Hierarchical Hybrid Sliced Wasserstein: A Scalable Metric for Heterogeneous Joint Distributions

Khai Nguyen

Department of Statistics and Data Sciences
The University of Texas at Austin
Austin, TX 78712
khainb@utexas.edu

Nhat Ho

Department of Statistics and Data Sciences
The University of Texas at Austin
Austin, TX 78712
minhnhat@utexas.edu

Abstract

Sliced Wasserstein (SW) and Generalized Sliced Wasserstein (GSW) have been widely used in applications due to their computational and statistical scalability. However, the SW and the GSW are only defined between distributions supported on a homogeneous domain. This limitation prevents their usage in applications with heterogeneous joint distributions with marginal distributions supported on multiple different domains. Using SW and GSW directly on the joint domains cannot make a meaningful comparison since their homogeneous slicing operator, i.e., Radon Transform (RT) and Generalized Radon Transform (GRT) are not expressive enough to capture the structure of the joint supports set. To address the issue, we propose two new slicing operators, i.e., Partial Generalized Radon Transform (PGRT) and Hierarchical Hybrid Radon Transform (HHRT). In greater detail, PGRT is the generalization of Partial Radon Transform (PRT), which transforms a subset of function arguments non-linearly while HHRT is the composition of PRT and multiple domain-specific PGRT on marginal domain arguments. By using HHRT, we extend the SW into Hierarchical Hybrid Sliced Wasserstein (H2SW) distance which is designed specifically for comparing heterogeneous joint distributions. We then discuss the topological, statistical, and computational properties of H2SW. Finally, we demonstrate the favorable performance of H2SW in 3D mesh deformation, deep 3D mesh autoencoders, and datasets comparison¹.

1 Introduction

Optimal transport [55, 47] is a powerful mathematical tool for machine learning, statistics, and data sciences. As an example, Wasserstein distance [47], defined as the optimal transportation cost between two distributions, has been used successfully in many areas of machine learning and statistics, such as generative modeling on images [2, 53], representation learning [35], vocabulary learning [57], and so on. Despite being accepted as an effective distance, Wasserstein distance has been widely known as a computationally expensive distance. In particular, when comparing two distributions that have at most n supports, the time complexity and the memory complexity of the Wasserstein distance scale with the order of $\mathcal{O}(n^3 \log n)$ [45] and $\mathcal{O}(n^2)$ respectively. In addition, the Wasserstein distance requires more samples to approximate a continuous distribution with its empirical distribution in high dimension since its sample complexity is of the order of $\mathcal{O}(n^{-1/d})$ [20], where n is the sample size and d is the number of dimensions. Therefore, Wasserstein distance is not statistically and computationally scalable, especially in high dimensions.

Along with entropic regularization [17] which can reduce the time complexity and memory complexity of computing optimal transport to $\mathcal{O}(n^2)$ and $\mathcal{O}(n^2)$ in turn, sliced Wasserstein (SW) distance [11]

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

¹Code for this paper is published at https://github.com/khainb/H2SW.

is one alternative approach for the original Wasserstein distance. The key benefit of the SW distance is that it scales the order $\mathcal{O}(n\log n)$ and $\mathcal{O}(n)$ in terms of time and memory respectively. The reason behind that fast computation is the closed-form solution of optimal transport in one dimension. To leverage that closed-form, sliced Wasserstein utilizes Radon Transform [23] (RT) to transform a high-dimensional distribution to its one-dimensional projected distributions, then the final distance is calculated as the average of all one-dimensional Wasserstein distance. By doing that, the SW distance has a very fast sample complexity i.e., $\mathcal{O}(n^{-1/2})$, which makes it computationally and statistically scalable in any dimension. Therefore, the SW distance has been applied successfully in various domains of applications including generative models [19], domain adaptation [32], clustering [27], 3D shapes [30, 29], gradient flows [34, 8], Bayesian inference computation [37, 58], texture synthesis [22], and many other tasks.

Despite being useful, the SW distance is not as flexible as the Wasserstein distance in terms of choosing the ground metric. In greater detail, the number of ground metrics in one dimension is limited, especially ground metrics that lead to the closed-form solution. As a result, the role of capturing the structure of distributions belongs to the slicing/projecting operators. To generalize RT to non-linear projection, generalized Radon Transform (GRT) is introduced in [3] with circular projection [28], polynomial projection [51], and so on. With GRT, Generalized Sliced Wasserstein (GSW) distance is proposed in [26]. In addition, there is a line of works on developing sliced Wasserstein variants on different manifolds such as hyper-sphere [6, 54, 49, 50], hyperbolic manifolds [7], the manifold of symmetric positive definite matrices [10], general manifolds and graphs [52]. In those works, special variants of GRT are proposed.

Although the SW has become more effective on multiple domains, no SW variant is designed specifically for heterogeneous joint distributions i.e., joint distributions that have marginals supported on different domains, except for the product of Hadamard manifolds [9]. It is worth noting that marginal domains of heterogeneous joint distributions can be any metric space and are not necessary manifolds. Heterogeneous joint distributions appear in many applications, e.g., domain adaptation domains [15, 4], comparing datasets with labels [1], 3D shape deformation [29], and so on. In this case, Wasserstein distance can be adapted by using a mixed ground metric, i.e., a weighted sum of metrics on domains [15, 1]. In contrast to the Wasserstein distance, the adaptation of SW has not been well-investigated. Using GSW directly with one type defining function for all marginals cannot separate the information within and among groups of arguments.

Contribution: In this work, we tackle the challenge of designing a sliced Wasserstein variant for heterogeneous joint distributions. In summary, our main contributions are three-fold:

- 1. We first extend the partial Radon Transform to the partial generalized Radon Transform (PGRT) to inject non-linearity into local transformation. We discuss the injectivity of PGRT for some choices of defining functions. We then propose a novel slicing operator for heterogeneous joint distributions, named Hierarchical Hybrid Radon Transform (HHRT). In particular, HHRT is a hierarchical transformation that first applies partial generalized Radon Transform with different defining functions on arguments of each marginal to gather marginal information, then applies partial Radon Transform on the joint transformed arguments to gather information among marginals. We show that HHRT is injective as long as the partial generalized Radon Transform is injective.
- 2. We propose Hierarchical Hybrid Sliced Wasserstein (H2SW) which is a novel metric for comparing heterogeneous joint distributions by utilizing the HHRT. Moreover, we investigate the topological properties, statistical properties, and computational properties of H2SW. In particular, we show that H2SW is a valid metric on the space of distribution over the joint space, H2SW does not suffer from the curse of dimensionality and enjoys the same computational scalability as SW distance.
- 3. A 3D mesh can be effectively represented by a point-cloud and corresponding surface normal vectors. Therefore, it can be seen as an empirical heterogeneous joint distribution. We conduct experiments on optimization-based 3D mesh deformation and deep 3D mesh autoencoder to show the favorable performance of H2SW compared to SW and GSW. Moreover, we also illustrate that H2SW can also provide a meaningful comparison for probability distributions on the product of Hadamard manifolds by conducting experiments on dataset comparison.

Organization. We first provide some preliminaries on SW distance, GSW distance, and joint Wasserstein distance in Section 2. We then define the hierarchical hybrid Radon transform and hierarchical hybrid sliced Wasserstein distance s in Section 3. Section 4 contains experiments on 3D mesh deformation, deep 3D mesh autoencoder, and datasets comparison. We conclude the paper in Section 5. Finally, we defer the proofs of key results, and additional materials in the Appendices.

2 Preliminaries

Wasserstein distance. For $p \geq 1$, the Wasserstein-p distance [55, 47] between two distributions $\mu \in \mathcal{P}(\mathcal{X})$ and $\nu \in \mathcal{P}(\mathcal{Y})$, where \mathcal{X} and \mathcal{Y} are subsets of \mathbb{R}^d and they share a ground metric $c: \mathcal{X} \times \mathcal{Y} \to \mathbb{R}^+$, is defined as:

$$\mathbf{W}_{p}^{p}(\mu,\nu;c) := \inf_{\pi \in \Pi(\mu,\nu)} \int_{\mathcal{X} \times \mathcal{Y}} c(x,y)^{p} d\pi(x,y), \tag{1}$$

where $\Pi(\mu,\nu) := \left\{ \pi \in \mathcal{P}(\mathcal{X} \times \mathcal{Y}) \right\} | \int_{\mathcal{Y}} d\pi(x,y) = \mu(x), \int_{\mathcal{X}} d\pi(x,y) = \nu(y) \right\}$. When μ and ν are discrete with at most n supports, the time complexity and the space complexity of the Wasserstein distance is $\mathcal{O}(n^3 \log n)$ and $\mathcal{O}(n^2)$ in turn which are very expensive. Therefore, sliced Wasserstein is proposed as an alternative solution. We first review the definition of Radon Transform.

Radon Transform [23] The *Radon Transform* $\mathcal{R}: \mathbb{L}_1(\mathbb{R}^d) \to \mathbb{L}_1\left(\mathbb{R} \times \mathbb{S}^{d-1}\right)$ is defined as:

$$(\mathcal{R}f)(t,\theta) = \int_{\mathbb{R}^d} f(x)\delta(t - \langle x, \theta \rangle)dx. \tag{2}$$

Radon Transform defines a linear bijection [23]. Given a projecting direction θ , $(\mathcal{R}f)(\cdot, \theta)$ is an one-dimensional function. With Radon Transform, we can now define the sliced Wasserstein distance.

Sliced Wasserstein distance. For $p \geq 1$, the *Sliced Wasserstein (SW)* distance [11] of p-th order between two distributions $\mu \in \mathcal{P}(\mathcal{X})$ and $\nu \in \mathcal{P}(\mathcal{Y})$ with an one-dimensional ground metric $c : \mathbb{R} \times \mathbb{R} \to \mathbb{R}^+$ is defined as follow:

$$SW_n^p(\mu, \nu; c) = \mathbb{E}_{\theta \sim \mathcal{U}(\mathbb{S}^{d-1})}[W_n^p(\mathcal{R}_\theta \sharp \mu, \mathcal{R}_\theta \sharp \nu; c)], \tag{3}$$

where $\mathcal{R}_{\theta}\sharp\mu$ and $\mathcal{R}_{\theta}\sharp\nu$ are the one-dimensional push-forward distributions created by applying Radon Transform (RT) [23] on the pdf of μ and ν with the projecting direction θ . The computational benefit of SW distance comes from the closed-form solution when the one-dimensional ground metric c(x,y)=h(x-y) for h is a strictly convex function:

$$\mathbf{W}_p^p(\mathcal{R}_\theta \sharp \mu, \mathcal{R}_\theta \sharp \nu; c) = \int_0^1 c \left(F_{\mathcal{R}_\theta \sharp \mu}^{-1}(z), F_{\mathcal{R}_\theta \sharp \nu}^{-1}(z) \right)^p dz,$$

where $F_{\mathcal{R}_{\theta}\sharp\mu}^{-1}$ and $F_{\mathcal{R}_{\theta}\sharp\nu}^{-1}$ are inverse CDF of $\mathcal{R}_{\theta}\sharp\mu$ and $\mathcal{R}_{\theta}\sharp\nu$ respectively. When μ and ν are discrete with at most n supports, the time complexity and the space complexity of the closed-form is $\mathcal{O}(n\log n)$ and $\mathcal{O}(n)$ respectively.

Generalized Radon Transform and Generalized Sliced Wasserstein distance. To generalize RT to non-linear operator, the *Generalized Radon Transform (GRT)* was proposed [3]. Given a defining function [26] $g: \mathbb{R}^d \times \Omega \to \mathbb{R}$, the Generalized Radon Transform [3] $\mathcal{GR}: \mathbb{L}_1(\mathbb{R}^d) \to \mathbb{L}_1$ ($\mathbb{R} \times \Omega$) is defined as:

$$(\mathcal{GR}f)(t,\theta) = \int_{\mathbb{R}^d} f(x)\delta(t - g(x,\theta))dx.$$

For example, we can have the circular function [28], i.e., $g(x,\theta) = \|x - r\theta\|_2$ for $r \in \mathbb{R}^+$ and $\theta \in \Omega := \mathbb{S}^{d-1}$, homogeneous polynomials with an odd degree [51] (m), i.e., $g(x,\theta) = \sum_{|\alpha| = m} \theta_\alpha x^\alpha$ with $\alpha = (\alpha_1, \ldots, \alpha_{d_\alpha}) \in \mathbb{N}^{d_\alpha}$, $|\alpha| = \sum_{i=1}^{d_\alpha} \alpha_i$, $x^\alpha = \prod_{i=1}^{d_\alpha} x_i^{\alpha_i}$, $\Omega = \mathbb{S}^{d_\alpha}$, and so on. Using GRT, the *Generalized Sliced Wasserstein (GSW)* distance is introduced in [26], which is formally defined as follow:

$$GSW_p^p(\mu,\nu;c,g) = \mathbb{E}_{\theta \sim \mathcal{U}(\mathbb{S}^{d-1})}[W_p^p(\mathcal{GR}_{\theta}^g \sharp \mu, \mathcal{GR}_{\theta}^g \sharp \nu;c)]. \tag{4}$$

It is worth noting that the injectivity of GRT is required to have the identity of indiscernible GSW.

