POT ought to be a local minimizer.

	$\lambda_1 = \lambda_2 = \lambda_3 = 1$	$\lambda_1 = \lambda_2 = \lambda_3 = 10^2$	$\lambda_1 = \lambda_2 = \lambda_3 = 10^4$
$C_0(\widehat{F}^*, \widehat{B}^*)$	0.136	0.386	0.390
$\ell_1(\widehat{F}^*, \widehat{B}^*)$	3.366×10^{-2}	1.034×10^{-3}	9.550×10^{-4}
$\ell_2(\widehat{F}^*, \widehat{B}^*)$	1.716×10^{-3}	5.758×10^{-5}	7.379×10^{-5}
$\ell_3(\widehat{F}^*, \widehat{B}^*)$	3.327×10^{-2}	2.689×10^{-4}	2.698×10^{-4}
$L_{\lambda_1, \lambda_2, \lambda_3}(\widehat{F}^*, \widehat{B}^*)$	0.205	0.522	13.377

Table 1: Minimum values of the Lagrangian form.

[38]. Letting $\mu = \frac{1}{m} \sum_{i=1}^m \delta_{x_i}$ and $\nu = \frac{1}{n} \sum_{j=1}^n \delta_{y_j}$ with $m = n = 30$, these bounds are computed as follows:

$$\begin{aligned} \text{FLB}^2(\mu, \nu) &= W_2^2 \left(\frac{1}{m} \sum_{i=1}^m \delta_{e_{\mathcal{X}}(x_i)}, \frac{1}{n} \sum_{j=1}^n \delta_{e_{\mathcal{Y}}(y_j)} \right) \approx 0.061, \\ \text{SLB}^2(\mu, \nu) &= W_2^2 \left(\frac{1}{m^2} \sum_{i, i'=1}^m \delta_{c_{\mathcal{X}}(x_i, x_{i'})}, \frac{1}{n^2} \sum_{j, j'=1}^n \delta_{c_{\mathcal{Y}}(y_j, y_{j'})} \right) \approx 0.135, \end{aligned}$$

where $e_{\mathcal{X}}(x_i) = \sqrt{\frac{1}{m} \sum_{i'=1}^m c_{\mathcal{X}}^2(x_i, x_{i'})}$ and $e_{\mathcal{Y}}(y_j) = \sqrt{\frac{1}{n} \sum_{j'=1}^n c_{\mathcal{Y}}^2(y_j, y_{j'})}$ are called the eccentricity; see [38] for details. These quantities, computed by using POT as well, are known to be lower bounds on $\text{GW}(\mu, \nu)^2$, hence

$$1 \leq \frac{\text{RGM}(\mu, \nu)}{\text{GW}(\mu, \nu)} \leq \sqrt{\frac{0.390}{0.135}} = 1.700.$$

Therefore, we can conclude that the ratio of the two distances is bounded by 1.7.

Convex formulation Next, we estimate the RGM distance based on the convex formulation; as in Theorem 5, we solve the convex optimization problem:

$$\min_{\substack{\mathbf{F}_{m,n} \in \mathbb{R}^{m \times n} \\ \mathbf{B}_{n,m} \in \mathbb{R}^{n \times m}}} \omega(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}), \quad (\text{E.2})$$