Heterogeneous joint distributions comparison. We are given two joint distributions $\mu(x_1, x_2) \in \mathcal{P}(\mathcal{X}_1 \times \mathcal{X}_2)$ and $\nu(y_1, y_2) \in \mathcal{P}(\mathcal{Y}_1 \times \mathcal{Y}_2)$ where X_1 are Y_1 share a ground metric $c_1 : \mathcal{X}_1 \times \mathcal{Y}_1 \to \mathbb{R}^+$ and X_2 are Y_2 share a ground metric $c_2 : \mathcal{X}_2 \times \mathcal{Y}_2 \to \mathbb{R}^+$ with $(c_1 \neq c_2)$. In this case, previous works utilize the joint distribution Wasserstein distance [15, 1] to compare μ and ν :

$$\mathbf{W}_{p}^{p}(\mu,\nu;c_{1},c_{2}) := \inf_{\pi \in \Pi(\mu,\nu)} \int_{\mathcal{X}_{1} \times \mathcal{X}_{2} \times \mathcal{Y}_{1} \times \mathcal{Y}_{2}} (c_{1}(x_{1},y_{1})^{p} + c_{2}(x_{2},y_{2})^{p}) d\pi(x_{1},x_{2},y_{1},y_{2}), \quad (5)$$

where
$$\Pi(\mu,\nu):=\Big\{\pi\in\mathcal{P}(\mathcal{X}_1\times\mathcal{X}_2\times\mathcal{Y}_1\times\mathcal{Y}_2)\}|\int_{\mathcal{Y}_1\times\mathcal{Y}_2}d\pi(x_1,x_2,y_1,y_2)=\mu(x_1,x_2),$$

 $\int_{\mathcal{X}_1 \times \mathcal{X}_2} d\pi(x_1, x_2, y_1, y_2) = \nu(y_1, y_2)$. We can easily extend the definition to joint distributions with more than two marginals (see Appendix B). In contrast to the Wasserstein distance, there is no variant of SW that is designed specifically for this case. SW variants can still be used by treating $\mathcal{X}_1 \times \mathcal{X}_2$ and $\mathcal{Y}_1 \times \mathcal{Y}_2$ as homogeneous spaces \mathcal{X} and \mathcal{Y} which share the same Radon Transform variant and one-dimensional ground metric c. However, that approach cannot differentiate the difference between \mathcal{X}_1 and \mathcal{X}_2 , and leverage the hierarchical structure, i.e., inside and among marginals.

3 Hierarchical Hybrid Sliced Wasserstein Distance

In this section, we propose the Hierarchical Hybrid Radon Transform (HHRT) which first applies P(G)RT on each marginal argument to gather each marginal information, then applies PRT on the joint transformed arguments from all marginals to gather information among marginals. After that, we introduce Hierarchical Hybrid Sliced Wasserstein distance by using HHRT as the slicing operator.

3.1 Hierarchical Hybrid Radon Transform

We first introduce the first building block in HHRT, i.e., Partial Generalized Radon Transform (PGRT).

Definition 1 (Partial Generalized Radon Transform). Given a defining function $g : \mathbb{R}^{d_1} \times \Omega \to \mathbb{R}$, Partial Generalized Radon Transform $\mathcal{PGR} : \mathbb{L}_1(\mathbb{R}^{d_1} \times \mathbb{R}^{d_2}) \to \mathbb{L}_1(\mathbb{R} \times \Omega \times \mathbb{R}^{d_2})$ is defined as:

$$(\mathcal{PGR}f)(t,\theta,y) = \int_{\mathbb{R}^{d_1}} f(x,y)\delta(t-g(x,\theta))dx. \tag{6}$$

When $g(x, \theta) = \langle x, \theta \rangle$, PGRT reverts into Partial Radon Transform (PRT) [33].

Proposition 1. For some defining function g such as linear, circular, and homogeneous polynomials with an odd degree; the Partial Generalized Radon Transform is injective, i.e., for any functions $f_1, f_2 \in \mathbb{L}^1(\mathbb{R}^d)$, $(\mathcal{PGR}f_1)(t,\theta,y) = (\mathcal{PGR}f_2)(t,\theta,y) \ \forall t,\theta,y$ implies $f_1 = f_2$.

The proof of Proposition 1 is given in Appendix A.1. The main idea to prove the injectivity of PGRT is to show that given a fixed y, the PGRT is the GRT of $f(\cdot, y)$.

Definition 2 (Hierarchical Hybrid Radon Transform). Given defining functions $g_1 : \mathbb{R}^{d_1} \times \Omega_1 \to \mathbb{R}$ and $g_2 : \mathbb{R}^{d_2} \times \Omega_2 \to \mathbb{R}$, Hierarchical Hybrid Radon Transform $\mathcal{HHR} : \mathbb{L}_1(\mathbb{R}^{d_1} \times \mathbb{R}^{d_2}) \to \mathbb{L}_1(\mathbb{R} \times \Omega_1 \times \Omega_2 \times \mathbb{S})$ is defined as:

$$(\mathcal{HHR}f)(t,\theta_1,\theta_2,\psi) = \int_{\mathbb{R}^{d_1} \times \mathbb{R}^{d_2}} f(x_1,x_2) \delta(t - \psi_1 g_1(x_1,\theta_1) - \psi_2 g_2(x_2,\theta_2)) \, dx_1 dx_2, \quad (7)$$

where $\psi = (\psi_1, \psi_2) \in \mathbb{S}$.

The reason for using PRT for the final transform is that the previous PGRTs are assumed to be able to transform the non-linear structure to a linear line. However, PGRT can still be used as a replacement for PRT. Definition 2 can be extended to more than two marginals (see Appendix B).

Proposition 2. For some defining functions g_1, g_2 such as linear, circular, and homogeneous polynomials with an odd degree; Hierarchical Hybrid Radon Transform is injective, i.e., for any functions $f_1, f_2 \in \mathbb{L}_1(\mathbb{R}^d)$, $(\mathcal{HHR}f_1)(t, \theta_1, \theta_2, \psi) = (\mathcal{HHR}f_2)(t, \theta_1, \theta_2, \psi) \ \forall t, \theta_1, \theta_2, \psi \text{ implies } f_1 = f_2.$

The proof of Proposition 2 is given in Appendix A.2. The main idea to prove the injectivity of HHRT is to show that HHRT is the composition of PRT and multiple PGRTs.

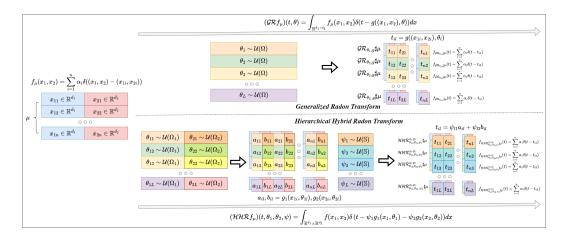


Figure 1: Generalized Radon Transform and Hierarchical Hybrid Radon Transform on a discrete distribution.

HHRT of discrete distributions. We are given $f(x) = \sum_{i=1}^n \alpha_i \delta((x_1, x_2) - (x_{1i}, x_{2i}))$ $(n \ge 1, \alpha_i \ge 0 \ \forall i)$. The HHRT of f(x) is $(\mathcal{HHR}f)(t, \theta_1, \theta_2, \psi) = \sum_{i=1}^n \alpha_i \delta(t - \psi_1 g_1(x_{i1}, \theta_1) - \psi_2 g_2(x_{i2}, \theta_2))$. For g_1 and g_2 that are the linear function and (or) the circular function, the time complexity of the transform is $\mathcal{O}(d_1 + d_2)$ which is the same as the complexity of using RT and GRT directly. However, HHRT has an additional constant complexity scaling linearly with the number of marginals, i.e., two marginals in Definition 2.

Example 1. In this paper, we focus on 3D shape data (mesh) with points and normals representation, i.e., shapes as points representation [46]. In particular, we can transform a 3D shape into a set of points and normals by sampling from the surface of the mesh. In addition, we can convert back to the 3D shape from points and normals with Poisson surface reconstruction [25] algorithm. In this setup, a shape is represented by a 6-dimensional vector $x = (x_1, x_2)$ where $x_1 \in \mathcal{X}_1 \in \mathbb{R}^3$ and $x_2 \in \mathcal{X}_2 \in \mathbb{S}^2$. For the set $\mathcal{X}_1 \in \mathbb{R}^3$, we can use directly the linear defining function $g_1(x_1, \theta_1) = \langle x_1, \theta_1 \rangle$ with $\theta_1 \in \mathbb{S}^2$. For the set $\mathcal{X}_2 \in \mathbb{S}^2$, we can utilize the circular defining function $g_2(x_2, \theta_2) = ||x_2 - r\theta_2||_2$ with $r \in \mathbb{R}^+$ and $\theta_2 \in \mathbb{S}^2$. As alternative options for \mathcal{X}_2 , we can use other defining functions from special cases of GRT including Vertical Slice Transform [49], Parallel Slice Transform [50], Spherical Radon Transform [6], and Stereographic Spherical Radon Transform [54].

Inversion. In Proposition 2, we show that HHRT is the composition of PRT and multiple PGRTs. Therefore, the inversion of HHRT is the composition of the inversion of multiple PGRT (invertibility of PGRT depends on the choice of defining functions [3, 28]) and the inversion of PRT [23].

3.2 Hierarchical Hybrid Sliced Wasserstein Distance

By using HHRT, we obtain a novel variant of SW which is specifically designed for comparing heterogeneous joint distributions.

Definitions. We now define the Hierarchical Hybrid Sliced Wasserstein (H2SW) distance.

Definition 3. For $p \geq 1$, defining functions g_1, g_2 , the hierarchical hybrid sliced Wasserstein-p (H2SW) distance between two distributions $\mu \in \mathcal{P}(\mathcal{X}_1 \times \mathcal{X}_2)$ and $\nu \in \mathcal{P}(\mathcal{Y}_1 \times \mathcal{Y}_2)$ with an one-dimensional ground metric $c : \mathbb{R} \times \mathbb{R} \to \mathbb{R}^+$ is defined as:

 $H2SW_p^p(\mu,\nu;c,g_1,g_2) = \mathbb{E}_{(\theta_1,\theta_2,\psi)\sim\mathcal{U}(\Omega_1\times\Omega_2\times\mathbb{S})}[W_p^p(\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\mu,\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\nu;c)], \quad (8)$ where $\mathcal{HHR}_{\theta_1,\theta_2,\psi}\sharp\mu$ and $\mathcal{HHR}_{\theta_1,\theta_2,\psi}\sharp\nu$ are the one-dimensional push-forward distributions created by applying HHRT.

Definition 3 can be easily extended to more than two marginals (see Appendix B)

Topological Properties. We first show that H2SW is a valid metric on the space of distributions on any sets $\mathcal{X} \times \mathcal{Y} \in \mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$ $(d_1, d_2 \ge 1)$.

Theorem 1. For any $p \ge 1$, ground metric c, and defining functions g_1, g_2 which lead to the injectivity of GRT, the hierarchical hybrid sliced Wasserstein $H2SW_p(\cdot,\cdot;c,g_1,g_2)$ is a metric on $\mathcal{P}(\mathbb{R}^{d_1}\times\mathbb{R}^{d_2})$ i.e., it satisfies the symmetry, non-negativity, triangle inequality, and identity of indiscernible.

The proof of Theorem 1 is given in Appendix A.3. It is worth noting that the identity of indiscernible property is proved by the injectivity of HHRT (Proposition 2). We now discuss the connection of H2SW to GSW and Wasserstein distance in some specific cases.

Proposition 3. For any $p \ge 1$, c(x,y) = |x-y|, and $\mu, \nu \in \mathcal{P}(\mathbb{R}^{d_1} \times \mathbb{R}^{d_2})$, we have: (i) $H2SW_p(\mu, \nu; c, g_1, g_2) \le GSW_p(\mu_1, \nu_1; c, g_1) + GSW_p(\mu_2, \nu_2; c, g_2)$, where $\mu_1(X) = \mu(X \times \mathbb{R}^{d_2})$ and $\mu_2(Y) = \mu(\mathbb{R}^{d_1} \times Y)$ (similar with ν_1 and ν_2).

- (ii) If g_1 , g_2 are linear defining functions, $H2SW_p(\mu, \nu; c, g_1, g_2) \le W_p(\mu_1, \nu_1; c) + W_p(\mu_2, \nu_2; c)$.
- (iii) If p=1, g_1 , g_2 are linear defining functions, $H2SW_1(\mu,\nu;c,g_1,g_2) \leq W_1(\mu,\nu;c)$.

The proof of Proposition 3 is given in Appendix A.4.

Sample Complexity. We now discuss the sample complexity of H2SW.

Proposition 4. For any $p \ge 1$, dimension $d_1, d_2 \ge 1$, q > p, c(x, y) = |x - y|, g_1, g_2 are linear defining functions or circular defining functions, and $\mu, \nu \in \mathcal{P}_q(\mathbb{R}^{d_1} \times \mathbb{R}^{d_2})$ with the corresponding empirical distributions μ_n and ν_n $(n \ge 1)$, there exists a constant $C_{p,q}$ depending on p,q such that:

$$\mathbb{E} |H2SW_p(\mu_n, \nu_n; c, g_1, g_2) - H2SW_p(\mu, \nu; c, g_1, g_2)|$$

$$\leq C_{p,q}^{\frac{1}{p}} \left(\sum_{i=0}^{q} q^{i} C_{g_{1},g_{2}}^{q-i}(M_{i}(\mu) + M_{i}(\nu)) \right)^{\frac{1}{p}} \begin{cases} n^{-1/2p} \text{ if } q > 2p, \\ n^{-1/2p} \log(n)^{\frac{1}{p}} \text{ if } q = 2p, \\ n^{-(q-p)/pq} \text{ if } q \in (p,2p), \end{cases}$$
(9)

where $M_q(\mu)$ and $M_q(\nu)$ are the q-th moments of μ and ν , C_{g_1,g_2} is a constant depends on g_1 , g_2 .

The proof of Proposition 4 is given in Appendix A.5. The rate in Proposition 4 is as good as the rate of SW in [38], however, it is slightly worse than the rate of SW in [44, 36, 5] due to the usage of the circular defining functions and simpler assumptions. To the best of our knowledge, the sample complexity of GSW has not been investigated.

Monte Carlo Estimation. Since the expectation in H2SW (Equation 8) is intractable, Monte Carlo estimation and Quasi-Monte Carlo approximation [39] can be used to form a practical evaluation of H2SW. Here, we utilize Monte Carlo estimation for simplicity. In particular, we sample $\theta_{11}, \ldots, \theta_{1L} \overset{i.i.d}{\sim} \mathcal{U}(\Omega_1), \, \theta_{21}, \ldots, \theta_{2L} \overset{i.i.d}{\sim} \mathcal{U}(\Omega_2), \, \text{and} \, \psi_1, \ldots, \psi_L \overset{i.i.d}{\sim} \mathcal{U}(\mathbb{S}).$ After that, we form the following estimation of H2SW:

$$\widehat{\text{H2SW}}_{p}^{p}(\mu,\nu;c,g_{1},g_{2},L) = \frac{1}{L} \sum_{l=1}^{L} W_{p}^{p}(\mathcal{HHR}_{\theta_{1l},\theta_{2l},\psi_{l}}^{g_{1},g_{2}} \sharp \mu, \mathcal{HHR}_{\theta_{1l},\theta_{2l},\psi_{l}}^{g_{1},g_{2}} \sharp \nu;c).$$
(10)

Proposition 5. For any $p \ge 1$, dimension $d_1, d_2 \ge 1$, and $\mu, \nu \in \mathcal{P}(\mathbb{R}^d_1 \times \mathbb{R}^{d_2})$, we have:

$$\mathbb{E}|\widehat{H2SW}_{p}^{p}(\mu,\nu;c,g_{1},g_{2},L) - H2SW_{p}^{p}(\mu,\nu;c,g_{1},g_{2})|$$

$$\leq \frac{1}{\sqrt{L}}Var\left[W_{p}^{p}(\mathcal{HHR}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu,\mathcal{HHR}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\nu;c)\right]^{\frac{1}{2}},$$
(11)

where the variance is with respect to $\mathcal{U}(\Omega_1 \times \Omega_2 \times \mathbb{S})$.

The proof of Proposition 5 is given in Appendix A.6. From the proposition, we see that the estimation error of H2SW is the same as SW which is $\mathcal{O}(L^{-1/2})$.

Computational Complexities. The time complexity and memory complexity of H2SW with linear and circular defining functions are $\mathcal{O}(Ln\log n + L(d_1 + d_2 + k)n)$ and $\mathcal{O}(Ln + (d_1 + d_2 + k)n)$ with k is the number of marginals i.e., 2. We can see that the complexities of H2SW are the same as those of SW in terms of the number of supports n and the number of dimensions d. We demonstrate the process of HHRT compared to GRT on a discrete distribution with L realization of θ_1, θ_2, ψ in Figure 1. Overall, the complexities of defining functions are often different in the number of dimensions, hence, H2SW is always scaled the same as SW in the number of supports i.e., $\mathcal{O}(n\log n)$.

Gradient Estimation. In applications, it is desirable to estimate the gradient $\nabla_{\phi} \text{H2SW}_{p}^{p}(\mu_{\phi}, \nu; c, g_{1}, g_{2})$. We can move the gradient operator to inside the expectation

Table 1: Summary of joint Wasserstein distances across time steps from deformation from the sphere mesh to the Armadillo mesh.

Distances	Step 100 (W $_{c_1,c_2}\downarrow$)	Step 300 (W $_{c_1,c_2}\downarrow$)	Step 500 (W $_{c_1,c_2}\downarrow$)	Step 1500 (W $_{c_1,c_2}\downarrow$)	Step 4000 (W $_{c_1,c_2}\downarrow$)	Step 5000 (W $_{c_1,c_2}\downarrow$)
SW L=10	1852.519±3.236	1436.686±3.056	1071.227±2.449	104.452±2.35	6.19±0.307	2.726±0.305
GSW L=10	1893.438±3.205	1535.737±3.363	1192.52±3.274	143.518±1.04	8.73±0.353	4.743±0.134
H2SW L=10	1840.73 ± 1.282	1422.667 ± 7.813	1058.171 ± 5.362	95.672 ± 4.376	6.326±0.151	2.602 ± 0.201
SW L=100	1847.572±0.303	1426.425±0.528	1059.127±1.106	89.693±0.793	4.453±0.22 7.905±0.373 4.61±0.431	1.171±0.056
GSW L=100	1889.312±0.883	1525.269±1.078	1179.1±2.052	122.618±1.175		3.226±0.388
H2SW L=100	1839.347 ± 1.986	1417.1 ± 3.677	1048.895 ± 4.008	86.078 ± 0.623		1.086 ± 0.177

Table 2: Summary of joint Wasserstein distances (multiplied by 100) across time steps from deformation from the sphere mesh to the Stanford Bunny mesh.

Distances	Step 100 ($W_{c_1,c_2} \downarrow$)	Step 300 (W $_{c_1,c_2}\downarrow$)	Step 500 ($W_{c_1,c_2}\downarrow$)	Step 1500 ($W_{c_1,c_2} \downarrow$)	Step 4000 (W $_{c_1,c_2}\downarrow$)	Step 5000 ($W_{c_1,c_2}\downarrow$)
SW L=10	26.868±0.579	4.46±0.195	1.52±0.081	0.623±0.024	0.221±0.023	0.14±0.018
GSW L=10	26.837±0.496	4.378±0.128	1.548±0.062	0.653±0.01	0.173 ± 0.018	0.146±0.013
H2SW L=10	23.283 ± 0.119	2.221 ± 0.124	1.452 ± 0.075	0.636±0.045	0.177±0.009	0.089 ± 0.022
SW L=100	26.678±0.168	4.109±0.138	$\substack{1.458 \pm 0.142\\1.375 \pm 0.049\\\textbf{1.358} \pm \textbf{0.051}}$	0.362±0.023	0.072±0.017	0.049±0.006
GSW L=100	26.795±0.202	4.084±0.109		0.372±0.026	0.048 ± 0.004	0.042±0.017
H2SW L=100	23.772±0.19	2.388 ± 0.009		0.488±0.026	0.064±0.01	0.038 ± 0.007

and then apply Monte Carlo estimation. The gradient $\nabla_{\phi} W_p^p(\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\mu_{\phi},\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\nu;c)]$ can be computed easily since the functions g_1,g_2 are usually differentiable.

Beyond uniform slicing distribution. H2SW is defined with the uniform slicing distribution in Definition 3, however, it is possible to extend it to other slicing distributions such as the maximal projecting direction [18], distributional slicing distribution [42], and energy-based slicing distribution [41]. Since the choice of slicing distribution is independent of the main contribution i.e., the slicing operator, we leave this investigation to future work.

H2SW for distributions on the product of Hadamard manifolds. A recent work [9] extends sliced Wasserstein on hyperbolic manifolds [7] and on the manifold of symmetric positive definite matrices [10] to Hadamard manifolds i.e., manifold non-positive curvature. The work discusses the extension of SW to the product of Hadamard manifolds. For the geodesic projection, the closed-form for the projection is intractable. For the Busemann projection, the Busemann projection on the product manifolds is the weighted sum of the Busemann projection with the weights belonging to the unit-sphere. In the work, the weights are a fixed hyperparameter i.e., Cartan-Hadamard Sliced-Wasserstein (CHSW) utilizes only one Busemann function to project the joint distribution. In contrast, H2SW utilizes the Radon Transform on the joint spaces of projections i.e., considering all distributed weighted combinations which is equivalent to considering all Busemann functions under a probability law. As a result, the H2SW is a valid metric as long as the Busemann projections can be proven to be injective (the injectivity of the Busemann projection has not been known at the moment) while Cartan-Hadamard Sliced-Wasserstein is only pseudo metric since the injectivity of a fixed weighted combination is not trivial to show. Moreover, H2SW does not only focus on the product of Hadamard manifolds i.e., H2SW is a generic distance for heterogeneous joint distributions in which marginal domains are not necessary manifolds e.g., images [40], functions [21], and so on. In the later experiments, we conduct experiments on comparing 3D shapes which are represented by a distribution on the product of the Euclidean space and the 2D sphere (not a Hadamard manifold).

4 Experiments

In this section, we first compare the performance of the proposed H2SW with SW and GSW in the 3D mesh deformation application. After that, we further evaluate the performance of H2SW in training a deep 3D mesh autoencoder compared to SW and GSW. Finally, we compare H2SW with SW and Cartan-Hadamard Sliced-Wasserstein (CHSW) in datasets comparison on the product of Hadamard manifolds. In the experiments, we use c(x,y) = |x-y| and p=2 for all SW variants.

4.1 3D Mesh Deformation

In this task, we would like to move from a source mesh to a target mesh. To represent those meshes, we sample 10000 points by Poisson disk sampling and their corresponding normal vectors of the mesh surface at those points. Let the source mesh be denoted as $X(0) = \{x_1(0), \ldots, x_n(0)\}$ and the target mesh be denoted as $Y = \{y_1, \ldots, y_n\}$. We deform X(0) to Y by integrating the ordinary differential equation $\dot{X}(t) = -n\nabla_{X(t)}\left[\mathcal{S}\left(\frac{1}{n}\sum_{i=1}^n\delta(x-x_i(t)), \frac{1}{n}\sum_{i=1}^n\delta(y-y_i)\right)\right]$, where \mathcal{S} denotes a SW

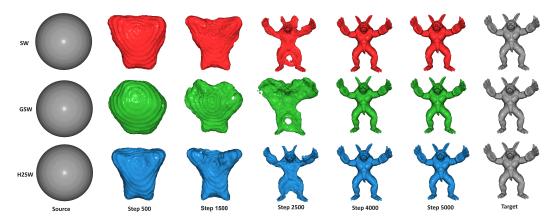


Figure 2: Visualization of deformation from the sphere mesh to the Armadillo mesh with L=10.