where

$$\begin{aligned} \omega(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}) &= \underbrace{\frac{1}{mn} \|\mathbf{K}_{\mathcal{Y}} \mathbf{B}_{n,m} \mathbf{K}_{\mathcal{X}} - \mathbf{K}_{\mathcal{Y}} \mathbf{F}_{m,n}^{\top} \mathbf{K}_{\mathcal{X}}\|^2}_{c_0(\mathbf{F}_{m,n}, \mathbf{B}_{n,m})} + \lambda_1 \cdot \underbrace{\left\| \frac{1}{m} \mathbf{K}_{\mathcal{X}}^{3/2} \mathbf{F}_{m,n} \mathbf{K}_{\mathcal{Y}}^{1/2} - \frac{1}{n} \mathbf{K}_{\mathcal{X}}^{1/2} \mathbf{B}_{n,m}^{\top} \mathbf{K}_{\mathcal{Y}}^{3/2} \right\|^2}_{m_1(\mathbf{F}_{m,n}, \mathbf{B}_{n,m})} \\ &\quad + \lambda_2 \cdot \underbrace{\left\| \mathbf{K}_{\mathcal{X}}^{1/2} \cdot \left(\frac{1}{m} \mathbf{1}_m - \mathbf{B}_{n,m}^{\top} \mathbf{K}_{\mathcal{Y}} \frac{1}{n} \mathbf{1}_n \right) \right\|^2}_{m_2(\mathbf{F}_{m,n}, \mathbf{B}_{n,m})} + \lambda_3 \cdot \underbrace{\left\| \mathbf{K}_{\mathcal{Y}}^{1/2} \cdot \left(\frac{1}{n} \mathbf{1}_n - \mathbf{F}_{m,n}^{\top} \mathbf{K}_{\mathcal{X}} \frac{1}{m} \mathbf{1}_m \right) \right\|^2}_{m_3(\mathbf{F}_{m,n}, \mathbf{B}_{n,m})} \end{aligned}$$

as derived in Section C. It should be noted that the minimum of (E.2) is a lower bound on the minimum of the Lagrangian, that is,

$$\min_{\substack{\mathbf{F}_{m,n} \in \mathbb{R}^{m \times n} \\ \mathbf{B}_{n,m} \in \mathbb{R}^{n \times m}}} \omega(\mathbf{F}_{m,n}, \mathbf{B}_{n,m}) \leq \min_{(F,B) \in \mathcal{F} \times \mathcal{B}} L_{\lambda_1, \lambda_2, \lambda_3}(F, B).$$

To see this, first note that the RHS is exactly (3.5).¹⁴ Then, recall from Section 3.4 that (3.5) = (3.10) is relaxed to (3.11) and is further relaxed to the convex problem (E.2) due to Theorem 5.

¹⁴Since μ and ν are discrete, the RHS is the same as its empirical estimate (3.5).

	$\lambda_1 = \lambda_2 = \lambda_3 = 1$	$\lambda_1 = \lambda_2 = \lambda_3 = 10^2$	$\lambda_1 = \lambda_2 = \lambda_3 = 10^4$
$c_0(\mathbf{F}_{m,n}^*, \mathbf{B}_{n,m}^*)$	0.044	0.106	0.108
$m_1(\mathbf{F}_{m,n}^*, \mathbf{B}_{n,m}^*)$	0.011	4.521×10^{-6}	6.369×10^{-10}
$m_2(\mathbf{F}_{m,n}^*, \mathbf{B}_{n,m}^*)$	0.012	2.592×10^{-6}	2.626×10^{-12}
$m_3(\mathbf{F}_{m,n}^*, \mathbf{B}_{n,m}^*)$	0.001	1.050×10^{-7}	2.461×10^{-10}
$\omega(\mathbf{F}_{m,n}^*, \mathbf{B}_{n,m}^*)$	0.068	0.107	0.108

Table 2: Minimum values of the convex problem (E.2) obtained by CVXPY.

Table 2 shows the results obtained by a convex optimization tool called CVXPY [16]. The minimum $\omega(\mathbf{F}_{m,n}^*, \mathbf{B}_{n,m}^*)$ of (E.2) is always smaller than $\text{GW}^2(\mu, \nu) \approx 0.171$ and is between the two lower bounds: $\text{FLB}^2(\mu, \nu) \approx 0.061$ and $\text{SLB}^2(\mu, \nu) \approx 0.135$. Also, the MMD terms vanish if we use the large Lagrangian multipliers, indicating that the constraints (represented via the MMD terms) are met. That said, we can see that the gap between the minimum and the RGM distance can be large. Therefore, finding conditions under which this gap vanishes would be interesting future work.

F Details of the Experiments in Section 4

Here, we provide implementation details of the experiments in Section 4.

F.1 Gaussian

In the Gaussian experiment in Section 4, we minimize (3.5) using Adam for 3000 iterations. The learning rate at the initial iteration is 0.1 and we halve it after every 500 iterations. Figure 8(a) shows the training curve.