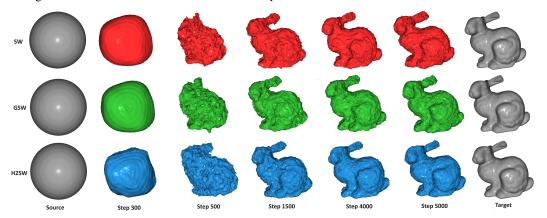


Figure 3: Visualization of deformation from the sphere mesh to the Stanford Bunny mesh with L=10.

variant. We utilize the Euler discretization scheme with step size 0.01 and 5000 steps. The normal vectors are projected back to the sphere after taking an Euler step. For evaluation, we use the joint Wasserstein distance in Equation 5 with the mixed distance from the Euclidean distance and the great circle distance. We use the circular defining function for GSW, and use the linear defining function and the circular defining function for H2SW. We vary the number of projections $L \in \{10, 100\}$ for all variants. For H2SW and GSW, we select the best hyperparameter of the circular defining function $r \in \{0.5, 0.7, 0.8, 0.9, 1, 5, 10, 50, 100\}$.

Results. We compare H2SW with GSW and SW by deforming the sphere mesh to the Armadillo mesh [59]. We report the quantitative results in Table 1 after 3 independent runs and the qualitative result for L=10 in Figure 2 and L=100 in Figure 6 in Appendix D. From Table 1, we observe that H2SW helps the deformation convergence faster at the beginning and better at the end in terms of the joint Wasserstein distance, especially for a small value of the number of projections i.e., L=10. The result for L=100 is better than L=10 which is consistent with Proposition 5. The qualitative results in Figure 2 and Figure 6 also reinforce the favorable performance of H2SW since they are visually consistent with quantitative scores. We also conduct deformation to the Stanford Bunny mesh [16, 59] in Table 2, Figure 3, and Figure 7 in Appendix D and we observe the same phenomenon that H2SW is the best variant for 3D meshes. From those experiments, H2SW has shown the benefit of the HHRT in transforming a joint distribution over the product of the Euclidean space and the 2D sphere compared to the conventional RT of SW and the conventional GRT of GSW.

4.2 Training deep 3D mesh autoencoder

We utilize the processed ShapeNet dataset [12] from [46], then sample 2048 points and the corresponding normal vectors from each shape in the dataset. Formally, we would like to train an autoencoder that contains an encoder f_{ϕ} that maps a mesh $X \in \mathbb{R}^{2048 \times 6}$ to a latent code $z \in \mathbb{R}^{1024}$, and a decoder g_{ψ} that maps the latent code z back to the reconstructed mesh $\tilde{X} \in \mathbb{R}^{2048 \times 6}$. We adopt

Table 3: Joint Wasserstein distance reconstruction errors (multiplied by 100) from three different runs of autoencoders trained by SW, GSW, and H2SW with the number of projections L=100 and L=1000.

Distance	Epoch 500 $W_{c_1,c_2}(\downarrow)$	Epoch 1000 $W_{c_1,c_2}(\downarrow)$	Epoch 2000 $W_{c_1,c_2}(\downarrow)$
SW L=100 GSW L=100	136.87 ± 1.18 136.51 ± 0.20	133.24 ± 0.50 133.36 ± 0.24	131.10 ± 0.34 130.80 ± 0.46
H2SW L=100	$egin{array}{c} 135.51 \pm 0.20 \ 135.54 \pm 0.72 \ \end{array}$	132.24 ± 0.36	130.24 ± 0.47
SW L=1000	135.85 ± 0.92	132.88 ± 0.27	130.93 ± 0.11
GSW L=1000	136.40 ± 0.10	133.02 ± 0.98	130.76 ± 0.26
H2SW L=1000	$\boldsymbol{135.47 \pm 0.64}$	$\boldsymbol{132.17 \pm 0.10}$	$\boldsymbol{129.87 \pm 0.44}$

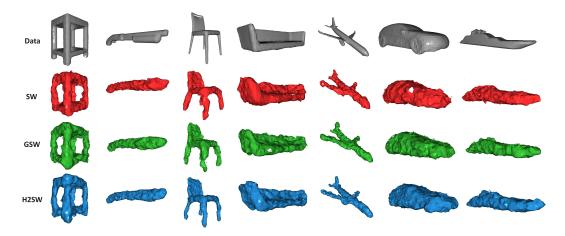


Figure 4: Visualization of some randomly selected reconstruction meshes from autoencoders trained by SW, GSW, and H2SW in turn with the number of projections L=100 at epoch 2000.

Point-Net [48] architecture to construct the autoencoder. We want to train the encoder f_{ϕ} and the decoder g_{ψ} such that $\tilde{X} = g_{\psi}(f_{\phi}(X)) \approx X$ for all shapes X in the dataset. To do that, we solve the following optimization problem:

$$\min_{\phi,\gamma} \mathbb{E}_{X \sim \mu(X)} [\mathcal{S}(P_X, P_{g_{\gamma}(f_{\phi}(X)))}],$$

where S is a sliced Wasserstein variant, and $P_X = \frac{1}{n} \sum_{i=1}^n \delta(x-x_i)$ denotes the empirical distribution over the point cloud $X = (x_1, \dots, x_n)$. We train the autoencoder for 2000 epochs on the training set of the ShapeNet dataset using an SGD optimizer with a learning rate of 1e-3, and a batch size of 128. For evaluation, we also use the joint Wasserstein distance in Equation 5 with the mixed distance from the Euclidean distance and the great circle distance to measure the average reconstruction loss on the testing set of the ShapeNet dataset. We use the circular defining function for GSW, and use the linear defining function and the circular defining function for H2SW. For H2SW and GSW, we select the best hyperparameter of the circular defining function $r \in \{0.5, 0.7, 0.8, 0.9, 1, 5, 10\}$. For more details such as the neural network architectures, we refer the reader to Appendix B.

Results. We report the joint Wasserstein reconstruction errors (measured in three independent times) on the testing set in Table 3 with trained autoencoder at epoch 500, 1000, and 2000 from SW, GSW, and H2SW with the number of projections L=100 and L=1000. In addition, we show some randomly reconstructed meshes for epoch 2000in Figure 4 and for epoch 500 in Figure 8 in Appendix D. From Table 3, we observe that H2SW yields the lowest reconstruction errors for both L=100 and L=1000. Moreover, we see that the reconstruction errors are lower with L=1000 than ones with L=100 for all SW variants. The qualitative reconstructed meshes in Figure 4 and Figure 8reflect the same relative comparison. It is worth noting that both the qualitative and the qualitative performance of autoencoders can be improved by using more powerful neural networks. Since we focus on comparing SW, GSW, and H2SW, we only use a light neural network i.e., Point-Net [48] architecture. The trained autoencoders can be further used to reduce the size of 3D meshes for data compression and for dimension reduction, however, such downstream applications are not our focus in the current investigation of the paper.

Table 4: Relative error to the joint Wasserstein distance of SW, CHSW, and H2SW.

			Di	star	ices		L = 100					L = 500						L = 1000				L = 2000						
			SV	SW				4.618 ± 0.744				4.253 ± 0.398					4	4.235 ± 0.310					4.198 ± 0.238					
			CI	CHSW			4.449 ± 0.497				4.063 ± 0.254				4	4.059 ± 0.167				4.035 ± 0.145								
			H	H2SW			$4.381 \pm 0.695 4.001 \pm 0$					0.2	67	$67 4.048 \pm 0.182$					$\boldsymbol{3.998 \pm 0.142}$									
																										_		
	HHIST	EMNST	Fashion/FMST	KPINIST	USFS			MWST	DHNST	FashiceMNST	KHNST	USPS		0.35	П	HMST	EMMST	FeshionHMST	KMNST	USPS			MWST	DHNST	FeshionHMST	KMNST	USPS	1200
MOUST		0.28±0.009	0.401±0.014	3241±0.006	0.99±0.002	0.4	MNIST		0.199±0.009	0.374±0.001	0.187±0.009	E 281±9.009	Ш		MUST					0.285±0.008	0.95	HMST		206.8			845.07	II
EMNST			0.051±0.013	3.225±0.007	0.488±0.013		EMMST	0.199±0.005	П	0.336±0.03	0.17±0.004	0.333±0.03		0.30	EMNST	8.201±0.006				0.34±0.01	0.30	SPMST	981		1278.35		891.65	- 1900
FashiceMNST				3.307±0.007	0.322+0.005	102	Fashior MNIST	0.878+0.031	0.336±0.01		0.236+0.007	0.15+0.004			PiceMNST	2.379+2.009	0.942+0.01			0.194+0.005		Peobler/MIST		1278.35		1353.50	607.44	600
KMNST			0.307±0.007		0.245±0.005		KMMIST	0.187±0.005	8.37±0.064	0.236±0.007		0.379±0.004		0.35	KMNST	8.389±8.805	0.173±0.005	0.245±0.007		0.007±0.004	0.15	KHMIST			1253.58		849.23	+400
USPS		0.485±0.003	0.22210.005	0.245±0.005			usps	0.281+0.009	8 333+8 EE	0.19+0.004	0.179+0.004			0.05	USPS		0.36+0.01		0.187+0.004		0.05	uses			607.44	843.23		200
	SW						CHSW							H2SW							Jo	oint	sser	stei	n			

Figure 5: Cost matrices between datasets from SW, CHSW, and H2SW with L=2000.

4.3 Comparing Datasets on The Product of Hadamard Manifolds

We follow the same experimental setting from [9]. Here, we have datasets as sets of feature-label pairs which are embedded in the space of $\mathbb{R}^{d_1} \times \mathbb{L}^{d_2}$ where \mathbb{L}^{d_2} denotes a Lorentz model of d_2 dimension (a hyperbolic space). We uses MNIST [31] dataset, EMNIST dataset [14], Fashion MNIST dataset [56], KMNIST dataset [13], and USPS dataset [24]. For CHSW, we use Busemann projection on the product space of Euclidean and the Lorentz model. For H2SW, we use the linear defining function and the Busemann function on the Lorentz model. We refer the reader to Appendix B for greater detail on Busemann functions and experimental setups. We compare SW, CHSW, and H2SW by varying $L \in \{100, 500, 1000, 2000\}$. For evaluation, we use the joint Wasserstein distance in [1] as the ground truth. In particular, let C_W be the cost matrix from the joint Wasserstein distance and C be a given cost matrix, we use $|C/\max(C) - C_W/\max(C_W)|$ as the relative error.

Results. We report the relative errors from SW, CHSW, and H2SW in Table 4 after 100 independent runs. In addition, we show the cost matrices from SW, CHSW, H2SW. and joint Wasserstein distance with L=2000 in Figure 5. Cost matrices for L=100, L=500, and L=1000 are given in Figure 9-11 in Appendix D. From Table 4, we see that H2SW gives a lower relative error than CHSW and SW. Therefore, using H2SW for comparing datasets is the most equivalent to the joint Wasserstein distance in terms of the relative error. We also observe that increasing the value of the number of projections also reduces the relative errors for all SW variants. Again, we would like to recall that H2SW can be used for heterogeneous joint distributions beyond the product of Hadamard manifolds as shown in previous experiments.

5 Conclusion

We have presented Hierarchical Hybrid Sliced Wasserstein (H2SW) distance, a novel sliced probability metric for heterogeneous joint distributions i.e., joint distributions have marginals on different domains. The key component of H2SW is the proposed hierarchical hybrid Radon Transform (HHRT) which is the composition of partial Radon Transform and multiples proposed partial generalized Radon Transform. We then discuss the injectivity of the proposed transforms and theoretical properties of H2SW including topological properties, statistical properties, and computational properties. On the experimental side, we show that H2SW has favorable performance in applications of 3D mesh deformation, training deep 3D mesh autoencoder, and datasets comparison. In those applications, heterogeneous joint distributions appear in the form of joint distributions on the product of Euclidean space and 2D sphere, and the product of Hadamard manifolds. In the future, we will extend the application of H2SW to more complicated heterogeneous joint distributions.

Acknowledgements

NH acknowledges support from the NSF IFML 2019844 and the NSF AI Institute for Foundations of Machine Learning.