Figure 8: Training curves for the experiments: (a) the Gaussian experiment, (b) the MNIST experiment with $\mathcal{X} = \mathbb{R}^2$.

F.2 MNIST

\mathbb{R}^2 and MMDs For numerical stability, we encode a variant of empirical estimate (3.5) as our loss:

$$\begin{aligned} & \lambda_1 \cdot \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n (c_{\mathcal{X}}(x_i, B(y_j)) - c_{\mathcal{Y}}(F(x_i), y_j))^2 + \text{MMD}_{K_{\mathcal{X} \otimes \mathcal{K}_{\mathcal{Y}}}}^2((\text{Id}, F)_{\#} \hat{\mu}_m, (B, \text{Id})_{\#} \hat{\nu}_n) \\ & + \lambda_2 \cdot \text{MMD}_{K_{\mathcal{X}}}^2(\hat{\mu}_m, B_{\#} \hat{\nu}_n) + \lambda_3 \cdot \text{MMD}_{K_{\mathcal{Y}}}^2(F_{\#} \hat{\mu}_m, \hat{\nu}_n). \end{aligned} \quad (\text{F.1})$$

We choose tuning parameters $(\lambda_1, \lambda_2, \lambda_3) = (0.01, 1, 1)$. For fully connected neural networks F and B , we apply the rectified linear unit (ReLU) activation function $\sigma(x) = \max(x, 0)$ to all three hidden layers of F and B , and an additional tangent hyperbolic (tanh) function $\tanh(x) = (e^x - e^{-x}) / (e^x + e^{-x})$ to the output layer of F , both of which are elementwise activation functions. To put it explicitly, $y = F(x)$ is defined as

$$\begin{aligned} h_0 &= x, \quad x \in \mathbb{R}^2 \\ h_l &= \sigma(W_l h_{l-1} + b_l), \quad l = 1, 2 \\ y &= \tanh(W_3 h_2 + b_3) \end{aligned}$$

with $W_1 \in \mathbb{R}^{50 \times 2}$, $W_2 \in \mathbb{R}^{50 \times 50}$, $W_3 \in \mathbb{R}^{784 \times 50}$, $b_1, b_2 \in \mathbb{R}^{50 \times 1}$, $b_3 \in \mathbb{R}^{784 \times 1}$. Similarly, $\tilde{x} = B(\tilde{y})$ is defined as

$$\begin{aligned} \tilde{h}_0 &= \tilde{y}, \quad \tilde{y} \in \mathbb{R}^{784} \\ \tilde{h}_l &= \sigma(\tilde{W}_l \tilde{h}_{l-1} + \tilde{b}_l), \quad l = 1, 2 \\ \tilde{x} &= \tilde{W}_3 \tilde{h}_2 + \tilde{b}_3 \end{aligned}$$

with $\tilde{W}_1 \in \mathbb{R}^{50 \times 784}$, $\tilde{W}_2 \in \mathbb{R}^{50 \times 50}$, $\tilde{W}_3 \in \mathbb{R}^{2 \times 50}$, $\tilde{b}_1, \tilde{b}_2 \in \mathbb{R}^{50 \times 1}$, $\tilde{b}_3 \in \mathbb{R}^{2 \times 1}$. The training set is randomly divided into minibatches of size 256, for which we run Adam again for 1000 iterations. The learning rate at the initial iteration is 0.005 and we halve it after every 500 iterations. Figure 9 shows the generated images during the training process.


\mathbb{R}^4 and Sinkhorn divergences We implement the Sinkhorn divergence with squared Euclidean cost by using GeomLoss [20]. Concretely, we first define the entropic regularized Kantorovich problem between $\hat{\mu}_m$ and $B_{\#} \hat{\nu}_n$ on some Euclidean space \mathcal{X}

$$W_{2,\epsilon}^2(\hat{\mu}_m, B_{\#} \hat{\nu}_n) = \min_{\gamma \in \Pi(\hat{\mu}_m, B_{\#} \hat{\nu}_n)} \sum_{i=1}^m \sum_{j=1}^n \gamma_{ij} (\|x_i - B(y_j)\|^2 + \epsilon \log(\gamma_{ij}))$$