References

- [1] D. Alvarez-Melis and N. Fusi. Geometric dataset distances via optimal transport. *Advances in Neural Information Processing Systems*, 33:21428–21439, 2020.
- [2] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223, 2017.
- [3] G. Beylkin. The inversion problem and applications of the generalized Radon transform. *Communications on pure and applied mathematics*, 37(5):579–599, 1984.
- [4] B. Bhushan Damodaran, B. Kellenberger, R. Flamary, D. Tuia, and N. Courty. Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 447–463, 2018.
- [5] M. T. Boedihardjo. Sharp bounds for the max-sliced Wasserstein distance. *arXiv preprint* arXiv:2403.00666, 2024.
- [6] C. Bonet, P. Berg, N. Courty, F. Septier, L. Drumetz, and M.-T. Pham. Spherical sliced-Wasserstein. *International Conference on Learning Representations*, 2023.
- [7] C. Bonet, L. Chapel, L. Drumetz, and N. Courty. Hyperbolic sliced-Wasserstein via geodesic and horospherical projections. In *Topological, Algebraic and Geometric Learning Workshops* 2023, pages 334–370. PMLR, 2023.
- [8] C. Bonet, N. Courty, F. Septier, and L. Drumetz. Efficient gradient flows in sliced-Wasserstein space. *Transactions on Machine Learning Research*, 2022.
- [9] C. Bonet, L. Drumetz, and N. Courty. Sliced-Wasserstein distances and flows on Cartan-Hadamard manifolds. *arXiv preprint arXiv:2403.06560*, 2024.
- [10] C. Bonet, B. Malézieux, A. Rakotomamonjy, L. Drumetz, T. Moreau, M. Kowalski, and N. Courty. Sliced-Wasserstein on symmetric positive definite matrices for m/eeg signals. In *International Conference on Machine Learning*, pages 2777–2805. PMLR, 2023.
- [11] N. Bonneel, J. Rabin, G. Peyré, and H. Pfister. Sliced and Radon Wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 1(51):22–45, 2015.
- [12] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [13] T. Clanuwat, M. Bober-Irizar, A. Kitamoto, A. Lamb, K. Yamamoto, and D. Ha. Deep learning for classical japanese literature. *arXiv preprint arXiv:1812.01718*, 2018.
- [14] G. Cohen, S. Afshar, J. Tapson, and A. Van Schaik. Emnist: Extending mnist to handwritten letters. In 2017 international joint conference on neural networks (IJCNN), pages 2921–2926. IEEE, 2017.
- [15] N. Courty, R. Flamary, A. Habrard, and A. Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. In *Advances in Neural Information Processing Systems*, pages 3730–3739, 2017.
- [16] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996.
- [17] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, 2013.
- [18] I. Deshpande, Y.-T. Hu, R. Sun, A. Pyrros, N. Siddiqui, S. Koyejo, Z. Zhao, D. Forsyth, and A. G. Schwing. Max-sliced Wasserstein distance and its use for GANs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10648–10656, 2019.

- [19] I. Deshpande, Z. Zhang, and A. G. Schwing. Generative modeling using the sliced Wasserstein distance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3483–3491, 2018.
- [20] N. Fournier and A. Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162:707–738, 2015.
- [21] R. C. Garrett, T. Harris, B. Li, and Z. Wang. Validating climate models with spherical convolutional Wasserstein distance. *arXiv preprint arXiv:2401.14657*, 2024.
- [22] E. Heitz, K. Vanhoey, T. Chambon, and L. Belcour. A sliced Wasserstein loss for neural texture synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9412–9420, 2021.
- [23] S. Helgason. The Radon transform on r n. In *Integral Geometry and Radon Transforms*, pages 1–62. Springer, 2011.
- [24] J. J. Hull. A database for handwritten text recognition research. *IEEE Transactions on pattern analysis and machine intelligence*, 16(5):550–554, 1994.
- [25] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, page 0, 2006.
- [26] S. Kolouri, K. Nadjahi, U. Simsekli, R. Badeau, and G. Rohde. Generalized sliced Wasserstein distances. In *Advances in Neural Information Processing Systems*, pages 261–272, 2019.
- [27] S. Kolouri, G. K. Rohde, and H. Hoffmann. Sliced Wasserstein distance for learning Gaussian mixture models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3427–3436, 2018.
- [28] P. Kuchment. Generalized transforms of Radon type and their applications. In *Proceedings of Symposia in Applied Mathematics*, volume 63, page 67, 2006.
- [29] T. Le, K. Nguyen, S. Sun, K. Han, N. Ho, and X. Xie. Diffeomorphic mesh deformation via efficient optimal transport for cortical surface reconstruction. *International Conference on Learning Representations*, 2024.
- [30] T. Le, K. Nguyen, S. Sun, N. Ho, and X. Xie. Integrating efficient optimal transport and functional maps for unsupervised shape correspondence learning. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2024.
- [31] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [32] C.-Y. Lee, T. Batra, M. H. Baig, and D. Ulbricht. Sliced Wasserstein discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10285–10295, 2019.
- [33] Z.-P. Liang and D. C. Munson. Partial Radon transforms. *IEEE transactions on image processing*, 6(10):1467–1469, 1997.
- [34] A. Liutkus, U. Simsekli, S. Majewski, A. Durmus, and F.-R. Stöter. Sliced-Wasserstein flows: Nonparametric generative modeling via optimal transport and diffusions. In *International Conference on Machine Learning*, pages 4104–4113. PMLR, 2019.
- [35] M. Luong, K. Nguyen, N. Ho, R. Haf, D. Phung, and L. Qu. Revisiting deep audio-text retrieval through the lens of transportation. In *The Twelfth International Conference on Learning Representations*, 2024.
- [36] T. Manole, S. Balakrishnan, and L. Wasserman. Minimax confidence intervals for the sliced Wasserstein distance. *Electronic Journal of Statistics*, 16(1):2252–2345, 2022.
- [37] K. Nadjahi, V. De Bortoli, A. Durmus, R. Badeau, and U. Şimşekli. Approximate Bayesian computation with the sliced-Wasserstein distance. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5470–5474. IEEE, 2020.
- [38] K. Nadjahi, A. Durmus, L. Chizat, S. Kolouri, S. Shahrampour, and U. Simsekli. Statistical and topological properties of sliced probability divergences. *Advances in Neural Information Processing Systems*, 33:20802–20812, 2020.
- [39] K. Nguyen, N. Bariletto, and N. Ho. Quasi-monte carlo for 3d sliced Wasserstein. In *The Twelfth International Conference on Learning Representations*, 2024.

- [40] K. Nguyen and N. Ho. Revisiting sliced Wasserstein on images: From vectorization to convolution. *Advances in Neural Information Processing Systems*, 2022.
- [41] K. Nguyen and N. Ho. Energy-based sliced Wasserstein distance. Advances in Neural Information Processing Systems, 2023.
- [42] K. Nguyen, N. Ho, T. Pham, and H. Bui. Distributional sliced-Wasserstein and applications to generative modeling. In *International Conference on Learning Representations*, 2021.
- [43] K. Nguyen, T. Ren, H. Nguyen, L. Rout, T. Nguyen, and N. Ho. Hierarchical sliced Wasserstein distance. *International Conference on Learning Representations*, 2023.
- [44] S. Nietert, R. Sadhu, Z. Goldfeld, and K. Kato. Statistical, robustness, and computational guarantees for sliced Wasserstein distances. *Advances in Neural Information Processing Systems*, 2022.
- [45] O. Pele and M. Werman. Fast and robust earth mover's distances. In 2009 IEEE 12th International Conference on Computer Vision, pages 460–467. IEEE, September 2009.
- [46] S. Peng, C. Jiang, Y. Liao, M. Niemeyer, M. Pollefeys, and A. Geiger. Shape as points: A differentiable poisson solver. Advances in Neural Information Processing Systems, 34:13032– 13044, 2021.
- [47] G. Peyré and M. Cuturi. Computational optimal transport, 2020.
- [48] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [49] M. Quellmalz, R. Beinert, and G. Steidl. Sliced optimal transport on the sphere. *Inverse Problems*, 39(10):105005, 2023.
- [50] M. Quellmalz, L. Buecher, and G. Steidl. Parallelly sliced optimal transport on spheres and on the rotation group. *arXiv preprint arXiv:2401.16896*, 2024.
- [51] F. Rouviere. Nonlinear Radon and Fourier transforms, 2015.
- [52] R. M. Rustamov and S. Majumdar. Intrinsic sliced Wasserstein distances for comparing collections of probability distributions on manifolds and graphs. In *International Conference* on *Machine Learning*, pages 29388–29415. PMLR, 2023.
- [53] I. Tolstikhin, O. Bousquet, S. Gelly, and B. Schoelkopf. Wasserstein auto-encoders. In *International Conference on Learning Representations*, 2018.
- [54] H. Tran, Y. Bai, A. Kothapalli, A. Shahbazi, X. Liu, R. D. Martin, and S. Kolouri. Stereographic spherical sliced Wasserstein distances. *International Conference on Machine Learning*, 2024.
- [55] C. Villani. Optimal transport: old and new, volume 338. Springer Science & Business Media, 2008.
- [56] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [57] J. Xu, H. Zhou, C. Gan, Z. Zheng, and L. Li. Vocabulary learning via optimal transport for neural machine translation. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 7361–7373, 2021.
- [58] M. Yi and S. Liu. Sliced Wasserstein variational inference. In Fourth Symposium on Advances in Approximate Bayesian Inference, 2021.
- [59] Q.-Y. Zhou, J. Park, and V. Koltun. Open3D: A modern library for 3D data processing. *arXiv*:1801.09847, 2018.

Supplement to "Hierarchical Hybrid Sliced Wasserstein: A Scalable Metric for Heterogeneous Joint Distributions"

We first provide skipped proofs in the main paper in Appendix A. We then provide some additional materials including additional background and extended definitions in Appendix B. After that, we discuss some related works in Appendix C. We report additional experimental results in Appendix D. Finally, we report computational infrastructure in Appendix E.

A Proofs

A.1 Proof of Proposition 1

For any t, θ, y , we are given $(\mathcal{PGR}f_1)(t, \theta, y) = (\mathcal{PGR}f_2)(t, \theta, y)$. By Definition 1, we have:

$$\int_{\mathbb{R}^{d_1}} f_1(x,y)\delta(t-g(x,\theta))dx = \int_{\mathbb{R}^{d_1}} f_2(x,y)\delta(t-g(x,\theta))dx.$$

For any $\varepsilon \in \mathbb{R}^{d_2}$, we have:

$$\int_{\mathbb{R}^{d_2}} \int_{\mathbb{R}^{d_1}} f_1(x,y) \delta(t-g(x,\theta)) e^{-i2\pi \langle \varepsilon,y \rangle} dx dy = \int_{\mathbb{R}^{d_2}} \int_{\mathbb{R}^{d_1}} f_2(x,y) \delta(t-g(x,\theta)) e^{-i2\pi \langle \varepsilon,y \rangle} dx dy.$$

Applying the Fubini's theorem, we have:

$$\int_{\mathbb{R}^{d_1}} f_1(x,y) \int_{\mathbb{R}^{d_2}} e^{-i2\pi \langle \varepsilon,y \rangle} dy \delta(t-g(x,\theta)) dx = \int_{\mathbb{R}^{d_1}} \int_{\mathbb{R}^{d_2}} f_2(x,y) e^{-i2\pi \langle \varepsilon,y \rangle} dy \delta(t-g(x,\theta)) dx,$$

which is

$$\left(\mathcal{GR}\int_{\mathbb{R}^{d_2}}f_1(x,y)e^{-i2\pi\langle\varepsilon,y\rangle}dy\right)=\left(\mathcal{GR}\int_{\mathbb{R}^{d_2}}f_2(x,y)e^{-i2\pi\langle\varepsilon,y\rangle}dy\right).$$

By the injectivity of GRT, we have:

$$\int_{\mathbb{R}^{d_2}} f_1(x,y) e^{-i2\pi\langle \varepsilon, y \rangle} dy = \int_{\mathbb{R}^{d_2}} f_2(x,y) e^{-i2\pi\langle \varepsilon, y \rangle} dy.$$

Then, for any $\epsilon \in \mathbb{R}^{d_1}$, we have

$$\int_{\mathbb{R}^{d_1}} \int_{\mathbb{R}^{d_2}} f_1(x,y) e^{-i2\pi\langle \varepsilon,y\rangle} e^{-i2\pi\langle \varepsilon,x\rangle} dy dx = \int_{\mathbb{R}^{d_1}} \int_{\mathbb{R}^{d_2}} f_2(x,y) e^{-i2\pi\langle \varepsilon,y\rangle} e^{-i2\pi\langle \varepsilon,x\rangle} dy dx.$$

which is $(\mathcal{F}f_1(x,y)) = (\mathcal{F}f_2(x,y))$ with \mathcal{F} denotes the Fourier transform. By the injectivity of the Fourier Transform, we have $f_1(x,y) = f_2(x,y)$ for any x,y, which concludes the proof.

A.2 Proof of Proposition 2

We first show that HHRT is the composition of PGRT and PRT. We have

$$\begin{split} &(\mathcal{PR}(\mathcal{PGR}(\mathcal{PGR}f)))(t,\theta_{1},\theta_{2},\psi) \\ &= \int_{\mathbb{R}^{2}} \int_{\mathbb{R}^{d_{1}}} \int_{\mathbb{R}^{d_{2}}} f(x,y) \delta(t_{1} - g_{1}(x,\theta_{1})) \delta(t_{2} - g_{2}(y,\theta_{2})) \delta(t - \psi_{1}t_{1} - \psi_{2}t_{2}) dx dy dt_{1} dt_{2} \\ &= \int_{\mathbb{R}^{d_{1}}} \int_{\mathbb{R}^{d_{2}}} f(x,y) \int_{\mathbb{R}^{2}} \delta(t_{1} - g_{1}(x,\theta_{1})) \delta(t_{2} - g_{2}(y,\theta_{2})) \delta(t - \psi_{1}t_{1} - \psi_{2}t_{2}) dt_{1} dt_{2} dx dy \\ &= \int_{\mathbb{R}^{d_{1}}} \int_{\mathbb{R}^{d_{2}}} f(x,y) \delta(t - \psi_{1}g_{1}(x,\theta_{1}) - \psi_{2}g_{2}(y,\theta_{2})) dx dy \\ &= (\mathcal{H}\mathcal{H}\mathcal{R}f)(t,\theta_{1},\theta_{2},\psi). \end{split}$$

For any $t, \theta_1, \theta_2, \psi$, we are given $(\mathcal{HHR}f_1)(t, \theta_1, \theta_2, \psi) = (\mathcal{HHR}f_2)(t, \theta_1, \theta_2, \psi)$, which is equivalent to:

$$(\mathcal{PR}(\mathcal{PGR}(\mathcal{PGR}f_1)))(t,\theta_1,\theta_2,\psi) = (\mathcal{PR}(\mathcal{PGR}(\mathcal{PGR}f_2)))(t,\theta_1,\theta_2,\psi).$$

By the injectivity of the PRT and the PGRT, we obtain $f_1(x,y) = f_2(x,y)$ for any x,y which completes the proof.

A.3 Proof of Theorem 1

To prove that the hierarchical hybrid sliced Wasserstein $H2SW_p(\cdot,\cdot;c,g_1,g_2)$ is a metric on the space of distributions on $\mathcal{P}(\mathbb{R}^{d_1}\times\mathbb{R}^{d_2})$ for any $p\geq 1$, ground metric c, and defining functions g_1,g_2 , we need to show that it satisfies non-negativity, symmetry, triangle inequality, and identity of indiscernible.

Non-Negativity. Since $W_p^p(\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\mu,\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\nu;c)\geq 0$ [47] for any θ_1,θ_2,ψ , we have:

$$\mathbb{E}_{(\theta_1,\theta_2,\psi)\sim\mathcal{U}(\Omega_1\times\Omega_2\times\mathbb{S})}[\mathbf{W}_p^p(\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\mu,\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\nu;c)]\geq 0,$$

which means that $H2SW_p(\mu, \nu; c, g_1, g_2) \ge 0$ for any μ and ν .

Symmetry. Since we have the symmetry of the Wasserstein distance $W_p^p(\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\mu,\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\nu;c)=W_p^p(\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\nu,\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\nu;t)$ [47] for any θ_1,θ_2,ψ , we have:

$$\begin{split} &\mathbb{E}_{(\theta_1,\theta_2,\psi)\sim\mathcal{U}(\Omega_1\times\Omega_2\times\mathbb{S})}[\mathbf{W}_p^p(\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\mu,\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\nu;c)]\\ &=\mathbb{E}_{(\theta_1,\theta_2,\psi)\sim\mathcal{U}(\Omega_1\times\Omega_2\times\mathbb{S})}[\mathbf{W}_p^p(\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\nu,\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\mu;c)], \end{split}$$

which means that $H2SW_p(\mu, \nu; c, g_1, g_2) = H2SW_p(\nu, \mu; c, g_1, g_2)$ any μ and ν .

Triangle Inequality. Given c to be a valid metric on \mathbb{R} , we can use the triangle inequality of the Wasserstein distance. For any distributions $\mu_1, \mu_2, \mu_3 \in \mathcal{P}(\mathbb{R}^{d_1} \times \mathbb{R}^{d_2})$, we have:

$$\begin{aligned} \operatorname{H2SW}_{p}(\mu_{1},\mu_{2};c,g_{1},g_{2}) &= \left(\mathbb{E}_{(\theta_{1},\theta_{2},\psi)\sim\mathcal{U}(\Omega_{1}\times\Omega_{2}\times\mathbb{S})}[\operatorname{W}_{p}^{p}(\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu_{1},\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu_{2};c)]\right)^{\frac{1}{p}} \\ &\leq \left(\mathbb{E}_{(\theta_{1},\theta_{2},\psi)\sim\mathcal{U}(\Omega_{1}\times\Omega_{2}\times\mathbb{S})}[(\operatorname{W}_{p}(\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu_{1},\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu_{3};c)\right. \\ &\left. + \operatorname{W}_{p}(\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu_{3},\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu_{2};c))^{p}]\right)^{\frac{1}{p}} \\ &\leq \left(\mathbb{E}_{(\theta_{1},\theta_{2},\psi)\sim\mathcal{U}(\Omega_{1}\times\Omega_{2}\times\mathbb{S})}[\operatorname{W}_{p}^{p}(\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu_{1},\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu_{3};c)]\right)^{\frac{1}{p}} \\ &+ \left(\mathbb{E}_{(\theta_{1},\theta_{2},\psi)\sim\mathcal{U}(\Omega_{1}\times\Omega_{2}\times\mathbb{S})}[\operatorname{W}_{p}^{p}(\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu_{3},\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu_{2};c)]\right)^{\frac{1}{p}} \\ &= \operatorname{H2SW}_{p}(\mu_{1},\mu_{3};c,q_{1},q_{2}) + \operatorname{H2SW}_{p}(\mu_{3},\mu_{2};c,q_{1},q_{2}), \end{aligned}$$

where the final inequality is due to Minkowski's inequality. Therefore, we complete the proof for the triangle inequality of the hierarchical hybrid sliced Wasserstein.

Identity of indiscernible. For any $p\geq 1$, ground metric c, and g_1,g_2 , when $\mu=\nu$, we have $\mathcal{H}\mathcal{H}\mathcal{R}^{g_1,g_2}_{\theta_1,\theta_2,\psi}\sharp\mu=(\mathcal{H}\mathcal{H}\mathcal{R}^{g_1,g_2}_{\theta_1,\theta_2,\psi}\sharp\nu)$. Therefore, we have $W^p_p(\mathcal{H}\mathcal{H}\mathcal{R}^{g_1,g_2}_{\theta_1,\theta_2,\psi}\sharp\mu_1,\mathcal{H}\mathcal{H}\mathcal{R}^{g_1,g_2}_{\theta_1,\theta_2,\psi}\sharp\mu_2;c)=0$ which leads to $\mathrm{H2SW}_p(\mu,\nu;c,g_1,g_2)=0$. Now, assume that $\mathrm{H2SW}_p(\mu,\nu;c,g_1,g_2)=0$, then $W^p_p(\mathcal{H}\mathcal{H}\mathcal{R}^{g_1,g_2}_{\theta_1,\theta_2,\psi}\sharp\mu_1,\mathcal{H}\mathcal{H}\mathcal{R}^{g_1,g_2}_{\theta_1,\theta_2,\psi}\sharp\mu_2;c)=0$ for almost everywhere $\theta_1\in\Omega_1,\theta_2\in\Omega_2,\psi\in\mathbb{S}$. By applying the identity property of the Wasserstein distance, we have $\mathcal{H}\mathcal{H}\mathcal{R}^{g_1,g_2}_{\theta_1,\theta_2,\psi}\sharp\mu=(\mathcal{H}\mathcal{H}\mathcal{R}^{g_1,g_2}_{\theta_1,\theta_2,\psi}\sharp\nu)$ for almost everywhere $\theta_1\in\Omega_1,\theta_2\in\Omega_2,\psi\in\mathbb{S}$. Since the HHRT is injective (proved in Proposition 2), we obtain $\mu=\nu$.

A.4 Proof of Proposition 3

(i) For any
$$p\geq 1$$
, $c(x,y)=|x-y|$, and $\mu,\nu\in\mathcal{P}(\mathbb{R}^{d_1}\times\mathbb{R}^{d_2})$, we have:
$$\text{H2SW}_p(\mu,\nu;c,g_1,g_2)$$

$$\begin{split} &= \left(\mathbb{E}_{(\theta_1,\theta_2,\psi) \sim \mathcal{U}(\Omega_1 \times \Omega_2 \times \mathbb{S})} [\mathbf{W}_p^p(\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2} \sharp \mu, \mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2} \sharp \nu; c)] \right)^{\frac{1}{p}} \\ &= \left(\mathbb{E}\left[\inf_{\pi \in \Pi(\mu,\nu)} \int |\psi_1(g_1(\theta_1,x_1) - g_1(\theta_1,y_1)) + \psi_2(g_2(\theta_2,x_2) - g_1(\theta_2,y_2))|^p d\pi(x_1,x_2,y_1,y_2)\right] \right)^{\frac{1}{p}} \end{split}$$

By applying the Cauchy-Schwartz inequality, we have: $\text{H2SW}_p(\mu, \nu; c, g_1, g_2)$

$$\begin{split} &\leq \left(\mathbb{E}\left[\inf_{\pi \in \Pi(\mu,\nu)} \int (\sqrt{\psi_1^2 + \psi_2^2})^p (\sqrt{(g_1(\theta_1,x_1) - g_1(\theta_1,y_1))^2 + (g_2(\theta_2,x_2) - g_2(\theta_2,y_2))^2})^p d\pi(x_1,x_2,y_1,y_2)\right]\right)^{\frac{1}{p}} \\ &\leq \left(\mathbb{E}\left[\inf_{\pi \in \Pi(\mu,\nu)} \int (|g_1(\theta_1,x_1) - g_1(\theta_1,y_1)| + |g_2(\theta_2,x_2) - g_2(\theta_2,y_2)|)^p d\pi(x_1,x_2,y_1,y_2)\right]\right)^{\frac{1}{p}} \\ &\leq \left(\mathbb{E}\left[\inf_{\pi \in \Pi(\mu,\nu)} \int |g_1(\theta_1,x_1) - g_1(\theta_1,y_1)|^p d\pi(x_1,x_2,y_1,y_2)\right]\right)^{\frac{1}{p}} \\ &+ \left(\mathbb{E}\left[\inf_{\pi \in \Pi(\mu,\nu)} \int |g_2(\theta_2,x_2) - g_2(\theta_2,y_2)|^p d\pi(x_1,x_2,y_1,y_2)\right]\right)^{\frac{1}{p}} \\ &= \left(\mathbb{E}\left[\inf_{\pi \in \Pi(\mu_1,\nu_1)} \int |g_1(\theta_1,x_1) - g_1(\theta_1,y_1)|^p d\pi(x_1,y_1)\right]\right)^{\frac{1}{p}} \\ &+ \left(\mathbb{E}\left[\inf_{\pi \in \Pi(\mu_2,\nu_2)} \int |g_2(\theta_2,x_2) - g_2(\theta_2,y_2)|^p d\pi(x_2,y_2)\right]\right)^{\frac{1}{p}} \\ &= \mathrm{GSW}_p(\mu_1,\nu_1;g_1,c) + \mathrm{GSW}_p(\mu_2,\nu_2;g_2,c), \end{split}$$

where the last inequality is due to the Minkowski's inequality.