where γ is a coupling matrix and γ_{ij} denotes its (i, j) element. Then the Sinkhorn divergence between two empirical measures is defined by $s_{\epsilon, \mathcal{X}}(\hat{\mu}_m, B_{\#} \hat{\nu}_n) = W_{2,\epsilon}^2(\hat{\mu}_m, B_{\#} \hat{\nu}_n) - \frac{1}{2} W_{2,\epsilon}^2(\hat{\mu}_m, \hat{\mu}_m) - \frac{1}{2} W_{2,\epsilon}^2(B_{\#} \hat{\nu}_n, B_{\#} \hat{\nu}_n)$. We encode our loss by replacing MMDs in (F.1) with Sinkhorn divergences


$$\begin{aligned} & \lambda_1 \cdot \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n (c_{\mathcal{X}}(x_i, B(y_j)) - c_{\mathcal{Y}}(F(x_i), y_j))^2 + s_{\epsilon, \mathcal{X} \times \mathcal{Y}}((\text{Id}, F)_{\#} \hat{\mu}_m, (B, \text{Id})_{\#} \hat{\nu}_n) \\ & + \lambda_2 \cdot s_{\epsilon, \mathcal{X}}(\hat{\mu}_m, B_{\#} \hat{\nu}_n) + \lambda_3 \cdot s_{\epsilon, \mathcal{Y}}(F_{\#} \hat{\mu}_m, \hat{\nu}_n), \end{aligned}$$

and choose tuning parameters $(\lambda_1, \lambda_2, \lambda_3) = (1, 1, 1)$. The Sinkhorn parameter ϵ is set to be 0.0001 for all three discrepancy measures. Again, $F: \mathbb{R}^2 \rightarrow \mathbb{R}^{784}$ and $B: \mathbb{R}^{784} \rightarrow \mathbb{R}^2$ are parametrized by fully connected neural networks with three hidden layers, whose activation functions are same as the MMD case. The rest of the setups, including the choice of optimizer, number of iterations, and batchsize, are same as the MMD case above.



(a) Before training

(b) After 20 iterations



(c) After 50 iterations

(d) After 1000 iterations

Figure 9: Generated images on MNIST data for digits 2, 4, 6, 7 during the training process, under the experimental setup with \mathbb{R}^2 and MMD discrepancy.

Comparison between RGM and GW Lastly, let us estimate the gap between the RGM distance and the GW distance in the MNIST example (\mathbb{R}^2 and MMDs) based on the discussions in Section E. Recall that we have obtained a minimizer (\hat{F}, \hat{B}) of (3.5) over $\mathcal{F} \times \mathcal{B}$ using samples $\{x_i\}_{i=1}^{20000}$ and $\{y_j\}_{j=1}^{20000}$ from $\mu = N(0, I_2)$ and $\nu =$ the distribution of the four digits, respectively; though this is a local minimizer as the optimization problem may be nonconvex, we will assume that this is indeed a global minimizer as in Section E. Letting $\hat{\mu}_m$ and $\hat{\nu}_n$ be the empirical measures constructed by $\{x_i\}_{i=1}^{20000}$ and $\{y_j\}_{j=1}^{20000}$, respectively ($m = n = 20000$), we obtain the following quantities:

$$\begin{aligned} \hat{C}_0(\hat{F}, \hat{B}) &\stackrel{\text{Section B}}{=} \int (c_{\mathcal{X}}(x, \hat{B}(y)) - c_{\mathcal{Y}}(\hat{F}(x), y))^2 d\hat{\mu}_m \otimes \hat{\nu}_n \approx 0.348, \\ m_1 &:= \text{MMD}_{K_{\mathcal{X} \otimes \mathcal{K}_{\mathcal{Y}}}}^2((\text{Id}, \hat{F})_{\#} \hat{\mu}_m, (\hat{B}, \text{Id})_{\#} \hat{\nu}_n) \approx 2.192 \times 10^{-3}, \\ m_2 &:= \text{MMD}_{K_{\mathcal{X}}}^2(\hat{\mu}_m, \hat{B}_{\#} \hat{\nu}_n) \approx 8.261 \times 10^{-5}, \\ m_3 &:= \text{MMD}_{K_{\mathcal{Y}}}^2(\hat{F}_{\#} \hat{\mu}_m, \hat{\nu}_n) \approx 1.437 \times 10^{-3}. \end{aligned}$$

Hence, $C(\hat{\mu}_m, \hat{\nu}_n, \hat{F}, \hat{B}) = \hat{C}_0(\hat{F}, \hat{B}) + \sum_{k=1}^3 \lambda_k m_k = 0.719$, where C is defined in (3.6). Recall from

Section B that

$$\begin{aligned} |C(\widehat{\mu}_m, \widehat{\nu}_n, \widehat{F}, \widehat{B}) - C(\mu, \nu, \widehat{F}, \widehat{B})| &\leq \sup_{(F, B) \in \mathcal{F} \times \mathcal{B}} |C(\widehat{\mu}_m, \widehat{\nu}_n, F, B) - C(\mu, \nu, F, B)| \\ &\lesssim \mathcal{M}(\mathcal{F}, \mathcal{B}, m, n, \delta) \end{aligned}$$

holds with probability at least $1 - \delta$, where $\mathcal{M}(\mathcal{F}, \mathcal{B}, m, n, \delta)$ is defined in Theorem 4. Assuming that this complexity measure is sufficiently small for $m = n = 20000$, we may roughly say $C(\mu, \nu, \widehat{F}, \widehat{B}) \approx C(\widehat{\mu}_m, \widehat{\nu}_n, \widehat{F}, \widehat{B}) = 0.719$. In the same vein, Theorem 4 indicates

$$\inf_{(F, B) \in \mathcal{F} \times \mathcal{B}} C(\mu, \nu, F, B) \approx C(\mu, \nu, \widehat{F}, \widehat{B}) \approx 0.719 .$$

Now, we combine this result with the discussion in Section E. First, by definition,

$$\inf_{(F, B) \in \mathcal{F} \times \mathcal{B}} C(\mu, \nu, F, B) = \inf_{(F, B) \in \mathcal{F} \times \mathcal{B}} L_{\lambda_1, \lambda_2, \lambda_3}(F, B) .$$

We have derived in Section E that

$$\text{RGM}(\mu, \nu)^2 = C_0(F^*, B^*) = L_{\lambda_1, \lambda_2, \lambda_3}(F^*, B^*) = \inf_{(F, B) \in \mathcal{F} \times \mathcal{B}} L_{\lambda_1, \lambda_2, \lambda_3}(F, B)$$

if $\mathcal{I}(\mu, \nu) \subseteq \mathcal{F} \times \mathcal{B}$ and the minimizer (F^*, B^*) defined in (E.1) satisfies $(F^*, B^*) \in \mathcal{I}(\mu, \nu)$. Therefore, under these assumptions, we can roughly estimate

$$\text{RGM}^2(\mu, \nu) \approx 0.719 .$$

Lastly, let us estimate the lower bounds, FLB and SLB, on the GW distance as in Section E.2. Due to the computational complexity, we will use subsets of the training data, say $\{x_i\}_{i=1}^{2000}$ and $\{y_j\}_{j=1}^{2000}$, to construct plug-in estimators $\text{FLB}(\widehat{\mu}_{2000}, \widehat{\nu}_{2000})$ and $\text{SLB}(\widehat{\mu}_{2000}, \widehat{\nu}_{2000})$; using POT again, we have

$$\begin{aligned} \text{FLB}^2(\mu, \nu) &\approx \text{FLB}^2(\widehat{\mu}_{2000}, \widehat{\nu}_{2000}) \approx 0.006 , \\ \text{SLB}^2(\mu, \nu) &\approx \text{SLB}^2(\widehat{\mu}_{2000}, \widehat{\nu}_{2000}) \approx 0.148 . \end{aligned}$$

Therefore, we can roughly conclude

$$1 \leq \frac{\text{RGM}(\mu, \nu)}{\text{GW}(\mu, \nu)} \leq \sqrt{\frac{0.719}{0.148}} = 2.204 .$$