(ii) From (i), we have $\mathrm{H2SW}_p(\mu,\nu;c,g_1,g_2) \leq \mathrm{GSW}_p(\mu_1,\nu_1;g_1,c) + \mathrm{GSW}_p(\mu_2,\nu_2;g_2,c)$. When, g_1,g_2 , and c(x,y)=|x-y| are linear defining functions, we have:

$$\begin{aligned} \text{GSW}_{p}(\mu_{1}, \nu_{1}; g_{1}, c) &= \left(\mathbb{E} \left[\inf_{\pi \in \Pi(\mu_{1}, \nu_{1})} \int (|\theta^{\top} x_{1} - \theta^{\top} y_{1}|^{p} d\pi(x_{1}, y_{1}) \right] \right)^{\frac{1}{p}} \\ &\leq \left(\mathbb{E} \left[\inf_{\pi \in \Pi(\mu_{1}, \nu_{1})} \int (||\theta||_{2} ||x_{1} - y_{1}||_{2}^{p} d\pi(x_{1}, y_{1}) \right] \right)^{\frac{1}{p}} \\ &\leq \left(\mathbb{E} \left[\inf_{\pi \in \Pi(\mu_{1}, \nu_{1})} \int (||x_{1} - y_{1}||^{p} d\pi(x_{1}, y_{1}) \right] \right)^{\frac{1}{p}} \\ &= \left(\inf_{\pi \in \Pi(\mu_{1}, \nu_{1})} \int (||x_{1} - y_{1}||^{p} d\pi(x_{1}, y_{1}) \right)^{\frac{1}{p}} \\ &= W_{p}(\mu_{1}, \nu_{1}; c). \end{aligned}$$

Similarly, we have $\mathrm{GSW}_p(\mu_2,\nu_2;g_1,c) \leq W_p(\mu_2,\nu_2;c)$. Therefore, we obtain the proof of $\mathrm{H2SW}_p(\mu,\nu;c,g_1,g_2) \leq W_p(\mu_1,\nu_1;c) + W_p(\mu_1,\nu_1;c)$.

(iii) When g_1 , g_2 are linear defining functions, we have:

 $\text{H2SW}_p(\mu, \nu; c, g_1, g_2)$

$$\leq \left(\mathbb{E} \left[\inf_{\pi \in \Pi(\mu,\nu)} \int (|\theta_{1}^{\top}x_{1} - \theta_{1}^{\top}y_{1})| + |\theta_{2}^{\top}x_{2} - \theta_{2}^{\top}y_{2}|)^{p} d\pi(x_{1}, x_{2}, y_{1}, y_{2}) \right] \right)^{\frac{1}{p}} \\
\leq \left(\mathbb{E} \left[\inf_{\pi \in \Pi(\mu,\nu)} \int (|\theta_{1}^{\top}x_{1} - \theta_{1}^{\top}y_{1})| + |\theta_{2}^{\top}x_{2} - \theta_{2}^{\top}y_{2}|)^{p} d\pi(x_{1}, x_{2}, y_{1}, y_{2}) \right] \right)^{\frac{1}{p}} \\
\leq \left(\mathbb{E} \left[\inf_{\pi \in \Pi(\mu,\nu)} \int (|x_{1} - y_{1})| + |x_{2} - y_{2}|)^{p} d\pi(x_{1}, x_{2}, y_{1}, y_{2}) \right] \right)^{\frac{1}{p}} \\
= \left(\inf_{\pi \in \Pi(\mu,\nu)} \int (|x_{1} - y_{1})| + |x_{2} - y_{2}|)^{p} d\pi(x_{1}, x_{2}, y_{1}, y_{2}) \right)^{\frac{1}{p}}$$

When p = 1, we obtain:

$$\begin{aligned} \text{H2SW}_1(\mu,\nu;c,g_1,g_2) &\leq \left(\inf_{\pi \in \Pi(\mu,\nu)} \int (|x_1-y_1|) + |x_2-y_2|) d\pi(x_1,x_2,y_1,y_2)\right)^{\frac{1}{p}} \\ &= W_1(\mu,\nu;c,c), \end{aligned}$$

which completes the proof.

A.5 Proof of Proposition 4

Let $p \ge 1$, c(x,y) = |x-y|, $\mu \in \mathcal{P}(\mathbb{R})$ with the corresponding empirical distribution μ_n , we assume that there exists q > p such that the q-th order moment of μ i.e, $M_q(\mu) = \int_{\mathbb{R}} |x|^q d\mu(x)$, is bounded by $B < \infty$. From Theorem 1 in [20], there exists a constant $C_{p,q}$ such that:

$$\mathbb{E}\left[W_p^p(\mu_n, \mu; c)\right] \le C_{p,q} B \begin{cases} n^{-1/2} \text{ if } q > 2p, \\ n^{-1/2} \log(n)^{\frac{1}{p}} \text{ if } q = 2p, \\ n^{-(q-p)/q} \text{ if } q \in (p, 2p). \end{cases}$$

We show that $\mathcal{HHR}^{g_1,g_2}_{\theta_1,\theta_2,\psi}\sharp\mu$ has finite bounded moments. In particular, we have:

$$M_{k}(\mathcal{HHR}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu) = \int_{\mathbb{R}} |t|^{k} d(\mathcal{HHR}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu)(t)$$

$$= \int_{\mathbb{R}^{d_{1}}\times\mathbb{R}^{d_{2}}} |\psi_{1}g_{1}(\theta_{1},x_{1}) + \psi_{2}g_{2}(\theta_{2},x_{2})|^{k} d\mu(x_{1},x_{2})$$

$$\leq \int_{\mathbb{R}^{d_{1}}\times\mathbb{R}^{d_{2}}} (\psi_{1}^{2} + \psi_{2}^{2})^{k/2} (g_{1}(\theta_{1},x_{1})^{2} + g_{2}(\theta_{2},x_{2})^{2})^{k/2} d\mu(x_{1},x_{2})$$

$$\leq \int_{\mathbb{R}^{d_{1}}\times\mathbb{R}^{d_{2}}} (|g_{1}(\theta_{1},x_{1})| + |g_{2}(\theta_{2},x_{2})|)^{k} d\mu(x_{1},x_{2}),$$

where the first inequality is due to the Cauchy-Schwarz inequality and the second inequality is due to the fact that $\|x\|_2 \leq |x|$. For the linear defining functions $g(\theta,x) = \theta^\top x$, we have $|g(\theta,x)| = |\theta^\top x| \leq \|x\|_1$. For the circular defining functions $g(\theta,x) = \|x - r\theta\|_2 \leq \|x - r\theta\|_1 \leq \|x\|_1 + \|r\theta\|_1 \leq \|x\|_1 + r$. Therefore, we have:

$$\begin{split} M_k(\mathcal{HHR}_{\theta_1,\theta_2,\psi}^{g_1,g_2}\sharp\mu) &\leq \int_{\mathbb{R}^{d_1}\times\mathbb{R}^{d_2}} (|x_1|+|x_2|+C_{g_1,g_2})^k d\mu(x_1,x_2) \\ &= \int_{\mathbb{R}^{d_1}\times\mathbb{R}^{d_2}} \sum_{i=0}^k k^i (|x_1|+|x_2|)^i C_{g_1,g_2}^{k-i} d\mu(x_1,x_2) \\ &= \sum_{i=0}^k k^i C_{g_1,g_2}^{k-i} \int_{\mathbb{R}^{d_1}\times\mathbb{R}^{d_2}} (|x_1|+|x_2|)^i d\mu(x_1,x_2) \\ &\leq \sum_{i=0}^k k^i C_{g_1,g_2}^{k-i} M_i(\mu), \end{split}$$

where $C_{g_1,g_2}=0$ if g_1,g_2 are linear, $C_{g_1,g_2}=r$ if g_1 and g_2 are linear and circular respectively (exchangeable), and $C_{g_1,g_2}=2r$ if both g_1 and g_2 are circular.

Now, using the triangle inequality of H2SW (Theorem 1), we have:

$$\begin{split} \mathbb{E} & \left| \mathsf{H2SW}_p(\mu_n, \nu_n; c, g_1, g_2) - \mathsf{H2SW}_p(\mu, \nu; c, g_1, g_2) \right| \\ & \leq \mathbb{E} \left| \mathsf{H2SW}_p(\mu, \mu_n; c, g_1, g_2) + \mathsf{H2SW}_p(\nu, \nu_n; c, g_1, g_2) \right| \\ & \leq \mathbb{E} \left| \mathsf{H2SW}_p(\mu, \mu_n; c, g_1, g_2) \right| + \mathbb{E} \left| \mathsf{H2SW}_p(\nu, \nu_n; c, g_1, g_2) \right| \\ & \leq \left(\mathbb{E} \left| \mathsf{H2SW}_p^p(\mu, \mu_n; c, g_1, g_2) \right| \right)^{\frac{1}{p}} + \left(\mathbb{E} \left| \mathsf{H2SW}_p^p(\nu, \nu_n; c, g_1, g_2) \right| \right)^{\frac{1}{p}}, \end{split}$$

where the last inequality is due to Holder's inequality. Combining with previous results, we obtain:

$$\begin{split} \mathbb{E} \left| \mathsf{H2SW}_p(\mu_n, \nu_n; c, g_1, g_2) - \mathsf{H2SW}_p(\mu, \nu; c, g_1, g_2) \right| \\ & \leq C_{p,q}^{\frac{1}{p}} \left(\sum_{i=0}^q q^i C_{g_1, g_2}^{q-i}(M_i(\mu) + M_i(\nu)) \right)^{\frac{1}{p}} \begin{cases} n^{-1/2p} \text{ if } q > 2p, \\ n^{-1/2p} \log(n)^{\frac{1}{p}} \text{ if } q = 2p, \\ n^{-(q-p)/pq} \text{ if } q \in (p, 2p), \end{cases} \end{split}$$

which completes the proof.

A.6 Proof of Proposition 5

For any $p \ge 1$, and $\mu, \nu \in \mathcal{P}(\mathbb{R}^{d_1} \times \mathbb{R}^{d_2})$, using the Holder's inequality, we have:

$$\begin{split} & \mathbb{E}|\widehat{\mathbf{H2SW}}_{p}^{p}(\mu,\nu;c,g_{1},g_{2},L) - \mathbf{H2SW}_{p}^{p}(\mu,\nu;c,g_{1},g_{2})| \\ & \leq \left(\mathbb{E}|\widehat{\mathbf{H2SW}}_{p}^{p}(\mu,\nu;c,g_{1},g_{2},L) - \mathbf{H2SW}_{p}^{p}(\mu,\nu;c,g_{1},g_{2})|^{2}\right)^{\frac{1}{2}} \\ & = \left(\mathbb{E}\left|\frac{1}{L}\sum_{l=1}^{L}\mathbf{W}_{p}^{p}(\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1l},\theta_{2l},\psi_{l}}^{g_{1},g_{2}}\sharp\mu,\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1l},\theta_{2l},\psi_{l}}^{g_{1},g_{2}}\sharp\nu;c) - \mathbb{E}\left[\mathbf{W}_{p}^{p}(\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1},\theta_{2},\psi}^{g_{1},g_{2}}\sharp\mu,\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1l},\theta_{2l},\psi_{l}}^{g_{1},g_{2}}\sharp\nu;c)\right]\right|^{\frac{1}{2}} \\ & = \left(Var\left[\frac{1}{L}\sum_{l=1}^{L}\mathbf{W}_{p}^{p}(\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1l},\theta_{2l},\psi_{l}}^{g_{1},g_{2}}\sharp\mu,\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1l},\theta_{2l},\psi_{l}}^{g_{1},g_{2}}\sharp\nu;c)\right]\right)^{\frac{1}{2}} \\ & = \frac{1}{\sqrt{L}}Var\left[\mathbf{W}_{p}^{p}(\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1l},\theta_{2l},\psi_{l}}^{g_{1},g_{2}}\sharp\mu,\mathcal{H}\mathcal{H}\mathcal{R}_{\theta_{1l},\theta_{2l},\psi_{l}}^{g_{1},g_{2}}\sharp\nu;c)\right]^{\frac{1}{2}}, \end{split}$$

which completes the proof.

B Additional Materials

HHRT with more than two marginals. We now extend the definition of HHRT to K>2 mariginals. **Definition 4** (Hierarchical Hybrid Radon Transform). Given $K\geq 2$, given defining functions $\{g_k:\mathbb{R}^{d_k}\times\Omega_i\to\mathbb{R}\}_{i=k}^K$, the Hierarchical Hybrid Radon Transform $\mathcal{HHR}:\mathbb{L}_1(\mathbb{R}^{d_1}\times\ldots\times\mathbb{R}^{d_K})\to\mathbb{L}_1\left(\mathbb{R}\times\Omega_1\ldots\times\Omega_K\times\mathbb{S}^{K-1}\right)$ is defined as:

$$(\mathcal{HHR}f)(t,\theta_1,\dots,\theta_K,\psi) = \int_{\mathbb{R}^{d_1}\times\dots\times\mathbb{R}^{d_K}} f(x_1,\dots,x_K)\delta\left(t - \sum_{k=1}^K \psi_k g_k(x_k,\theta_k)\right) dx_1\dots dx_K.$$
 (12)

H2SW with more than two marginals. From the new definition of HRRT on K > 2 mariginals, we now can define H2SW between joint distributions with K mariginals.

Definition 5. For $p \geq 1, K \geq 2$, defining functions g_1, \ldots, g_K , the hierarchical hybrid sliced Wasserstein-p (H2SW) distance between two distributions $\mu \in \mathcal{P}(\mathcal{X}_1 \times \ldots \times \mathcal{X}_K)$ and $\nu \in \mathcal{P}(\mathcal{Y}_1 \times \ldots \times \mathcal{Y}_K)$ with an one-dimensional ground metric $c : \mathbb{R} \times \mathbb{R} \to \mathbb{R}^+$ is defined as:

$$H2SW_{p}^{p}(\mu,\nu;c,g_{1},\ldots,g_{K})$$

$$= \mathbb{E}_{(\theta_{1},\ldots,\theta_{K},\psi)\sim\mathcal{U}(\Omega_{1}\times\ldots\times\Omega_{K}\times\mathbb{S}^{K-1})}[W_{p}^{p}(\mathcal{HHR}_{\theta_{1},\ldots,\theta_{K},\psi}^{g_{1},\ldots,g_{K}}\sharp\mu,\mathcal{HHR}_{\theta_{1},\ldots,\theta_{K},\psi}^{g_{1},\ldots,g_{K}}\sharp\nu;c)], \qquad (13)$$

where $\mathcal{HHR}^{g_1,\ldots,g_K}_{\theta_1,\ldots,\theta_K,\psi}\sharp\mu$ and $\mathcal{HHR}^{g_1,\ldots,g_K}_{\theta_1,\ldots,\theta_K,\psi}\sharp\nu$ are the one-dimensional push-forward distributions created by applying HHRT.

Lorentz Model and Busemann function. The Lorentz model $\mathbb{L}^d \in \mathbb{R}^{d+1}$ of a d-dimensional hyperbolic space is [7]:

$$\mathbb{L}^{d} = \left\{ (x_{1}, \dots, x_{d}) \in \mathbb{R}^{d+1}, -x_{0}y_{0} + \sum_{i=1}^{d} x_{i}y_{i} = -1, x_{0} > 0 \right\}.$$

Given a direction $\theta \in T_{x_0} \mathbb{L}^d \cap \mathbb{S}^d$, $x \in \mathbb{L}^d$, the Busemann function is:

$$B(x, \theta) = \log(-\langle x, x_0 + \theta \rangle).$$

Busemann function on product Hadamard manifolds. For distributions supports on the product of $K \ge 2$ Hadamard manifolds with the corresponding Busemann functions B_1, \ldots, B_K , we have a Busemann function of the product manifolds is:

$$B(x_1, \dots, x_K, \theta_1, \dots, \theta_K) = \sum_{k=1}^K \lambda_k B_k(x_k, \theta_k),$$

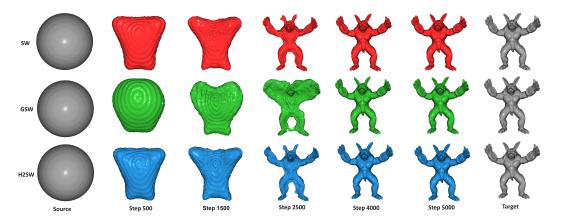


Figure 6: Visualization of deformation from the sphere mesh to the Armadillo mesh with L=100.

for $(\lambda_1,\ldots,\lambda_K)\in\mathbb{S}^{K-1}$. The Cartan-Hyperbolic Sliced-Wasserstein distance use a fixed value of $(\lambda_1,\ldots,\lambda_K)$ e.g., $(\lambda_1,\ldots,\lambda_K)=(1/\sqrt{K},\ldots,1/\sqrt{K})$ (see 2). In our proposed H2SW, we treat $(\lambda_1,\ldots,\lambda_K)$ as a random variable follows $\mathcal{U}(\mathbb{S}^{K-1})$ and the value of H2SW is defined as the mean of such random variable.

C Related Works

HHRT and Generalized Radon Transform. HHRT can be also seen as a special case of GRT [3] with the defining function $g(x,\theta) = \psi_1 g_1(x_1,\theta_1) + \psi_2 g_2(y,\theta_2)$ with $x = (x_1,x_2)$ and $\theta = (\theta_1,\theta_2,\psi)$ ($\Omega = \Omega_1 \times \Omega_2 \times \mathbb{S}$). However, without approaching via the hierarchical construction, the injectivity of the transform might be a challenge to obtain.

HHRT and Hierarchical Radon Transform. Hierarchical Radon Transform (HRT) [43] is the composition of Partial Radon Transform and Overparameterized Radon Transform, which is designed specifically for reducing projection complexity when using Monte Carlo estimation. Moreover, HRT is introduced with linear projection and does not focus on the problem of comparing heterogeneous joint distributions. In contrast to HRT, the proposed HHRT is the composition of multiple partial Generalized Radon Transform and Partial Random Transform, which is suitable for comparing heterogeneous joint distributions.

HHRT and convolution slicers. Convolution slicers [40] are introduced to project an image into a scalar. It can be viewed as a Hierarchical Partial Radon Transform i.e., small parts of the image are transformed first, then be aggregated later. Although convolution slicers can separate global and local information as HHRT, they focus on the domain of images only and have not been proven to be injective. Again, HHRT is designed to compare heterogeneous joint distributions and is proven to be injective in Proposition 2. As a result, H2SW is a valid metric while convolution sliced Wasserstein [40] is only a pseudo metric. Moreover, H2SW can also use convolution slicers when having marginal domains as images.

D Additional Experiments

3D Mesh Deformation. As mentioned in the main text, we present the deformation visualization to the Armadillo mesh with L=100 in Figure 6, and the deformation visualization to the Stanford Bunny o mesh with L=10 and L=100 in Figure 3-7 in turn. The quantitative result for the Armadillo mesh is given in Table 2. Here, we set the step size to 0.1. From these results, we see that the proposed H2SW gives the best flow deformation flow in general. The performance gap is especially larger when L=10 i.e., having a small number of projections.

²https://github.com/clbonet/Sliced-Wasserstein_Distances_and_Flows_on_ Cartan-Hadamard_Manifolds/blob/0eb05450e7f9f27586d0ddb1ce6e58f07eb75786/Experiments/xp_otdd/OTDD_SW.ipynb

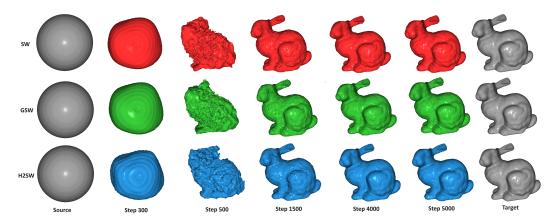


Figure 7: Visualization of deformation from the sphere mesh to the Stanford Bunny mesh with L=100.

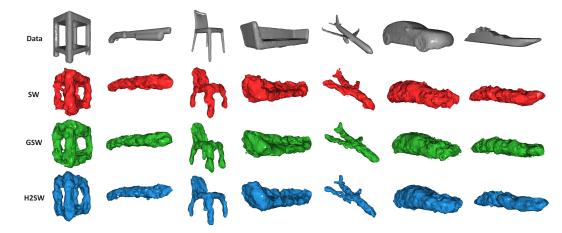


Figure 8: Visualization of some randomly selected reconstruction meshes from autoencoders trained by SW, GSW, and H2SW in turn with the number of projections L=100 at epoch 500.

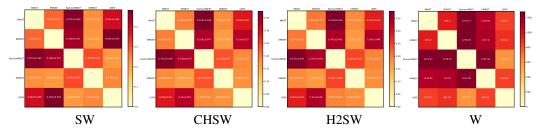


Figure 9: Cost matrices between datasets from SW, CHSW, and H2SW with L=100.

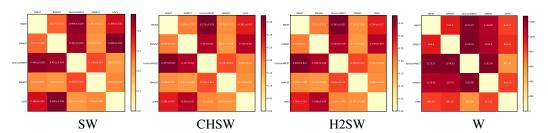


Figure 10: Cost matrices between datasets from SW, CHSW, and H2SW with L=500.

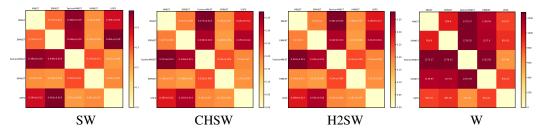


Figure 11: Cost matrices between datasets from SW, CHSW, and H2SW with L = 1000.

Deep 3D mesh autoencoder. We first report the neural network architectures that we use in the experiments.

- The encoder: Conv1d(6,64,1) \rightarrow BatchNorm1d \rightarrow LeakyReLU(0.2) \rightarrow Conv1d(64, 128, 1) \rightarrow BatchNorm1d \rightarrow LeakyReLU(0.2) \rightarrow Conv1d(128, 256, 1) \rightarrow BatchNorm1d \rightarrow LeakyReLU(0.2) \rightarrow Conv1d(256, 512, 1) \rightarrow BatchNorm1d \rightarrow LeakyReLU(0.2) \rightarrow Conv1d(512, 1024, 1) \rightarrow BatchNorm1d \rightarrow LeakyReLU(0.2) \rightarrow Max-Pooling \rightarrow Linear(1024, 1024).
- The decoder: Linear(1024, 1024) → BatchNorm1d → LeakyReLU(0.2) → Linear(1024, 2048) → BatchNorm1d → LeakyReLU(0.2) → Linear(2048, 4096) → BatchNorm1d → LeakyReLU(0.2) → Linear(2048, 2048*6). The output of the decoder is the concatenation of the location and normal vector. We normalize the normal vector to the unit-sphere.

As mentioned in the main text, we report the reconstruction of randomly selected meshes for L=100 at epoch 500 in Figure 8. We see that the reconstructed meshes at epoch 500 are visually worse than the reconstructed meshes at epoch 2000. Therefore, the joint Wasserstein distances in Table 3 are consistent with the qualitative results.

Dataset Comparison. We follow the same procedure in Section 6.2 in [9]. We refer the reader to the reference for a detailed description. Here, we show the cross-dataset cost matrices with the number of projections L=100 in Figure 9, L=500 in Figure 10, and L=1000 Figure 11.

E Computational Infrastructure

For the non-deep-learning experiments, we use a HP Omen 25L desktop for conducting experiments. For 3D mesh autoencoder experiments, we use a single NVIDIA A100 GPU.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: As in the abstract and introduction, we focus on designing a sliced Wasserstein variant for heterogeneous joint distributions.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
 contributions made in the paper and important assumptions and limitations. A No or
 NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The proposed hierarchical hybrid Radon transform costs slightly more computation as discussed in Section 3.1. Also, the injectivity of the hierarchical hybrid Radon transform depends on the injectivity of its partial generalized Radon Transform component as discussed in Section 3.1.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We state all assumptions for our theoretical results in the paper.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Ouestion: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We report all experimental settings for our experiments in Section 4 and Appendices.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
- (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We submitted the anonymized code for experiments in the paper.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/ public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https: //nips.cc/public/guides/CodeSubmissionPolicy) for more details.

- The authors should provide instructions on data access and preparation, including how
 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We report the training and test details in the experimental parts of the paper in Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We run our experiments at least three independent times and report the error bars.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We report the computational devices that we use in Appendix E

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: We follow the NeurIPS Code of Ethics when conducting the research.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We propose a new metric for comparing heterogeneous joint distributions. As shown in the paper, the proposed metric can improve applications of 3D mesh and datasets comparison. We believe that there are no direct negative societal impacts of the work.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not collect any data in the paper.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite and credit all used assets in the paper.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the
 package should be provided. For popular datasets, paperswithcode.com/datasets
 has curated licenses for some datasets. Their licensing guide can help determine the
 license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We provide anonymized code for the paper with instructions for running the code.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.

- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We do not use crowdsourcing experiments and research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.